Progress in Engineering Science 3 (2026) 100313

Contents lists available at ScienceDirect

Progress in Engineering Science

journal homepage: www.sciencedirect.com/journal/progress-in-engineering-science

Inverse unit-cell design for truss lattices with percolation-inspired peak-stress tailoring☆

Salvatore Annunziata, Luca Lomazzi*, Marco Giglio, Andrea Manes

Department of Mechanical Engineering, Politecnico di Milano, Via la Masa, 1, 20156, Milan, Italy

ARTICLE INFO

Keywords:
Lattice structures
Inverse modelling
Percolation
Impact absorption
Neural network
High strain-rate

ABSTRACT

Metamaterials offer strong potential for lightweight impact mitigation because their architecture can be tailored to control stiffness, energy absorption, and transmitted stress. However, their practical design remains limited by the cost of generating large numerical databases and by the difficulty of solving inverse design problems, where unit-cell parameters must be identified from target macroscopic responses. This study presents an automated finite element workflow for generating a dataset of periodic truss lattices and training a Unit-cell Optimization Network (UON) for inverse unit-cell design. Kelvin and octagonal beam-based lattices are generated parametrically, simulated under dynamic compression, and post-processed to extract four mechanical descriptors: initial stiffness, plateau stress, densification strain, and absorbed energy. After cleaning, 1376 valid simulations are retained for UON training and evaluation. The inverse model combines regression for the continuous strut thickness with classification for the discrete unit-cell size, achieving a test RMSE of 0.012571 mm for thickness prediction and a cell-size classification accuracy of 92.7536%. In parallel, a percolation-inspired study is performed separately from the UON training dataset by randomly removing beams according to an assigned absence probability. The results show an approximately linear decrease in plateau stress and absorbed energy with increasing absence probability, while peak stress follows an exponential-type reduction. These findings demonstrate that automated dataset generation, mixed regression-classification inverse modelling, and controlled connectivity reduction can support efficient tailoring of lattice energy absorption performance.

1. Introduction

The demand for lightweight structures capable of dissipating large amounts of energy while limiting transmitted loads is increasing in transportation, aerospace, and protective engineering. In these applications, crashworthiness and impact/blast mitigation require a careful balance between stiffness, strength, and stable deformation. Functionally graded thin-walled components and cellular architectures are widely used to tune collapse modes and improve energy absorption [1]. Under impact-like loading, the dynamic response becomes particularly relevant because early stress peaks and rate-dependent collapse can dominate transmitted loads. Hopkinson-bar experiments have shown that architected lattices can distribute impact loading over time and reduce peak stress, although strong rate effects remain [2]. Complementary nonlinear finite element studies indicate that beam-based models can capture key trends and help distinguish intrinsic material rate effects from structural rate sensitivity arising from topology and layer-by-layer collapse mechanisms [3]. More broadly, comparisons among lattice families confirm that the dominant deformation mode, either stretching- or bending-dominated, governs failure initiation and strongly affects specific stiffness, strength, and energy absorption [4]. Mechanical metamaterials, defined as architected cellular solids whose macroscopic response is mainly controlled by microstructural geometry, provide a powerful route for tailoring properties beyond those of the base material. Their foundations are rooted in the mechanics of natural cellular solids [5], where effective properties depend on relative density, cell-wall behaviour, and architecture [6]. Additive manufacturing has accelerated the development of periodic architectures and programmable deformation modes, but has also highlighted that high energy absorption requires careful multiscale design [7]. Hierarchical concepts, such as dual-scale plate-lattice structures, show how combining architectures at different length scales can improve specific performance under compression [8]. Recent perspectives further emphasize that data-driven design and optimization will be essential for the next generation of metamaterial systems [9]. Among architected materials,

☆ This article is part of a Special issue entitled: 'PES_IGF28-MedFract3' published in Progress in Engineering Science.

* Corresponding author.
E-mail address: luca.lomazzi@polimi.it (L. Lomazzi).

https://doi.org/10.1016/j.pes.2026.100313
Received 28 February 2026; Received in revised form 14 May 2026; Accepted 1 June 2026
Available online 8 June 2026

periodic truss lattices are especially attractive because their behaviour can be controlled through a compact set of design variables, including unit-cell topology, cell size, strut slenderness, relative density, and connectivity. Classic analyses show that lattice connectivity governs the transition between bending- and stretching-dominated behaviour, with stretching-dominated architectures often providing higher weight efficiency in stiffness and strength [10]. The octet-truss represents a canonical stretching-dominated topology, with well-established effective properties and collapse mechanisms, including the role of strut buckling [11]. Kelvin-type lattices remain widely used in energy absorption studies, and experiments confirm their sensitivity to architectural parameters such as infill and relative density [12]. However, performance is not governed by topology alone. Unit-cell size, specimen dimensions, and layout can modify compressive behaviour and strength convergence, introducing scale and manufacturing-resolution effects that are relevant for benchmarking and optimization [13]. Distinct Kelvin-based responses have also been reported in tension and compression [14], while octagonal-based unit cells have been investigated for crashworthiness-oriented absorbers, showing how geometry-driven load paths can tune energy absorption [15]. To convert these design variables into actionable maps, numerical models must balance physical fidelity with computational scalability. Both three-dimensional continuum models and beam-based finite element formulations can reproduce experimental compression trends when geometry, boundary conditions, contacts, and constitutive assumptions are carefully defined [16]. Validated procedures also show that accurate material calibration and faithful reproduction of the test configuration are essential to predict stiffness, peak stress, and post-yield response [17]. Since large design spaces are expensive to explore with high-fidelity models alone, reduced-order descriptions that preserve connectivity, such as lattice and spring-network models, are useful for studying rigidity and connectivity-driven mechanics [18]. More recently, dataset-generation frameworks have enabled large metamaterial databases under explicit geometric constraints [19], while multi-fidelity strategies combining low-cost simulations with higher-fidelity validation inside Bayesian optimization loops have been proposed to accelerate architected-material discovery [20]. A further requirement for scalable design pipelines is the robust extraction of macroscopic descriptors from simulated or experimental stress–strain curves. Lattice responses are commonly reduced to initial stiffness, plateau stress, densification strain, and absorbed energy, but these quantities can be sensitive to oscillations, transient peaks, and discrete collapse events. Controlled smoothing and regularization can improve the stability of post-processing and parameter identification [21]. Densification is particularly delicate, and energy-absorption-efficiency methods provide consistent estimates while distinguishing onset and densification strains [22]. Under dynamic loading, inertia, specimen size, and base-material rate dependence can further modify the apparent crushing stress and complicate the interpretation of plateau behaviour [23]. These aspects motivate standardized modelling and post-processing routines so that descriptors remain comparable across large simulation datasets. Once stable descriptors are available, lattice design can be formulated as an optimization or inverse-design problem constrained by performance and manufacturability. Topology optimization provides a rigorous route for architected-material design [24], while database-driven comparisons show that unit-cell choice and arrangement are strongly application-dependent [25]. Hybrid solid-lattice strategies address component-level requirements and additive-manufacturing constraints by combining dense regions with lattice infills [26]. Since iterative search may be computationally expensive, Bayesian optimization has been used to guide designs towards tailored responses [27]. In parallel, machine-learning-assisted genetic algorithms have enabled inverse design targeting customized loading curves [28] and improved buckling resistance in truss lattices [29]. Forward surrogates are also increasingly important for rapid screening, as highlighted in recent reviews [30]; examples include CNN-based prediction of effective properties from limited data [31], end-to-end workflows connecting finite element datasets to lattice design [32], and surrogate studies extracting insight from simulation databases [33]. Current models are also moving towards richer outputs, including sequence-aware prediction of full compressive stress–strain curves [34], graph neural networks that encode node-edge connectivity and predict deformation or stress–strain behaviour [35], physics-informed learning for improved data efficiency [36], and hypernetwork-driven constitutive surrogates adaptable across metamaterial families at large deformation [37]. Beyond deterministic geometric parameterization, connectivity can itself be treated as a design variable. Percolation theory describes the emergence or loss of system-spanning clusters under stochastic bond or site removal, linking local connectivity to macroscopic load transfer near critical thresholds [38]. These thresholds depend on lattice type and coordination number [39]. Percolation concepts have been applied to microstructural networks [40] and random fiber systems [41], while stochastic lattice architectures and defect-tolerant designs show that connectivity can control stiffness, strength, and deformation-mode transitions [42]. More generally, defect engineering and mesoscale heterogeneity can delay localization and improve post-yield robustness [43]. Simulation-driven approaches targeting prescribed stress–strain signatures further support the idea of enlarging the design space beyond pure geometric scaling [44]. Despite this progress, several gaps remain. Automated dataset-generation and multi-fidelity optimization strategies are available [19,20], but end-to-end workflows that jointly enforce admissible geometry, standardized finite element evaluation, and robust descriptor extraction are still uncommon at scale [22]. Iterative optimization also remains costly when design variables are mixed, namely continuous and discrete [24,27]. In addition, percolation-inspired connectivity reduction has rarely been investigated as a controlled design lever for peak-stress tailoring in energy-absorbing lattices, despite its direct connection with load-path continuity and defect-tolerant mechanics [38,39,42,43]. Finally, simulation-driven strategies targeting complete or partial stress–strain signatures [44] are still rarely coupled with reproducible high-throughput workflows and mixed-variable inverse tools. This work addresses these gaps by presenting an end-to-end workflow for impact-oriented lattice design that integrates automated geometry generation, high-throughput finite element evaluation, standardized post-processing of crushing descriptors, and mixed inverse learning. The proposed Unit-cell Optimization Network recovers unit-cell parameters from target mechanical descriptors by combining regression for continuous strut thickness and classification for discrete unit-cell size. In addition, a percolation-inspired connectivity-reduction strategy based on probabilistic beam removal is introduced to tailor early stress peaks and quantify the associated trade-off in plateau stress and absorbed energy. Selected higher-fidelity three-dimensional simulations are used to check whether the observed peak-stress reduction persists beyond the beam formulation.

2. Materials and methods

The workflow developed in this study is implemented in MATLAB and is conceived as an end-to-end routine for (i) parametrically generating beam-based lattice metamaterials, (ii) exporting simulation-ready finite element (FE) input files in series, (iii) post-processing solver outputs into compact performance descriptors, and (iv) training an inverse surrogate (Unit-cell Optimization Network, UON) that maps target mechanical descriptors back to unit-cell design parameters. Concerning the part of the workflow dedicated to the generation of the beam-based lattice metamaterials, two branches are available: a periodic branch, which constructs lattices by repeating a prescribed unit cell, and a stochastic branch, which produces spatially random networks via a Voronoi–Delaunay procedure (Fig. 1); however, the present work focuses on periodic truss metamaterials based on Kelvin and octagonal unit cells. In the periodic routine, the unit cell is defined by its nodal coordinates and beam connectivity (edge list), while the geometric

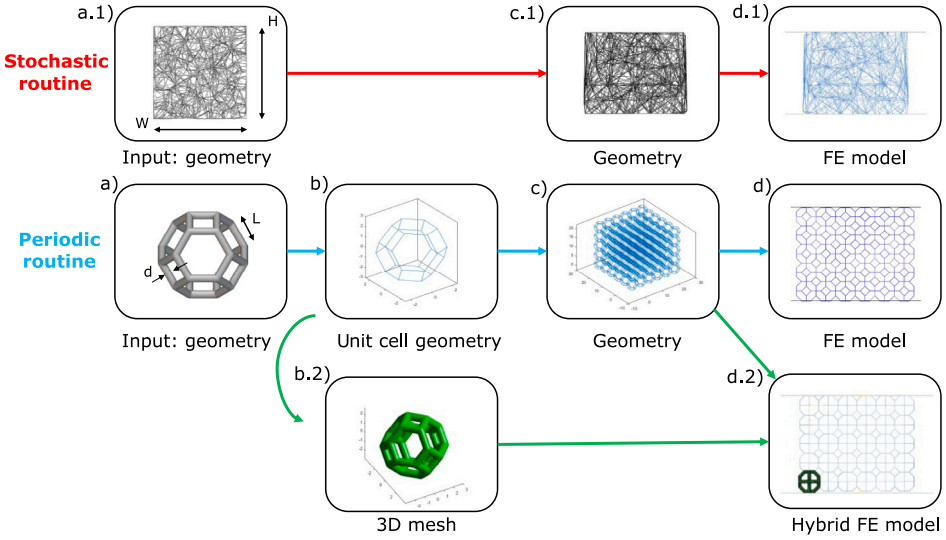

Fig. 1. Framework workflow. Periodic routine: (a) The user selects the type of unit cell and defines its geometric parameters (b) The developed code generates an input for the selected unit cell. (c) The unit cell input file is the input to generate the corresponding lattice structure input file. (d) Finite element analyses are performed. (b.2) A custom algorithm generates a .stl file building iso-surface around the cell created at step (b). (d.2) The hybrid 1D+3D model is generated. Stochastic routine: (a.1) The user selects the geometric parameters of the stochastic lattice structure. (c.1) The corresponding lattice structure input file is generated. (d.1) Finite element analyses are performed. The hybrid 1D–3D branch is retained in the workflow as an optional capability of the framework; in the present study, the UON dataset is generated from beam-based simulations [45].

Table 1
Main finite element settings adopted for the beam-based simulations used in dataset generation.

Item	Setting
Software and solver	Abaqus/Explicit 2023, dynamic explicit analysis
Lattice element type	B31 beam elements, circular cross-section
Compression plates	Rigid plates discretized with R3D4 elements
Specimen dimensions	$30 \times 30 \times 25$ mm
Material	PA12 elastoplastic material with rate dependence
Density	9.19×10^{-10} tonne/mm^3
Elastic constants	$E = 1300$ MPa, $\nu = 0.33$
Plasticity	Tabulated stress-plastic strain curve
Rate dependence	Johnson–Cook rate term, $C = 0.05$, $\dot{\varepsilon}_0 = 0.02$ s^{-1}
Contact	General contact, hard normal behaviour, rough tangential behaviour
Boundary conditions	Lower plate fixed; upper plate laterally/rotationally constrained
Loading	Prescribed vertical velocity on upper plate
Dataset velocity	35 m/s
Step time	6.2×10^{-4} s for the 35 m/s case
Output extraction	Upper-plate displacement, lower-plate reaction force

degrees of freedom considered for dataset generation include the unit-cell characteristic length (selected from a discrete set of admissible cell sizes) and the beam cross-section size (or equivalently, a target relative density). The full lattice is obtained by periodic replication of the unit cell along the three directions to match user-prescribed global dimensions (width, depth, height); when required, the unit-cell geometry is scaled to ensure an integer number of repetitions within the prescribed domain; the corresponding beam thickness is then assigned consistently with the selected design variable or with the target relative-density condition, depending on the specific generation mode. The beam lattice is then discretized with one-dimensional elements using a user-defined element size, with automatic assignment of section properties and material data, and the resulting FE model is exported to the selected solver in a standardized format together with a structured metadata record (unique identifier, design parameters, mesh size, boundary-condition flags, and analysis settings).

For clarity and reproducibility, the main numerical settings used for the beam-based FE simulations are summarized in Table 1. The table refers to the high-throughput simulations used for dataset generation; selected 3D simulations used in the percolation study follow the same boundary and loading scheme, but replace the beam discretization with solid elements, as discussed in the corresponding section. The beam mesh density was defined by subdividing each lattice beam into a prescribed number of elements. To verify that the extracted macroscopic descriptors were not dominated by the selected beam discretization, a mesh-density sensitivity check was performed on representative Kelvin and octagonal configurations in Ref. [46] confirming that the adopted mesh density is adequate for the high-throughput descriptor extraction performed in this study.

For applications where local stress fields are required at limited computational cost, an optional hybrid modelling capability is included for periodic lattices: a subset of the unit cells can be represented with three-dimensional solid elements while the rest of the structure is retained as a beam network. In the present manuscript, this hybrid 1D–3D strategy is retained in the workflow description only as an optional capability of the broader framework. The high-throughput dataset used for UON training is generated from fully beam-based models, while the selected 3D simulations discussed in the percolation section are used only as higher-fidelity numerical checks of the peak-stress trend. The detailed formulation and validation of the hybrid 1D–3D approach are reported in Ref. [47]. In parallel, the stochastic Voronoi–Delaunay routine, reported for completeness, builds random lattices by first seeding

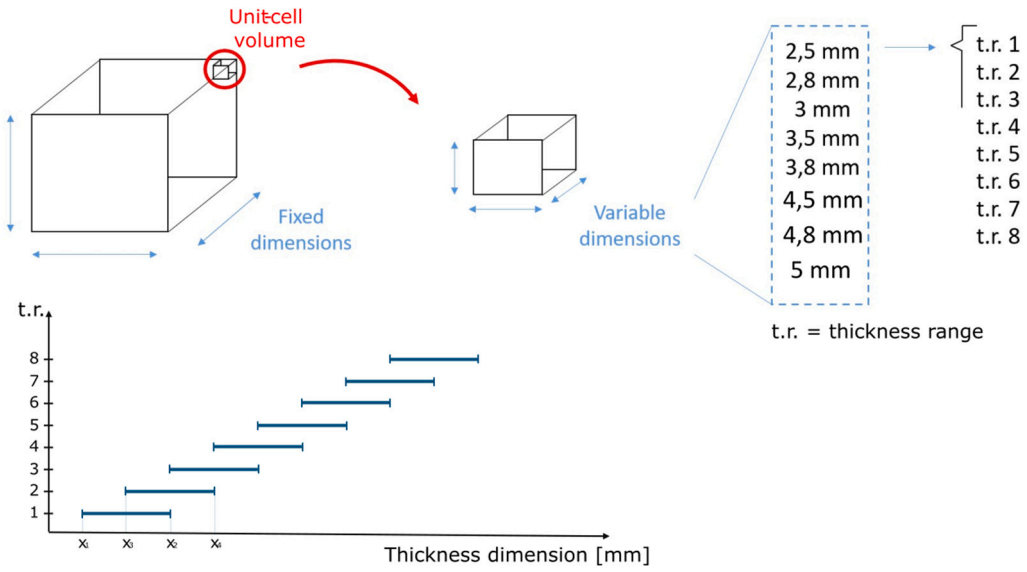

Fig. 2. Schematic representation of the dataset-generation strategy based on discrete unit-cell sizes and cell-size-dependent beam-thickness ranges. The lattice specimen has fixed global dimensions, while the unit-cell size is selected from an admissible discrete set. For each cell-size class, a dedicated thickness range (t.r.) is computed to preserve physically meaningful thickness-to-length ratios across the dataset.

points on the top and bottom faces with a minimum-spacing constraint, supplementing them with a three-dimensional Poisson-disk sampling within the volume, and finally generating a 3D Delaunay triangulation whose tetrahedral edges define the beam connectivity; the procedure returns a spatially random yet well-conditioned network controlled by lattice size, minimum beam length, and target relative density.

Dataset generation is performed by repeatedly sampling the admissible design space, generating the corresponding lattice geometry, meshing it, and exporting one FE input file per configuration; this production in series is orchestrated by a batch driver that logs each configuration into a parameter matrix and associates it with a unique simulation identifier to enable traceable database assembly and systematic exclusion of failed or non-physical realizations (e.g., invalid geometries, solver divergence, or out-of-bound derived properties). For periodic Kelvin and octagonal lattices, particular attention is devoted to ensuring that the sampled beam thickness remains geometrically consistent across different unit-cell sizes, since naive thickness sampling would produce unrealistic slenderness ratios (either excessively thin beams for large cells or overly stocky beams for small cells) and would bias both the FE response and the learning stage. To address this, thickness bounds are generated through an explicit scaling procedure embedded in the dataset routine: for each candidate unit cell and each discrete cell size, a specific function computes the length of every strut in the unit cell from the nodal coordinates and connectivity, and scales these lengths according to the selected cell size, returning a vector of characteristic edge lengths representative of the resized unit cell. These scaled edge lengths are then passed to thickness range generator, which converts geometric information into an admissible thickness interval expressed as a fraction of a characteristic strut length, thereby enforcing thickness-to-length ratios that remain within physically meaningful bounds when the cell size changes. In practice, this step provides a cell-size-dependent thickness sampling window used by the batch driver to draw the continuous thickness variable for each realization, ensuring that the generated models exhibit reasonable proportions and comparable manufacturability/structural plausibility across all cell-size classes; when global lattice dimensions are not commensurate with the nominal cell size, the same scaling logic remains consistent with the subsequent adjustment of unit-cell scaling to fit an integer number of replicas and the corresponding update of beam thickness required to preserve the targeted relative density. The unit-cell size is treated as a discrete design variable rather than as a continuous parameter. This choice reflects the way the lattice-generation routine constructs the specimens: the external dimensions of the lattice are kept fixed at $30 \times 30 \times 25$ mm, while the unit-cell size is selected from a finite set of admissible values. For each selected nominal cell size, the routine determines the corresponding number of cell repetitions along the three spatial directions and applies the necessary scaling to ensure compatibility with the fixed specimen dimensions. Therefore, the cell size represents a class label associated with a valid lattice layout, rather than an unrestricted continuous variable. This also motivates the mixed formulation adopted in the UON. The unit-cell size is predicted through multi-class classification because only predefined size classes are admissible in the generated dataset, whereas the beam thickness is predicted through regression because it is sampled as a continuous variable within each cell-size class. The thickness range is not fixed globally; instead, it is computed separately for each cell size from the scaled unit-cell edge-length distribution, so that the generated structures maintain physically reasonable thickness-to-length ratios across the entire dataset. The complete logic of this procedure is schematically illustrated in Fig. 2.

Once the FE analyses are completed, solver outputs are post-processed to extract the macroscopic stress–strain response under the prescribed loading condition, and a compact set of performance descriptors is computed and stored in an output matrix; these descriptors include the initial elastic stiffness (E_0), representative plateau stress levels (σ_{pl}), densification strain (ε_{dens}), and absorbed energy measures (En_{abs}), thereby transforming full curves into consistent scalar targets suitable for learning and inverse design. All four quantities are extracted from the macroscopic stress–strain curve associated with each simulation, which is obtained from the solver reaction force and plate displacement and converted into nominal stress and engineering strain as $\varepsilon = |u|/H_0$ and $\sigma = F/A_0$, where u is the imposed axial displacement, H_0 is the initial specimen height, F is the axial reaction force, and A_0 is the initial loaded area. Prior to feature extraction, the raw curve is processed by an adaptive pre-trimming step aimed at removing the initial transient portion in which contact establishment and numerical noise may dominate. This step relies on a lightly smoothed stress signal $\sigma_s(\varepsilon)$ (Savitzky–Golay filtering) and on its strain derivative $d\sigma_s/d\varepsilon$, and detects the first index at which either the stress or its derivative exceeds robust baseline thresholds computed on the initial portion of the curve (median and median absolute deviation estimators), with an additional mild hysteresis check to avoid spurious triggers; trimming is

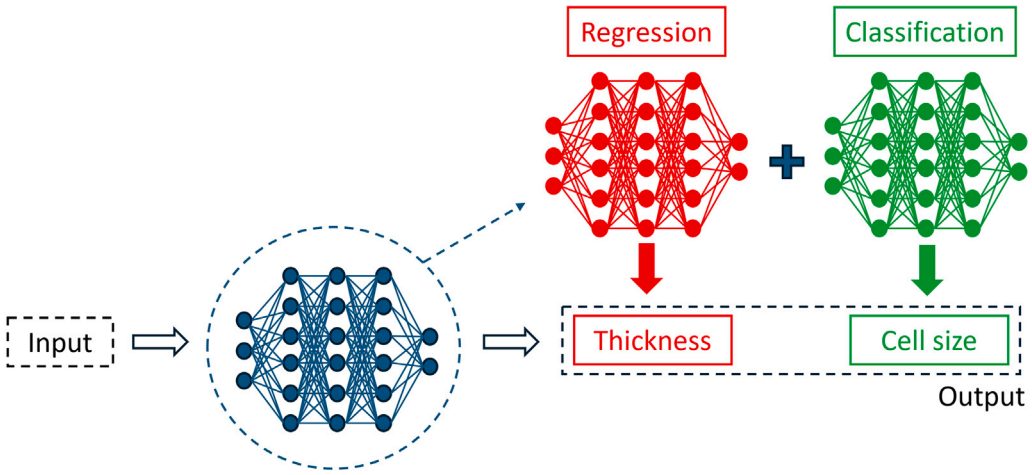

Fig. 3. Schematic representation of the Unit-cell Optimization Network (UON) used for inverse design. The mechanical descriptors extracted from the stress–strain curve, together with the one-hot encoding of the unit-cell topology, are used as common input features for two independent neural networks. The first network is trained as a regression model to predict the continuous beam thickness t, whereas the second network is trained as a multi-class classifier to predict the discrete unit-cell size s. The two networks share the same input vector and the same general multilayer perceptron architecture, but they are trained separately and do not share weights.

Table 2
Summary of the UON architecture and training settings. The UON is composed of two independent neural networks trained on the same normalized input descriptors: one regression network for beam-thickness prediction and one classification network for unit-cell-size prediction.

Item	Setting
Input descriptors	$[E_0, \epsilon_{dens}, En_{abs}, \sigma_{pl}, t_K, t_O]$
Topology encoding	One-hot encoding of Kelvin and octagonal unit cells
Regression model output	Beam thickness t
Classification model output	Unit-cell size class s
Network coupling	Two independent networks, no shared weights
Hidden layers	3 fully connected layers per network
Neurons per hidden layer	128
Activation function	ReLU
Weight initialization	He initialization
Batch normalization	After the first two ReLU blocks
Dropout	$p = 0.05$ before the output layer
Regression loss	Mean-squared error
Classification loss	Cross-entropy loss
Input normalization	Min–max scaling using training-set statistics
Thickness scaling	Min–max scaling using training-set statistics
Optimizer	ADAM
Initial learning rate	10^{-3}
L2 regularization	10^{-3}
Maximum epochs	600
Data split	70/20/10 training/validation/test split
Evaluation metrics	RMSE for thickness, accuracy/confusion matrix for cell size

capped to a small fraction of the total samples to prevent excessive loss of early data. After trimming, the stress curve is filtered again using an adaptive Savitzky–Golay window (or a moving-average fallback for short sequences) to obtain a smooth representation $\sigma_f(\varepsilon)$ employed in slope- and plateau-related computations, while the unfiltered curve $\sigma(\varepsilon)$ is retained for energy integration in order to avoid systematic bias in the absorbed-energy estimate.

The initial elastic stiffness E_0 is estimated from the early quasi-linear segment of the response, restricting the analysis to a small strain interval ($\varepsilon \leq 0.04$) and to a limited number of points (typically 8–20 samples) to focus on the elastic regime while maintaining numerical robustness. Within this window, a monotonicity-based selection retains only the first approximately increasing segment of $\sigma_f(\varepsilon)$ (small stress drops within a tolerance band are allowed), thus preventing local oscillations from corrupting the slope. Two independent stiffness estimators are then computed and combined in a conservative manner: (i) a derivative-based estimator E_A obtained as the median value of $d\sigma_f/d\varepsilon$ over the retained early window (computed through a Savitzky–Golay derivative when possible, otherwise through a finite-difference gradient), and (ii) a constrained least-squares estimator E_B obtained by fitting a straight line anchored at the first point of the selected segment, i.e. $\sigma_f(\varepsilon) \approx E_B(\varepsilon - \varepsilon_0) + \sigma_0$. The final stiffness is taken as $E_0 = \max(E_A, E_B)$ to avoid underestimation due to smoothing and small early nonlinearities, and is further capped using a robust upper bound derived from the distribution of local slopes in the initial segment (e.g., $E_0 \leq 1.3, P_{90}\Delta\sigma_f/\Delta\varepsilon$) to prevent unrealistically large values induced by residual noise or discretization artefacts. Representative plateau stress is computed on the filtered curve as the mean stress within a prescribed mid-strain interval, $\sigma_{pl} = \langle\sigma_f(\varepsilon)\rangle_{\varepsilon\in[0.20,0.50]}$, which captures the characteristic crushing level while reducing sensitivity to localized fluctuations; when the curve does not span the full interval, a fallback evaluation around $\varepsilon \approx 0.30$ is adopted to ensure consistent definition across samples.

Densification strain is determined using an energy-absorption-efficiency criterion, which provides a robust indicator of the transition

Fig. 4. As the absence probability p increases, the retained connectivity decreases. Below the critical absence threshold, a system-spanning load path can still persist, whereas above it, the network tends to fragment and the load-transfer capability is progressively lost.

from progressive crushing to rapid stiffening [48]. The absorbed energy per unit volume up to a strain level ε is computed as shown in Eq. (1).

$$En(\varepsilon) = \int_0^{\varepsilon} \sigma(\xi)\, d\xi \tag{1}$$

The corresponding energy absorption efficiency is defined as the ratio between the cumulative absorbed energy and the energy absorbed by an ideal energy absorber when both produce the same peak stress σ_{peak}(Eq. (2)). $\varepsilon_{ideal} = 1$ refers to 100% deformation of the ideal absorber.

$$\eta(\varepsilon) = \frac{En(\varepsilon)}{\sigma_{peak} * 1} = \frac{\int_0^{\varepsilon} \sigma(\xi)\, d\xi}{\sigma_{peak} * 1} \tag{2}$$

The densification strain is identified as the strain at which $\eta(\varepsilon)$ reaches its maximum over a physically meaningful post-yield interval, here restricted to $\varepsilon \geq 0.35$ to exclude the elastic–early plastic regime and to improve robustness in the presence of small initial oscillations (Eq. (3)):

$$\varepsilon_{dens} = \arg \max_{\varepsilon \geq 0.35} \eta(\varepsilon) \tag{3}$$

This definition is complemented in the code by a safeguard that defaults to the last available strain value if the detected maximum is ill-posed (e.g., too few points), thereby avoiding empty or non-physical outputs. Finally, the absorbed energy used as a scalar descriptor is computed as the area under the stress–strain curve up to densification, as shown in Eq. (4).

$$En_{abs} = En(\varepsilon_{dens}) = \int_0^{\varepsilon_{dens}} \sigma(\xi)\, d\xi \tag{4}$$

Numerically, this is evaluated via trapezoidal integration on the raw stress signal. In addition to the efficiency-based criterion, a secondary densification indicator based on the onset of rapid stiffening can be computed from the slope of the filtered curve at high strains (e.g., for $\varepsilon \geq 0.65$), by detecting the first index where $d\sigma_f/d\varepsilon$ exceeds a fraction of its high-strain maximum; this auxiliary criterion is used as a consistency check and as a fallback in edge cases.

The resulting dataset of scalar descriptors is therefore assembled into a matrix whose columns correspond to $(E_0, \sigma_{pl}, \varepsilon_{dens}, En_{abs})$ and whose rows correspond to the simulations retained after the automated filtering and cleaning stages. The matrix is then employed to train the Unit-cell Optimization Network (UON) inverse model, which is formulated as a mixed learning problem reflecting the metamaterial parameterization: beam thickness is continuous and is therefore predicted via regression, whereas unit-cell size is selected from a discrete set of admissible values and is thus predicted via multi-class classification, as schematized in Fig. 3.

Accordingly, the UON is implemented as a mixed inverse-design model composed of two independent neural networks trained on the same input features. The first model is a regression network used to predict the continuous strut thickness t, whereas the second model is a classification network used to predict the unit-cell size class s. This separation reflects the different mathematical nature of the two design variables: t is continuous, while s is constrained to a discrete set of admissible values by the lattice construction and size-compatibility requirements, as schematized in Fig. 3. The input vector is built from the mechanical descriptors extracted from each stress–strain curve and from a topology-conditioning term, i.e.

$$\mathbf{x} = [E_0, \varepsilon_{dens}, En_{abs}, \sigma_{pl}, t_K, t_O],$$

where t_K and t_O are a one-hot encoding of the unit-cell family. In the adopted convention, Kelvin and octagonal lattices are encoded by assigning the corresponding binary topology indicators. For the classification task, the cell size is mapped to the closest value in the admissible set, and samples with cell sizes outside the considered set are discarded to ensure consistent class labels. Before training, the input features are normalized using min–max scaling computed only on the training subset, and the thickness target is scaled with the same min–max strategy to improve numerical conditioning. The corresponding scaling parameters are stored and reused during inference. Both networks adopt the same multilayer perceptron architecture, but they are trained separately and do not share weights. Each model uses three fully connected hidden layers with 128 neurons, ReLU activation functions, He weight initialization, batch normalization after the first two ReLU blocks, and dropout before the output layer. The regression network ends with a single linear output neuron followed by a regression layer, providing the predicted thickness $\hat{t}$. The classification network ends with $|S|$ output neurons, where $|S|$ is the number of admissible cell-size classes, followed by a softmax layer and a classification layer, returning the probability distribution over the cell-size classes. The regression network is trained by minimizing the mean-squared error on the normalized thickness target, whereas the classification network is trained by minimizing the cross-entropy loss associated with the cell-size labels. Training is performed with the ADAM optimizer using an initial learning rate of 10^{-3}, L2 regularization equal to 10^{-3}, and a fixed training/validation/test split with controlled random seed for reproducible evaluation. Model performance is quantified on held-out test data through regression metrics for thickness, such as the de-normalized RMSE, and classification metrics for cell size, such as accuracy and confusion matrix. The main training settings are summarized in Table 2.

To extend the design space beyond purely continuous geometric variables, a percolation-inspired defect-generation module is implemented for periodic truss lattices by introducing controlled, stochastic connectivity reductions. Starting from a fully connected reference lattice, the beam network is represented through its edge list, and each

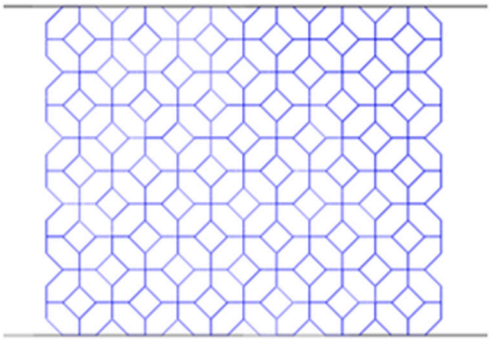

(a) Kelvin lattice structure.

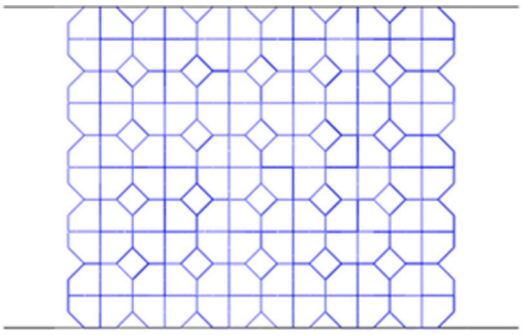

(b) Octagonal lattice structure.

Fig. 5. Periodic lattice topologies included in the UON dataset. Beam thickness and unit-cell size are treated as design variables, whereas the overall lattice dimensions are fixed. All models feature an internal lattice between two rigid plates and are subjected to compression at the same constant velocity.

beam is assigned an independent absence probability p. The defective topology is then generated through a set of independent Bernoulli trials, in which each edge is retained with probability $1 - p$ and removed with probability p, yielding a family of realizations whose connectivity is progressively reduced while preserving the underlying spatial embedding and periodic reference architecture (Fig. 4). This procedure directly mirrors classical bond-percolation on graphs, but it is applied here to a mechanical lattice, enabling systematic exploration of defective configurations without changing nodal positions or re-defining unit-cell topology. Controlled realizations are generated reproducibly by fixing the pseudo-random seed and storing it alongside the sample identifier, so that each defective lattice can be exactly reconstructed from $(p, seed)$. In practical use, additional validity checks can be introduced after beam removal (e.g., removal of isolated nodes, or verification of a load-path connecting top and bottom faces) to reject trivially disconnected configurations when required by the loading scenario; when such checks are adopted, the acceptance/rejection outcome is stored as metadata to ensure full traceability.

3. Results

The workflow described above was employed to generate and analyse beam-based lattice metamaterials with periodic topology, focusing on Kelvin (Fig. 5(a)) and octagonal (Fig. 5(b)) unit cells to produce large-scale dataset and train a surrogate model. The lattice specimens present overall dimensions of 30 × 30 × 25 mm (width, depth, height) subjected to constant-velocity dynamic compression. The relative density of 13.7% refers to the representative benchmark configuration used for model verification and topology comparison. In the dataset used for UON training, the strut thickness is varied within cell-size-dependent admissible bounds, and therefore the relative density varies consistently with the selected thickness and unit-cell size.

Before the generation of the full dataset, the automated FE modelling chain was verified against a reference Kelvin lattice case available in the literature. In particular, the geometry, global dimensions, unit-cell size, material description, boundary conditions, and loading conditions were selected to reproduce the configuration reported by Habib et al. [46]. The numerical settings follow the beam-based formulation summarized in Table 1, with B31 beam elements for the lattice struts, rigid compression plates, a fixed lower reference point, and a prescribed velocity applied to the upper reference point along the compression direction. Fig. 6 compares the macroscopic stress–strain response obtained with the present workflow with the reference result. The comparison shows that the proposed automated model generation and post-processing procedure is able to reproduce the main features of the reference response, including the initial stiffness and the subsequent crushing regime. Small deviations are mainly associated with differences in numerical implementation, stress-transmission delay, and curve-processing conventions. This validation case is reported here to provide a self-contained verification of the FE workflow before its use for high-throughput dataset generation; a more extended comparison of the automatic lattice-generation framework is provided in Ref. [45].

Following verification and topology benchmarking, large-scale dataset generation was carried out for periodic lattices under high strain-rate loading, and the resulting database was used to train the UON inverse model. In particular, the dataset was populated through FE simulations representative of impact-rate conditions (35 m/s), and surrogate training was limited to fully populated periodic Kelvin and octagonal lattices (i.e., percolated lattices were analysed separately and were not included in the UON training set). All dataset specimens retained the same external dimensions adopted in the verification stage, to ensure consistency of the extracted macroscopic descriptors. The unit-cell size was selected from the discrete admissible set s [2.5, 2.8, 3.0, 3.5, 3.8, 4.5, 4.8] mm, where the available values were chosen to (i) guarantee geometric compatibility with the target specimen size (integer replication and controlled scaling), and (ii) remain within additive-manufacturing feasibility constraints, ensuring a sufficient number of unit cells within the specimen and avoiding feature sizes below typical printable limits. For each cell size, strut thickness values were sampled within size-dependent bounds computed automatically as described in the previous section: the scaled unit-cell edge-length distribution was first evaluated and then converted into an admissible thickness interval through the dedicated thickness-range generator, enforcing physically reasonable thickness-to-length ratios across all cell-size classes.

A typical macroscopic response obtained from the high strain-rate compression simulations is shown in Fig. 7. The curve exhibits the characteristic sequence observed in dynamically crushed cellular materials, namely an initial elastic stage followed by a peak, a plateau-like regime associated with progressive collapse, and a final densification stage marked by a steep stress increase. Starting from the raw stress–strain history (blue curve), the post-processing routine applies the smoothing and trimming procedures described in the previous section (red curve) and automatically extracts the key performance descriptors used throughout this work. In particular, the initial elastic stiffness E_0 is evaluated from the slope of the early quasi-linear segment, the representative plateau stress σ_{pl} is computed over a prescribed mid-strain interval, the densification strain ε_{dens} is identified through the energy-absorption-efficiency criterion, and the absorbed energy En_{abs} is obtained by integrating the stress–strain curve up to ε_{dens} (shaded area). This automated feature extraction is repeated for every simulated configuration, thereby populating the dataset consistently for each combination of unit-cell type (Kelvin or octagonal), cell size, and strut thickness under the selected loading condition. All these parameters are collected through the automated database assembly driven by the

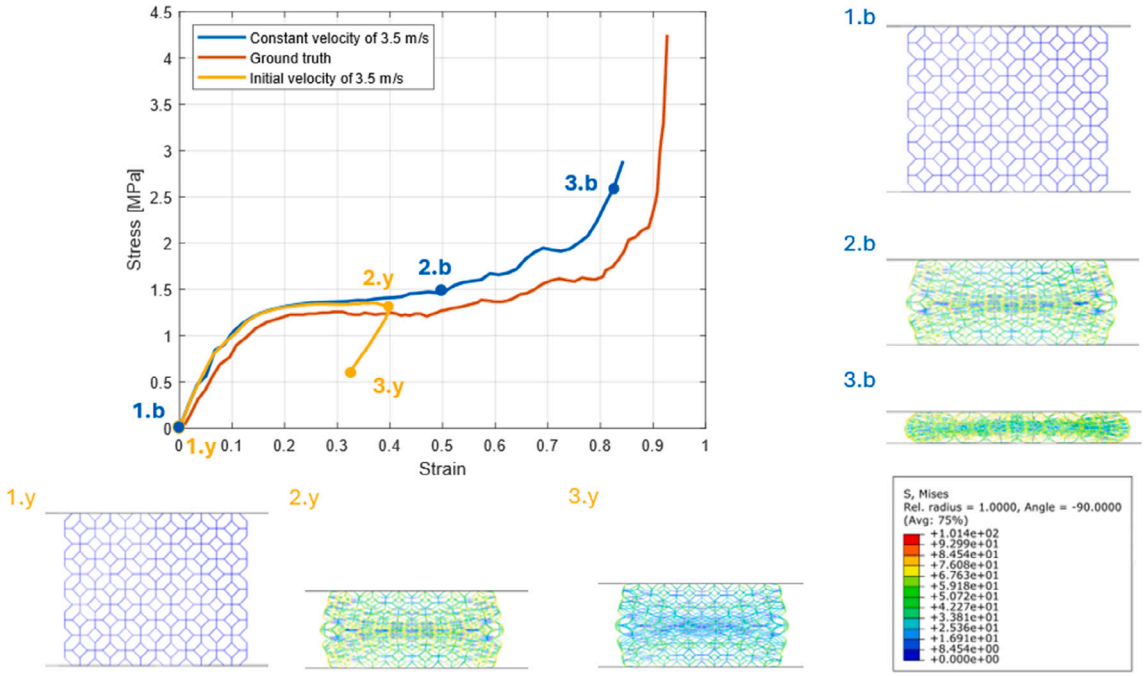

Fig. 6. Verification of the automated FE modelling chain against the reference Kelvin lattice compression case reported by Habib et al. [46]. The comparison is used as a representative validation case before large-scale dataset generation [45].

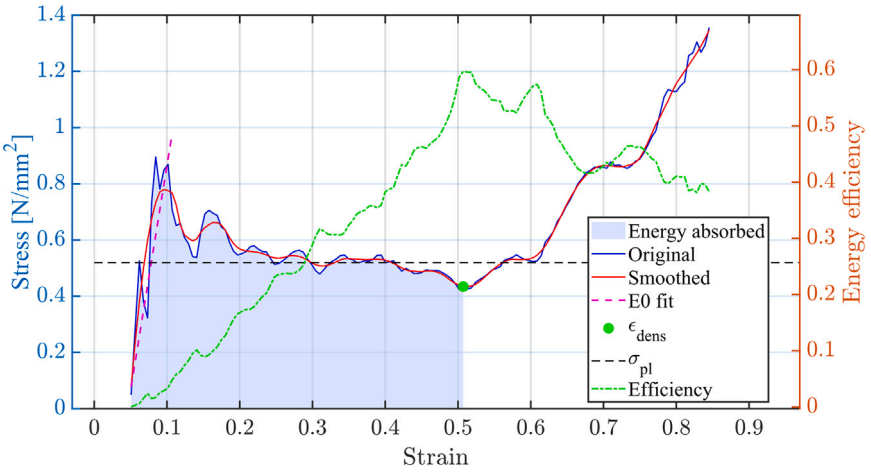

Fig. 7. Typical stress–strain curve of a lattice structure under high strain-rate compression (octagonal topology, 35 m/s), extracted from the dataset. The raw response (blue) and the filtered curve (red) are shown together with the procedures used to extract the performance descriptors. The red dashed line indicates the initial stiffness E_0, the black dashed line represents the plateau stress σ_{pl}, and the green dashed curve denotes the energy-absorption efficiency used to determine the densification strain ϵ_{dens} (green marker). The absorbed energy is computed as the area under the raw stress–strain curve up to ϵ_{dens} (shaded region).

Table 3
Composition of the cleaned dataset used for UON training and evaluation.

Cell size [mm]	Kelvin samples	Octagonal samples	Total samples
2.5	26	94	120
2.8	90	105	195
3.0	105	114	219
3.5	92	109	201
3.8	111	113	224
4.5	117	83	200
4.8	124	93	217
Total	665	711	1376

batch generation/export routine described in the Materials and Methods section: each parameter sample produces one FE input file and an associated metadata record, enabling systematic traceability between parameter space, solver outputs, and post-processed descriptors.

The distribution of generated designs and responses is summarized through dedicated dataset-visualization plots. Specifically, Fig. 8 summarizes the distribution of the Kelvin-lattice dataset in terms of the two design inputs (unit-cell size and beam thickness) and the corresponding macroscopic descriptors extracted from the 35 m/s simulations. Each marker represents one simulated configuration; beam thickness and unit-cell size define the horizontal axes, while the vertical axis reports the extracted descriptor, with the colour map used as an additional visual cue for the same quantity. The plots highlight clear and smooth trends across the sampled design space: the initial stiffness and the plateau stress increase systematically with increasing thickness, with

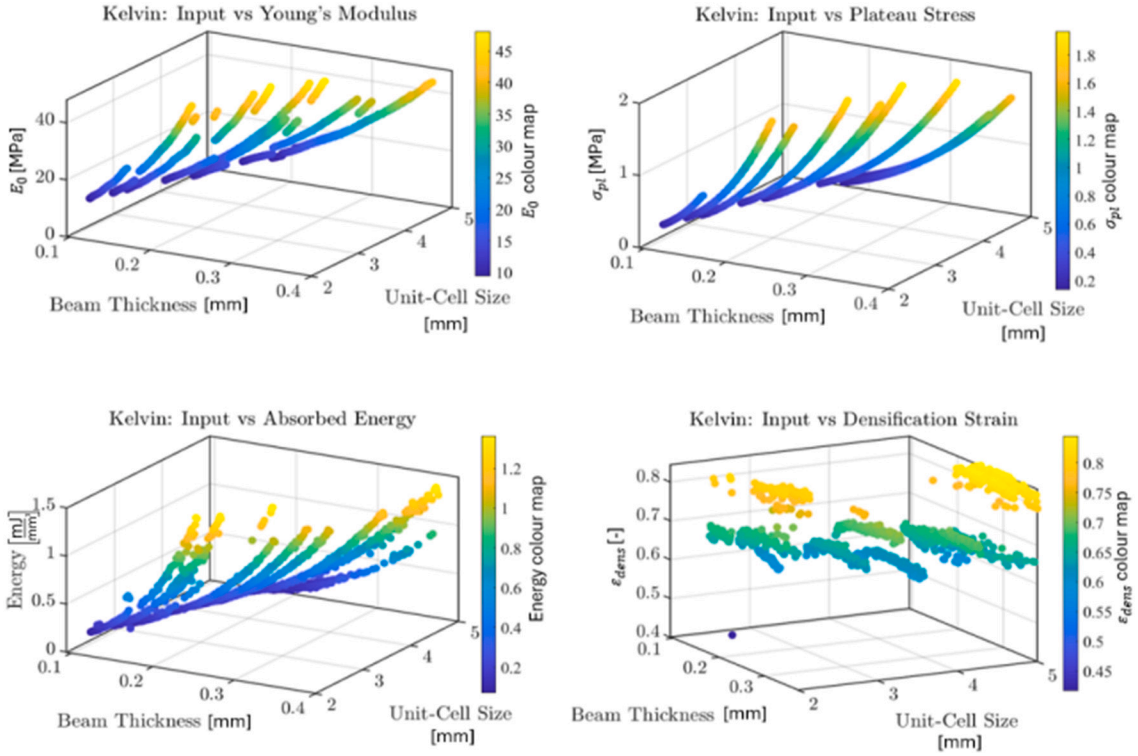

Fig. 8. Dataset visualization. Each performance parameter is plotted in a diagram which presents the value of the parameter as the vertical axis, the beam thickness and unit cell size as the horizontal axes. Each marker represents one simulated configuration. The reported descriptors are extracted from compression simulations performed at 35 m/s. Each marker represents a simulated model.

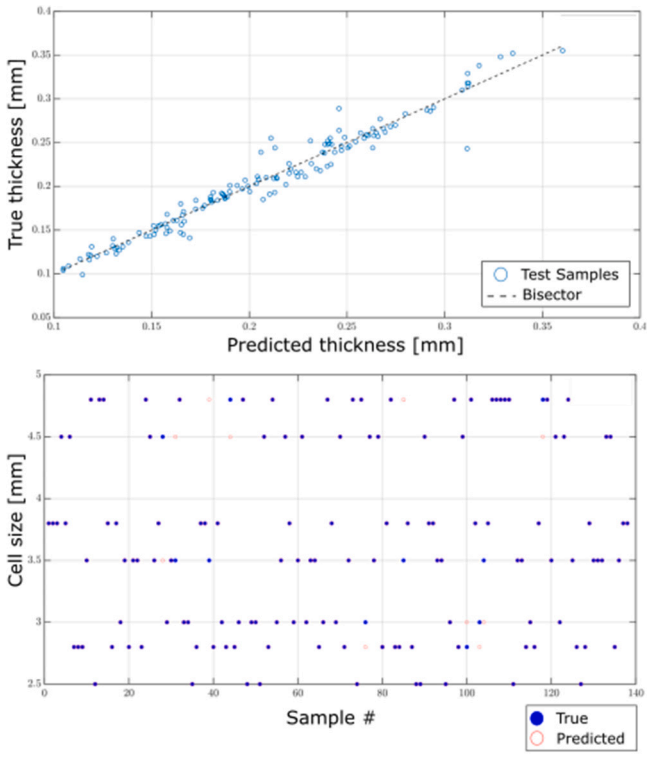

Fig. 9. Thickness and unit-cell size parity plots. For thickness, the y-axis reports the ground-truth values from the dataset, while the x-axis reports the corresponding model predictions. For unit-cell size, each sample is shown by comparing the ground-truth class (blue markers) and the predicted class (red circles). These plots are valid for both kelvin and octagonal unit cell types, subjected to 35 m/s compression.

a consistent shift across cell-size levels, whereas the absorbed energy exhibits a similar monotonic dependence, reflecting the combined effect of stiffness and plateau level on the area under the stress–strain curve. In contrast, the densification strain shows a more bounded response, with values clustered in bands that depend primarily on unit-cell size and only secondarily on thickness, indicating that the onset of densification is less sensitive to thickness variations than stiffness- and stress-related descriptors within the explored domain. To ensure that the learning stage is informed only by physically meaningful and numerically reliable samples, the raw simulation set was processed through an automated cleaning pipeline. This pipeline filters out incomplete or solver-failed analyses, removes samples with missing values or corrupted curves, and enforces admissible-domain constraints derived from the intended design space. After cleaning, 1376 simulations were retained as valid samples for surrogate training and evaluation. The composition of the cleaned dataset is reported in Table 3, where the retained samples are grouped by unit-cell topology and cell-size class.

Concerning the UON, both the regression and classification models use the same multilayer perceptron architecture, consisting of three fully connected hidden layers with 128 neurons, ReLU activation functions, batch normalization, and a dropout layer with $p = 0.05$ before the output layer. The two networks are trained independently on the same normalized input descriptors. Training is performed with the ADAM optimizer using an initial learning rate of 10^{-3}, L2 regularization of 10^{-3}, and a fixed 70/20/10 split for training, validation, and testing, with a controlled random seed for reproducible evaluation. The maximum number of epochs is 600 for both networks, and validation data are used to monitor generalization. Figs. 9, 10 and 11 summarize the predictive performance of the UON, reported through both regression and classification metrics to reflect the mixed nature of the inverse design task.

The reported metrics and plots refer to the complete held-out test set, including both Kelvin and octagonal unit-cell designs. The unit-cell topology is provided to the UON as a one-hot conditioning input;

Table 4
Sensitivity of UON predictions to Gaussian perturbations applied to the input mechanical descriptors on the test set. Values are reported as mean ± standard deviation over 100 noise realizations.

Noise level	Thickness RMSE [mm]	Cell-size accuracy [%]
0%	0.01257 ± 0.00000	92.754 ± 0.000
1%	0.01731 ± 0.00119	82.166 ± 2.516
2.5%	0.03109 ± 0.00399	62.549 ± 3.603

therefore, the results should be interpreted as the performance of a topology-conditioned inverse model rather than as the performance on one specific lattice family. Since the unit-cell size is selected from a discrete admissible set, cell-size identification is treated as a multi-class classification problem, whereas strut thickness is treated as a continuous regression target. Thickness regression is reported through the parity plot in Fig. 9, where predicted thickness is compared against the true value on the test set; the proximity of the samples to the bisector provides a direct visual measure of agreement across the explored thickness range. Quantitatively, the regression network achieves a test RMSE equal to 0.012571 mm. Cell-size prediction is reported in complementary forms: Fig. 9 shows the true (blue) and predicted (red) cell-size labels for each test sample, providing a sample-by-sample visualization of classification consistency, while Fig. 10 reports the corresponding confusion matrix, where diagonal entries represent correct classifications and off-diagonal terms quantify misclassifications between neighbouring size classes. Separate inspection of the two topology subsets showed the same qualitative behaviour: thickness predictions remained consistent across the explored range, while cell-size misclassifications were mainly concentrated between neighbouring size classes.. In addition, Fig. 11 reports the classification accuracy broken down by cell-size class, highlighting the class-dependent performance of the classifier and supporting the interpretation of the confusion matrix. Overall, the classification network achieves an accuracy of 92.7536% on the test set.

A descriptor-noise sensitivity analysis was also performed to evaluate the robustness of the inverse model to small uncertainties in the extracted mechanical descriptors. The four mechanical inputs were perturbed on the test set as

$$\tilde{x}_i = x_i(1 + \delta_i),$$

where δ_i is a zero-mean Gaussian perturbation with standard deviation equal to 1% and 2.5% of the corresponding descriptor value. For each noise level, the perturbation was repeated over 100 realizations and the mean and standard deviation of the thickness RMSE and cell-size classification accuracy were computed. The results, summarized in Table 4, show the sensitivity of the UON predictions to descriptor-level uncertainty.

In the high strain-rate compression simulations (35 m/s), the macroscopic stress–strain response of fully populated periodic lattices typically exhibits an initial peak that can be noticeably higher than the subsequent plateau level. Since the initial peak is directly associated with the maximum transmitted stress during the early loading stage, controlling its magnitude is relevant for impact-mitigation applications, where excessive early stress transmission may be undesirable even when the plateau behaviour and the overall absorbed energy remain acceptable. For this reason, a percolation-inspired strategy is investigated as an additional design lever, aimed at selectively reducing lattice connectivity to tailor the early-stage response and, in particular, mitigate the peak-to-plateau contrast.

The percolation campaign is performed by introducing an absence probability p for beam elements in the Kelvin lattice and generating multiple stochastic realizations for each value of p. For the 1D beam formulation, the effect of increasing absence probability on the macroscopic response is illustrated by the family of stress–strain curves in Fig. 12. Each curve corresponds to a distinct defective realization, with the

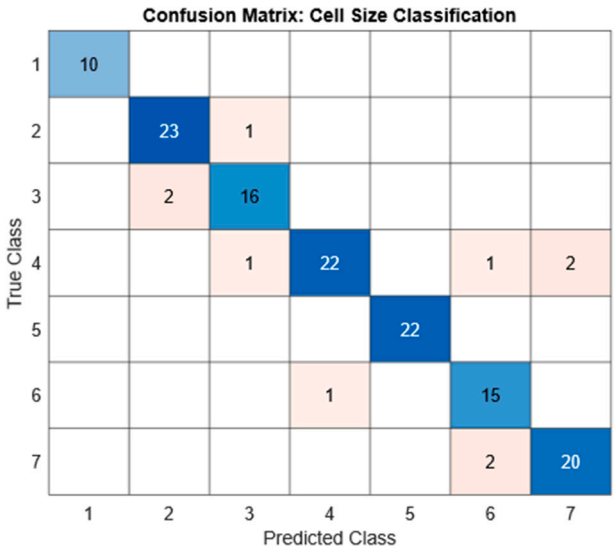

Fig. 10. Cell size confusion matrix.This plot is valid for both kelvin and octagonal unit cell types, subjected to 35 m/s compression.

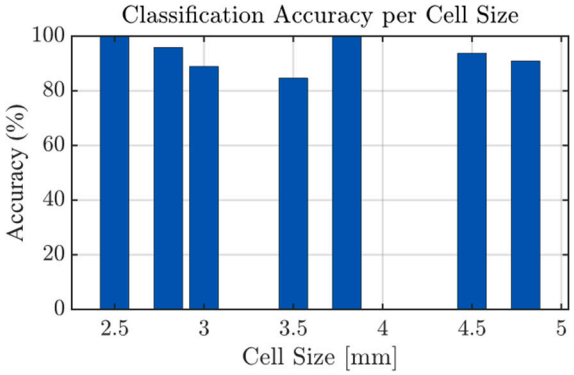

Fig. 11. Predictions accuracy percentage for each cell size class. This plot is valid for both kelvin and octagonal unit cell types, subjected to 35 m/s compression.

colourmap indicating the associated value of p. The figure highlights a progressive reduction of the initial peak stress as p increases, while the subsequent crushing regime evolves more smoothly. The same post-processing pipeline used for the dataset generation is applied to each curve, ensuring that peak stress, plateau stress, densification strain and absorbed energy are extracted consistently across all realizations. The evolution of key performance quantities, that is, plateau stress and absorbed energy, with increasing absence probability is summarized in Figs. 13 and 14, respectively.

In particular, the plateau stress decreases approximately linearly with p over the explored interval, and the linear fit reported in the figure (95% confidence band shown) captures the trend with high coefficient of determination ($R^2 \approx 0.99$) Similarly, the absorbed energy density (computed as the area under the stress–strain curve up to densification, and therefore expressed in units of stress) decreases approximately linearly with p, again with a tight confidence band and $R^2 \approx 0.99$. These two plots provide a compact representation of how connectivity reduction influences the global crushing level and the cumulative energy dissipation, while keeping the analysis at the level of macroscopic descriptors rather than individual curve features.

In contrast to plateau stress and absorbed energy, the initial peak stress shows a markedly stronger sensitivity to p. Fig. 15 reports the dependence of peak stress on absence probability for the 1D simulations,

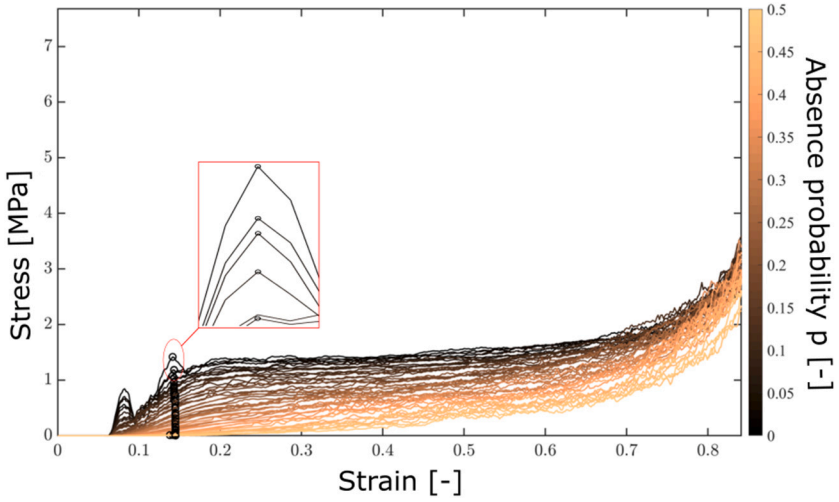

Fig. 12. Stress–strain curves related to the 1D Kelvin model subjected to 35 m/s compression, applying percolation. In the picture, the peak stress values are highlighted.

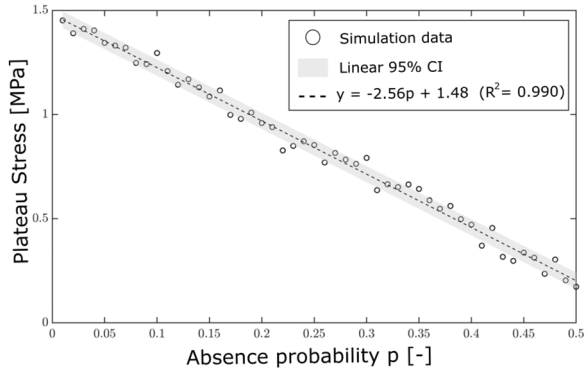

Fig. 13. Variation of plateau stress with beam absence probability p for the 1D Kelvin lattice subjected to compression at 35 m/s.

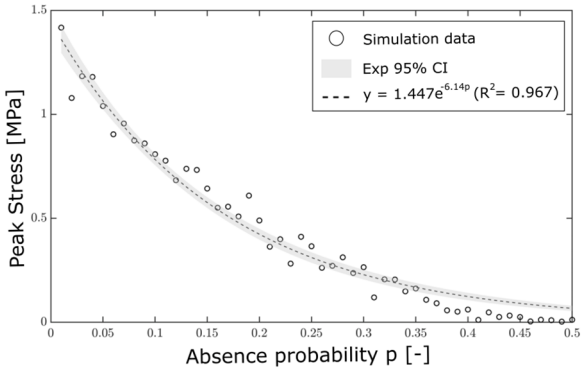

Fig. 15. Variation of peak stress with beam absence probability p for the 1D Kelvin lattice subjected to compression at 35 m/s.

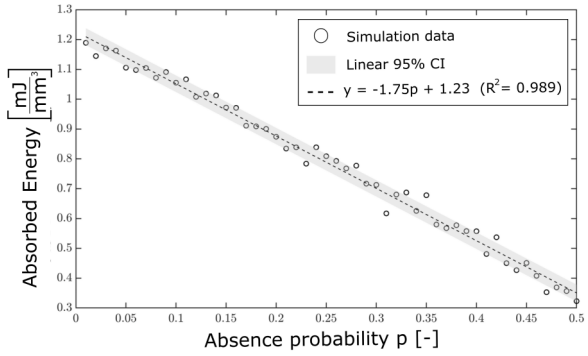

Fig. 14. Variation of absorbed energy with beam absence probability p for the 1D Kelvin lattice subjected to compression at 35 m/s.

together with an exponential-type empirical fit and the corresponding confidence band. The fit shown in the plot shows that peak stress is well represented, over the investigated range of p, by an exponential-type reduction, consistent with the qualitative trend observed in Fig. 12. This representation is particularly useful because it summarizes, with a single curve, the systematic reduction of the early stress peak obtained by progressively removing load-carrying members.

To verify that the same qualitative behaviour persists under a higher-fidelity discretization, selected simulations are repeated using a 3D solid formulation. Fig. 16 reports representative stress–strain curves obtained for increasing absence probability, showing the same overall tendency: larger p values correspond to lower initial peak stress and a globally reduced stress level during crushing, while densification is still characterized by a final steep stress increase. The corresponding peak-stress trend extracted from the 3D simulations is summarized in Fig. 17, where peak stress is plotted against p (reported as 1%–50% in the figure) and represented through the same exponential-type empirical fit adopted for the 1D trend. The fitted curve provides a compact description of the peak-stress reduction in the 3D case and supports direct comparison with the 1D percolation trends using the same macroscopic metric.

It is important to clarify that the beam-removal procedure adopted in the percolation study is not intended to reproduce the most frequent defect morphology observed in additively manufactured lattices. In high-quality printed lattices, imperfections more commonly appear as surface roughness, local porosity, diameter deviations, or partial weakening of struts rather than as complete member absence. In the present work, binary strut removal is therefore used as an idealized representation of severe connectivity loss and, more importantly, as a deliberate topology-tailoring variable. From this perspective, the absence probability p provides a controlled way to investigate how interruption of load paths affects the initial peak stress and the subsequent crushing response.

The different trends observed for plateau stress, absorbed energy, and peak stress can be interpreted from the distinct mechanical role

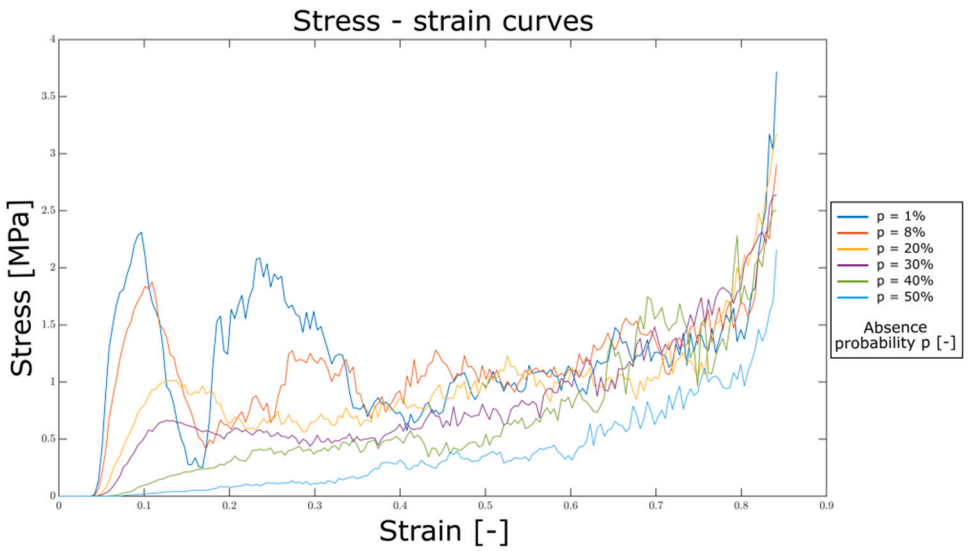

Fig. 16. Stress–strain curves obtained from the 3D Kelvin model subjected to compression at 35 m/s for different beam absence probabilities p.

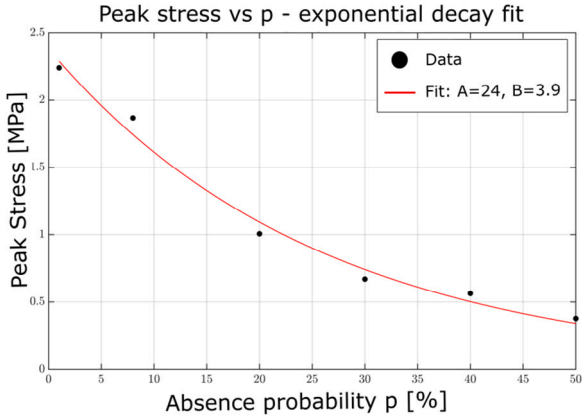

Fig. 17. Variation of peak stress with beam absence probability p for the 3D Kelvin lattice subjected to compression at 35 m/s. The solid line represents the fitted exponential-type trend.

of these quantities during compression. Plateau stress and absorbed energy are global descriptors of the progressive crushing stage. They are mainly governed by the average amount of load-bearing material and by the number of active load paths that remain available after beam removal. Since each beam is independently retained with probability $1 - p$, the expected number of retained beams is proportional to $1 - p$. Therefore, quantities associated with distributed crushing, such as σ_{pl} and En_{abs}, are expected to decrease approximately linearly over the investigated range, provided that the global deformation mode does not change abruptly.

The initial peak stress is controlled by a different mechanism. It occurs during the early loading stage, before the development of a stable crushing regime, and is strongly affected by the initial stiffness of the lattice and by the presence of continuous, highly connected load paths between the loading plates. Removing even a limited number of beams can interrupt some of these direct load paths, promote a more progressive engagement of the structure, and reduce the intensity of the stress transmitted in the initial phase. From a graph-based viewpoint, the probability that a load path composed of n beams remains fully intact scales as

$$P_{\text{path}} = (1 - p)^n.$$

For moderate values of p, this relation provides a physical rationale for using an exponential-type empirical dependence over the investigated probability range. The exponential fit should therefore be interpreted as a compact description of the explored range, rather than as a universal law independent of topology, geometry, or loading conditions. This explains why peak stress exhibits a stronger nonlinear reduction than plateau stress and absorbed energy.

Selected 3D solid simulations were then used as higher-fidelity numerical checks of this trend. These simulations were not used to populate the UON dataset and were not intended as a complete quantitative cross-validation of the beam formulation over the full percolation space. Their purpose was to verify whether the peak-stress reduction persists when local stress states are resolved more explicitly. As shown in Figs. 16 and 17, the 3D simulations confirm the same qualitative reduction of the initial peak stress observed with the 1D beam formulation. Therefore, they support the numerical robustness of the percolation trend, while the quantitative values remain dependent on the adopted element formulation, material model, and damage assumptions.

4. Conclusions

This work presented an automated workflow for the generation, simulation, post-processing, and inverse design of periodic truss lattice metamaterials under dynamic compression. The framework enables the parametric construction of Kelvin and octagonal lattices, the export of simulation-ready finite element models, and the extraction of compact macroscopic descriptors from the resulting stress–strain curves. The dataset-generation strategy includes a cell-size-dependent thickness-bounding procedure, which preserves physically meaningful thickness-to-length ratios across the discrete cell-size classes and supports consistent high-throughput database construction. The cleaned dataset, composed of 1376 valid simulations, was used to train a Unit-cell Optimization Network (UON) for inverse design. Since the design variables have different mathematical characteristics, the inverse problem was formulated as a mixed learning task: regression for the continuous strut thickness and classification for the discrete unit-cell size. The UON achieved a test RMSE of 0.012571 mm for thickness prediction and a cell-size classification accuracy of 92.7536%, showing that the selected descriptors $(E_0, \sigma_{pl}, \varepsilon_{dens}, En_{abs})$ are informative enough to recover feasible unit-cell parameters within the investigated design space. However, the inverse mapping is not necessarily unique. Different combinations of beam thickness and unit-cell size can produce

similar scalar descriptors, especially between neighbouring cell-size classes. Therefore, the UON prediction should be interpreted as one feasible solution within the trained design domain, rather than as proof of a unique physical inverse solution. A percolation-inspired strategy was also investigated as a complementary design lever for peak-stress tailoring. This analysis was performed separately from the UON training dataset. By assigning an independent absence probability p to lattice beams, controlled reductions in connectivity were introduced into the reference architecture. The results showed that plateau stress and absorbed energy decrease approximately linearly with increasing p, whereas the initial peak stress exhibits a stronger exponential-type reduction over the investigated probability range. Selected 3D simulations confirmed the same qualitative peak-reduction trend observed with the 1D beam model, supporting the numerical robustness of the connectivity-based tailoring strategy. These simulations should be interpreted as preliminary higher-fidelity checks of the observed trend, rather than as a complete quantitative cross-validation of the beam formulation over the full percolation design space. The present study remains primarily numerical and is restricted to the considered topologies, specimen dimensions, material model, and loading conditions. Although the modelling chain was verified against a reference Kelvin lattice case and selected 3D checks were performed, beam-based simulations cannot fully resolve joint-level stress concentrations, cross-section distortion, manufacturing-induced defects, or progressive strut fracture. These effects may influence the quantitative agreement with experiments, especially when local damage develops before or during the plateau regime. However, the use of PA12, a nylon-based polymer with comparatively ductile and deformable behaviour, supports the assumption that the global response is mainly governed by lattice bending, buckling, and progressive collapse rather than by brittle strut fracture. This makes the beam formulation suitable for high-throughput prediction of macroscopic stress–strain descriptors, while not replacing local damage modelling. Higher-fidelity aspects were investigated in a previous contribution by the authors through hybrid 1D–3D simulations [47], and selected 3D models were also used here to check the percolation-induced peak-stress trend. Future work should extend the dataset to additional lattice families, broader density ranges, and different loading conditions, include percolation variables directly in the inverse-design loop, and introduce strategies to handle non-unique inverse solutions, such as additional response descriptors, full-curve information, probabilistic outputs, or forward-model verification of multiple candidate designs. A more systematic multi-fidelity validation campaign, including beam, hybrid 1D–3D, and fully solid models across different absence probabilities and random realizations, should also be performed to quantify the agreement between modelling fidelities. Targeted experimental validation on manufactured lattices under dynamic compression will finally be required to assess the influence of local failure, manufacturing tolerances, and material rate effects.

CRediT authorship contribution statement

Salvatore Annunziata: Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Luca Lomazzi:** Writing – review & editing, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Marco Giglio:** Supervision. **Andrea Manes:** Supervision.

Data and code availability

The codes and datasets generated and used in this study, including the parameter-sampling routines, post-processing scripts, UON training configuration, and representative metadata fields, are available from the corresponding author upon reasonable request. Full FE simulation outputs are not included as supplementary material due to file-size and software-specific limitations, but can be shared upon explicit request where possible.

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.

Data availability

Data will be made available on request.

References

[1] F. Xu, X. Zhang, H. Zhang, A review on functionally graded structures and materials for energy absorption, Eng. Struct. 171 (2018) 309–325.
[2] Z. Ozdemir, E. Hernandez-Nava, A. Tyas, J.A. Warren, S.D. Fay, R. Goodall, I. Todd, H. Askes, Energy absorption in lattice structures in dynamics: Experiments, Int. J. Impact Eng. 89 (2016) 49–61.
[3] Z. Ozdemir, A. Tyas, R. Goodall, H. Askes, Energy absorption in lattice structures in dynamics: Nonlinear FE simulations, Int. J. Impact Eng. 102 (2017) 1–15.
[4] N. Jin, F. Wang, Y. Wang, B. Zhang, H. Cheng, H. Zhang, Failure and energy absorption characteristics of four lattice structures under dynamic loading, Mater. Des. 169 (2019) 107655.
[5] L.J. Gibson, Biomechanics of cellular solids, J. Biomech. 38 (2005) 377–399.
[6] E.W. Andrews, L.J. Gibson, M.F. Ashby, The creep of cellular solids, Acta Mater. 47 (10) (1999) 2853–2863.
[7] S. Yuan, C.K. Chua, K. Zhou, 3D-printed mechanical metamaterials with high energy absorption, Adv. Mater. Technol. 4 (3) (2019) 1800419.
[8] R. Xue, X. Cui, P. Zhang, K. Liu, Y. Li, W. Wu, H. Liao, Mechanical design and energy absorption performances of novel dual scale hybrid plate-lattice mechanical metamaterials, Extrem. Mech. Lett. 40 (2020) 100918.
[9] P. Jiao, J. Mueller, J.R. Raney, et al., Mechanical metamaterials and beyond, Nat. Commun. 14 (2023) 6004.
[10] V.S. Deshpande, M.F. Ashby, N.A. Fleck, Foam topology: bending versus stretching dominated architectures, Acta Mater. 49 (2001) 1035–1040.
[11] V.S. Deshpande, N.A. Fleck, M.F. Ashby, Effective properties of the octet-truss lattice material, J. Mech. Phys. Solids 49 (2001) 1747–1769.
[12] R. Daya Karthic, J. Joy, G. Sakthivel, R. Nadimpalli, Mechanical characterization of 3D-printed kelvin cell with varying infill densities, Mater. Today: Proc. 84 (2023) 41–46.
[13] J. Park, Y. Roh, S. Lee, Effects of the unit-cell size and arrangement on the compressive behaviors of lattice structures in powder bed fusion additive manufacturing, Results Mater. 22 (2024) 100587.
[14] Q. Hu, G. Lu, K.M. Tse, Tensile mechanical behaviors of re-entrant and kelvin cell lattice structures, JOM 76 (2024) 387–396.
[15] M. Sadeghzade, et al., Experimental study of energy absorption capability in the lattice structures based on the octagonal bipyramid unit cell, J. Braz. Soc. Mech. Sci. Eng. 45 (9) (2023) 460.
[16] M. Smith, Z. Guan, W.J. Cantwell, Finite element modelling of the compressive response of lattice structures manufactured using the selective laser melting technique, Int. J. Mech. Sci. 67 (2013) 28–41.
[17] S. Ruiz de Galarreta, J.R.T. Jeffers, S. Ghouse, A validated finite element analysis procedure for porous structures, Mater. Des. 189 (2020) 108546.
[18] M. Ostoja-Starzewski, Lattice models in micromechanics, Appl. Mech. Rev. 55 (1) (2002) 35–59.
[19] S. Abu-Mualla, H. Huang, A dataset generation framework for symmetry-induced mechanical metamaterials, J. Mech. Des. 147 (4) (2025) 041705.
[20] S. Mo, P. Perdikaris, J.R. Raney, Accelerated design of architected materials with multifidelity Bayesian optimization, J. Eng. Mech. 149 (6) (2023) 04023032.
[21] C.-S. Liu, Smoothing elastoplastic stress–strain curves obtained by a critical modification of conventional models, Int. J. Solids Struct. 40 (2003) 2121–2145.
[22] Q.M. Li, I. Magkiriadis, J.J. Harrigan, Compressive strain at the onset of densification of cellular solids, J. Cell. Plast. 42 (2006) 371–392.
[23] G.W. Ma, Z.Q. Ye, Z.S. Shao, Modeling loading rate effect on crushing stress of metallic cellular materials, Int. J. Impact Eng. 36 (2009) 775–782.
[24] M. Osanov, J.K. Guest, Topology optimization for architected materials design, Annu. Rev. Mater. Res. 46 (1) (2016) 211–233.
[25] K.M. Park, K.S. Min, Y.S. Roh, Design optimization of lattice structures under compression: Study of unit cell types and cell arrangements, Materials 15 (1) (2022) 97.
[26] G. Dong, Y. Tang, D. Li, Y.F. Zhao, Design and optimization of solid lattice hybrid structures fabricated by additive manufacturing, Addit. Manuf. 33 (2020) 101116.
[27] H.M. Sheikh, T. Meier, B. Blankenship, Z. Vangelatos, N. Zhao, P.S. Marcus, C.P. Grigoropoulos, Systematic design of Cauchy symmetric structures through Bayesian optimization, Int. J. Mech. Sci. 236 (2022) 107741.

[28] Y. Wang, Q. Zeng, J. Wang, Y. Li, D. Fang, Inverse design of shell-based mechanical metamaterial with customized loading curves based on machine learning and genetic algorithm, Comput. Methods Appl. Mech. Engrg. 401 (2022) 115571.
[29] M. Maurizi, C. Gao, F. Berto, Inverse design of truss lattice materials with superior buckling resistance, Npj Comput. Mater. 8 (2022) 247.
[30] J. Song, J.H. Lee, N. Kim, K. Min, Artificial intelligence in the design of innovative metamaterials: A comprehensive review, Int. J. Precis. Eng. Manuf. 25 (1) (2024) 225–244.
[31] M.C. Messner, Convolutional neural network surrogate models for the mechanical properties of periodic structures, J. Mech. Des. 142 (2) (2020) 024503.
[32] A.E. Gongora, C. Friedman, D.K. Newton, et al., Accelerating the design of lattice structures using machine learning, Sci. Rep. 14 (2024) 13703.
[33] P. Suttakul, D. Vo, T. Fongsamootr, et al., The role of machine learning for insight into the material behavior of lattices: a surrogate model based on data from finite element simulation, Results Eng. 23 (2024) 102547.
[34] L. Xiao, G. Shi, W. Song, Machine learning predictions on the compressive stress–strain response of lattice-based metamaterials, Int. J. Solids Struct. 300 (2024) 112893.
[35] F. Xia, G. Shi, Z. Gao, J. Li, S. Xu, D. Ruan, Predicting deformation and stress-strain behaviour of lattice truss structures under compression using dual graph neural network, Compos. Struct. 372 (2025) 119557.
[36] K. Chen, X.-J. Dong, P. Gao, C. Qian, Physics-informed neural networks for topological metamaterial design and mechanical applications, Int. J. Mech. Sci. 301 (2025) 110489.
[37] L. Zheng, D.M. Kochmann, S. Kumar, HyperCAN: Hypernetwork-driven deep parameterized constitutive models for metamaterials, Extrem. Mech. Lett. 72 (2024) 102243.
[38] A. Aharony, Percolation, fractals, and anomalous diffusion, J. Stat. Phys. 34 (5–6) (1984) 931–939.
[39] A.S. Balankin, M.A. Martinez-Cruz, M.D. Alvarez-Jasso, et al., Effects of ramification and connectivity degree on site percolation threshold on regular lattices and fractal networks, Phys. Lett. A 383 (2019) 957–966.
[40] D.T. Fullwood, J.A. Basinger, B.L. Adams, Lattice-based structures for studying percolation in two-dimensional grain networks, Acta Mater. 54 (2006) 1033–1043.
[41] D.T. Casey, S. Bou Jawde, J. Herrmann, et al., Percolation of collagen stress in a random network model of the alveolar wall, Sci. Rep. 11 (2021) 16654.
[42] S. Kechagias, R.N. Oosterbeek, M.J. Munford, S. Ghouse, J.R.T. Jeffers, Controlling the mechanical behaviour of stochastic lattice structures: The key role of nodal connectivity, Addit. Manuf. 54 (2022) 102730.
[43] M.-S. Pham, C. Liu, I. Todd, J. Lertthanasarn, Damage-tolerant architected materials inspired by crystal microstructure, Nature 565 (7739) (2019) 305–311.
[44] B. McDonnell, E.M. O'Hara, N.M. Harrison, Simulation-driven-design of metal lattice structures for a target stress–strain curve, Mater. Des. 237 (2024) 112543.
[45] S. Annunziata, E. Shehu, L. Lomazzi, A. Manes, Validation of an automatic lattice generation framework and comparative analysis of lattice-based metamaterials for impact protection, Adv. Sci. Technol. 171 (2025) 167–175.
[46] F. Habib, P. Iovenitti, S. Masood, M. Nikzad, Fabrication of polymeric lattice structures for optimum energy absorption using multi jet fusion technology, Mater. Des. 155 (2018) 86–98.
[47] S. Annunziata, L. Lomazzi, A. Manes, A hybrid 1D–3D computational framework for dynamic analysis of lattice structures for impact protection, Eng. Proc. 131 (1) (2026) 7.
[48] J. Miltz, O. Ramon, Energy absorption characteristics of polymeric foams used as cushioning materials, Polym. Eng. Sci. 30 (2) (1990) 129–133.