TU Delft

Delft University of Technology

Document Version
Final published version

Licence
CC BY-NC-ND

Citation (APA)
Kazemi, P., Entezami, A., Turrin, M., Andriotis, C., Mariani, S., & Ghisi, A. (2026). Tall buildings with diagrids: design exploration and assessment of the influence of the architectural form on the structural behavior. *Structures*, *93*, Article 113161. https://doi.org/10.1016/j.istruc.2026.113161

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.

Structures 93 (2026) 113161



Contents lists available at ScienceDirect

Structures

journal homepage: www.elsevier.com/locate/istruc



Tall buildings with diagrids: design exploration and assessment of the influence of the architectural form on the structural behavior

Pooyan Kazemi [a], Alireza Entezami [a], Michela Turrin [b], Charalampos Andriotis [b], 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

HIGHLIGHTS

- Detailed workflow for the numerical generation of a tall building database.
- Database of 1000 tall buildings with diagrids available online.
- Seismic analysis in the time domain for seven moderate earthquakes.
- Assessment of the influence of architectural form on the structural behavior.
- Geometrical, structural and environmental (i.e., cost, and carbon footprint) features.

ARTICLE INFO

Keywords:
AI in building design
High-rise buildings
Diagrid
Structural database

ABSTRACT

Tall buildings with diagrids offer several advantages over alternative structural solutions, including the possibility of reducing the dimensions of interior cores and freeing up interior space. However, the interplay between structural systems and architectural form is not always adequately considered during the early stages of design to support quantitative decision-making. In this study, a database comprising 1000 tall buildings has been generated using parametric modeling. The buildings range in height from 140 to 300 meters. These buildings have then been subjected to seven ground motions selected from the NGA-West2 database, for a total of 7000 simulations. Data related to geometric, structural, and environmental variables will subsequently be processed to support investigations based on artificial intelligence-driven generative approaches aimed at producing alternative optimized solutions. In this paper, the data are described in terms of specific 22 input (10 architectural, 12 related to the ground motion) and 99 output (69 structural, 23 environmental, 7 geometry-dependent) features, demonstrating the extent to which the database has been expanded with respect to the original design space. The dataset is made available through a repository for the architecture, engineering, and construction community.

1. Introduction

Structural systems with large outer steel frames arranged in a triangular module, known as diagrids, present an opportunity for designers of tall buildings, also in view of the sustainability advantages over other systems [1–3]. By carrying the gravity and the lateral loads mostly with diagonal, axially-working members, diagrid structures allow interesting complex and curved forms [4], attracting architects who can also more freely exploit the interior space due to a reduced need for a core. While the first examples of such structures date back to the mid 20th century, see for example the IBM building in Pittsburgh (13 stories, built in 1960), these systems have received significant interest in recent years; see, e.g., the Swiss Re Tower in London (41 stories, 2004), the Hearst Tower in New York (46 stories, 2006), the Tornado Tower in Qatar (52 stories, 2008). The structural behavior of these buildings has been studied in terms of the optimal inclination angles for the diagrids to carry lateral loads [5], including wind-induced and earthquake-related ones. The simplicity of the structural behavior of the individual elements, which basically work as trusses, makes these structures a natural target for ingenious solutions, from the informal approaches attempted in the

* Corresponding author.
Email addresses: seyedpooyan.kazemi@polimi.it (P. Kazemi), alireza.entezami@polimi.it (A. Entezami), M.Turrin@tudelft.nl (M. Turrin), C.Andriotis@tudelft.nl (C. Andriotis), stefano.mariani@polimi.it (S. Mariani), aldo.ghisi@polimi.it (A. Ghisi).

https://doi.org/10.1016/j.istruc.2026.113161
Received 19 May 2026; Received in revised form 24 September 2026; Accepted 25 September 2026
Available online 3 October 2026
2352-0124/© 2026 The Author(s). Published by Elsevier Ltd on behalf of Institution of Structural Engineers. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

1960s to the optimization methods exploiting the improved computing power of the 1990s [6,7] and, more recently, refined metaheuristic techniques [8] or multi-objective optimization with genetic algorithms [9]; interested readers are referred to the reviews in [10,11].

Works by Moon and coworkers [1,12] offered guidelines for the design of diagrid structures of tall buildings, based on the stiffness of regular grids. Recently, the authors have compared the diagrid system with other structural configurations [2,3], also expanding their study to non-regular patterns in [13]. In [14], an optimization method has been described for complex and non conventional diagrid patterns, showing that non-regular forms become more feasible.

Recent experimental testing on diagrid systems has confirmed the interest in this building typology. In [15], the probability of failure of a portion of a diagrid structure has been investigated under cyclic displacement-controlled loading, mimicking a seismic one, with the induced stiffness reduction. In [16], polygonal planes with sides from 3 to 12, connected by a vertical diagrid, have instead been examined experimentally, to obtain a simple relationship between lateral deformation and stiffness. Moreover, the overall architectural form obviously also affects a aerodynamic behavior of the tall building, see, e.g., [17].

Framing the optimization of the structural and architectural requirements of tall buildings naturally leads to a multi-objective problem, where economic and sustainability issues also come into play, see, e.g., [18,19]. The complexity of the problem requires many competences [20], fruitful interaction among the stakeholders, scientific and technical rigor, and a distinct sensitivity to aesthetics. On the other hand, parametric design tools [21] make it relatively easy to create complex and curved architectural forms by handling a small number of control parameters, and can also automate the generation of data for subsequent structural analysis. Therefore, there is an opportunity to compare different architectural forms by varying the aforementioned parameters within a large design space [22].

With generative design, see, e.g., [23,24], algorithms and parametric modeling can be combined in an iterative loop to optimize the design aims under pre-defined constraints. As pointed out in a number of reviews, see, e.g., [25], the lack of reliable freely available data hinders the diffusion and development of this approach. As a large amount of data would be unsustainable for a purely human-based analysis, the advent of artificial intelligence (AI) is considered to be the only feasible solution.

The computer-aided inspection of a database of alternative building forms can in fact exploit the advances in machine learning (ML) and deep learning (DL), see, e.g., [26,27]. Based on prescribed building and/or structure features, designs can be classified using supervised learning frameworks or clustered with unsupervised approaches. In [28], human and artificial intelligence both contribute to finding non-standard structural forms in static equilibrium, exploiting several ML algorithms. In [29], ML tools have been adopted to classify the ambient-induced vibrations of an existing tall building, showing that physical and spurious modes can be effectively distinguished and the fundamental vibration frequencies can also be correctly captured in the monitored period. In [30], eight ML models have been compared to assess damage from nonlinear dynamical analyses selected from a database of 469 steel frames and 240 ground motions, finding that random forest outperformed the other investigated methods. In [31], several ML methods have been shown to correctly estimate the response of composite laminated timber tall structures to seismic loads, upon choosing the appropriate features.

In the review paper [26], applications of ML to different types of structures have been collected while structural applications have only been described in [32], again underscoring the wide range of feasible solutions offered by AI tools. These contributions paved the way to allow for architectural and structural perspectives in unprecedented ways, either through a formal mathematical formulation [33] or through image analysis [34]. Additionally, generative ML or DL can naturally provide surrogate models to enlarge the initial database of building geometries, see, e.g., [35].

While the seismic behavior of diagrid structures and the application of AI/ML tools to structural engineering problems have each received considerable attention, as reviewed above, the two streams of research have so far developed largely in parallel, excluding very recent exceptions [36]. Studies on diagrid structural behavior have primarily focused on regular or near-regular grid geometries and on optimizing member sizing and diagrid angle for a given architectural form [5,8,9,14,37], while ML applications to tall buildings have mostly targeted damage classification or response prediction on databases of comparatively simple, non-diagrid structural typologies [31,38]. Quantitative studies that jointly vary architectural form parameters, such as tapering, twisting, curvilinearity, and polygonal cross-sections, and assess their combined effect on seismic performance metrics at the early design stage remain limited. This gap is precisely what the present database is intended to address, by pairing a systematically sampled architectural design space with a consistent seismic simulation pipeline applicable across all generated forms.

In the present work, through a set of architectural forms and the results of structural analyses, quantitative support is provided for the creative stage of the design of tall buildings with diagrids. This solution is specifically able to account for the influence of the architectural form on the structural behavior under seismic loads. The proposed solution is a follow up of a former one, wherein the feasibility of ML approaches for the classification of tall buildings with diagrids has been explored using a database of 144 buildings subjected to static loads equivalent to moderate earthquakes [39]. To overcome the issues linked to the small size of the dataset, data augmentation strategies have been implemented, see [40,41], showing reasonable outcomes, provided that building features and responses are carefully selected within supervised learning frameworks [40]. With respect to previous work, the present study offers four main contributions. First, the scale and diversity of the generated dataset: 1000 models spanning a 10-dimensional architectural design space, compared to the authors' earlier dataset of 144 buildings subjected to static-equivalent loads only [39], and to other structural databases for ML purposes reported in the literature, such as the 469-steel-frame/240-ground-motion database used in [30]. Second, the use of full linear dynamic time-history analysis, rather than static-equivalent loads, across the entire dataset, providing a more realistic characterization of the seismic response. Third, the joint availability, for the same 1000 models, of structural, environmental (embodied carbon and cost), and geometry-dependent responses, enabling multi-objective analyses that are seldom supported by existing datasets. Fourth, the public release of the complete dataset and generation code through a dedicated repository, supporting reproducibility and reuse by the wider architecture, engineering, and construction research community. As this contribution aims to demonstrate the feasibility of a workflow for the early design of tall buildings that incorporates parametric generation of structures and the quantification of structural responses in relation to architectural form choices, comparisons with actual buildings will be left for future tests. In fact, the database and workflow should be useful during the idealization phase, enabling designers to recognize the nuances of architectural choices before pursuing subpar options later as a consequence of poor initial decisions.

To manage the resulting data, an accurate organization of input data and output features along with data storage is required: these details are provided in Section 2. The post-processing of data is discussed next in Section 3. The generated building characteristics are discussed in Section 4, by considering constraints in terms of structural behavior, cost and embodied carbon footprint. This work has several limitations, as explicitly pointed out in Section 5. Conclusions of this work and a look into the future are finally gathered in Section 6. All data used in this study are available in the relevant GitHub repository Tall-buildings-with-diagrids-design-exploration.

Table 1
Parameters defining the diagrid shape.

	Code	Description	Number of options	Values
	X1	Number of sides (top plan). Polygonal edges of the upper plan.	10	[3,4,...,12]
	X2	Number of sides (bottom plan). Defines the polygonal edges of the lower plan.	10	[3,4,...,12]
	X3	Relative orientation. Alignment of upper and lower plans (corner–corner = 0, corner–edge = 1).	2	[0,1]
	X4	Story number. Total number of building levels.	6	[40,44,48,...,60]
	X5	Height per story. Vertical spacing between consecutive floors (m).	3	[3.5,4,4.5]
	X6	Vertical morphing strategy. Rule by which top and bottom plans are linked.	3	[0,1,2]
If X6 = 0	X7	Tapering ratio. Scale factor applied to the area of the top plan.	6	[50%,60%,...,100%]
If X6 = 1	X8	Twisting. Rotation applied to the top plan (degrees).	5	[30,45,60,75,90]
If X6 = 2	X9	Control floor position. Location of curvilinear deformation along the height.	3	[0.25,0.5,0.75]
	X10	Control floor intensity. Magnitude of curvilinear scaling effect.	2	[0.85,1.15]
Total combinations			**61,200**	

2. Architectural modeling

To obtain a comprehensive database including 1000 models of tall buildings featuring diagrid façades, alternative transformation rules have been adopted for the vertical connection between the top and bottom floor plans. For the geometry of the top and bottom plans, polygons have been chosen with a number of sides varying between 3 and 12. Along the height of the building, the top and bottom plan geometries have been connected with tapered, twisted or more generic curvilinear side faces. Table 1 lists all the geometric options considered in the present work, each with a short description: the 10 parameters X1–X10 specify the design space that allows the generation of architectural forms. Each feature X_i, $i = 1, \ldots, 10$, can take discrete values in the design range as listed in the last column of Table 1. It is also worth mentioning that there are interdependencies between some parameters: X6, which is the vertical transformation method, acts as a switch that may affect the values of X7–X10. For example, if a tapered shape is specified (X6 = 0), parameters X8–X10 are not applicable and are therefore skipped in the analysis.

As shown in Fig. 1(a), X6 = 0 specifies the typology of the tapering transformation, while X7 quantifies the tapering intensity. A value of X7 of 50% indicates a highly tapered form where the top plan is scaled down to 50%, while a value of 100% corresponds to no tapering. For twisting transformations (X6 = 1), see Fig. 1(b), where X8 represents the relative rotational angle between the top and bottom plans. If present, the minimum twisting angle is considered to be 30°, while the maximum is 90°. Curvilinearity (X6 = 2) in Fig. 1(c), i.e., how top and bottom plans are connected by a curve instead of a straight line, is characterized by parameters X9 and X10. The former, see Fig. 2 (left), defines the position of the control floor, which is pivotal for the definition of the curve along the building height; this control floor can be positioned in the first, second, or third quarter of the building's total height. In Fig. 2 (right), X10 instead specifies the intensity of the curvilinearity, by scaling the control floor plan: for instance, a scale factor of 0.85 produces a concave curvilinearity, whereas 1.15 produces a convex one, as depicted in Fig. 2.

The number of generated tall building models, as detailed in Table 1, is a compromise between the constraints provided by the available computational resources and the full exploration of the design space. Accordingly, 61,200 possibilities are guaranteed by the choices shown in the table to maintain a balanced representation across the different variables. Three sample tall building models are presented in Fig. 1, along with their corresponding parameters. Out of the 61,200 building models, 1000 have been included in the dataset as described in the following.

Independently of the number of stories, the ratio between the core area and the total floor area plays an important role in the structural response to lateral loads. While other works have focused on supertall buildings (e.g., [42] refers to 300–600 m height buildings) or have inspected the nonlinear behavior of diagrid structures with cores (as in [4]), here the maximum building height is set to 300 m. Heights near the design space limits, namely 140 m and 270 m, have been selected, referring also to buildings of heights close to those in [43], to adopt their mean core ratio values. A linear relationship for core ratios across the different heights has been established and implemented in Grasshopper™, so that a ratio of 19.3% for 140 m based on five existing buildings, and a ratio of 23.8% for 270 m based on seven existing buildings are here considered.

The selection of the core geometry has presented specific challenges. Maintaining identical core areas across all models would lead to issues in the case of a triangular plan, which would inherently have the largest width. This could result in overlap with the top plan geometry in cases of extreme tapering, or absent overlap, significantly limit the usable space on the top floors. Since the core has a significant role in the building's load-bearing capacity, to provide a fair comparison among all the solutions its geometry has been standardized to a polygon of 12 sides.

2.1. Generation of architectural models

To sample the handled building models from the possible 61,200 solutions in the design space listed in Table 1, the Latin hypercube algorithm was adopted.

The sampling strategy begins by defining the number of parameters (10) and the number of samples (1000): the samples generated the $[0 - 1]$ range by the pyDOE2 Python library are then mapped to the actual parameter ranges. A subsequent validation stage has been adopted to ensure a well-distributed sample set within the design space, and to also make sure that a possible subsequent surrogate model will be capable of making predictions for tall buildings not included in the initial sample set. The representativeness of the sampling has been checked with two metrics: the space-filling index and the coverage index. The former computes the distance to the nearest neighbor for each point in the sample, with higher average values indicating more even distribution. The latter generates a large number of random points in the unit hypercube and finds the nearest sample for each random point; smaller maximum distances (fill distance) indicate better coverage of the space. For the X6–X10 variables, which govern the possible shapes along the building heights, the 1000 data points have been equally divided in practice. Instead, the metrics have been computed for the variables [X1, X2, X4, X5] (variable X3 has also been divided into two equal parts). The sampling results for 1000 cases have been compared with the outcomes for the full possible database of 61,200 cases: the differences have been found to be 7% and 14% for the space-filling and the coverage index, respectively, which are deemed to be acceptable.

The subsequent step consists of generating the corresponding building models by way of the Colibri Grasshopper™ plugin. In this way, the architectural model (i.e., the geometry), the structural model (i.e., a finite element model), and a Python file containing the seismic simulation data are all generated. Data are supplemented by a picture of the model, and information collected in a CSV file in terms of total weight, gross area, and façade area.

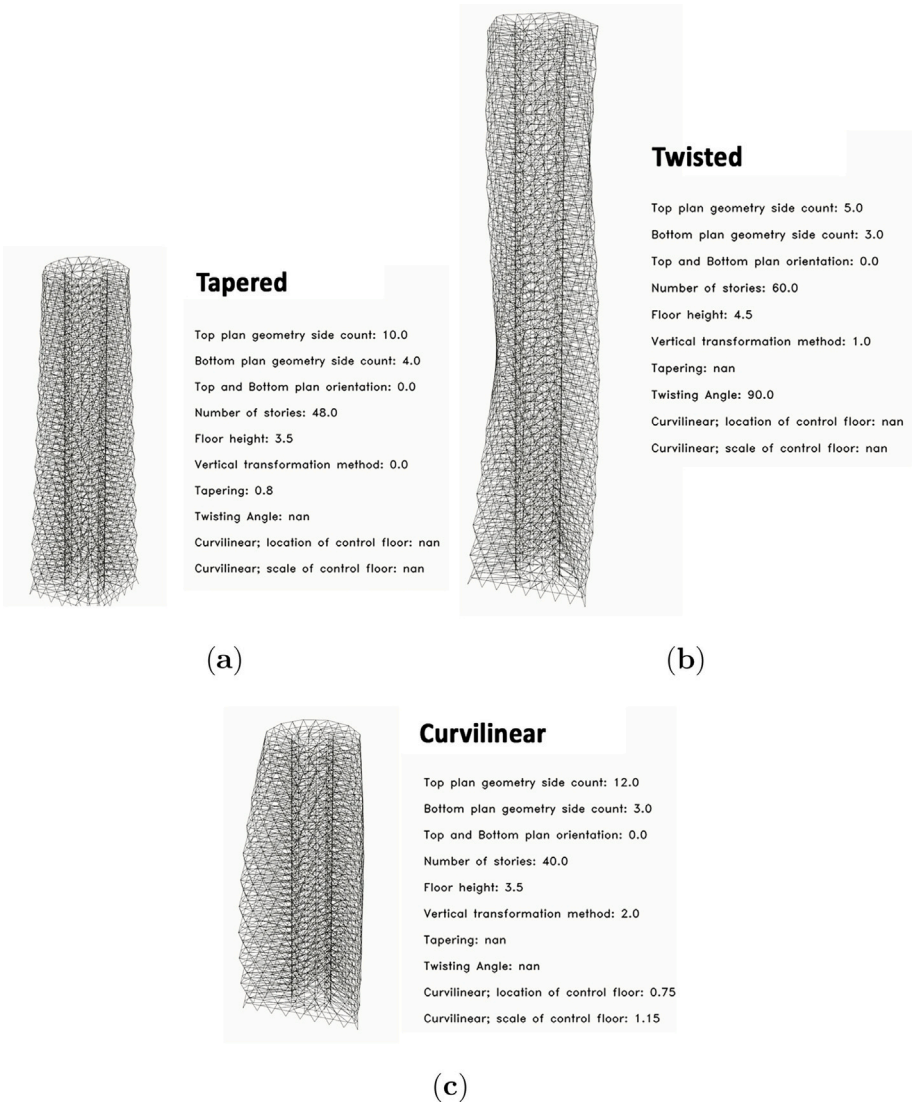


Fig. 1. Samples selected from the 1000 generated ones are shown along with their main characteristics: (a) tapered model with vertical transformation X6 = 0; (b) twisted model with X6 = 1; (c) curvilinear model with X6 = 2.

2.2. Structural modeling

This work exploits Alpaca4d, a simulation plugin for Grasshopper™, to generate OpenSees [44] input files that have been exploited to run dynamic simulations. From Alpaca4d, 1000 Python scripts for execution in OpenSeesPy have been generated.

All the structural systems have been assumed to comprise a diagrid and a braced core, interconnected by radial beams. Floor slabs, which consist of steel decks with a structural concrete topping that provides proper shear stud connections to the perimeter and collector beams, have been modeled as rigid diaphragms, with their inertial effects concentrated at floor centroids.

Static loads due to the weight of structural elements and permanent loads have first been considered. The gravity load and area loads, accounting for concrete floor slabs and live loads considering an office function, have been distributed among floor radial beams, horizontal beams of the diagrid, and horizontal beams of the core. Overall, the dead loads range from 7.5 kN/m^2 to 14.5 kN/m^2 and the live load is assumed to be 2.5 kN/m^2. No reduction of the live load is applied along building elevation. Load combinations according to the Eurocodes 1 and 3 have been considered with the characteristic values of permanent actions, variable actions, and seismic actions, as described in the following and formally expressed by the ultimate limit state load combination:

$$E_d = \gamma_G \cdot G_k + \gamma_Q \cdot Q_{k1} + \gamma_Q \cdot \psi_{02} \cdot Q_{k2} + \cdots + \gamma_A \cdot A_k \tag{1}$$

where E_d represents the design effect (e.g., bending moment, shear force, axial force); γ_G, γ_Q, and γ_A are partial factors for permanent actions, variable actions, and seismic actions, respectively, assumed to be 1.35, 1.5, and 1, according to Eurocodes; ψ_0 are factors for the combination value of a variable action, see Eurocode 1; G_k, Q_k, and A_k are the characteristic values of permanent actions, variable actions, and seismic actions, respectively.

A static analysis has first been carried out with the purpose of pre-dimensioning the cross-sections of the structural members, using the cross-section optimizer algorithm of Karamba, another plugin for Grasshopper™. Zero displacement has been imposed on all nodes at the base of the buildings for this pre-dimensioning. Loads statically equivalent to the seismic ones have been adopted with the parameters detailed in [39] within an ultimate limit state combination. To do that, the optimizer has been provided the following information:

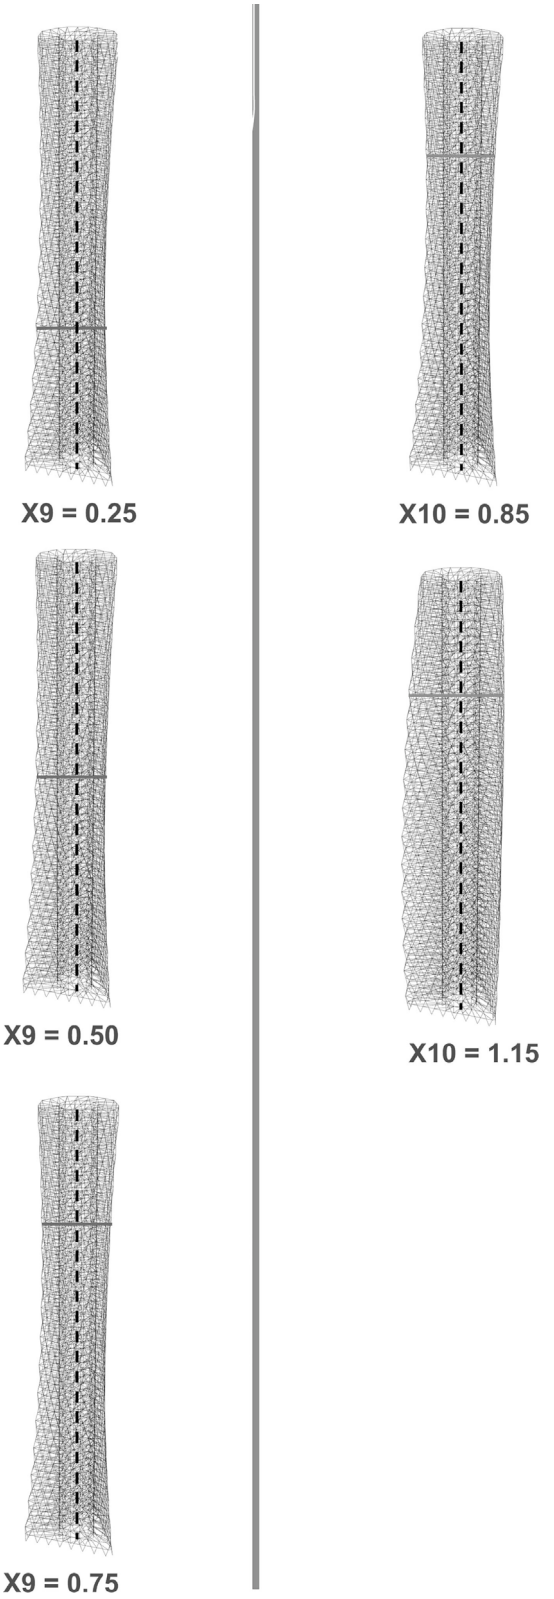


Fig. 2. Curvilinear vertical transformation adopted in this work, with (left) X9 to define the position of curvilinearity, and (right) X10 to define the intensity of curvilinearity.

1. Cross-section lists. The algorithm has been fed with a list of 2065 cross-sections ranging from EN 10220 ROR10.2/0.5 to ROR2540/65, including standard tube sections.
2. Optimization groups. Cross-sections have been divided into groups, to mandate identical sizes for members within each one. For example, diagonal diagrid members, core columns, bracings, beams, and radial beams have been grouped by story, ensuring uniformity.
3. Displacement ratio. A threshold of $h/500$, h being the building height, has been enforced for the lateral deflection for each building.
4. Utilization ratio, defined as the ratio of resultant stress to the allowable one for a structural member in tension/compression, has been set to 95% in line with the research objectives. By balancing between conservative (70 to 85%) and aggressive (100%) approaches, this value has been adopted for both the cross-section optimizer and the dynamic simulator.

The output of the cross-section optimizer is therefore a tall building model with optimized cross-sections, as shown in Fig. 3. Validation checks have been performed to ensure that the designed cross-sections are adequate for static and seismic loadings. It would not be possible to list all the cross-sections obtained from the initial design phase for each building, as there are too many of them. Even if they are grouped together, there are still tens of them. Moreover, they differ depending on the geometry of the building. However, all this data is present in the OpenSees input files, which can be found in the GitHub repository mentioned at the end of the Introduction.

To improve the clarity and reproducibility of the structural model, Fig. 4 shows a representative floor plan for a reference tall building model, including the diagrid perimeter, the radial beams connecting the perimeter nodes to the core, and the core geometry, standardized to a 12-sided polygon and positioned centrally within the plan. Typical dimensions of the tubular cross sections after pre-dimensioning are on the order of 800 mm in diameter for the diagonal diagrid, 500 mm for the horizontal diagrid, 700 mm for the core column and 450 mm for the core beam. Story heights range from 3.5 to 4.5 m depending on the sampled value of parameter X5, and radial beam spans vary with the plan geometry (number of sides and building height) as governed by parameters X1–X5.

It is important to emphasize that two distinct structural analysis stages are involved in the workflow. The aforementioned equivalent-static analysis is used exclusively for the preliminary sizing (pre-dimensioning) of member cross-sections, via the Karamba cross-section optimizer, under an ultimate-limit-state combination including seismic-equivalent static loads, as detailed in [39]. Subsequently, the actual seismic performance of each of the 1000 pre-dimensioned models is assessed through linear dynamic time-history analysis in OpenSeesPy, as described in the following.

Regarding seismic loading, the PEER NGA-West2 web tool [45] has been adopted, as it accommodates a broad period domain, up to 10 s, suitable for tall buildings, albeit without the capability for geolocation-specific searches. The selection procedure comprises three steps. First, a design (target) spectrum is defined by specifying a magnitude range, a distance range, and a soil/site class representative of stiff soil to rock conditions (with the time averaged shear wave velocity at a depth of 30 m, Vs30, in the range of the selected records, 361–550 m/s); no site-specific seismic hazard analysis is performed, since the present study is intentionally generic and comparative in nature, aimed at ranking the relative seismic performance of 1000 alternative architectural forms under a consistent, moderate-intensity scenario rather than at the seismic design of a specific building at a specific site. Second, among the candidate records returned by the tool, seven have been selected to ensure a diversified set in terms of Arias intensity and 5–95% significant duration (D5-95), explicitly avoiding the selection of multiple records from the same seismic event; fault mechanism and Joyner–Boore/rupture (Rjb/Rrup) distances have also been considered (Table 2). Seven ground motions, shown in Fig. 5, have been adopted in line with the Eurocode 8 minimum requirement for

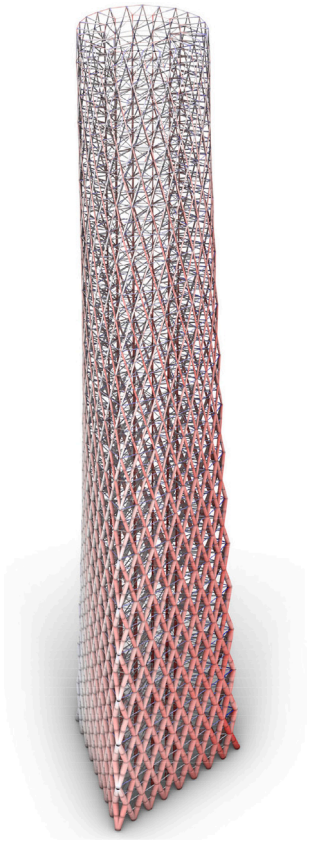


Fig. 3. Sample tall building model provided by the cross-section optimization stage, showcasing the structural framework.

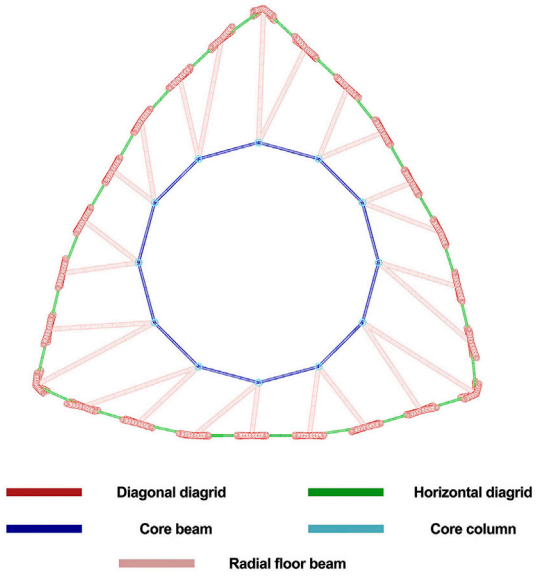


Fig. 4. Planar view of an exemplary floor. The outer surface side measures approximately 39 meters.

time-history analysis; we note that this is a less demanding requirement than those of guidelines such as The Tall Buildings Initiative and The Los Angeles Tall Buildings Structural Design Council, which recommend larger suites including long-period and long-duration records, and which have not been specifically targeted in the present selection.

Table 2
Characteristics of the selected ground motions from the NGA-West2 database [45]. Entries highlighted in red are considered in this work's database.

Ground motion ID	Spectral ordinate	Record sequence number	Mean squared error	Scale factor	5-75 % duration (s)	5-95 % duration (s)	Arias intensity (m/s)	Earthquake name	Year	Station name	Magnitude	Mechanism	Rjb (km)	Rrup (km)	Vs30 (m/s)	Lowest usable frequency (Hz)	Initial-search scale factor
1	RotD50	28	0.6007	0.6676	14.7	29	0.1	"Parkfield"	1966	"Cholame - Shandon Array #12"	6.19	Strike slip	17.64	17.64	408.93	0.125	0.7636
2	RotD50	57	0.2348	0.3512	10.6	16.8	1	"San Fernando"	1971	"Castaic - Old Ridge Route"	6.61	Reverse	19.33	22.63	450.28	0.375	0.3792
3	RotD50	313	0.2005	0.239	5.4	15.4	0.9	"Corinth_ Greece"	1981	"Corinth"	6.6	Normal Oblique	10.27	10.27	361.4	0.25	0.258
4	RotD50	359	0.3823	0.3239	3.8	10.9	0.4	"Coalinga-01"	1983	"Parkfield - Vineyard Cany 1E"	6.36	Reverse	24.83	26.38	381.27	0.2875	0.3497
5	RotD50	815	0.3757	24.139	4.4	11.6	0	"Griva_ Greece"	1990	"Kilkis"	6.1	Normal	26.75	29.2	454.56	0.75	26.062
6	RotD50	954	0.2037	0.5873	6.1	10.7	0.5	"Northridge-01"	1994	"Big Tujunga_ Angeles Nat F"	6.69	Reverse	19.1	19.74	550.11	0.4125	0.6341
7	RotD50	1643	0.0933	10.132	4.7	13.8	0.1	"Sierra Madre"	1991	"LA - City Terrace"	5.61	Reverse	23.67	25.69	365.22	0.2	1.094

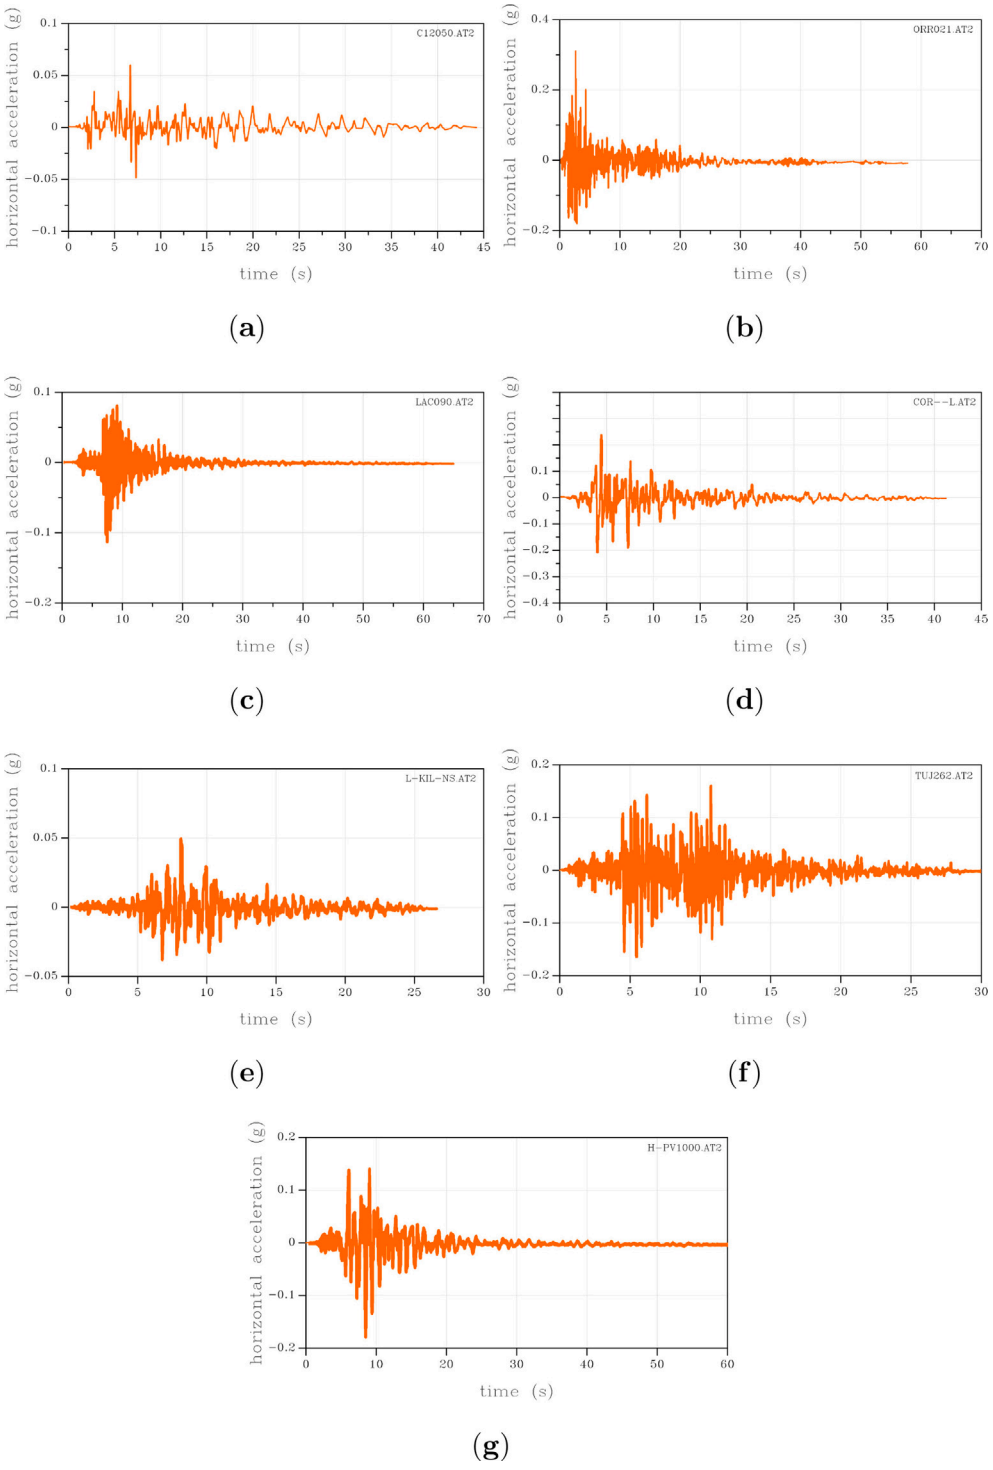


Fig. 5. Visualization of the seven selected ground motions before scaling: (a) Parkfield; (b) San Fernando; (c) Corinth_Greece; (d) Coalinga-01; (e) Griva_Greece; (f) Northridge-01; (g) Sierra Madre.

Third, each selected record has been scaled so that the average response spectrum fits within the bounds of the design spectrum (Fig. 6).

As the analyses presented here are in the linear elastic regime, additional calculations should carefully account for nonlinear constitutive laws in the time-domain analysis in the case of significant earthquake excitations that might initiate inelastic effects and even (partial) collapse mechanisms. The current dataset does not account for this; however, the same workflow can be followed, as it has been designed to allow for this, despite the significant increase in computational cost.

Due to the limited plan size compared to building elevations, all nodes restrained to the ground are simultaneously excited by the adopted ground motions, both in the $X-$ and $Y-$directions. Due to the assumed moderate excitations, the vertical ($Z-$direction) excitation has not been considered in this work, but it can easily be included in future developments.

For transient analysis, Rayleigh damping has been imposed. The procedure described in [46] has been adopted, with a 5% ($\xi = 0.05$) damping ratio related to the first and 10th modes ($\xi = \xi_{\max}$); between

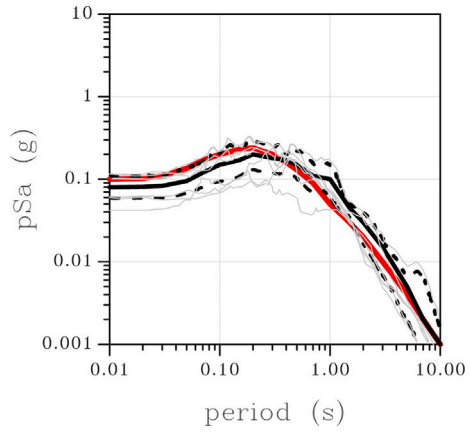


Fig. 6. Spectra of the seven selected ground motions. In the graph, the black dashed lines represent the upper and lower bounds of the design spectrum, the black and red continuous lines stand for the mean and target spectra, while the gray lines indicate the scaled ground motion spectra. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

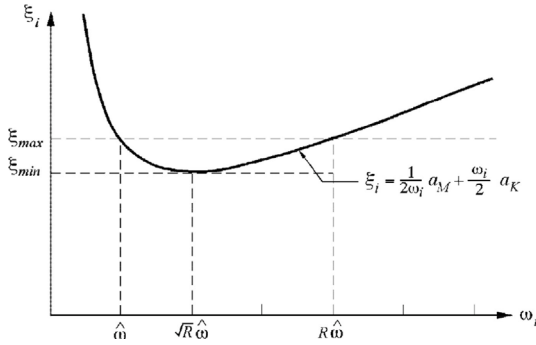


Fig. 7. Representation of the damping ratio ξ_i illustrating coefficients α and β relative to target frequencies $\omega_1 = \hat{\omega}$ and $\omega_2 = R\hat{\omega}$ ($R > 1$) (modified from [46]).

the mentioned modes, ξ varies as shown in Fig. 7. Accordingly, with the damping force given by $\boldsymbol{F}_d = \alpha \boldsymbol{M}\ddot{\boldsymbol{u}} + \beta \boldsymbol{K}\boldsymbol{u}$, where $\boldsymbol{M}$ is the mass matrix, $\ddot{\boldsymbol{u}}$ the acceleration vector, $\boldsymbol{K}$ the stiffness matrix, and $\boldsymbol{u}$ the displacement vector, the coefficients α and β read:

$$\alpha = \frac{2\xi(\omega_1\omega_2)}{(\omega_1 + \omega_2)} \qquad \beta = \frac{2\xi}{(\omega_1 + \omega_2)} \tag{2}$$

ω_1 and ω_2 being the circular frequencies of the 1st and 10th modes, respectively.

An eigen-analysis has first been carried out to elucidate the dynamic properties, and to inspect how the building modes relate to the chosen earthquake signals. The first 10 vibration modes of an exemplary model are detailed in Fig. 8. It can be noted that modal deformations range from those akin to a cantilever beam, to torsional ones and others featuring a spring-like deformation along the longitudinal axis of the structure. To bound the information provided by the transient simulations, a selective recording strategy has been implemented in OpenSeesPy. This procedure seems compulsory to adhere to simulation time constraints and post-processing, as each tall building features a number of nodes ranging from 9000 to 12,000.

Given that the selected ground motions are defined with a 10 ms time step, aiming for computational efficiency, the analysis time step Δt is set to 30 ms. A Newmark time integration method has been adopted with parameters set to guarantee unconditional stability and reliable simulation outcomes.

The criteria for node output selection have been informed by both structural significance and computational feasibility, prioritizing nodes that are expected to yield the most insightful data regarding the building's seismic response. Aiming as said at computational efficiency, seven node output types, also known as node "recorders", have therefore been defined for each model, as listed in Table 3.

While recording all nodal results within the database could generate output files exceeding 300 MB in size, becoming impractical for large-scale simulations, three nodes per floor, in addition to the centroid, have been selected for recording, to maximize the spatial distribution of the collected information across the structural geometry. These nodes are chosen for their maximal distance from one another, ensuring an even representation of the floor's dynamic response. To validate this selection methodology, a comprehensive test has been conducted on a sample tall building. This has involved collecting nodal results for the entire structure, subsequently analyzing the data to confirm whether the nodes exhibiting the highest displacement values corresponded to the pre-selected nodes. For the monitored nodes, the response variables to record, as listed in Table 3, have been selected knowing that not all response types are equally relevant for this study. Some, like velocity, have been considered redundant when the focus is on linear elastic structural behavior. Others, like inter-story drift, are not directly available within OpenSees and therefore require a post-processing phase; see Section 2.4.

The selection of element outputs mirrors the rationale employed for node selection, with an emphasis on quantities critical to understanding the seismic behavior of the structure. As a result, nine element output positions have been established capturing critical force interactions, as listed in Table 4. These elements have also been chosen for their pivotal roles in the structural integrity of the building. Notably, element force responses chosen for every element position account for both endpoints of an element, encapsulating a comprehensive set of twelve values per element detailing axial force N, shear forces V_y and V_z, torque M_x, and bending moments M_y and M_z at each endpoint. Time-dependent dynamic simulations have then been carried out: an isotropic, linear elastic constitutive law for the materials has been adopted, under the hypothesis of small displacements. A rigid diaphragm has been defined for each floor, constraining the motion of all nodes on the i-th floor to be contained in a plane.

2.3. Implementation details

Python libraries such as VFO and Opsvis have been employed, as they facilitate the visualization of node and element recorders. The controlling Python files have required the development of additional Python scripts and the use of GhPython components within Grasshopper™. This process has been automated by integrating Colibri into the architectural modeling phase, allowing for the systematic storage of Python files alongside snapshots and model information for each simulated tall building, and ultimately streamlining the simulation workflow. The dataset has been split into batches of 100 models, and the tasks have been distributed between two workstations featuring slightly different computing capabilities, specifically, one with Intel Core i7 10,700 CPU, Nvidia Quadro P620 graphics, and 32 GB of RAM; and one with Intel Core i9 13,900 H CPU, Nvidia GeForce RTX 4070 graphics and 32 GB of RAM.

To run multiple instances of Rhinoceros™ and Grasshopper™ scripts concurrently, leading to a processing time of about 3 h, the post-processing stage has been organized as follows:

1. Each model folder has included a Python script for execution on a Linux cluster;
2. Inside each model folder, four CSV files have been collected for the subsequent post-processing stage, containing node and element tags, diagrid and core support node information and tags;
3. A separate folder has been generated to contain the seven ground motion datasets;

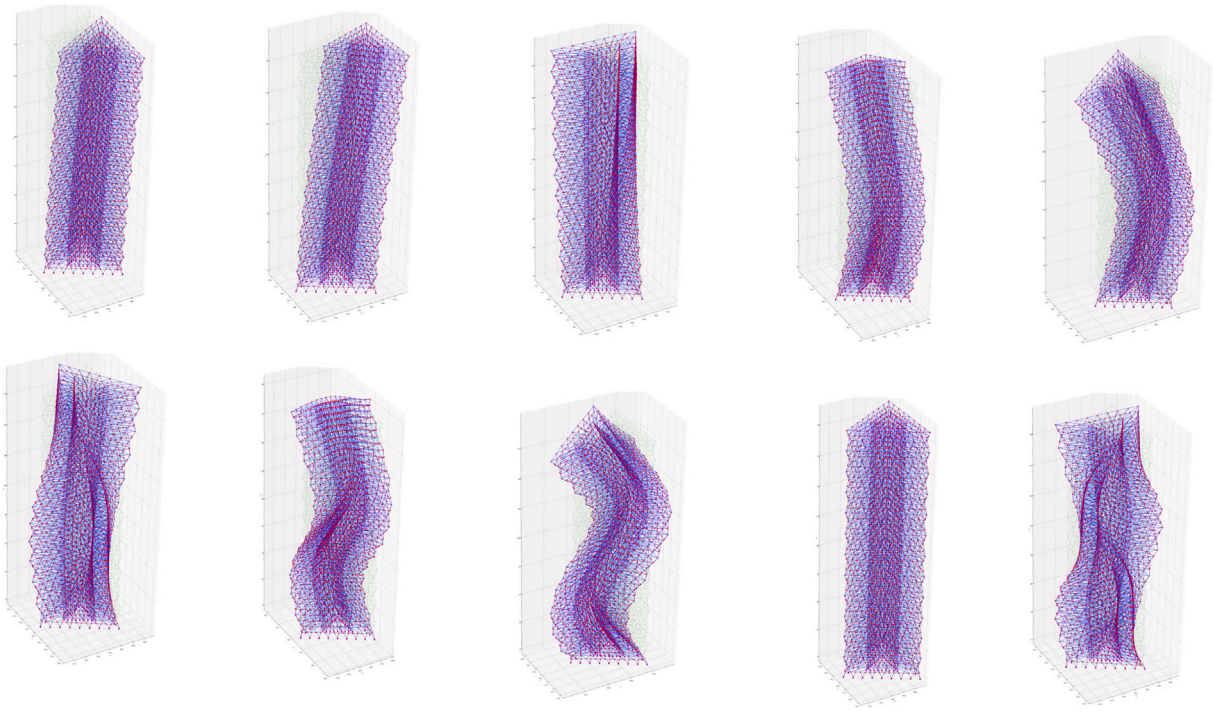


Fig. 8. First 10 vibration modes for a sample tall building model.

Table 3
Node output description.

No.	Parameter	Location
1	Reaction forces at restrained core nodes	Base
2	Reaction forces at restrained diagrid nodes	Base
3	Displacement at centroid node	Each floor
4	Acceleration at centroid node	Each floor
5	Displacement at corner node 1	Each floor
6	Displacement at corner node 2	Each floor
7	Displacement at corner node 3	Each floor

Table 4
Positions of elements used for output.

No.	Element	Location
1	Core support columns	Base (support)
2	Core bracing	Base (support)
3	Diagrid support node elements	Base (support)
4	Diagrid diagonal member	Corner 1
5	Diagrid horizontal member	Corner 1
6	Diagrid diagonal member	Corner 2
7	Diagrid horizontal member	Corner 2
8	Diagrid diagonal member	Corner 3
9	Diagrid horizontal member	Corner 3

4. A sequence of 1000 images, each representing a model, has also been created;
5. A single CSV file listing the parameters X1–X10 and providing essential information on each model has been produced.

Opting for this approach has not only allowed for easy scripting and automation, but also resulted in an efficient data management system.

The seismic simulations of the 1000 tall buildings subjected to seven different ground motions, have then been carried out via the Delft Blue Supercomputer, whose specifications are collected in Table 5. To conveniently work within the cluster job scheduling rules, the 7000 simulations have been split into batches of 1000 runs. The choice of the finite element solver is pivotal for the simulation efficiency of thousands of analyses: for relatively large-scale models such as tall buildings, the UmfPack solver which efficiently handles sparse matrices, allows for a computing time of about 30 min per analysis with the hardware system described above. This efficiency is indeed attributed to UmfPack's features for sparse matrix factorization, making it particularly suited for the structural matrices characteristic of the considered structural models.

Table 5
Summary of the main Delft Blue system features.

Metrics	Value
Intel Cascade Lake (Xeon E5-6248R and E5-6448Y) CPUs (total)	676
Compute cores	17,842
GPUs (total)	80
Tensor cores	42,880
CUDA cores	481,280

2.4. Post-processing

OpenSees, even though robust in conducting seismic simulations, does not directly provide outputs in terms of critical structural metrics, such as member stresses and inter-story drifts which are used to assess the structural performance and seismic resilience of buildings; hence, they must be computed during the post-processing phase. However, as mentioned above, simulation and post-processing have been integrated into a Python script developed within Grasshopper™ to avoid storing extensive volumes of raw simulation data. In this way, relevant quantities have been computed at every step of the analysis for each model, namely: inter-story drifts, i.e., relative lateral displacements between stories; normal, shear, and von Mises stresses for each structural element; acceleration vectors and torsion of the floor centroids; total shear forces and bending moments at the building base, distinguishing between the contributions from the core and the diagrid systems.

3. Data preparation and organization for AI-driven data analytics

A key part of the proposed methodology involves linking critical data from the outputs, as provided by the post-processing phase, to the model parameters X1–X10. Although the dataset includes comprehensive time-series results, this work on purpose exploits only the overall maxima of the time-varying output quantities, which look significant from a structural perspective. The entire dataset made available provides an opportunity for future research, to explore in more depth alternate strategies; it can also be explored with several AI algorithms to generate new geometries, as shown in [47].

A Python script, utilized to systematically process the data files for each model and each ground motion, has been crafted to recognize specific patterns and structures within each output file, and to facilitate the compilation of a CSV file that categorizes each variable recorded.

3.1. Dataset: inputs and outputs

The dataset is constructed to include the following information:

1. Overall maxima provided by the designated recorders, see Tables 3 and 4.
2. Model parameters X1 to X10, and additional information such as total mass, façade area, aspect ratio, gross area, and diagrid angles.
3. Environmental metrics, with further discussion to follow.
4. Characteristics of ground motion, focusing on the 12 parameters from the PEER database highlighted in Table 2.

More specifically, the dataset consists of two categories of input features considered as independent variables for the ML model, which are:

- Building features (10), encompassing X1 to X10 and therefore related to the input parameters of tall building models.
- Ground motion features (12), as listed in Table 2, to effectively characterize the ground motions; irrelevant features are removed, so that only the features highlighted in red in the table are stored in the database.

The output features consist of 99 responses, which are subdivided into the following categories:

- Environmental responses (23), are metrics related to the environmental impact of the building, see Section 3.2.
- Structural responses (69), are data linked to the structural integrity and performance under varying conditions as listed in Table 6, see also Section 3.3.
- Geometry-dependent responses (7), are metrics governed by the geometrical characteristics of the building, see Section 3.4.

3.2. Environmental responses

The environmental impact of each building is synthesized using variables for embodied carbon calculation and others related to cost assessment. They focus on quantifying the carbon footprint associated with material production, transportation, and construction processes, as well as evaluating the overall cost implications of different materials used in construction; they are listed in Table A.10 and relevant details are given in Appendix A.

3.3. Structural responses

The structural responses primarily focus on the overall maximum values characterizing the output of the simulations, as previously mentioned. The aforementioned 69 structural responses extracted encompass a broad spectrum of performance metrics, including but not limited to overall acceleration, displacement, inter-story drift, stress levels in critical structural components of the diagrids, and comprehensive assessments of torsional behavior and reaction forces. Each of these metrics provides insights into the structural integrity and resilience of the modeled buildings under various simulated conditions. The complete list of these responses is collected in Table 6.

Table 6
Structural responses.

Category	Parameters
Acceleration	- Overall maximum (magnitude) - Acceleration in X and Y directions - Time of occurrence, story level
Displacement	- Maximum displacement (magnitude) - Displacement in X and Y directions - Time of occurrence, story level
Inter-story drift	- Overall maximum drift (magnitude) - Drift in X and Y directions - Time of occurrence, story level
Diagrid stresses	- Maximum tensile/compression stress - Maximum shear stress, von Mises stress - Time and story of occurrence
Total structural stresses	- Overall maximum tension, compression, shear, von Mises - Time and story of occurrence
Torsional response	- Overall maximum torsion - Time of occurrence, floor level
Diagrid reaction forces and moments	- Maximum reaction force (magnitude) - Reaction forces in X and Y directions (R_y, R_z) - Maximum magnitude of moments - Bending in X and Y directions (M_y, M_z) - Time of occurrence
Core reaction forces and moments	- Maximum reaction force (magnitude) - Reaction forces in X and Y directions (R_y, R_z) - Maximum magnitude of moments - Bending in X and Y directions (M_y, M_z) - Time of occurrence
Total reaction forces and moments	- Combined maximum reaction force (magnitude) - Combined bending in X and Y directions (M_y, M_z) - Combined maximum magnitude of moments - Combined bending in X and Y directions (M_y, M_z) - Diagrid contribution in maximum overall reaction force - Core contribution in maximum overall reaction force - Diagrid contribution in maximum bending - Core contribution in maximum bending - Time of occurrence

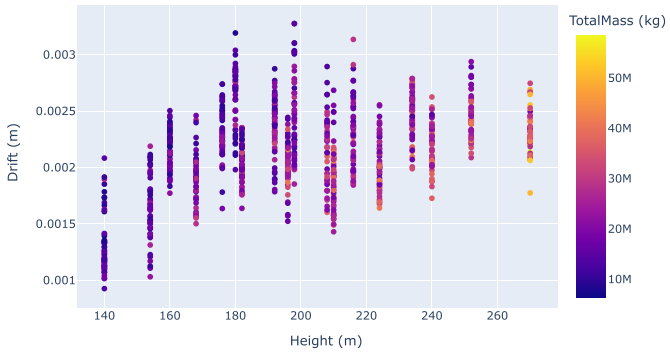
3.4. Geometry-dependent responses

Geometry-dependent responses emerge from the application of the Colibri tool within the Grasshopper™ environment, based on the design parameters X1–X10. Unlike structural responses, which are directly derived from the whole simulation procedure, geometry-dependent responses are ascertained post-definition of the input parameters, i.e., they are not a direct result of the dynamical simulation processes even though they are determined in the pre-dimensioning phase. This underlines the rationale for their separate categorization from structural responses, as they provide insights into the architectural and geometric aspects of the buildings shaped by the foundational design inputs.

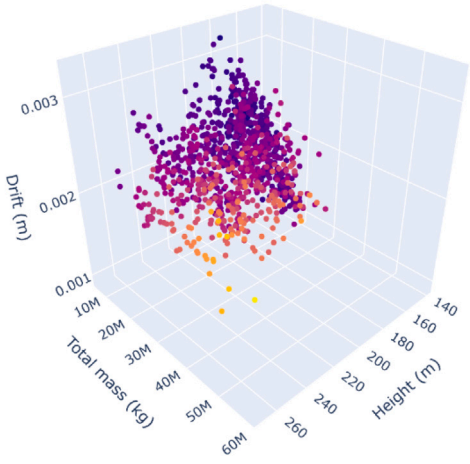
Among the geometry-dependent responses, all but one are non-structural, focusing on aspects such as the building's architectural geometry and façade area. The exception within this category is the total weight of the structure, which is determined through the application of the cross-section optimizer. In total, the 7 parameters identified within this category elucidate how initial design decisions influence the final architectural and geometric characteristics of the building according to Table 7.

Table 7
Geometry-dependent responses.

Parameter	Description
Total gross area	Total built-up floor area of the building (m^2)
Aspect ratio	Ratio of building height to its base width (slenderness indicator)
Total façade area	Overall external envelope surface area (m^2)
Diagrid angle (bottom)	Inclination angle of diagrid members at the base
Diagrid angle (top)	Inclination angle of diagrid members at the top
Diagrid angle (average)	Mean inclination angle of diagrid members along the height
Total mass	Total structural mass of the building (kg or tonnes)



(a)



(b)

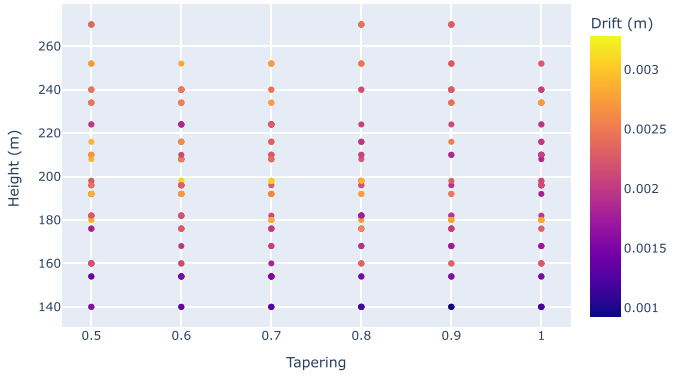
Fig. 9. Ground motion 1. Maximum inter-story drift versus height as a function of the total mass of the building. The 3D plot can be inspected interactively with the HTML file available at the GitHub repository in the Visualization folder.

4. Results

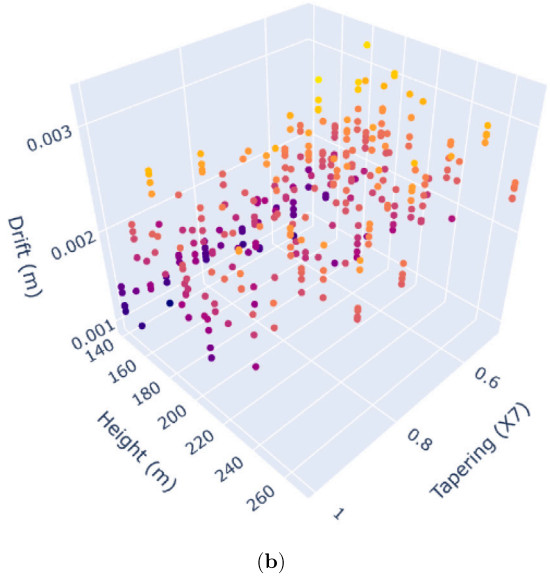
This section provides information on the input and output variables of the dataset, as linked to a specific earthquake or to all the ground motions. The aim is to help to understand possible limitations of the parameter space explored in this database, and to guide further steps in the application of AI-based procedures or to envisage possible data augmentation strategies. HTML versions of the 3D graphs reported in the following are also available and can be viewed interactively in the GitHub repository Tall-buildings-with-diagrids-design-exploration.

Fig. 9(a) plots the relationship between the maximum inter-story drift induced by ground motion 1 (ground IDs are defined in Table 2) against building height; each dot is colored according to the mass of the corresponding building in millions of kg. The same data are also depicted in a 3D plot in Fig. 9(b); these results are taken as a reference for all other cases subsequently discussed. As expected, the tallest buildings (250+ m) often have the largest masses (yellow), while the shortest buildings (160 m or less) have the smallest masses (dark blue); however, there are also several tall buildings with relatively lower masses at 250 m. The latter class often exhibits relatively higher inter-story drifts, but in all cases there is a limited yet significant set of buildings that perform well in terms of inter-story drift and have a moderate mass and cost, since these two factors are mostly correlated. The same distribution of results can also be seen for the shortest buildings, where a small mass often combines with large drift, but not in all cases. For buildings of average height, between 160 and 250 m, high inter-story drifts are often associated with smaller masses, but there are some exceptions. Moreover, there are also examples of buildings with a relatively high mass (shown in red) that exhibit a moderate drift, showcasing an inefficient use of the structures. As the results combine expected and unexpected outcomes, the soundness of the approach is confirmed, as it highlights intriguing combinations of traits with interesting behavior, even if they are not optimized.

To inspect the influence of the building shape on the results, Fig. 10(a) plots height against the tapering ratio X7, with color

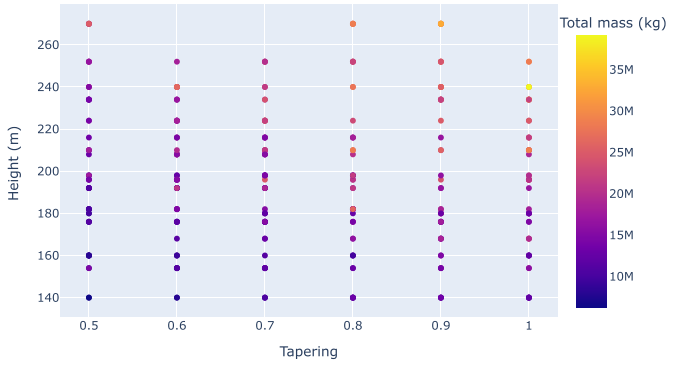


(a)

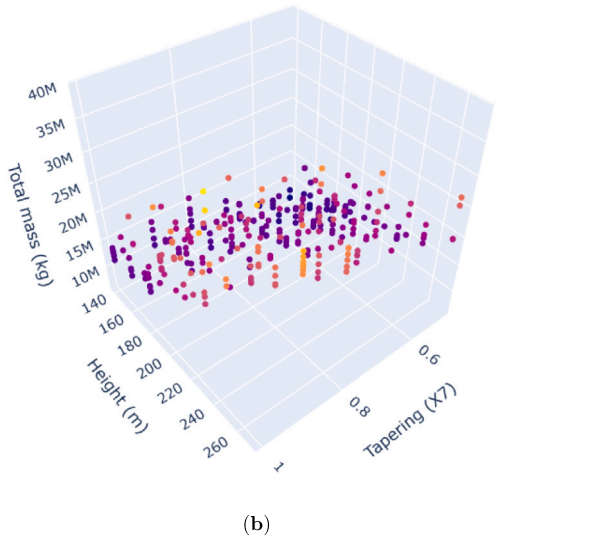


(b)

Fig. 10. Ground motion 1. Height vs tapering ratio (X7) as a function of the maximum interstory drift. The 3D plot can be inspected interactively using the HTML file available in the GitHub repository in the Visualization folder.

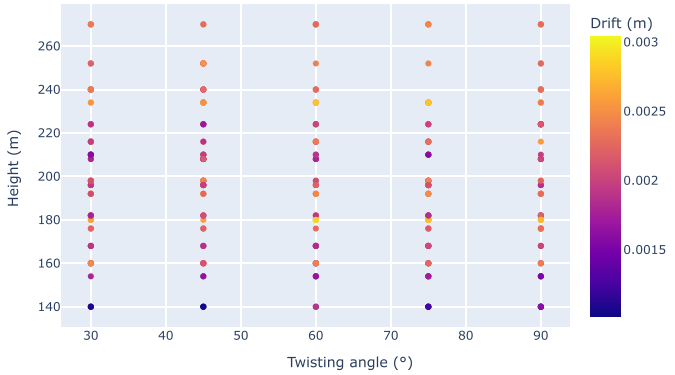


(a)

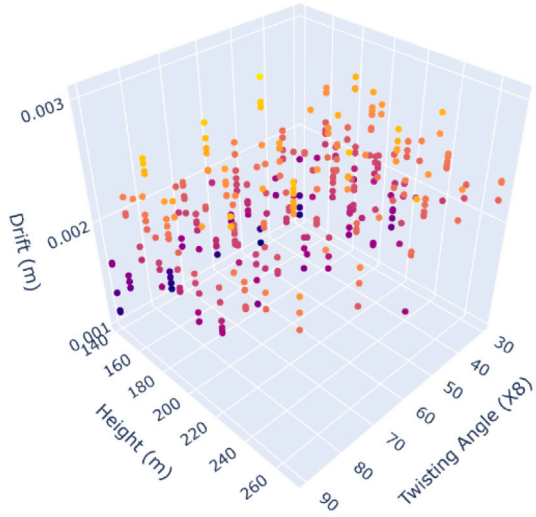


(b)

Fig. 11. Ground motion 1. Height vs tapering ratio (X7) as a function of the total mass of the building. The 3D plot can be inspected interactively with the HTML file available in the GitHub repository in the Visualization folder.



(a)

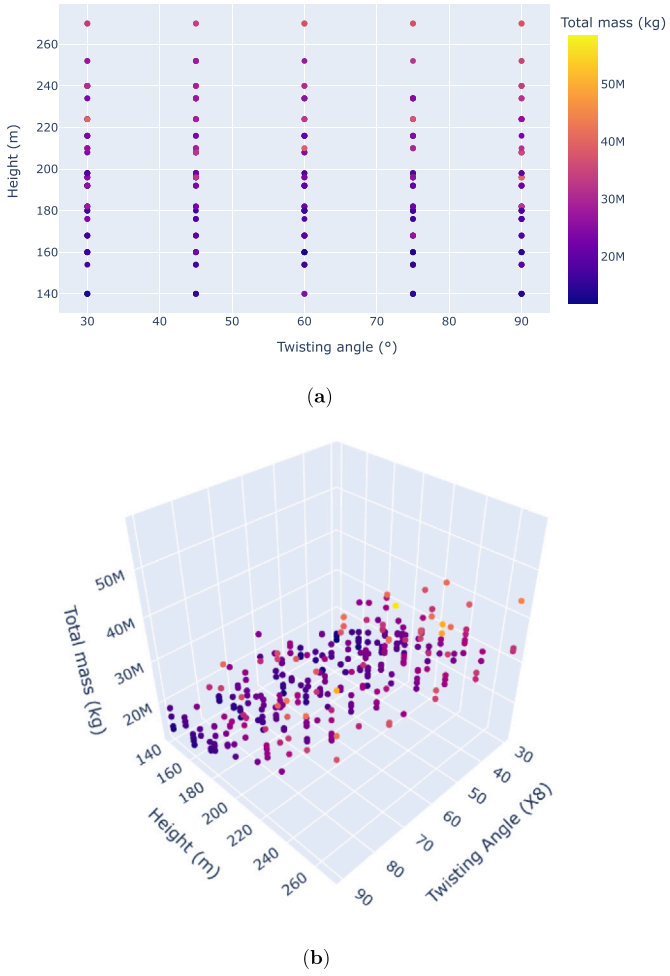


(b)

Fig. 12. Ground motion 1. Height vs twisting angle as a function of the maximum inter-story drift. The 3D plot can be inspected interactively with the HTML file available at the GitHub repository in the Visualization folder.

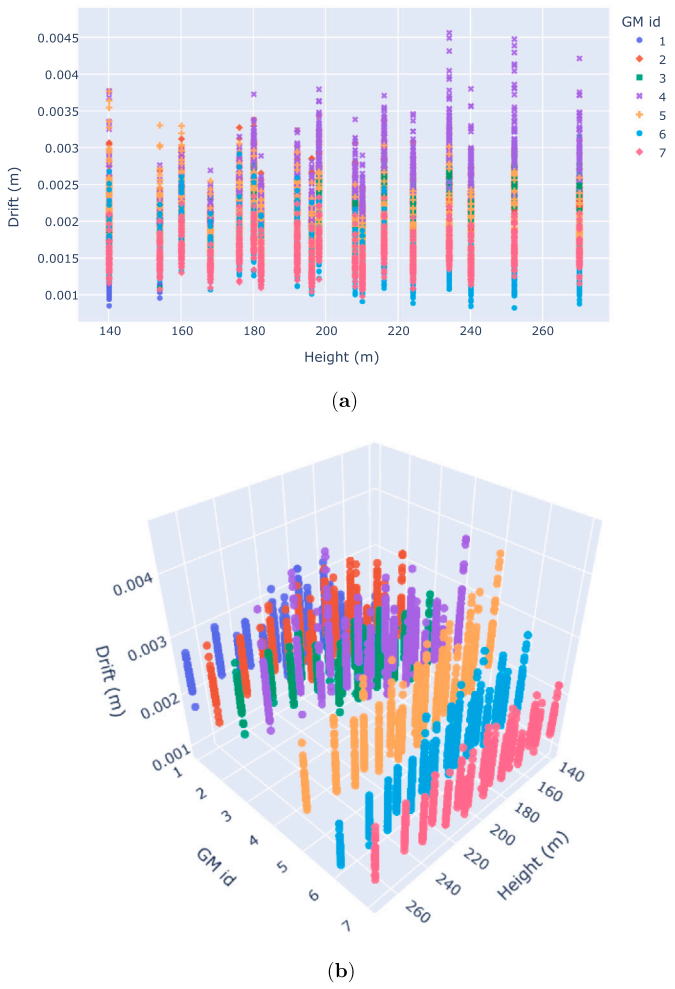
indicating the maximum inter-story drift. It is worth noting that 100% tapering means no reduction of the top plan area with respect to the bottom one; conversely, a ratio lower than 1 applies the corresponding reduction to the top floor area, see Table 1 and Fig. 1(a). The results in Fig. 10(a) and (b) show that the tapering has a relatively limited effect on the maximum drift, independently of the building height. Only in some cases, in the moderate height range (around 180–210 m), tapering can affect the stiffness and increase the drift. It is however worth underscoring that it is difficult to recognize them by manually inspecting the database; therefore, AI algorithms appear a natural choice to identify the buildings satisfying appropriate design criteria. In Fig. 11(a) and (b), the same data points shown in Fig. 10 are plotted with a color instead representing the building mass. It is easily checked that, when the tapering ratio is equal to 50%, the mass is conveniently reduced; moreover, taller buildings usually have larger masses. However, in this plot some counter-intuitive options are also visible, such as buildings with significant mass and a 180–210 m height, featuring a tapering ratio of around 70–80%.

In the dataset, some buildings present a significant twisting angle of the rising floor plans with respect to the bottom one as linked to X8, see Fig. 1(b), ranging from 30° to 90°. In Fig. 12, results as a function of twisting angle and height are reported with colors referring to the maximum inter-story drift: the effect of this architectural choice on the stiffness is not straightforward, since there are several cases with relatively high inter-story drifts both at moderate heights and at low or high twisting angles. In the considered database, the twisting angle is thus an architectural choice available to the designers as its influence on the building stiffness can be mild. Height and twisting angle are also comparatively represented in Fig. 13, where the point color identifies the mass. As mentioned above, taller buildings typically have the largest masses and, conversely, shorter buildings the smallest masses; however, there are several cases at moderate heights (180–230 m) where lighter solutions combine with varying twisting angles, without any clear indication of the role of the latter parameter. Therefore, these results confirm again that twisting angle can be left, within some limits, to the architect's sensitivity, without affecting the structural stiffness too much at least for the seismic loading here considered. The influence of the seven ground motions is investigated in Fig. 14. In this plot, the maximum inter-story drift and the building height are compared; it is shown that ground motion 2 (red diamonds) gives more conservative results in terms of inter-story drift, while ground motion 4 (purple crosses) leads to the largest inter-story drifts for the tallest buildings. If the focus is moved to the maximum acceleration



(a) (b)

Fig. 13. Ground motion 1. Height vs twisting angle as a function of the total mass of the building. The 3D plot can be inspected interactively with the HTML file available at the GitHub repository in the Visualization folder.



(a) (b)

Fig. 14. All ground motions. Maximum inter-story drift vs height as a function of the given seismic input. The 3D plot can be inspected interactively with the HTML file available on the GitHub repository in the Visualization folder.

generated by the ground motions, Fig. 15 reports that ground motion 7 prevails at almost all heights.

Fig. 16 shows the effects of the ratio of the number of polygon sides (top vs. bottom floor plan sides) and twisting angle on the inter-story drift for ground motion 1. The database contains a limited but significant number of buildings with a top/bottom ratio greater than 1, meaning there are more sides on the top than on the bottom floor. In these cases, the inter-story drift tends to be relatively high; conversely, when the ratio is smaller than 1, i.e., when the top floor has fewer sides than the bottom floor, the inter-story drift tends to be smaller. However, there are several buildings whose behavior contradicts these observations for all twisting angles, which is a further confirmation of the freedom left in this geometric choice.

The building responses can also be characterized in terms of the maximum von Mises stress for each structural element. Appropriate predimensioning usually prevents exceeding the acceptable (i.e., design stress) limits for common steel. However, as can be seen in Fig. 17(a), there are some outliers, i.e., stresses larger than 200 MPa, which should therefore be treated with care. It can be noted that exceeding stress values are usually limited to a few structural elements, so a redesign could help even if section verifications are passed in most cases. However, these outliers are kept to exemplify which design choices could be problematic within the considered design space. Fig. 17(b) focuses instead on buildings whose stress is lower than 200 MPa, distinguishing the ground motion by color. It can be seen that some earthquakes tend to generate higher stresses in the models, while others seem less dangerous. It would be interesting to combine this information with that on deformability, i.e., inter-story drift or maximum displacement, or earthquake characteristics, using automatic tools that exploit AI to assess which architectural forms are both lightweight and tough.

The effect of curvilinearity, whose meaning is shown in Figs. 1(c) and 2, is considered in Fig. 18 in terms of three potential locations for implementing this geometric choice linked to the parameter X9, see again Table 1 and Fig. 2. This geometric effect aims to either expand or narrow the cross-section along the height of the building by creating a curvilinear guideline for sweeping the base floor; the two alternatives, to expand or narrow, are represented in Fig. 18(a) by square or dot symbols, and by blue and yellow colors in Fig. 18(b), respectively. In the database, the location of the curvilinearity affects the inter-story drift only when it is applied at half height (i.e., when the curvilinearity location is equal to 0.5), where it slightly increases the drift. When the curvilinearity location is away from this position, the influence on the stiffness turns out to be small. The role of the shape induced by the curvilinearity in increasing or decreasing the cross-section area is also not identifiable in a straightforward way, so architectural design can conveniently

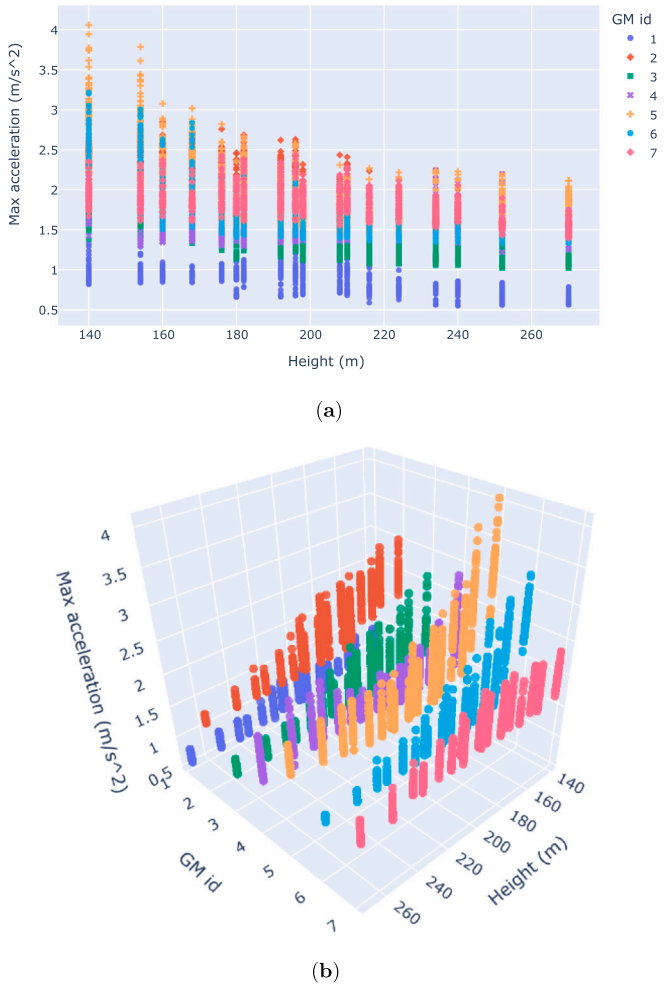


Fig. 15. All ground motions. Maximum acceleration vs height as a function of the given seismic input. The 3D plot can be inspected interactively with the HTML file available in the GitHub repository in the Visualization folder.

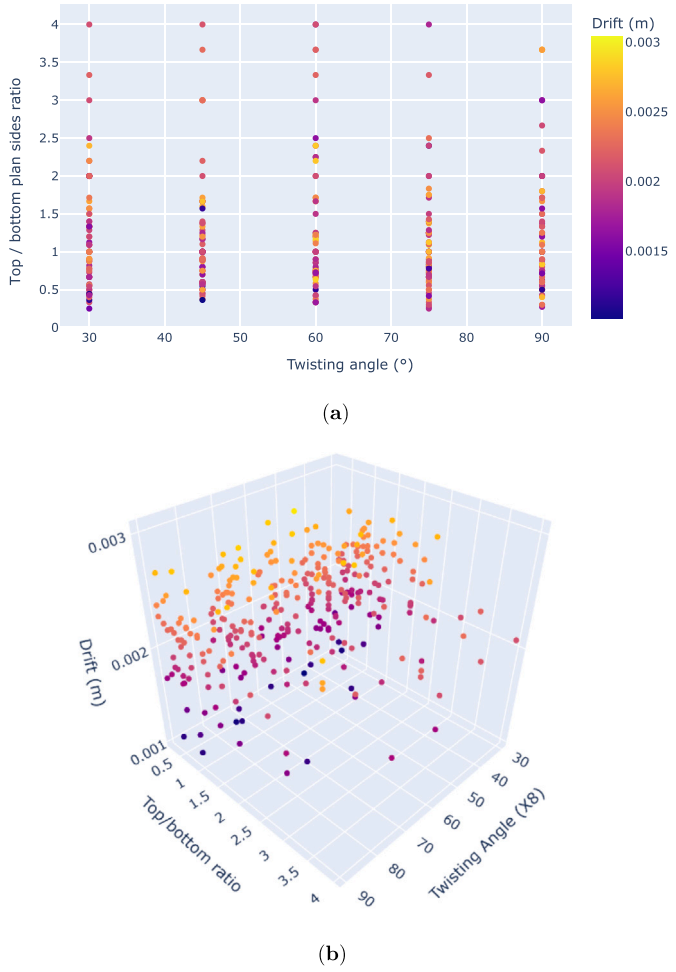


Fig. 16. Ground motion 1. Top/bottom plan side ratio vs twisting angle as a function of the maximum interstory drift. The 3D plot can be inspected interactively with the HTML file available in the GitHub repository in the Visualization folder.

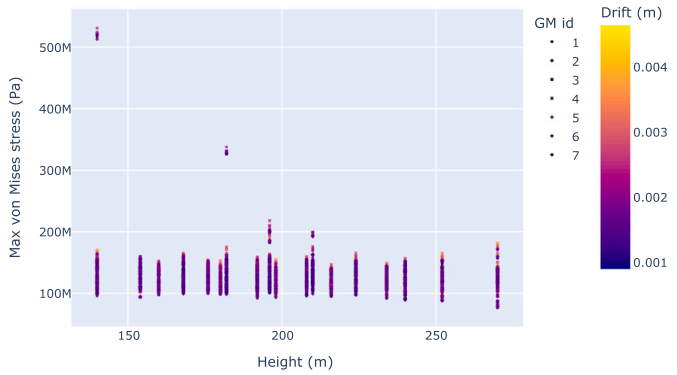
explore these options. Again, the effect of these choices is not easily captured by visual inspection, and the use of AI tools can therefore be beneficial.

These results allow a direct link to be drawn between each vertical transformation type and its effect on structural performance. For curvilinear models, a higher curvilinearity intensity (scale factor further from 1, i.e., more pronounced concave or convex forms) combined with a control-floor position closer to mid-height is generally associated with larger maximum interstory drift for a given building height, as shown in Fig. 18. For tapered models, a more pronounced tapering ratio (greater reduction of the top plan area relative to the bottom one) tends to reduce lateral mass at the top of the building, generally lowering drift for a given height, though this trend interacts with the number of stories and story height. For twisted models, larger twisting angles introduce additional coupling between lateral and torsional response, which can increase peak drift and torsional demand, particularly for taller, more slender buildings. Regarding the polygon side count of the top and bottom plans (parameters X1 and X2), buildings with a higher number of sides approach a more circular plan, generally leading to a more uniform distribution of diagrid stiffness around the perimeter, whereas lower-sided (triangular to hexagonal) plans exhibit more direction-dependent stiffness and consequently more variable drift between the X and Y directions.

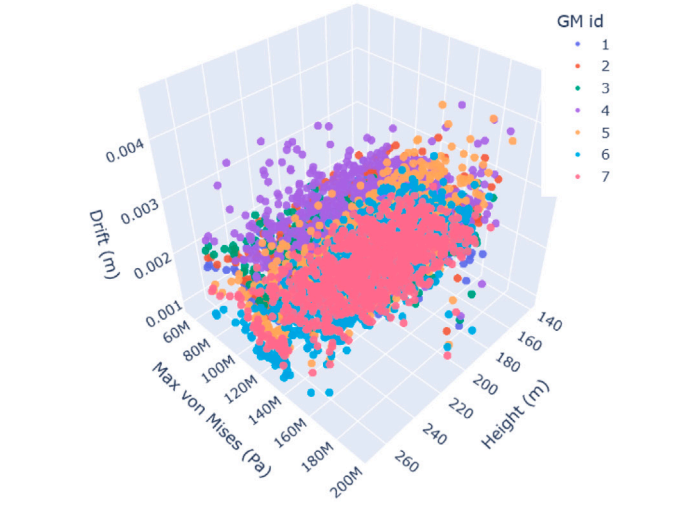
Finally, in Fig. 19 all the outputs are compared in terms of their maximum observed displacement, height and cost. Different symbols are used to identify each ground motion in Fig. 19(a), while colors are used instead in Fig. 19(b). By comparing Fig. 19(b) with Fig. 17(b), it is possible to better appreciate the effect of each ground motion on the buildings: for example, models subjected to ground motion 4 (purple) often show large displacements but small stresses, showcasing flexible structures or the possible contribution of different vibration modes. For ground motion 2 (red) larger stresses and moderate displacements are observed, suggesting a more balanced design in structural terms. The cost appears generally lower for buildings whose height is lower (i.e., dark color points in Fig. 19a) and higher for the tallest buildings, but there are both a significant number of relatively low cost cases around 250 m and a smaller number of expensive solutions for low heights. The impact of a specific ground motion is less clear, since exceptions in terms of cost/height are reported for every ground motion and AI therefore seems particularly suitable to dig into the database and identify valuable building models.

5. Limitations of this study

This study has several limitations, which should be considered when interpreting its results and could be improved in future work. (i) All

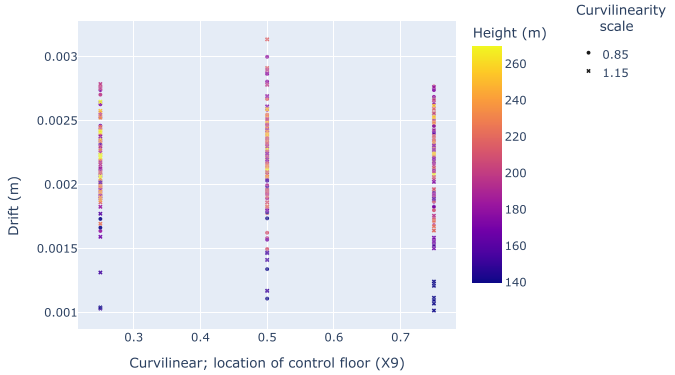


(a)

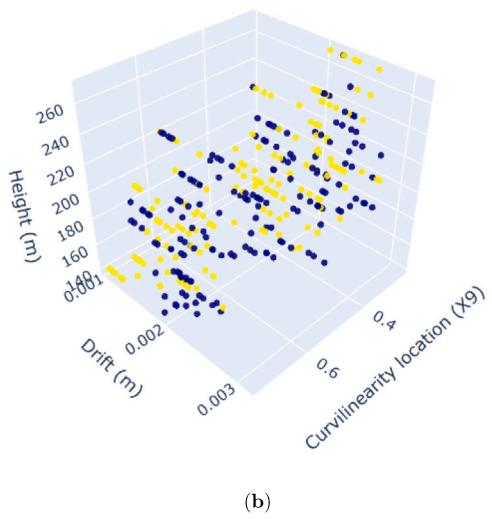


(b)

Fig. 17. All ground motions. Maximum von Mises stress vs height as a function of the maximum inter-story drift and earthquake ID. The 3D plot can be inspected interactively with the HTML file available at the GitHub repository in the Visualization folder.



(a)



(b)

Fig. 18. All ground motions. Maximum inter-story drift vs the location of applied curvilinearity as a function of height and earthquake ID. The 3D plot can be inspected interactively with the HTML file available at the GitHub repository in the Visualization folder.

structural responses are derived from linear-elastic dynamic time-history simulations; for the more demanding combinations of architectural form and ground motion within the dataset, some degree of inelastic behavior cannot be excluded, and nonlinear material and element models have not been employed due to the computational cost of running 7000 simulations. (ii) No site-specific seismic hazard analysis has been performed: the target spectrum is generic, reflecting the comparative rather than site-specific aim of the study. (iii) Seven ground motions have been adopted, consistent with the Eurocode 8 minimum requirement for time-history analysis, but fewer and less duration-diverse than those recommended by guidelines such as the Tall Buildings Initiative and the Los Angeles Tall Buildings Structural Design Council; long-period and long-duration records have not been specifically targeted. (iv) Several modeling simplifications have been adopted for tractability across 1000 models, including rigid floor diaphragms, no explicit *P*-delta effects, no wind loading, no vertical seismic excitation, and a core geometry standardized to a 12-sided polygon across all models. (v) The dataset comprises 1000 of the 61,200 possible parameter combinations, a trade-off dictated by the available computational resources. (vi) The present analysis exploits only the overall maxima of the time-varying structural responses, while the complete time-series data, made available in the repository, remain to be more fully exploited.

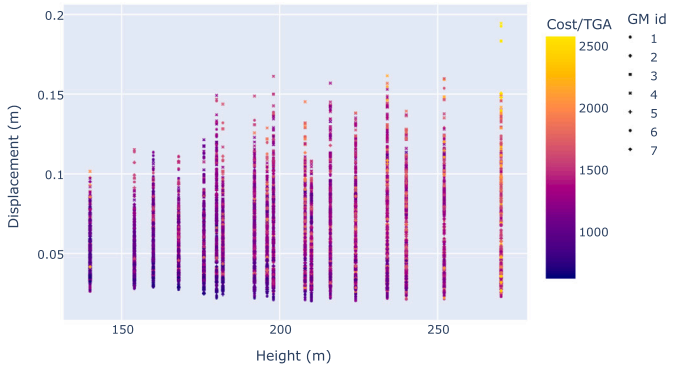


Fig. 19. All ground motions. Maximum building displacement vs height as a function of the total cost and earthquake ID. The 3D plot can be inspected interactively with the HTML file available at the GitHub repository in the Visualization folder.

6. Conclusions and future work

A dataset consisting of 1000 models of tall buildings with diagrids has been generated, and the structural responses to seven ground motions have been computed. The geometry of the buildings features curvilinear and tapered forms, with the top and bottom floors characterized by 3- to 12-sided polygons, and a twisting angle allowed along the building height. Dynamical analyses in the time domain have provided structural responses that have been supplemented by geometrical and environmental variables, to quantitatively characterize the behavior of each building. The database can be conveniently explored by the community, either to test AI tools effectiveness in the context of architecture and building construction, or to extract useful quantitative information regarding the influence of architectural choices. As an example, these data can be used for classification or regression, to train artificial neural networks or generative adversarial networks, or to implement generative design strategies.

Beyond the description of the dataset, the results point to several engineering-relevant trends. Taller, lighter buildings (particularly around 250 m and above) tend to coincide with higher interstory drift, while a limited but significant subset of tall models achieves favorable drift at moderate mass and cost, indicating that architectural form choices can substantially mitigate the seismic penalty typically associated with height. Buildings of intermediate height (160–250 m) display the least consistent relationship between mass and drift, suggesting that architectural form has a comparatively larger influence on seismic performance in this height range than at the extremes of the design space. The influence of the specific ground motion on the observed cost/height and drift/height trends is comparatively weak, since exceptions are found for every ground motion considered, suggesting that no single architectural rule of thumb dominates and reinforcing the case for AI-based exploration of the database to identify valuable building configurations that would not be apparent from simple trend analysis alone.

The procedure and dataset presented here can be extended in several directions. Incorporating wind loads would require an aerodynamic or code-based wind-load model compatible with the parametric geometry-generation pipeline, together with additional environmental input features; given the sensitivity of tall building aerodynamic behavior to architectural form, this is expected to interact strongly with the existing geometric parameters. Introducing material nonlinearity would require nonlinear fiber-section elements and calibrated steel constitutive laws, at a substantially higher computational cost per simulation, likely requiring either a reduction in the number of sampled models or a shift toward surrogate-assisted sampling strategies to keep the overall computational burden tractable. Extending the architectural design space with additional geometric features would require a corresponding extension of the sampling strategy and of the surrogate model architecture. Beyond the overall maxima considered here as a first proof of concept, the retained time-series data could be further exploited to extract response spectra and dominant-frequency content per building (to study form-dependent dynamic amplification), cumulative damage or fatigue-oriented indices such as the number of drift-threshold exceedances, and time-to-peak/phase relationships between core and diagrid contributions, potentially revealing load-redistribution patterns not visible from maxima alone.

Finally, by building a surrogate model and exploiting generative design, a natural-language design-assistant tool could be developed to let designers query the surrogate model directly. This development would require (i) translating a designer's prompt into constraints or targets over the parameters X1–X10, (ii) querying a surrogate model, e.g., a trained neural network [47], to estimate the relevant responses in near-real time, and (iii) optionally invoking a generative step to propose compliant alternative designs, with the surrogate model's prediction uncertainty communicated to the designer to avoid overconfident feedback, and with any final proposed solution verified against full time-history simulation before use at later design stages. On this issue, the main technical challenges lie in mapping qualitative design intents onto the quantitative design space, in communicating prediction uncertainty to avoid overconfident feedback, and in maintaining a human-in-the-loop verification of any proposed solution against full time-history simulation before its use at later design stages.

CRediT authorship contribution statement

Pooyan Kazemi: Writing – original draft, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Alireza Entezami:** Writing – review & editing, Supervision, Conceptualization. **Michela Turrin:** Supervision, Resources. **Charalampos Andriotis:** Supervision, Resources. **Stefano Mariani:** Writing – review & editing, Conceptualization. **Aldo Ghisi:** Writing – review & editing, Writing – original draft, Supervision, Investigation, Conceptualization.

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. Embodied carbon and cost calculation

The methodology for calculating embodied carbon follows the guidelines and data primarily derived from [48]. The formula used to compute the embodied carbon in kilograms of CO_2-equivalent ($kgCO_2e$) is:

$$\text{Embodied carbon } (kgCO_2e) = \text{Material quantity } (kg) \times \text{Carbon factor } (kgCO_2e/kg) \tag{A.1}$$

Eq. (A.1) incorporates life cycle stages A1–A3 (raw material supply, transport, and manufacturing), A4–A5 (transport and construction and installation), and C2–C4 (end of life) according to the nomenclature used in [48], and is here adopted to obtain comprehensive carbon factors for each material, as detailed in Table A.8.

For structural components, such as columns and beams, and for the flooring systems the embodied carbon factors $E_{\text{structure}}$ and E_{floor} are computed by allowing for the relevant weights. The analysis considers multiple material options, including typical steel, recycled steel, reused steel, typical concrete, and sustainable concrete variants. The calculations span six columns to cover all the combinations of materials, and six other columns with a special focus on embodied carbon per gross internal area (EC GIA) in $kgCO_2e/m^2$.

As for the financial assessment, it includes the cost of structural components, joint assemblies, and flooring systems. The calculations are underpinned by consolidated data in Table A.9. For the two primary materials, the analysis delineates the costs associated with concrete floor systems and the steel structural framework, encompassing core elements, diagrids, and radial beams as follows:

- Cost of structure. The structural cost is derived from the weight of the material W_{material} factored by the sum of unitary material cost UMC_{material} and unitary labor cost ULC_{material}:

$$\text{Cost of structure} = W_{\text{material}} \times (UMC_{\text{material}} + ULC_{\text{material}}) \tag{A.2}$$

- Cost of joints. Given the structural reliance on joint integrity, the joint cost is a function of the structural cost adjusted by the proportion of fixed to total joints and pinned joints to total joints, factoring their respective cost coefficients according to:

$$\text{Cost of joints} = \text{Cost of structure} \times \left(\frac{\text{Fixed joints}}{\text{Total joints}} \times 0.6 + \frac{\text{Pinned joints}}{\text{Total joints}} \times 0.2 \right) \tag{A.3}$$

To align this value with the structural analysis assumption where all joints are considered fixed, the cost is amplified by a factor equal to 1.6 to reflect the design choice.

Table A.9
Cost analysis.

Item	Cost	Unit
Steel material cost (UMC) [53]	1.5	€/kg
Steel labor cost (ULC) [53]	0.9	€/kg
Precast RC slab material cost [54]	50	€/m^2
Precast RC slab labor cost [54]	25	€/m^2
Fixed joints cost coefficient [55]	1.6	–
Pinned joints cost coefficient [55]	1.2	–

Table A.10
Environmental responses, with material IDs listed in Table A.8.

Category	Parameters
Economic costs	- Cost of structure - Cost of floors - Cost of land - Total project cost - Cost per Total Gross Area (TGA) - Marketable price
Embodied Carbon (EC)	- EC Steel option 1 [ID 1] - EC Steel option 2 [ID 2] - EC Steel option 3 [ID 3] - EC Concrete option 1 [ID 4] - EC Concrete option 2 [ID 5] - EC total Steel: ID 1 + Concrete: ID 4 - EC total Steel: ID 2 + Concrete: ID 4 - EC total Steel: ID 3 + Concrete: ID 4 - EC total Steel: ID 1 + Concrete: ID 5 - EC total Steel: ID 2 + Concrete: ID 5 - EC total Steel: ID 3 + Concrete: ID 5
Embodied Carbon per Gross Internal Area (EC GIA)	- EC GIA Steel: ID 1 + Concrete: ID 4 - EC GIA Steel: ID 2 + Concrete: ID 4 - EC GIA Steel: ID 3 + Concrete: ID 4 - EC GIA Steel: ID 1 + Concrete: ID 5 - EC GIA Steel: ID 2 + Concrete: ID 5 - EC GIA Steel: ID 3 + Concrete: ID 5

- Cost of floor. The flooring cost calculation leverages the floor area and aggregates the unitary material and labor costs as:

$$\text{Cost of floor} = \text{Area} \times (UMC_{\text{floor}} + ULC_{\text{floor}}) \tag{A.4}$$

The cost analysis thus accounts for the nuanced financial implications of material selection, labor, and construction methodologies, providing a granular even if estimated view of the project's economic demands. In the dataset, an item is dedicated to the total cost while another delineates the total cost per unit area.

The complete list of environmental responses is listed in Table A.10.

Table A.8
Carbon coefficient factor ($kgCO_2e/kg$) for different steel and concrete types. A5 values are extracted from [48], while A1–A4, and C2–C4 values are obtained from the reported sources.

ID	Material	Source	WF	A1–A3	A4	A5	C2–C4	Total carbon coefficient
1	Structural steel "typical", global average [49]	ICE	0.01	1.55	0.005	0.0175	0.018	1.5905
2	Recycled structural steel sections, EA furnace (Beltrame SPA) [50]	EPD	0.01	0.3693	0.005	0.0057	0.018	0.398
3	Reused structural steel sections (EMR Group) [51]	EPD	0.01	0.0466	0.005	0.0025	0.018	0.0721
4	RC30 reinforced concrete CEM I, 0% GGBS [49]	ICE	0.053	0.129	0.005	0.0081	0.018	0.1601
5	C28/35 CIIIB ready mix concrete, 12% GGBS (Hanson UK) [52]	EPD	0.053	0.0696	0.005	0.0049	0.018	0.0975

References

[1] Moon KS. Sustainable structural engineering strategies for tall buildings. Struct Des Tall Spec Build 2008;17:895–914. https://doi.org/10.1002/tal.475

[2] Moon KS. Comparative evaluation of structural systems for tapered tall buildings. Buildings 2018;8. https://doi.org/10.3390/buildings8080108. https://www.mdpi.com/2075-5309/8/8/108.

[3] Ali M, Moon K. Advances in structural systems for tall buildings: emerging developments for contemporary urban giants. Buildings 2018;8:104. https://doi.org/10.3390/buildings8080104

[4] Jalali A, Massumi A. Effect of type of internal core on the behavior of diagrid tube systems under lateral seismic loads. Structures 2024;63:106464. https://doi.org/10.1016/j.istruc.2024.106464. https://www.sciencedirect.com/science/article/pii/S2352012424006167.

[5] Terán-Gilmore A, Roeslin S, Tapia-Hernández E, Cuadros-Hipólito E. Displacement-based design of tall earthquake-resistant diagrid systems. J Build Eng 2021;35:102022. https://doi.org/10.1016/j.jobe.2020.102022. https://www.sciencedirect.com/science/article/pii/S2352710220336548.

[6] Adeli H, Park HS. Fully automated design of super-high-rise building structures by a hybrid AI model on a massively parallel machine. AI Mag 1996;17:87. https://doi.org/10.1609/aimag.v17i3.1233. https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/1233.

[7] Park HS, Adeli H. Distributed neural dynamics algorithms for optimization of large steel structures. ASCE J Struct Eng 1997;123:880–8. https://doi.org/10.1061/(ASCE)0733-9445(1997)123:7(880)

[8] Wang B, Liu C. Intelligent design of high-rise diagrid structures using the snow ablation optimizer. Autom Constr 2025;178:106402. https://doi.org/10.1016/j.autcon.2025.106402

[9] Liu C, Wang B, Ma Y, Wu Z. Multi-objective optimization design of high-rise diagrid structure based on non-dominated sorting genetic algorithm II. J Build Eng 2025;108:112801. https://doi.org/10.1016/j.jobe.2025.112801

[10] Aldwaik M, Adeli H. Advances in optimization of highrise building structures. Struct Multidiscip Optim 2014;50:1488–615. https://doi.org/10.1007/s00158-014-1148-1

[11] Scaramozzino D, Albitos B, Lacidogna G, Carpinteri A. Selection of the optimal diagrid patterns in tall buildings within a multi-response framework: application of the desirability function. J Build Eng 2022;54:104645. https://doi.org/10.1016/j.jobe.2022.104645

[12] Ali MM, Moon KS. Structural developments in tall buildings: current trends and future prospects. Archit Sci Rev 2007;50:205–23. https://doi.org/10.3763/asre.2007.5027

[13] Moon KS. Optimal grid geometry of diagrid structures for tall buildings. Archit Sci Rev 2008;51:239–51. https://doi.org/10.3763/asre.2008.5129

[14] Tomei V, Imbimbo M, Mele E. Optimization of structural patterns for tall buildings: the case of diagrid. Eng Struct 2018;171:280–97. https://doi.org/10.1016/j.engstruct.2018.05.043. https://www.sciencedirect.com/science/article/pii/S0141029617325695.

[15] Liu C, Luo X, Fang D, Shi C, Sarvar S, Zhao B. Study on flexural stiffness of diagrid non-stiffened node based on four-spring assemblage model. Eng Struct 2019;198:109500. https://doi.org/10.1016/j.engstruct.2019.109500. https://www.sciencedirect.com/science/article/pii/S014102961932111X.

[16] Fang D, Liu C. Mechanical characteristics and deformation calculation of steel diagrid structures in high-rise buildings. J Build Eng 2021;42:103062. https://doi.org/10.1016/j.jobe.2021.103062. https://www.sciencedirect.com/science/article/pii/S2352710221009207.

[17] Jafari M, Alipour A. Review of approaches, opportunities, and future directions for improving aerodynamics of tall buildings with smart facades. Sustain Cities Soc 2021;72:102979. https://doi.org/10.1016/j.scs.2021.102979. https://www.sciencedirect.com/science/article/pii/S2210670721002651.

[18] Asadi E, Salman AM, Li Y. Multi-criteria decision-making for seismic resilience and sustainability assessment of diagrid buildings. Eng Struct 2019;191:229–46. https://doi.org/10.1016/j.engstruct.2019.04.049

[19] Ekici B, Kazanasmaz ZT, Turrin M, Tasgetiren MF, Sariyildiz IS. Multi-zone optimisation of high-rise buildings using artificial intelligence for sustainable metropolises. Part 2: Optimisation problems, algorithms, results, and method validation. Sol Energy 2021;224:309–26. https://doi.org/10.1016/j.solener.2021.05.082

[20] Muthumanickam NK, Brown N, Duarte JP, Simpson TW. Multidisciplinary design optimization in architecture, engineering, and construction: a detailed review and call for collaboration. Struct Multidiscip Optim 2023;66. https://doi.org/10.1007/s00158-023-03673-y. https://www.scopus.com/inward/record.uri?eid=2-s2.0-85176043578&doi=10.1007/s00158-023-03673-y&partnerID=40&md5=60810ba36bd5f00642d11e39f1f1eed2., Cited by: 25; All Open Access, Green Open Access.

[21] Preisinger C. Linking structure and parametric geometry. Archit Des 2013;83:110–3. https://doi.org/10.1002/ad.1564

[22] Bertagna F, D'acunto P, Ohlbrock PO, Moosavi V. Holistic design explorations of building envelopes supported by machine learning. J Facade Des Eng 2021;9:31–46. https://doi.org/10.7480/jfde.2021.1.5423, Cited by: 7.

[23] BuHamdan S, Alwisy A, Bouferguene A. Generative systems in the architecture, engineering and construction industry: a systematic review and analysis. Int J Archit Comput 2021;19:226–49. https://doi.org/10.1177/1478077120934126

[24] Kookalani S, Parn E, Brilakis I, Dirar S, Theofanous M, Faramarzi A, Mahdavipour MA, Feng Q. Trajectory of building and structural design automation from generative design towards the integration of deep generative models and optimization: a review. J Build Eng 2024;97:110972. https://doi.org/10.1016/j.jobe.2024.110972

[25] Khan A, Chang S, Chang H. Generative AI approaches for architectural design automation. Autom Constr 2025;180:106506. https://doi.org/10.1016/j.autcon.2025.106506

[26] Sun H, Burton HV, Huang H. Machine learning applications for building structural design and performance assessment: state-of-the-art review. J Build Eng 2021;33:101816. https://doi.org/10.1016/j.jobe.2020.101816

[27] Farrar CR, Worden K. Structural health monitoring: a machine learning perspective. John Wiley & Sons Inc; 2012. https://doi.org/10.1002/9781118443118

[28] Saldana Ochoa K, Ohlbrock PO, D'acunto P, Moosavi V. Beyond typologies, beyond optimization: exploring novel structural forms at the interface of human and machine intelligence. Int J Archit Comput 2021;19:466–90. https://doi.org/10.1177/1478077120943062, Cited by: 27.

[29] Zhou K, Xie D-L, Xu K, Zhi L-H, Hu F, Shu Z-R. A machine learning-based stochastic subspace approach for operational modal analysis of civil structures. J Build Eng 2023;76:107187. https://doi.org/10.1016/j.jobe.2023.107187

[30] Nguyen HD, LaFave JM, Lee Y-J, Shin M. Rapid seismic damage-state assessment of steel moment frames using machine learning. Eng Struct 2022;252:113737. https://doi.org/10.1016/j.engstruct.2021.113737

[31] Junda E, Málaga-Chuquitaype C, Chawgien K. Interpretable machine learning models for the estimation of seismic drifts in CLT buildings. J Build Eng 2023;70:106365. https://doi.org/10.1016/j.jobe.2023.106365

[32] Thai H-T. Machine learning for structural engineering: a state-of-the-art review. Structures 2022;38:448–91. https://doi.org/10.1016/j.istruc.2022.02.003. https://www.sciencedirect.com/science/article/pii/S2352012422000947.

[33] Zhang W, Wang Y, Du Z, Liu C, Youn S-K, Guo X. Machine-learning assisted topology optimization for architectural design with artistic flavor. Comput Methods Appl Mech Eng 2023;413:116041. https://doi.org/10.1016/j.cma.2023.116041. https://www.sciencedirect.com/science/article/pii/S0045782523001652.

[34] Millán E, Belmonte M-V, Boned F-J, Gavilanes J, de-la Cruz J-LP, Díaz-López C. Using machine learning techniques for architectural design tracking: an experimental study of the design of a shelter. J Build Eng 2022;51:104223. https://doi.org/10.1016/j.jobe.2022.104223. https://www.sciencedirect.com/science/article/pii/S2352710222002364.

[35] Orhan T, Taşkın K. Automated topology design of high-rise diagrid buildings by genetic algorithm optimization. Struct Des Tall Spec Build 2021;30:e1853. https://doi.org/10.1002/tal.1853

[36] Jahangiri V, Akbarzadeh MR, Shahamat SA, Asgari A, Naeim B, Ranjbar F. Machine learning-based prediction of seismic response of steel diagrid systems. Structures 2025;80:109791. https://doi.org/10.1016/j.istruc.2025.109791

[37] Liu C, Li G, He B, Zhou C, Ma Y. Sustainable seismic performance of diagrid core-tube structure with replaceable steel coupling beam. Sustainability 2024;16. https://doi.org/10.3390/su16072690

[38] Nguyen S, Reiter S, Rigo P. A review on simulation-based optimization methods applied to building performance analysis. Appl Energy 2014;113:1043–58. https://doi.org/10.1016/j.apenergy.2013.08.061

[39] Kazemi P, Ghisi A, Mariani S. Classification of the structural behavior of tall buildings with a diagrid structure: a machine learning-based approach. Algorithms 2022;15. https://doi.org/10.3390/a15100349

[40] Kazemi P, Entezami A, Ghisi A. Machine learning techniques for diagrid building design: architectural–structural correlations with feature selection and data augmentation. J Build Eng 2024;86:108766. https://doi.org/10.1016/j.jobe.2024.108766. https://www.sciencedirect.com/science/article/pii/S2352710224003346.

[41] Kazemi P, Ghisi A, Entezami A. Artificial intelligence-powered computational strategies in selecting and augmenting data for early design of tall buildings with outer diagrids. Buildings 2024;14. https://doi.org/10.3390/buildings14041118. https://www.mdpi.com/2075-5309/14/4/1118.

[42] Ilgin HE. Space efficiency in contemporary supertall residential buildings. Architecture 2021;1:25–37. https://doi.org/10.3390/architecture1010004. https://www.mdpi.com/2673-8945/1/1/4.

[43] Demirel GN. Calculating the core area percentage for steel and concrete tall building. Private communication 2024.

[44] McKenna F, Fenves GL, Scott MH, et al. Open system for earthquake engineering simulation. Berkeley, CA: University of California; 2000.

[45] Seyhan E, Stewart JP, Ancheta TD, Darragh RB, Graves RW. NGA-West2 site database. Earthq Spectra 2014;30:1007–24. https://doi.org/10.1193/062913EQS180M

[46] Hall J. Problems encountered from the use (or misuse) of Rayleigh damping. Earthq Eng Struct Dyn 2006;35:525–45. https://doi.org/10.1002/eqe.541

[47] Kazemi P. A workflow for architectural and structural early design of tall buildings with outer diagrids through AI-driven and computational design strategies [Ph.D. thesis]. Doctoral School in Structural, Seismic and Geotehcnical Engineering, Politecnico Di Milano, Italy; 2024. https://hdl.handle.net/20.500.14242/206749.

[48] Kanyilmaz A, Hoi Dang V, Kondratenko A. How does conceptual design impact the cost and carbon footprint of structures? Structures 2023;58:105102. https://doi.org/10.1016/j.istruc.2023.105102. https://www.sciencedirect.com/science/article/pii/S2352012423011918.

[49] Gibbons OP, Orr J, Archer-Jones C, Arnold W, Green D. How to calculate embodied carbon. Institution of Structural Engineers (ISTRUCTE); 2022.

[50] Beltrame Group. Environmental product declaration. Technical report Beltrame Group; 2020.

[51] EMR. Reusable steel environmental product declaration programme. Technical report EMR; 2022.
[52] Hanson UK. Environmental product declaration. Environmental Product Declaration Series 44. 2017.
[53] Janisch A. Kostenfaktoren im Stahlbau - was kostet ein kg Stahlbau 2023? 2023. https://jactio.com/kostenfaktoren-im-stahlbau/.
[54] Rete Ferroviaria Italiana. Tariffa dei prezzi FA. Opere civili e industriali. Edizione 2024; 2023.
[55] Simões LMC. Optimization of frames with semi–rigid connections. Comput Struct 1996;60:531–9. https://doi.org/10.1016/0045-7949(95)00427-0. https://www.sciencedirect.com/science/article/pii/0045794995004270.