

Research paper

# Mechanics-guided optimization of hybrid periodic architectures

Luca Liu<sup>a, </sup>, Laura Maria Vergani<sup>a, b, \* </sup>, Federica Buccino<sup>a, b </sup><sup>a</sup> Department of Mechanical Engineering (DMEC), Politecnico di Milano, Via La Masa 1, 20156 Milan, Italy<sup>b</sup> IRCCS Galeazzi-Sant'Ambrogio, Via Cristina Belgioioso 173, 20157 Milan, Italy

## ARTICLE INFO

## Keywords:

Triply periodic minimal surfaces  
 Surrogate-assisted optimization  
 Architected materials  
 Stiffness-to-weight ratio  
 Bayesian optimization

## ABSTRACT

Lightweight architected materials require systematic design strategies to maximize stiffness-to-weight efficiency under realistic finite-domain loading conditions. However, optimization of hybrid triply periodic surface architectures remains challenging due to implicit geometric representations, high-dimensional parameter spaces, and the computational cost of mechanical evaluation.

This study presents a mechanics-guided optimization framework for hybrid triply periodic surface architectures. An 18-dimensional coefficient representation based on linearly independent periodic basis functions enables continuous hybridization across canonical surface families within a unified design space. Candidate geometries are evaluated as finite  $2 \times 2 \times 2$  multi-cell specimens using voxel-based finite element analysis under uniaxial compression. Exploration of the high-dimensional space is performed using a Gaussian process surrogate with expected improvement, combined with connectivity screening to ensure structural feasibility.

The surrogate-assisted Bayesian optimization achieves an effective modulus of  $E_{\text{eff}} = 216,536$  MPa, outperforming random search (10 seeds, equal budget) by 31.0%. The optimized hybrid design also exhibits approximately 112–114% higher effective stiffness than representative gyroid and Schwarz P structures evaluated under identical conditions. Gibson–Ashby scaling gives an exponent of  $n = 1.082 \pm 0.015$  ( $R^2 = 0.999$ ), compared with  $n = 2.2$ – $3.4$  for canonical TPMS surfaces, supporting a transition toward stretching-dominated load transfer. The improvement is associated with material reorganization into vertically continuous load paths that enhance axial force transmission and reduce bending-driven compliance.

The proposed framework provides a computationally efficient and mechanically interpretable approach for high-dimensional design of periodic surface architectures under finite-domain constraints.

## 1. Introduction

Lightweight architected materials provide a powerful route to surpass the intrinsic performance limits of bulk solids by exploiting geometry rather than chemistry alone [1,2]. Intrinsic material properties and failure mechanisms can often be investigated through smaller-scale computational approaches, such as atomistic or molecular simulations [3,4]. Architected materials, however, introduce an additional design principle: the same constituent material can exhibit substantially different effective properties when arranged into different spatial architectures. As a result, stiffness, strength, deformation mode, and energy absorption can be tuned through geometry even when the base material remains unchanged. Natural systems such as bone exemplify this principle: high stiffness and strength are achieved at low mass through hierarchically organized porous architectures that efficiently transmit load while preserving global connectivity [1,5,6]. Inspired by

such biological strategies, engineered lattice and cellular materials have been intensively investigated for applications spanning aerospace structures, mechanical metamaterials, biomedical implants, and energy absorption systems [7,8]. In these materials, mechanical performance is governed primarily by architecture, and small geometric variations can fundamentally alter stiffness, strength, and deformation mechanisms [9,10].

Among architected materials, triply periodic minimal surfaces (TPMS) constitute a particularly attractive class. Canonical TPMS geometries, such as the gyroid, Schwarz primitive, and diamond surfaces, form smooth, continuous networks without sharp nodes or stress-concentrating junctions. Compared with strut-based lattices, they exhibit enhanced load distribution, reduced stress localization, and favorable damage tolerance due to their curvature continuity [11–14]. Their inherent periodicity promotes efficient load transfer across unit cells and has motivated their adoption in lightweight structural

\* Corresponding author.

E-mail address: [laura.vergani@polimi.it](mailto:laura.vergani@polimi.it) (L.M. Vergani).<https://doi.org/10.1016/j.rineng.2026.111235>

Received 12 April 2026; Received in revised form 17 May 2026; Accepted 25 May 2026

Available online 26 May 2026

2590-1230/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

components, thermal devices, and porous scaffolds.

More broadly, architected metamaterials are increasingly being designed for multifunctional performance beyond stiffness alone, including coupled acoustic–mechanical response, mechanical wave attenuation, and energy absorption. Recent acoustic-mechanical metamaterials demonstrate how geometry can be used to couple load-bearing capacity with acoustic absorption or wave-control functionality, while dual-phase and bioinspired microlattice strategies further highlight the potential of architected materials for multifunctional mechanical performance [15–17]. These developments emphasize the broader design opportunity offered by architected matter, while the present work focuses specifically on deterministic, mechanics-guided TPMS hybridization for effective stiffness enhancement under finite-domain compression.

Despite these advantages, the mechanical response of TPMS architectures is highly sensitive to geometric parameters [12,18]. Modest changes in surface morphology, anisotropy, or relative density can produce substantial differences in stiffness-to-weight efficiency and deformation mode. However, most existing studies restrict the design space to a single canonical TPMS family with only a few tunable parameters, typically wall thickness or one or two shape coefficients [14, 19,20]. Such low-dimensional parameterizations inherently limit exploration and preclude systematic investigation of hybrid morphologies that combine geometric characteristics across multiple TPMS families. As a result, potentially superior configurations remain unexplored.

The challenge is compounded when considering finite multi-cell architectures rather than idealized infinite periodic unit cells [21]. In mechanical components, boundary effects and inter-cell interactions influence the global response, and stiffness trends observed in homogenized unit-cell analyses may not directly translate to finite domains. Discovering stiffness-efficient morphologies in this setting requires navigating a high-dimensional, nonconvex design landscape under computationally expensive mechanical evaluation.

This setting exposes fundamental limitations of conventional optimization strategies. The geometry is defined implicitly through level-set formulations [22], and the effective stiffness objective can only be evaluated via a full simulation pipeline. Voxel-based finite element analysis provides a robust and geometry-agnostic solver capable of handling arbitrarily complex topologies without mesh generation, yet each evaluation remains computationally costly and may exhibit discretization-induced numerical noise [20]. Gradient-based optimization is therefore poorly suited: analytical sensitivities are unavailable, numerical derivatives are unreliable, and local methods are prone to premature convergence in a rugged landscape.

Bayesian optimization (BO) provides a principled framework for global optimization under expensive black-box evaluations by iteratively constructing a probabilistic surrogate model and selecting new candidates via an acquisition function [23–26]. Nevertheless, BO is not inherently immune to the curse of dimensionality [27,28]. As dimensionality increases, surrogate fidelity degrades, acquisition landscapes become increasingly multimodal, and the number of evaluations required for reliable exploration grows rapidly.

Moreover, in geometric design problems, a significant portion of the design space can correspond to infeasible or disconnected configurations, causing optimization runs to waste computational budget on rejected candidates [29,30]. Effective application of Bayesian optimization in high-dimensional architected-material design therefore benefits from a well-structured parameter space, an informed initialization strategy, and practical feasibility handling.

In this work, we introduce a unified computational framework that addresses these challenges through three integrated components. First, we construct an 18-dimensional TPMS representation based on a set of linearly independent periodic basis functions. Unlike conventional family-specific parameterizations, this formulation enables continuous hybridization across multiple canonical TPMS types within a single

coherent design space, substantially enlarging the searchable morphology spectrum while promoting structurally connected morphologies. Second, we couple this representation with a voxel-based finite element evaluation of finite multi-cell specimens under uniaxial compression, thereby assessing mechanical performance in a physically consistent setting that captures boundary and inter-cell effects beyond single-cell homogenized approximations. Third, we employ Bayesian optimization to efficiently explore the resulting high-dimensional, expensive black-box objective. To limit wasted evaluations due to disconnected geometries or failed solves, we incorporate structure-aware initialization and connectivity pre-screening throughout; additionally, we evaluate an optional feasibility-aware BO variant that augments expected improvement with a probabilistic feasibility model.

The novelty of the present study lies in establishing a mechanically grounded, high-dimensional hybrid TPMS representation together with a robust simulation–optimization pipeline for implicitly defined design spaces. By integrating an expressive geometric parameterization, finite-domain mechanical evaluation, and surrogate-guided Bayesian optimization, the proposed framework enables systematic discovery of stiffness-efficient hybrid architectures that are difficult to access using canonical low-dimensional parameterizations or purely local optimization strategies. (Fig. 1)

## 2. Materials and Methods

### 2.1. TPMS geometry representation and design space

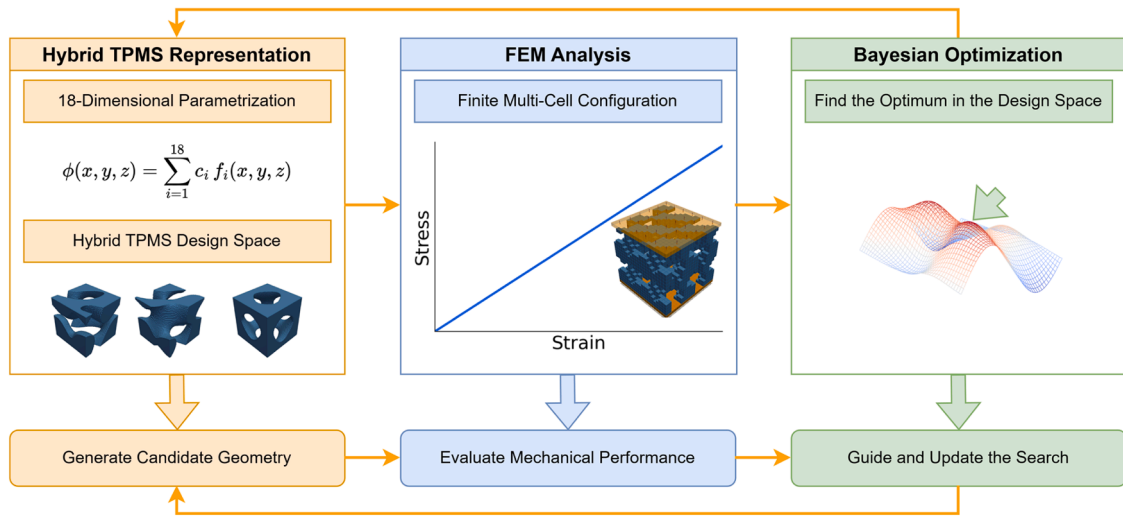
The architected materials investigated in this work are based on triply periodic minimal surface (TPMS) geometries described through an implicit level-set formulation. Instead of restricting the design space to a small set of canonical TPMS families (e.g., gyroid, Schwarz, or diamond), we adopt a unified and expressive parameterization built from a set of linearly independent periodic basis functions. This formulation enables the continuous generation of a wide spectrum of hybrid TPMS architectures within a unified design space.

The implicit field defining the TPMS geometry is expressed as:

$$\phi(x, y, z) = \sum_{i=1}^{18} c_i f_i(x, y, z) \quad (1)$$

where  $f_i(x, y, z)$  are periodic trigonometric basis functions with compatible spatial periodicity, and  $c_i$  are scalar coefficients serving as design variables. To retain expressive power while limiting redundancy, we adopt an 18-term linearly independent periodic basis grouped by geometric role (see Table 1), including Schwarz-type, gyroid-type, diamond-type, harmonic, and mixed-frequency modes. The resulting design space is 18-dimensional, with coefficients constrained to  $c_i \in [-3, 3]$ ,  $i = 1, \dots, 18$ .

The selected basis was designed to span hybridizations of the primary TPMS families considered in this work while maintaining a compact and tractable parameterization. Modes 1–4 include Schwarz-P-like and gyroid-like combinations, including rotated and sign-varied gyroid contributions. Modes 5–9 introduce diamond- and Neovius-like product terms that enrich the representation with additional surface symmetries and connectivity patterns. Modes 10–12 provide pairwise cosine-product terms, while modes 13–15 introduce second-harmonic single-axis terms ( $\cos 2x$ ,  $\cos 2y$ ,  $\cos 2z$ ) that enable modulation of wall-thickness distribution along the coordinate directions. Finally, modes 16–18 introduce mixed-frequency cross terms ( $\sin 2x \cos y$ ,  $\sin 2y \cos z$ ,  $\sin 2z \cos x$ ), enabling directionally biased and anisotropic hybridization. This structured selection expands the searchable morphology space beyond individual canonical TPMS families while limiting redundancy and preserving a manageable 18-dimensional design space for Bayesian optimization. It also helps reduce mode-cancellation effects and the occurrence of fragmented geometries; in preliminary tests, naïve combinations of periodic terms with arbitrary



**Fig. 1.** Overview of the proposed optimization framework. Hybrid TPMS geometries  $\phi$  are represented using an 18-dimensional implicit parameterization, where  $f_i(x, y, z)$  are periodic trigonometric basis functions with compatible spatial periodicity, and  $c_i$  are scalar coefficients serving as design variables. Candidate designs are generated within this design space and evaluated through finite multi-cell voxel-based finite element method (FEM) to obtain mechanical performance. The results are then used by a Bayesian optimization algorithm to iteratively guide the search toward improved stiffness-to-weight designs.

**Table 1**  
Linearly independent periodic basis functions used in the 18-dimensional TPMS representation.

Index	Basis function $f_i(x, y, z)$
1	$\cos x + \cos y + \cos z$
2	$\sin x \cos y + \sin y \cos z + \sin z \cos x$
3	$\sin x \cos z + \sin z \cos y + \sin y \cos x$
4	$\sin x \cos y - \sin y \cos z + \sin z \cos x$
5	$\sin x \sin y \sin z$
6	$\sin x \cos y \cos z$
7	$\cos x \sin y \cos z$
8	$\cos x \cos y \sin z$
9	$\cos x \cos y \cos z$
10	$\cos x \cos y$
11	$\cos y \cos z$
12	$\cos z \cos x$
13	$\cos 2x$
14	$\cos 2y$
15	$\cos 2z$
16	$\sin 2x \cos y$
17	$\sin 2y \cos z$
18	$\sin 2z \cos x$

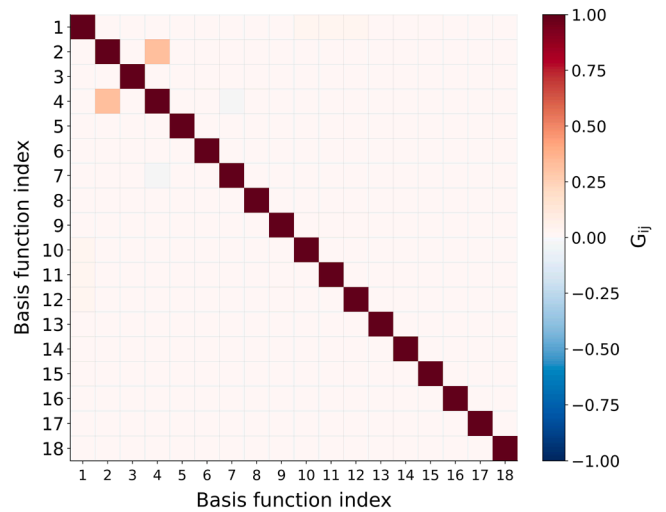
coefficients frequently produced disconnected material islands or interrupted bottom–top load paths, leading to a high proportion of candidates failing connectivity screening.

To verify numerical distinguishability after discretization, we performed a pairwise basis-function correlation analysis (Fig. 2). Each basis function was sampled on the same voxel grid used for geometry generation (resolution = 64,  $n_{\text{cells}} = 2$ ), vectorized, mean-centered, and normalized to unit Euclidean norm. The normalized inner-product matrix was defined as:

$$G_{ij} = u_i^T u_j, \text{ for } i, j = 1, \dots, 18.$$

where  $u_i$  denotes the mean-centered, unit-norm vectorized form of the  $i$ -th basis function. The maximum, mean, and median absolute off-diagonal correlations were 0.333, 0.004, and  $<0.001$ , respectively, supporting that the selected modes are numerically distinguishable on the discretized design domain and do not introduce strongly redundant coefficient directions.

Solid TPMS geometries are obtained by thresholding the implicit field  $\phi$  using a quantile-based criterion such that a prescribed target volume fraction is satisfied (the target volume fraction is 0.5 for all



**Fig. 2.** Normalized inner-product matrix of the 18 basis functions. Each basis function was sampled on the voxel grid used for geometry generation, vectorized, mean-centered, and normalized to unit Euclidean norm. Diagonal entries are equal to one by construction, while off-diagonal entries quantify normalized similarity between distinct basis-function fields. The maximum, mean, and median absolute off-diagonal values were 0.333, 0.004, and  $<0.001$ , respectively, indicating that the selected basis functions are numerically distinguishable on the discretized design domain and do not introduce strongly redundant coefficient directions.

designs considered in this study). This procedure enforces consistent mass across designs and enables meaningful comparison of stiffness-to-weight performance.

To enforce global structural connectivity, we apply a post-processing step based on connected-component analysis. Small, disconnected fragments are removed by discarding any isolated component whose volume is less than 1% of the total solid volume. If the geometry still contains more than one connected component after this pruning, it fails a connectivity pre-test, is labeled infeasible.

This connectivity pre-test is applied to all candidate designs (both the initialization phase and subsequent BO iterations) within the optimization workflow. Only feasible geometries are forwarded to finite element method (FEM) and used for surrogate-model training, whereas

infeasible candidates are rejected prior to simulation.

## 2.2. Multi-cell architectures and compression endcaps

All TPMS designs are evaluated as multi-cell architectures rather than single periodic unit cells. In the present study, the implicit field is instantiated over a  $2 \times 2 \times 2$  array of unit cells, allowing interactions between neighboring cells to influence the global mechanical response (see Fig. 3). This configuration better reflects the behavior of finite-sized architected components fabricated by additive manufacturing.

To emulate uniaxial compression tests and stabilize boundary effects, stiff endcaps are added to the top and bottom of each structure (see Fig. 3). The endcaps occupy a small fraction of the total height and are assigned a Young's modulus one order of magnitude larger than that of the TPMS body (2000 GPa vs 200 GPa). This promotes uniform load transfer while minimizing artificial stress concentrations at the loading interfaces.

## 2.3. Voxel-based FEM evaluation

Mechanical performance is evaluated using voxel-based FEM. TPMS geometries, including endcaps, are discretized on a Cartesian voxel grid, where each voxel is assigned a material phase (void, TPMS body, or endcap). To reduce computational cost, the voxel grid is coarsened by aggregating blocks of fine voxels into larger hexahedral finite elements, while averaging the local Young's modulus over solid voxels to retain the multi-material distribution. A coarse hexahedral element is retained as solid only if more than half of its constituent fine voxels are solid, otherwise it is treated as void.

Each active element is modeled as an eight-node linear hexahedron (Hex8) with isotropic linear elasticity. Element stiffness matrices are computed via standard  $2 \times 2 \times 2$  Gauss integration and assembled into a sparse global stiffness matrix.

Boundary conditions emulate uniaxial compression: all bottom surface nodes are constrained only in the loading direction ( $u_z = 0$ , z-axis is normal to the largest endcap surface (see Fig. 3), while three additional anchor constraints suppress rigid-body translations and rotation. A uniform pressure is applied on the top surface over solid faces, and the resulting sparse linear system is solved using a direct sparse solver.

The apparent mechanical properties are derived from the global force-displacement field. Let  $\delta_{\text{top}}$  denote the mean vertical displacement of the loaded (top) surface:

$$\delta_{\text{top}} = \frac{1}{N_{\text{top}}} \sum_{n \in \Gamma_{\text{top}}} u_z^{(n)} \quad (2)$$

where  $\Gamma_{\text{top}}$  is the set of top-surface nodes (restricted to solid regions) and  $N_{\text{top}}$  is their count. Here  $u_z^{(n)}$  denotes the z-component of the nodal

displacement at node  $n$ . Using the mean top-surface displacement makes the metric less sensitive to local displacement extrema than pointwise measures. The resultant applied force ( $F_{\text{app}}$ ) is computed directly from the assembled load vector as:

$$F_{\text{app}} = \left| \sum_i f_{z,i} \right| \quad (3)$$

where  $f_{z,i}$  is the z-component of the nodal force applied at top-surface node  $i$ . The apparent axial stiffness ( $k_{\text{app}}$ ) is then obtained as:

$$k_{\text{app}} = \frac{F_{\text{app}}}{\delta_{\text{top}}} \quad (4)$$

and the apparent Young's modulus ( $E_{\text{app}}$ ) is computed by normalizing with the specimen geometry:

$$E_{\text{app}} = \frac{F_{\text{app}} L}{A \delta_{\text{top}}} \quad (5)$$

where  $L$  is the specimen height and  $A = L^2$  is the nominal cross-sectional area of the cubic sample. The term 'apparent' is used because the TPMS architecture is a heterogeneous porous structure;  $E_{\text{app}}$  represents the equivalent homogeneous modulus that produces the same global force-displacement response under uniaxial loading.

To compare designs at fixed mass, we also define a normalized performance metric,  $E_{\text{eff}}$ , by dividing the apparent modulus by the realized solid volume fraction VF of the voxelized geometry:

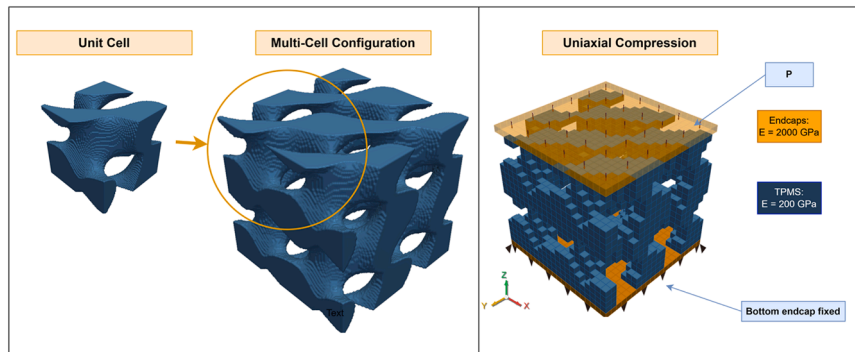
$$E_{\text{eff}} = \frac{E_{\text{app}}}{\text{VF}} \quad (6)$$

We refer to this quantity as the effective stiffness, which serves as an indicator of stiffness-to-weight performance. It is used as the scalar performance metric in the optimization. Here, VF is computed as the ratio of solid voxels (TPMS body + endcaps) to the total number of voxels in the full specimen.

The target volume fraction of 0.5 is enforced only for the TPMS body via quantile thresholding of the implicit field, whereas the final realized VF of the full specimen may slightly differ because: (i) the endcaps add a fixed amount of solid material, and (ii) voxelization and post-processing (e.g., removal of disconnected components) introduce small deviations from the target body fraction. Since the endcaps are identical across all designs, their contribution to VF is constant and does not affect the relative ranking of candidates.

## 2.4. Bayesian optimization framework

BO is used to explore the 18-dimensional TPMS coefficient space under expensive FEM-based evaluations. Each design is defined by the



**Fig. 3.** Finite-element simulation setup. A single TPMS unit cell is replicated to form a finite multi-cell specimen. Uniaxial compression is applied through rigid endcaps, with uniform pressure imposed on the top surface and the bottom surface fixed. The endcaps are assigned a high stiffness to ensure uniform load transfer, while the TPMS material is modeled according to the target elastic properties.

coefficient vector  $R^{18}$  and evaluated through the full geometry generation and voxel-FEM pipeline, yielding the effective stiffness  $E_{eff}$  as the scalar objective.

#### 2.4.1. Initialization

Before starting the BO iterations, an initial training dataset is constructed using a structure-aware (magnitude-controlled) sampling strategy. A fixed number of coefficient vectors are generated within a restricted region of the design space corresponding to moderate amplitudes. This produces smooth geometries dominated by low-frequency backbone contributions, while higher-frequency and anisotropic modes are introduced with smaller magnitudes to avoid highly irregular or fragmented structures at the start of the search.

All initial designs are subjected to the connectivity pre-test described in Section 2.1. Only feasible geometries are evaluated using FEM and included in the training dataset, whereas infeasible candidates are discarded prior to simulation.

#### 2.4.2. Surrogate model and acquisition

A Gaussian process (GP) surrogate is fit to the set of feasible FEM evaluations and updated sequentially. Candidate designs that fail pre-FEM feasibility checks or for which the FEM solve fails are treated as infeasible and excluded from GP training (equivalently assigned  $E_{eff} = -\infty$  for bookkeeping and filtered prior to fitting). New candidates are selected by maximizing an expected improvement (EI) acquisition function over the bounded domain  $[-3, 3]^{18}$ . Acquisition maximization is performed sequentially ( $q = 1$ ) using multi-start optimization.

All BO-proposed candidates are filtered using the same connectivity pre-test prior to FEM evaluation. Only feasible designs are simulated and

used to update the GP, while infeasible ones are rejected and excluded from surrogate training.

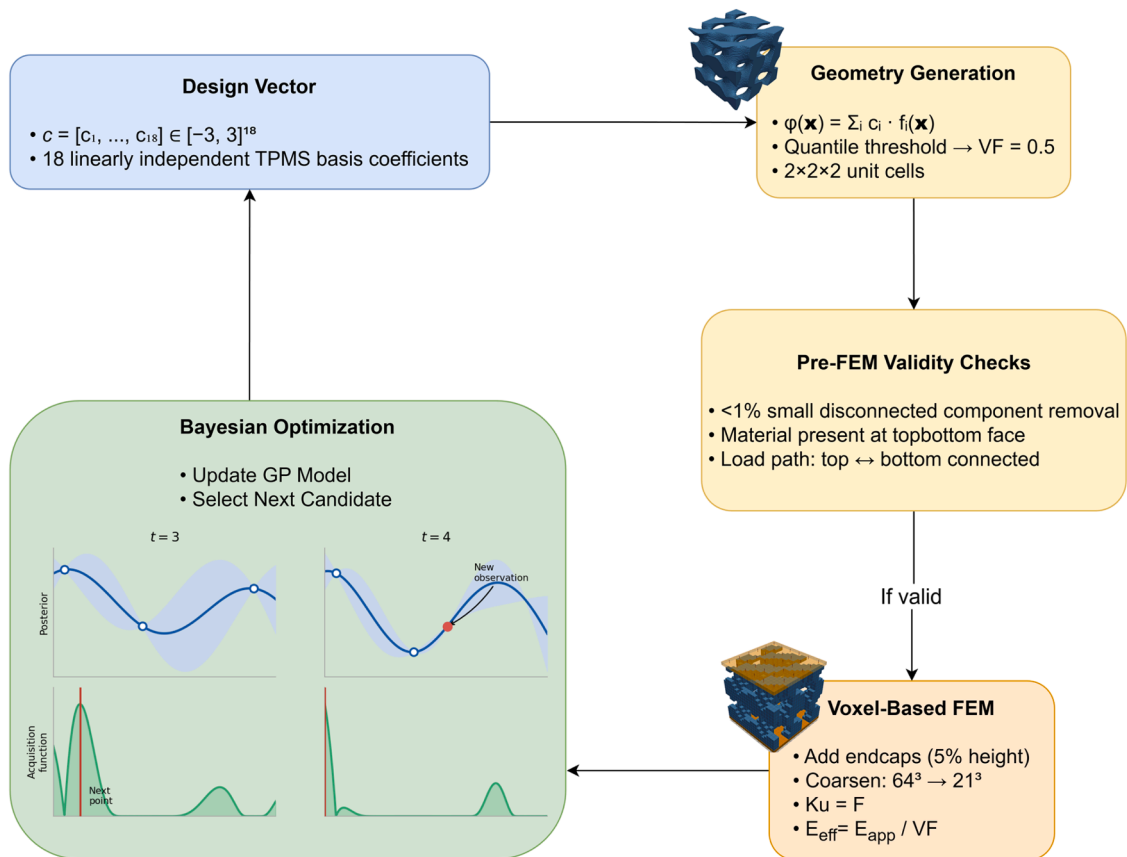
Two BO configurations are considered:

- 1) in the baseline formulation (or unconstrained formulation), the surrogate model is a BoTorch SingleTaskGP with default settings.
- 2) in the constraint-aware formulation, the objective surrogate uses a Matérn kernel ( $\nu = 2.5$ ) with automatic relevance determination (ARD) over the 18 coefficients with a Gamma prior on the lengthscales (Gamma(2, 4)). The ARD formulation allows the surrogate model to learn individual lengthscales for each coefficient, enabling data-driven identification of the most influential geometric modes and effectively adapting the search to the intrinsic dimensionality of the problem. To reduce wasted evaluations, a separate feasibility surrogate is trained on a binary feasibility indicator ( $y \in \{0, 1\}$ ) using the same GP form; the posterior mean, clamped to  $[0, 1]$ , is used as an estimate  $\hat{P}(\text{feasible}|\mathbf{c})$ . Once sufficient feasible and infeasible observations are available, the acquisition is modified as:

$$a(\mathbf{c}) = EI(\mathbf{c}) \hat{P}(\text{feasible}|\mathbf{c}) \quad (7)$$

This feasibility-aware acquisition favors candidates that are expected to improve the objective while remaining likely to pass geometric and numerical checks, reducing wasted evaluations when failures are costly.

Acquisition maximization is performed via multi-start gradient-based optimization (L-BFGS-B): the baseline uses 10 restarts and 100 raw samples, while the feasibility-aware variant uses 20 restarts and 200



**Fig. 4.** Overview of the optimization pipeline. An 18-dimensional design vector defines the hybrid TPMS implicit field, which is thresholded to a multi-cell body at the target volume fraction. Pre-FEM validity checks on the body (small-component removal, top-bottom connectivity, boundary material) filter candidates; valid designs are given endcaps and evaluated with voxel-based FEM to compute the stiffness-to-weight objective. A Gaussian Process surrogate is updated in a Bayesian optimization loop that iteratively proposes new candidates.

raw samples to handle the more rugged acquisition landscape induced by feasibility weighting.

Fig. 4 summarizes the end-to-end evaluation loop: starting from the 18D coefficient vector, we generate the TPMS geometry (body only), apply pre-FEM feasibility checks on that body (connectivity, load path, volume fraction, top/bottom material). If the checks pass, we add compression endcaps (5% height), run voxel-based FEM on the full structure to compute  $E_{app}$ , and evaluate the effective stiffness objective  $E_{eff} = E_{app}/VF$ . This observation updates the GP and drives the selection of the next candidate.

## 2.5. Computational implementation

All simulations and optimization procedures were implemented in Python (version 3.10). The Bayesian optimization loop was developed using PyTorch v.2.1 [31] for tensor computations and automatic differentiation during acquisition maximization. Gaussian process (GP) surrogate models and acquisition functions were implemented using BoTorch v0.16.1 [32], built on top of GPyTorch v1.14.3 [33], enabling numerically stable exact marginal likelihood training and scalable GP inference.

The voxel-based finite element solver was implemented using sparse matrix data structures and linear algebra routines provided by SciPy v1.16.3 [34]. Global stiffness matrices were assembled in compressed sparse row (CSR) format to reduce memory consumption, and the resulting linear systems were solved using direct sparse solvers available

within SciPy. Core numerical operations, including array manipulation, voxel aggregation, and element-wise computations, were performed using NumPy v2.3.5 [35] to ensure computational efficiency.

Three-dimensional rendering, geometry inspection, and post-processing of optimized architectures were performed using PyVista v0.46.5 [36]. This enabled verification of connectivity, boundary-condition application, and qualitative assessment of deformation patterns.

All simulations were executed on a single CPU (Intel(R) Core(TM) i5-14400F (2.50 GHz)) workstation without GPU acceleration. The modular implementation separates geometry generation, finite element evaluation, and surrogate-model updating, ensuring reproducibility and facilitating future extension to alternative solvers or optimization strategies.

## 3. Results and Discussion

### 3.1. Finite Element Model validation

The voxel-based FEM implementation was validated using an analytical benchmark and a suite of internal consistency checks. The main benchmark considers a homogeneous solid cube under uniaxial compression: a uniform pressure is applied on the top face ( $z = L$ ), while the bottom face ( $z = 0$ ) is constrained in the loading direction ( $u_z = 0$ ) (see Fig. 5a); a minimal set of additional anchor constraints is introduced only to eliminate rigid-body modes. For the analytical benchmark, a

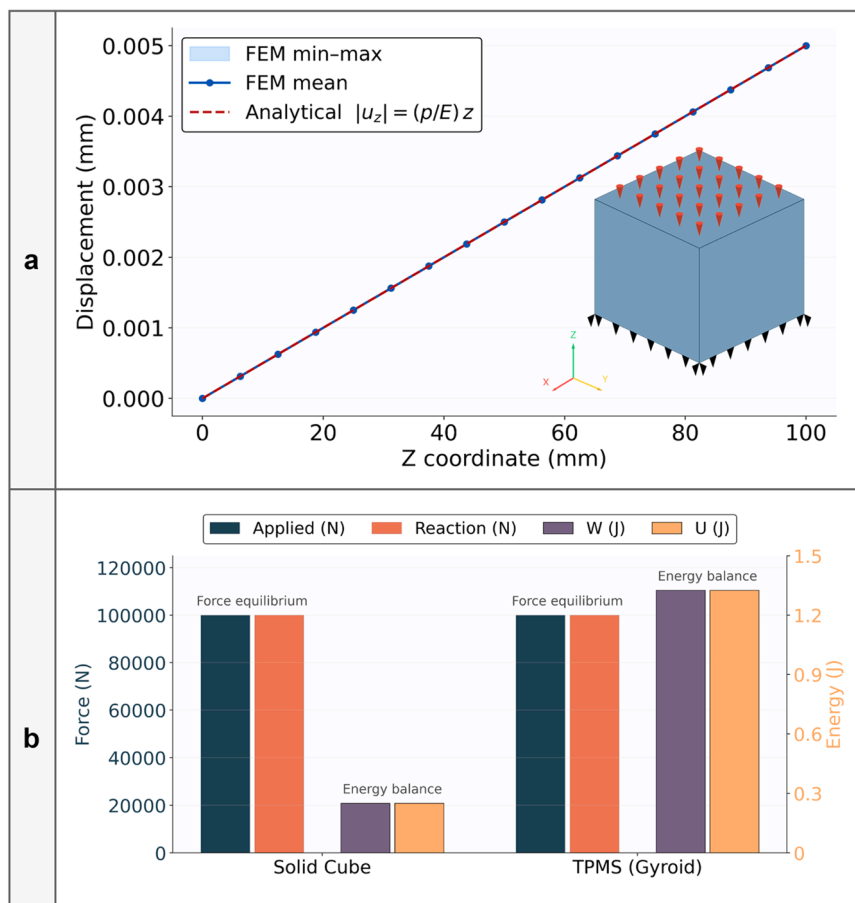


Fig. 5. Verification of the voxel-based FEM implementation. (a) Displacement profile along the loading direction for the solid-cube benchmark. The analytical solution (red dashed line) is compared with the FEM results, showing the layer-averaged vertical displacement (blue markers) and the corresponding min-max range within each layer (shaded band). (b) Global equilibrium and energy consistency checks. For both a solid cube and a representative TPMS (gyroid) geometry, the total applied load matches the sum of reaction forces, confirming force equilibrium. The external work  $W = \frac{1}{2} \mathbf{f}^T \mathbf{u}$  also coincides with the internal strain energy  $U = \frac{1}{2} \mathbf{u}^T \mathbf{K} \mathbf{u}$ , demonstrating correct stiffness assembly and load application.

cube of side length  $L = 100$  mm and Young's modulus  $E = 200$  GPa was subjected to a uniform compressive pressure  $= 10$  MPa. Assuming uniaxial stress with traction-free lateral surfaces (i.e., free lateral deformation), the analytical mean top displacement is  $\delta_{\text{top}}^{\text{analytical}} = pL/E = 0.005$  mm. Under these conditions, the solver matches the analytical mean top displacement (and the corresponding effective Young's modulus) to within numerical precision, showing excellent agreement with the analytical solution. Table 2 reports the comparison at the coarsest discretization ( $16^3$  voxels), and the same agreement is observed at higher resolutions.

To assess the displacement solution beyond a single scalar metric, we inspected the displacement field for the solid-cube benchmark (Fig. 5a). We report the layer-wise displacement profile along the loading direction, computed for each nodal layer  $z = z_k$  as the mean  $u_z(z_k) = \langle |u_z| \rangle_{z=z_k}$ , together with the corresponding in-plane min-max range  $[\min(|u_z|), \max(|u_z|)]$  within that layer. The FEM profile overlaps the analytical prediction  $|u_z(z)| = (p/E)z$  across the full height, confirming the expected linear displacement evolution from the constrained bottom surface ( $|u_z| \approx 0$ ) to the loaded top surface ( $|u_z| \approx 0.005$  mm). The min-max range within each layer is too small to be visually resolved at the plot scale, indicating an essentially uniform vertical displacement within each cross-section.

Global equilibrium and energetic consistency were verified through two internal checks. First, a global force-balance test was performed for the solid cube benchmark and a representative gyroid TPMS geometry by comparing the total applied load with the sum of reaction forces; the resulting net force is essentially zero (Fig. 5b), confirming equilibrium to numerical precision. Second, energetic consistency was by comparing the external work  $W = \frac{1}{2} \mathbf{f}^T \mathbf{u}$  with the internal strain energy  $U = \frac{1}{2} \mathbf{u}^T \mathbf{K} \mathbf{u}$ . As shown in Fig. 5b,  $W$  and  $U$  coincide for both the solid cube and the gyroid case, indicating a consistent stiffness assembly and load application.

These tests confirm that the FEM framework is physically consistent and suitable for comparative stiffness evaluation of optimized TPMS architectures. The analytical cube benchmark verifies that, under uniaxial compression with free lateral contraction, the solver reproduces the expected uniform stress state  $\sigma_{zz} = p$  and the associated linear displacement field  $u_z(z) = (p/E)z$ . The agreement of both the mean top displacement and the full displacement profile confirms that the discretized stiffness operator correctly transmits axial load without introducing artificial shear or bending modes. The negligible in-plane displacement variation within each horizontal layer further indicates that the voxel aggregation and element interpolation do not induce spurious localization or grid-alignment artifacts. The global force-balance check establishes that external tractions are exactly balanced by internal reaction forces, ensuring that the assembled stiffness matrix satisfies equilibrium conditions to machine precision. This is particularly important in voxel-based discretizations, where improper treatment of boundary voxels can otherwise generate artificial residual forces. The observed equilibrium confirms that boundary conditions and load application are implemented consistently. The energetic consistency test provides a stronger, integral-level validation of the formulation. For a linear elastic system, equality between the external work and the internal strain energy is a necessary condition for variational correctness of the discretized system. The coincidence of  $W$  and  $U$  for both the homogeneous cube and a representative porous TPMS geometry

**Table 2**

Validation of the voxel-based FEM model through comparison between the analytical solution and the simulated mean top displacement for a homogeneous solid cube under uniaxial compression at a resolution of  $16^3$  voxels.

Case	Resolution	$\delta_{\text{top}}^{\text{analytical}}$ (mm)	$\delta_{\text{top}}^{\text{FEM}}$ (mm)
Solid cubes	$16^3$	0.0050	0.0050

confirms that stiffness assembly, material interpolation, and multi-material treatment are internally consistent. Importantly, this validation on a heterogeneous porous structure demonstrates that the solver accurately captures stress redistribution across complex internal interfaces. From a mechanical standpoint, these results ensure that stiffness variations observed in the optimized TPMS architectures arise from geometric load-path differences rather than discretization artifacts, boundary inconsistencies, or numerical imbalance. The framework therefore provides a mechanically reliable basis for comparative effective stiffness evaluation across the high-dimensional hybrid design space.

To verify solver accuracy on porous geometries, a mesh-convergence study was performed for both a canonical gyroid and the unconstrained BO-optimized hybrid design, using the same coarsening factor  $\text{CF}=3$  (Fig. 6). The generation-grid resolution  $N$  was varied, and the finest feasible solution was used as the reference. Fig. 6 shows that both structures approach a stable response as the grid is refined. At the paper-default setting ( $N=64$ ,  $\text{CF}=3$ ), the relative error is 1.33% for the canonical gyroid and 0.19% for the optimized hybrid design with respect to their finest-reference solutions. These results support that the discretization used in the manuscript is within the converged response regime for both a canonical TPMS benchmark and the optimized hybrid geometry, while remaining computationally efficient.

### 3.2. Optimization convergence and mechanical performance

The optimization outcomes provide insight not only into algorithmic convergence but also into the underlying mechanical structure of the hybrid TPMS design space. By coupling high-dimensional geometric parameterization with finite element mechanical evaluation of multi-cell specimens, the framework enables systematic identification of stiffness-efficient morphologies under uniaxial compression. The following results analyze the evolution of the effective stiffness objective, the effect of feasibility-aware search, and the performance relative to canonical TPMS architectures. Particular attention is given to the mechanical mechanisms responsible for the observed performance gains, including load-path organization, deformation mode transition, and boundary-driven structural adaptation in finite-size specimens.

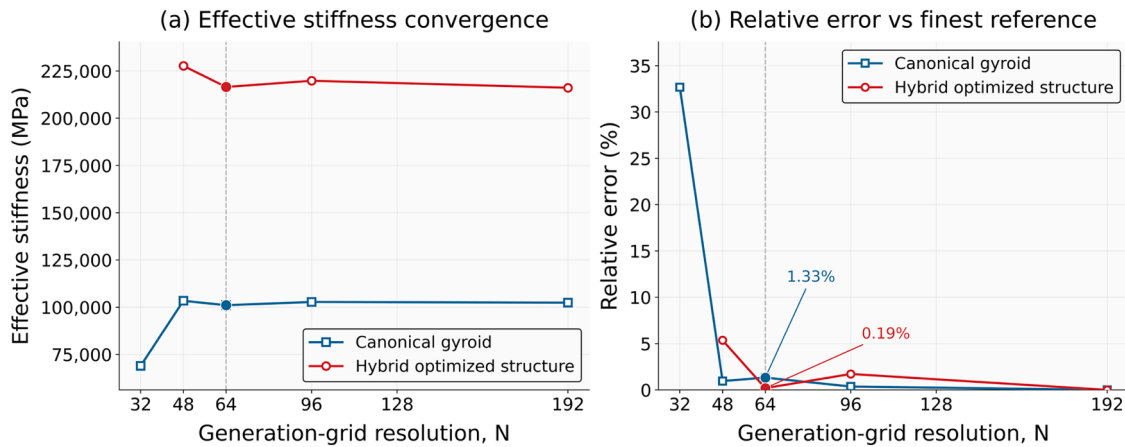
#### 3.2.1. Objective convergence

Fig. 7a shows the evolution of the effective stiffness objective  $E_{\text{eff}}$  as a function of the iteration number. After an initial exploration phase characterized by large variability in the evaluated designs, the best objective value increases monotonically and progressively saturates. This behavior is consistent with the expected exploitation-driven regime of Bayesian optimization once promising regions of the design space have been identified. The final best design achieves a effective stiffness value approximately three times larger than the median of the initial samples, indicating that the optimization is able to discover substantially improved TPMS architectures beyond the seeded configurations.

The gap between the "current" and "best-so-far" curves highlights the exploratory nature of the acquisition strategy: while many proposed candidates do not improve upon the current best solution, they contribute valuable information to the surrogate model, enabling sustained progress over hundreds of iterations.

For comparison, a random search baseline was constructed using uniform sampling of the 18-dimensional design space within the same parameter bounds and evaluated using the identical FEM pipeline, feasibility criteria and evaluation budget as Bayesian optimization (Fig. 7c). Random search is a widely used and often competitive baseline in high-dimensional optimization, as it can rapidly identify reasonable designs without model assumptions. To account for sampling variability, the random search baseline was evaluated over 10 independent seeds, each with 820 evaluations.

Fig. 7c compares Bayesian optimization with the multi-seed random search baseline. Random search achieved rapid early improvements but



**Fig. 6.** Mesh-convergence study for a canonical gyroid and the unconstrained BO-optimized hybrid design. (a) The effective modulus  $E_{eff}$  is plotted as a function of generation-grid resolution  $N$ . (b) Relative error with respect to the finest feasible reference solution for each structure. All cases were evaluated using the same coarsening factor ( $CF = 3$ ). The dashed vertical line marks the main-study setting ( $N = 64$ ), for which the relative errors are 1.33% for the canonical gyroid and 0.19% for unconstrained BO-optimized hybrid design. The results support the numerical adequacy of the selected discretization for the tested TPMS geometries.

saturated after approximately 80–200 evaluations. Across the 10 independent runs, the final random search best values had a mean of  $1.65 \times 10^5$  MPa and a standard deviation of  $5.33 \times 10^3$  MPa, with a 95% bootstrap confidence interval of  $[1.62 \times 10^5, 1.68 \times 10^5]$  MPa. The best random search run reached  $1.73 \times 10^5$  MPa. In contrast, Bayesian optimization continued to improve over the full evaluation budget and reached a final effective stiffness of  $2.17 \times 10^5$  MPa. This corresponds to an improvement of 31.0% over the random search mean and 25.4% over the best random search run. The Bayesian-optimization result exceeded all 10 random-search runs, placing it at the 100th percentile of the random-search distribution and approximately 9.6 standard deviations above the random search mean.

These results suggest that the surrogate-guided acquisition strategy provides a clear advantage over uniform exploration for identifying high-performance TPMS architectures under the tested fixed evaluation budget. In this comparison, Bayesian optimization surpassed the final best-so-far performance of the 10-seed random-search baseline by approximately 200 evaluations.

### 3.2.2. Feasibility and sample efficiency

Infeasible candidates may occur due to disconnected geometries and therefore consume computational budget without providing useful objective evaluations. Fig. 7b reports the cumulative number of feasible designs obtained during the unconstrained BO run. Over 820 attempted evaluations, 747 designs were feasible (91.1 leading to an almost linear growth of usable samples with iteration.

A short plateau is observed around iterations  $\sim 400 - 470$ , indicating a temporary phase in which several proposed candidates were rejected before the optimizer returned to predominantly feasible regions. Overall, the near-linear accumulation of feasible evaluations confirms that the optimization budget is largely converted into informative training data for the surrogate model, which is crucial given the high cost of voxel-based mechanical simulations.

To assess sensitivity to the fragment-removal threshold, the optimized design was re-evaluated at thresholds of 0.5%, 1.0%, 2.0%, and 5.0% of the total solid voxel count. The effective modulus remained unchanged at 216,536 MPa across all tested thresholds ( $\Delta E = 0.00\%$  relative to the 1% baseline), indicating that the optimized design is a single well-connected structure and that the default 1% threshold does not affect the reported optimum.

### 3.2.3. Feasibility-aware Bayesian optimization

Fig. 8 summarizes the behaviour of the feasibility-aware formulation in terms of both convergence and usable sample budget. As shown in

Fig. 8a, the best-so-far effective stiffness increases rapidly during the early iterations and then plateaus, reaching a final value of  $1.91 \times 10^5$  MPa. Fig. 8b reports the cumulative number of feasible designs, which grows almost perfectly linearly and reaches 814 feasible evaluations, indicating that only a small fraction of the computational budget is spent on rejected candidates.

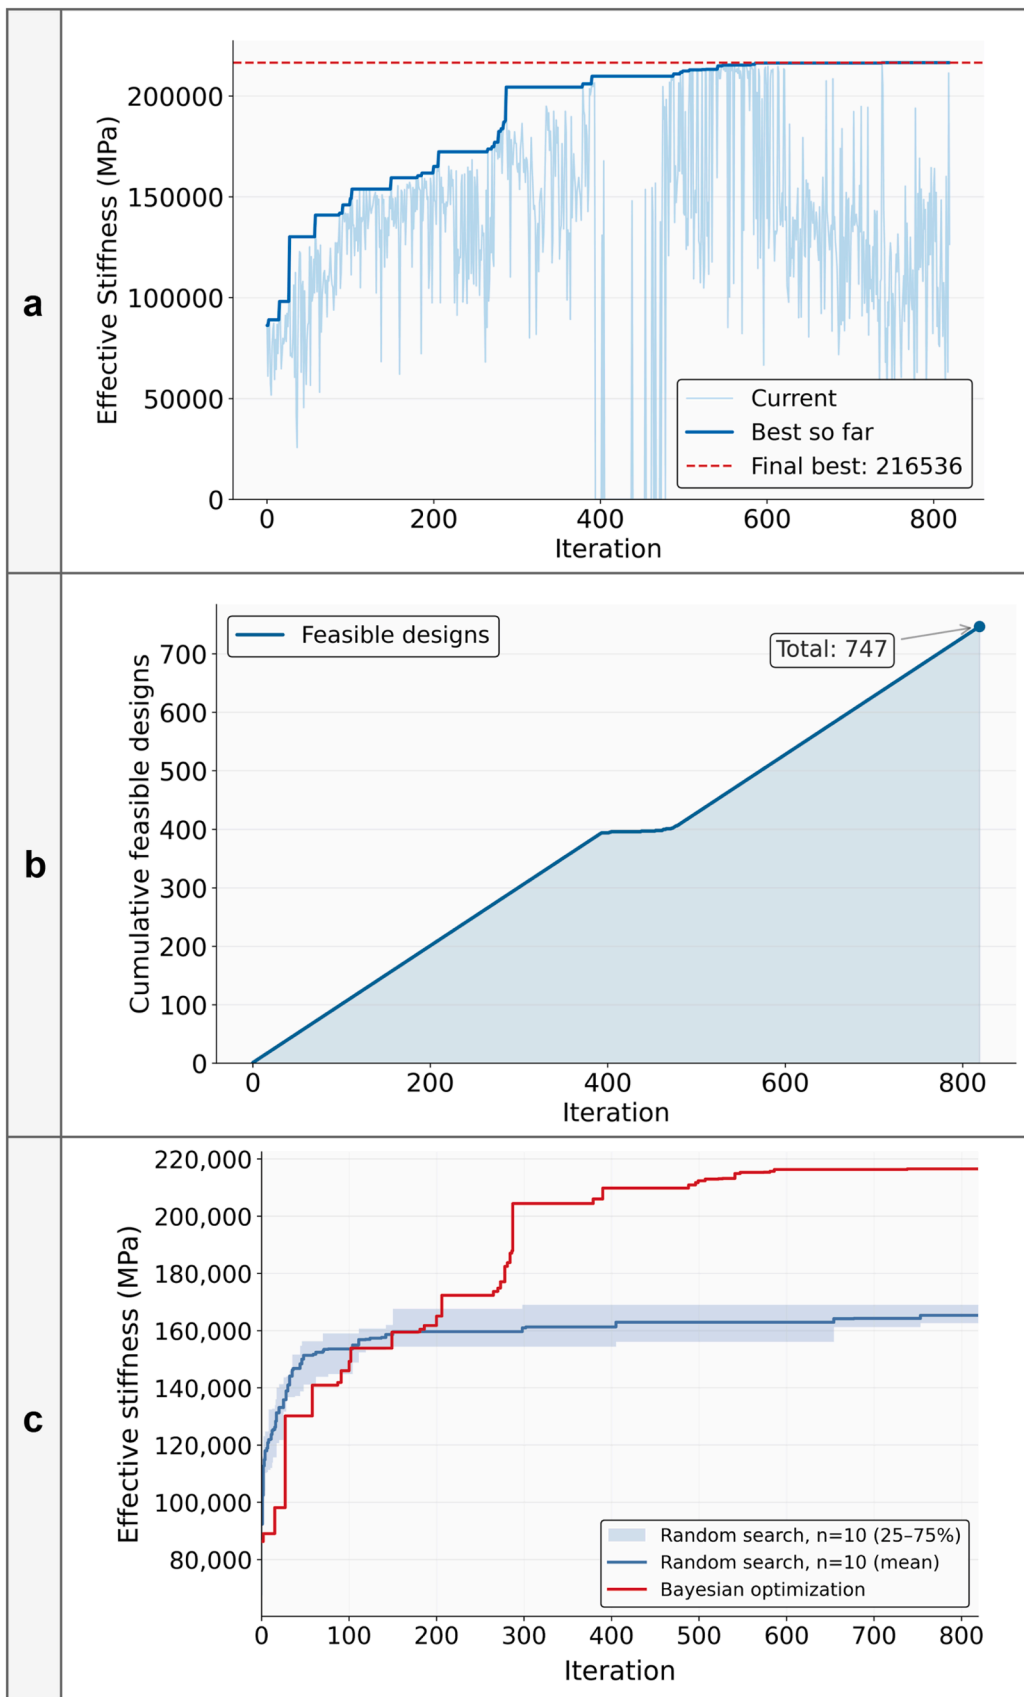
Despite the improved feasibility efficiency, the best objective achieved by the feasibility-aware strategy remains lower than that obtained with the unconstrained BO ( $2.17 \times 10^5$  MPa). This behavior suggests that high-performing designs may lie near the feasibility boundary, where small parameter changes can lead to geometric disconnections or numerical instability. By downweighting candidates with low predicted feasibility, the constraint-aware acquisition becomes more conservative, prioritizing robustness and evaluation efficiency at the expense of peak performance.

To quantify this trade-off, the unconstrained BO run evaluated 820 candidates, of which 747 were feasible and 73 were rejected as infeasible, corresponding to 8.9% wasted evaluations. In comparison, the feasibility-aware BO run evaluated 830 candidates, of which 824 were feasible and only 6 were rejected, reducing the wasted-evaluation fraction to 0.7%. This improved sample efficiency also led to faster early convergence: the feasibility-aware BO run reached 90% of its final best value in 66 evaluations, compared with 287 evaluations for the unconstrained BO run. However, this gain in evaluation efficiency came with an 11.8% reduction in peak performance, with the feasibility-aware BO run reaching  $E_{eff} = 190,899$  MPa compared with 216,536 MPa for the unconstrained BO run. These results indicate a practical trade-off between evaluation efficiency and peak stiffness: feasibility-aware acquisition substantially reduces wasted evaluations and accelerates early convergence, but may avoid high-performing designs near the feasibility boundary.

### 3.2.4. Performance relative to canonical TPMS

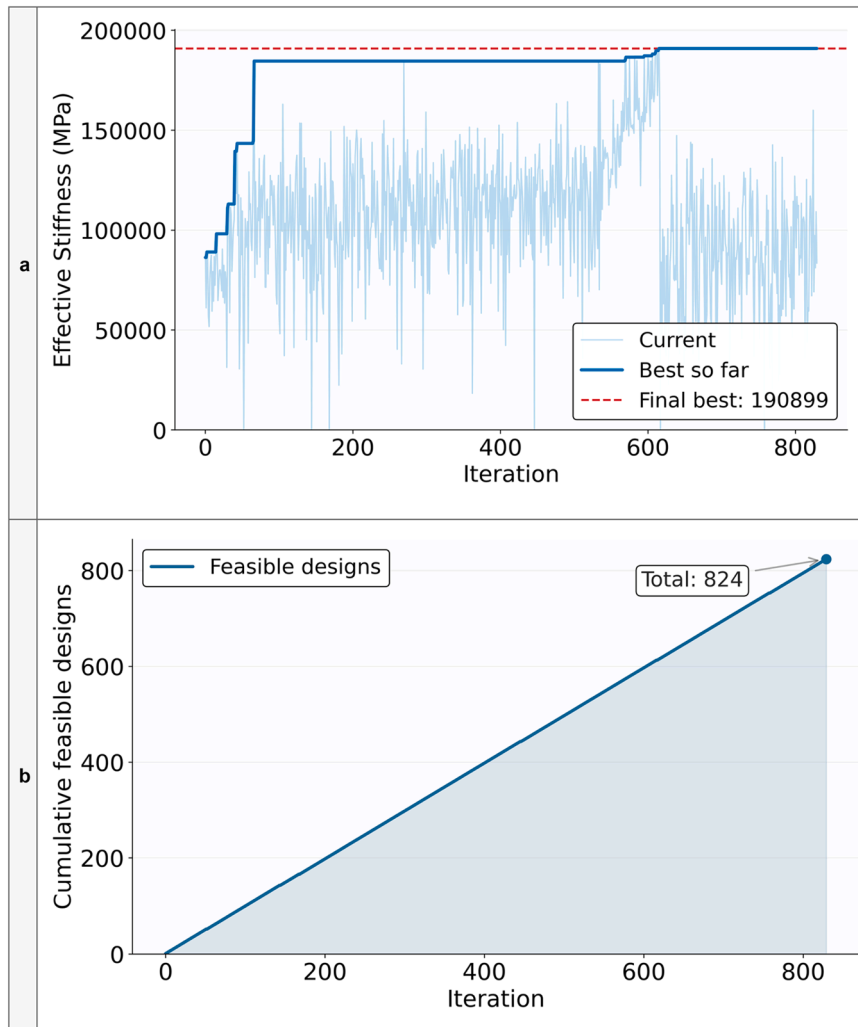
To place the obtained performance levels in context, the best designs identified by Bayesian optimization were compared with representative canonical TPMS architectures. Gyroid and Schwarz P structures were evaluated under the same geometric, boundary, and loading conditions, using identical multi-cell configurations. For a fair comparison, the canonical geometries were analyzed using the same voxel-based FEM solver, and the target volume fraction of the TPMS body was fixed to 0.5 in both cases.

As shown in Fig. 9, the optimized hybrid architectures significantly outperform the canonical designs in terms of effective stiffness. The unconstrained Bayesian optimization achieves a value of  $2.17 \times 10^5$



(caption on next page)

**Fig. 7.** Performance of the unconstrained Bayesian optimization and comparison with random search. (a) Optimization history showing the effective stiffness of each evaluated design (“current”), the best-so-far trajectory, and the final optimum of 216,536 MPa. (b) Cumulative number of feasible designs obtained during the Bayesian optimization run, reaching 747 feasible evaluations over the 820-evaluation budget. (c) Comparison between Bayesian optimization and random search under the same evaluation budget. The blue curve shows the mean best-so-far performance across 10 independent random-search runs, and the shaded region indicates the interquartile range (25th–75th percentile). Bayesian optimization continues improving after the random-search baseline has plateaued and achieves a final effective stiffness exceeding all random-search runs.



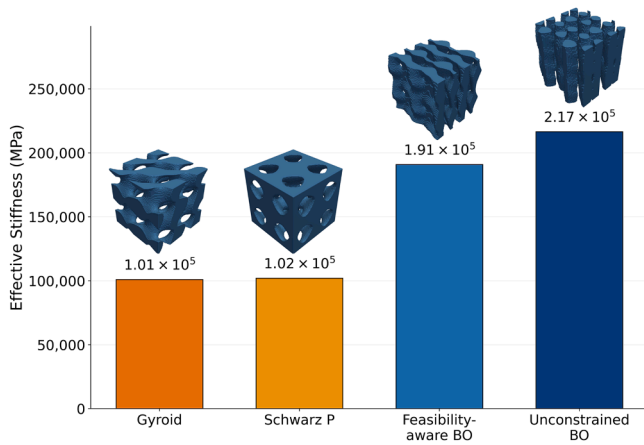
**Fig. 8.** Performance of the feasibility-aware Bayesian optimization. (a) Convergence history of the effective stiffness objective. The light curve shows the objective value of each evaluated design, while the dark blue line indicates the best-so-far performance. The red dashed line marks the final optimum ( $\approx 1.91 \times 10^5$  MPa). (b) Cumulative number of feasible designs obtained during the optimization. A total of 824 feasible designs were collected over the evaluation budget, exhibiting an almost linear growth.

MPa, compared with  $1.01 \times 10^5$  MPa for the gyroid and  $1.02 \times 10^5$  MPa for the Schwarz P structure, corresponding to an improvement of approximately 112–114%. The feasibility-aware formulation also yields a substantially higher effective stiffness ( $1.91 \times 10^5$  MPa) than the canonical references, despite its more conservative search strategy.

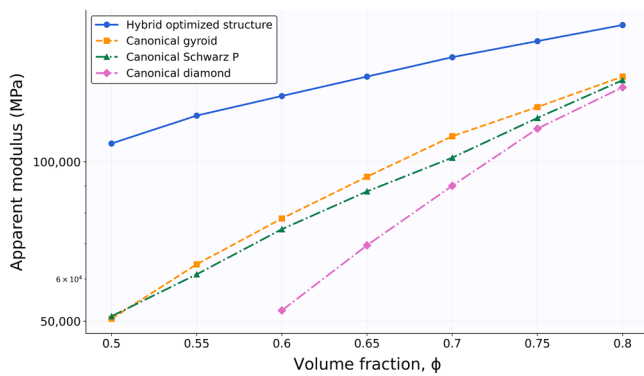
These results indicate that restricting the design to a single TPMS family may limit achievable performance, and that the unified coefficient-based representation enables the discovery of hybrid morphologies with more efficient load-bearing characteristics. The comparison also confirms that the performance gains achieved by the optimization correspond to meaningful improvements relative to commonly used TPMS architectures, rather than incremental variations within a standard design family.

To further compare the deformation regime of the optimized hybrid architecture with canonical TPMS surfaces, a Gibson–Ashby scaling

analysis was performed by fitting the relationship between apparent stiffness  $E_{app}$  and volume fraction. As shown in Fig. 10, the optimized hybrid design maintains higher apparent stiffness ( $E_{eff} = E_{app}/VF$ ) than the canonical TPMS surfaces across the tested volume-fraction range. Power-law fitting of  $E_{app} \sim \varphi^n$  gives  $n = 1.082 \pm 0.015$  ( $R^2 = 0.999$ ) for the optimized hybrid design, compared with  $n = 2.242 \pm 0.041$ ,  $n = 2.170 \pm 0.043$ , and  $n = 3.412 \pm 0.078$  for the canonical gyroid, Schwarz P, and diamond surfaces, respectively. The markedly lower exponent of the optimized hybrid design indicates a more stretching-dominated stiffness scaling, whereas the canonical TPMS surfaces exhibit bending-dominated scaling over the same volume-fraction range. This result provides quantitative support for the interpretation that the optimized hybrid architecture improves stiffness-to-weight performance through more efficient axial load transfer rather than incremental tuning of a single canonical TPMS family.



**Fig. 9.** Comparison of effective stiffness between representative canonical TPMS architectures and the optimized designs obtained with Bayesian optimization. Gyroid and Schwarz P structures (target volume fraction = 0.5) were evaluated using the same voxel-based FEM framework. Both optimization strategies (Unconstrained Bayesian Optimisation, UBO and Feasibility-aware Bayesian Optimisation, FBO) significantly outperform the canonical geometries, with the unconstrained Bayesian optimization achieving the highest performance, followed by the feasibility-aware variant.



**Fig. 10.** Gibson–Ashby scaling of the optimized hybrid design and canonical TPMS surfaces. Apparent stiffness  $E_{app}$  is plotted as a function of volume fraction  $\phi$ . Power-law fitting was performed using  $E_{app} \sim \phi^n$ . The optimized hybrid design gives  $n = 1.082 \pm 0.015$ , compared with  $n = 2.242 \pm 0.041$ ,  $n = 2.170 \pm 0.043$ , and  $n = 3.412 \pm 0.078$  for the canonical gyroid, Schwarz P, and diamond surfaces, respectively. The lower exponent supports a transition toward stretching-dominated load transfer in the optimized hybrid architecture.

To further assess whether the optimized hybrid architecture depends on a single dominant basis mode or on a distributed combination of modes, a local coefficient-ablation analysis was performed. Starting from the unconstrained BO-optimized design, each of the 18 coefficients was set to zero one at a time while the remaining coefficients were kept fixed, and the resulting geometry was reprocessed using the same thresholding, connectivity screening, and voxel-FEM pipeline. Among the feasible ablations, the change in effective stiffness remained within approximately  $\pm 1.1\%$ , indicating that no single coefficient alone dominates the optimized stiffness. Several coefficient removals led to infeasible geometries, indicating local topological sensitivity of the optimized design to those modes. Overall, the analysis supports that the optimized architecture arises from the combined nonlinear interaction of multiple basis functions rather than from a single canonical TPMS mode.

### 3.2.5. Optimized TPMS geometries

Fig. 11a visualizes the best-performing TPMS design obtained with

the unconstrained Bayesian optimization and Fig. 11b the best design obtained with the feasibility-aware variant. In both cases, the structures are rendered from the voxelized material field and post-processed with a surface-smoothing step to reduce staircase artifacts; the endcaps are shown separately for clarity. Additional high-ranking designs (e.g., second- and third-best) were visually similar to the corresponding optimum and are omitted for brevity.

Despite differences in the coefficient values, both optimization strategies converge to morphologies characterized by continuous, vertically oriented load-bearing pathways. The resulting architectures exhibit column-like regions aligned with the loading direction, which provide direct force transmission between the top and bottom surfaces while minimizing bending-dominated deformation. Compared with canonical TPMS geometries, which distribute material more isotropically, the optimized hybrid structures concentrate material along the principal load direction, leading to a more efficient axial load transfer and improved effective stiffness.

This consistent structural motif across different optimization settings suggests that the proposed framework is able to identify physically meaningful design trends rather than isolated high-performing configurations.

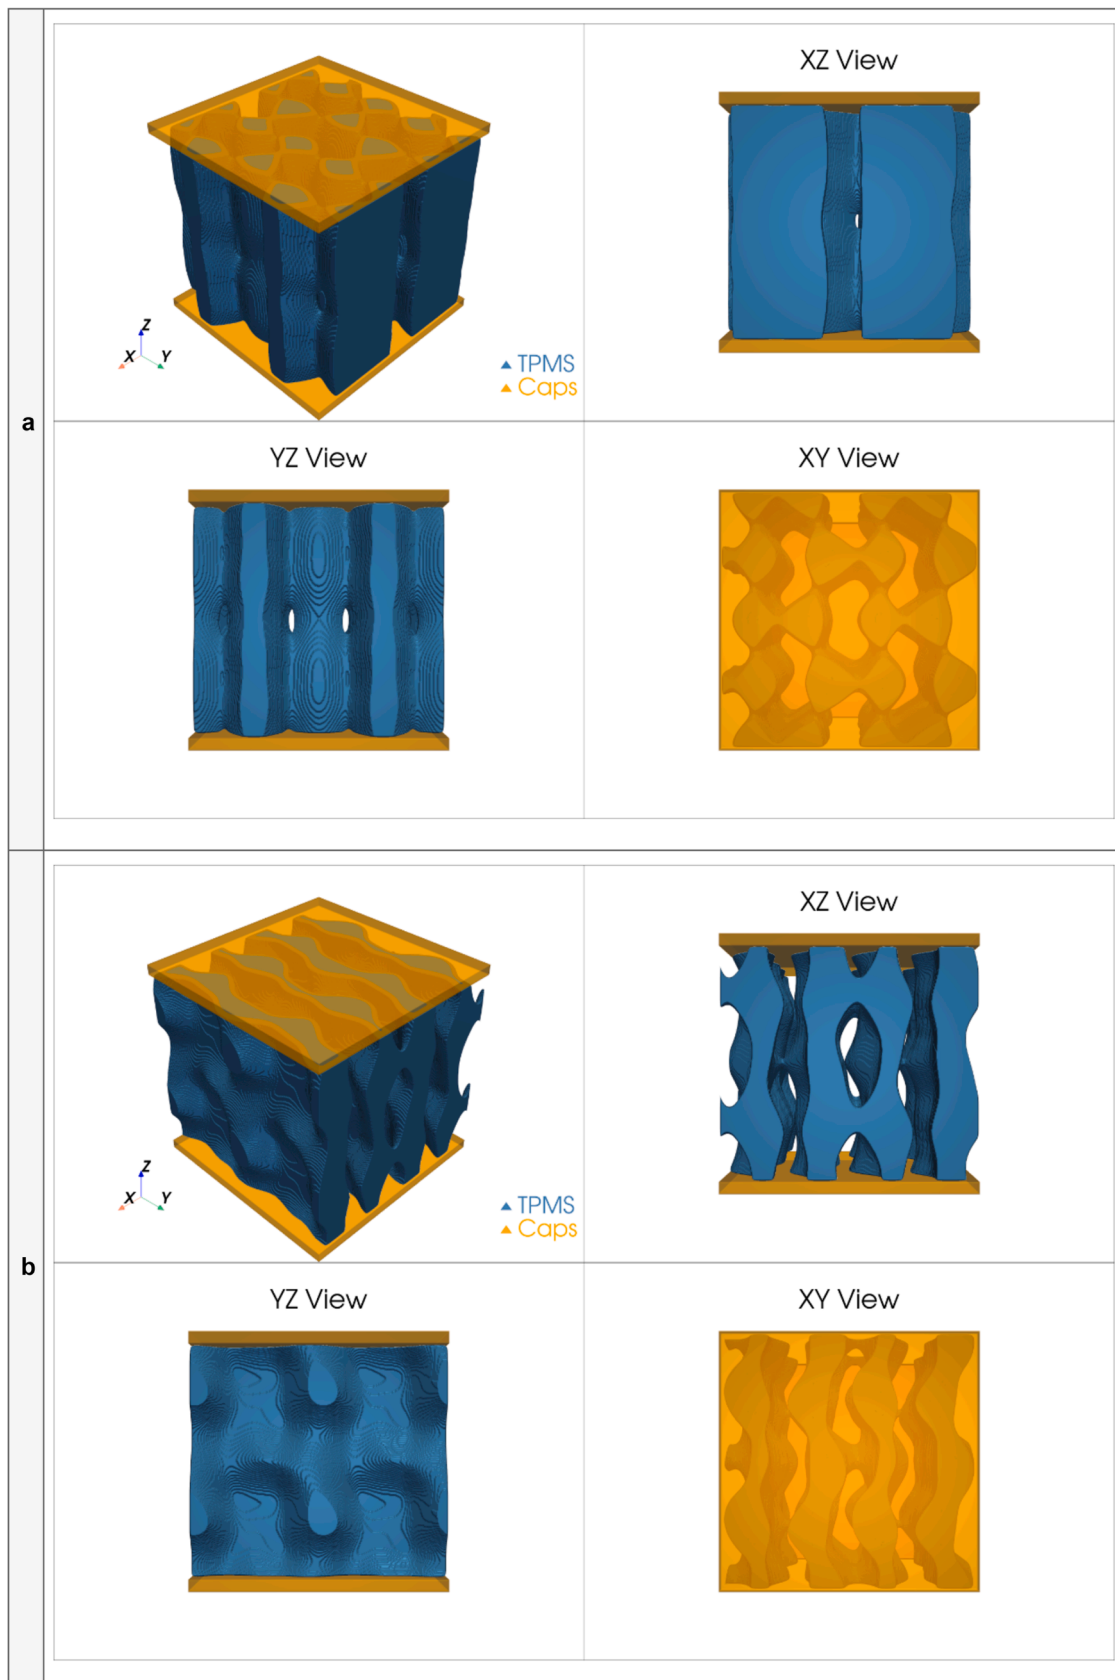
### 3.2.6. Mechanistic interpretation of the optimization results

From a mechanical standpoint, the observed convergence behavior and performance gains can be interpreted in terms of load-path optimization and deformation mode transition within the finite multi-cell architecture. Canonical TPMS geometries such as gyroid and Schwarz P distribute material relatively isotropically, promoting a combination of axial and bending-dominated deformation under uniaxial compression. While this isotropy provides multifunctionality and robustness, it does not maximize stiffness along a single loading direction at fixed volume fraction.

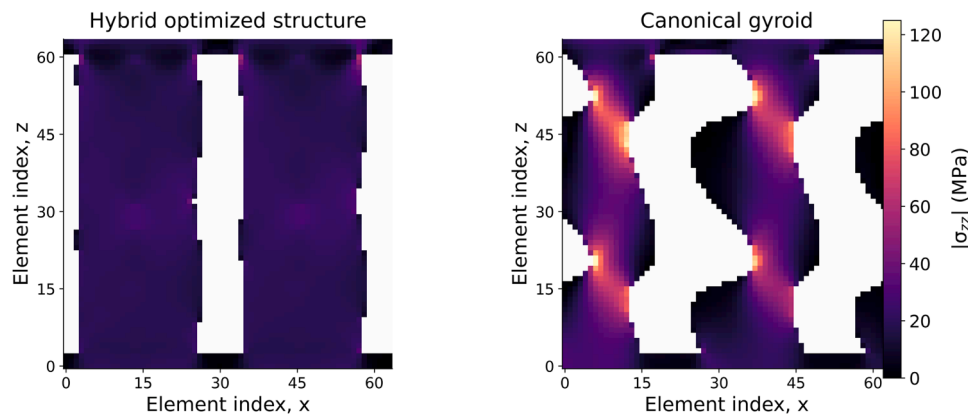
The optimized hybrid architectures, in contrast, systematically reorganize material toward vertically aligned, continuous load-bearing pathways spanning the specimen height. This reconfiguration effectively increases the fraction of material participating in axial stress transmission ( $\sigma_{zz}$ ) while reducing transverse bending and shear components. In structural mechanics terms, the optimization drives the architecture toward a more stretching-dominated regime along the loading direction, where deformation energy is stored primarily through axial strain rather than bending curvature. Since stretching-dominated mechanisms scale more favorably with relative density than bending-dominated ones, this shift directly enhances stiffness-to-weight efficiency.

The Gibson–Ashby analysis in Section 3.2.4 provides global evidence that the optimized hybrid design follows a more stretching-dominated stiffness scaling than the canonical TPMS surfaces. To examine whether this behavior is also reflected at the local stress-field level, the axial stress distribution was evaluated under identical uniaxial compression for the optimized hybrid design and the canonical gyroid. Fig. 12 shows mid-plane maps of the absolute axial stress  $|\sigma_{zz}|$ . The optimized hybrid structure exhibits lower peak axial stress concentration than the canonical gyroid, with a peak  $|\sigma_{zz}|$  of 60.1 MPa compared with 145.9 MPa for the gyroid. The mean absolute axial stress is also lower for the optimized hybrid design, 19.2 MPa compared with 25.9 MPa for the gyroid. These results indicate that the optimized morphology transfers compressive load with reduced local stress concentration, consistent with the stretching-dominated response inferred from the Gibson–Ashby exponent.

The fact that both unconstrained and feasibility-aware optimization runs converge to qualitatively similar column-like motifs indicates that this vertical load-path alignment represents a mechanically robust optimum within the admissible design space. The surrogate model is therefore not identifying arbitrary geometric irregularities but rather exploiting a consistent mechanical principle: concentration of material along the principal stress trajectory induced by the boundary conditions.



**Fig. 11.** Best-performing TPMS designs obtained from Bayesian optimization. (a) Optimal geometry from the unconstrained formulation, achieving an effective stiffness of  $\approx 2.17 \times 10^5$  MPa. (b) Optimal geometry from the feasibility-aware formulation, achieving  $\approx 1.91 \times 10^5$  MPa. For each case, the finite multi-cell specimen  $2 \times 2 \times 2$  is shown together with orthogonal cross-sectional views (XZ, YZ, and XY). Both designs exhibit continuous vertical load-bearing pathways connecting the top and bottom endcaps, indicating efficient axial force transmission.



**Fig. 12.** Axial stress distribution under uniaxial compression. Mid-plane maps of absolute axial stress  $|\sigma_{zz}|$  are shown for the optimized hybrid structure and the canonical gyroid under identical loading and boundary conditions. A common color scale is used for both panels. The optimized hybrid structure exhibits a more distributed axial stress field and lower peak stress concentration than the gyroid, supporting the interpretation of more continuous axial load transfer and a stretching-dominated deformation mechanism.

The difference between unconstrained and feasibility-aware BO further supports this interpretation. The highest-performing designs appear to lie near the feasibility boundary, where aggressive coefficient combinations produce slender but continuous vertical members that maximize axial stiffness. Small perturbations in these configurations can lead to geometric disconnections or numerical instability, explaining why the feasibility-aware strategy—by penalizing low predicted feasibility—converges to slightly more conservative architectures with marginally reduced peak stiffness. Mechanically, this reflects a trade-off between structural efficiency and geometric robustness: the most stiffness-efficient morphologies approach the limit of connectivity preservation.

The comparison with canonical TPMS structures clarifies that the performance improvement is not merely a parametric refinement within a known surface family. Instead, the unified coefficient representation enables redistribution of material density in a directionally biased manner that canonical minimal surfaces, by construction, do not favor. Whereas classical TPMS are constrained by their minimal-surface property (zero mean curvature), the hybrid representation relaxes this restriction and allows curvature redistribution that enhances axial load transfer in finite domains. This additional geometric freedom permits the emergence of quasi-columnar features that better align with the applied compression field.

Finally, the progressive convergence behavior of the surrogate-guided optimization reflects the underlying mechanical landscape of the design space. Early iterations explore broadly varying morphologies with diverse deformation modes, leading to high variability in stiffness response. As the surrogate identifies regions associated with stretching-dominated axial load paths, subsequent iterations refine coefficient combinations that strengthen vertical continuity and suppress bending-dominated branches. The monotonic improvement in best-so-far performance therefore corresponds to a gradual mechanical transition toward architectures that more efficiently convert applied compressive work into uniform axial strain energy.

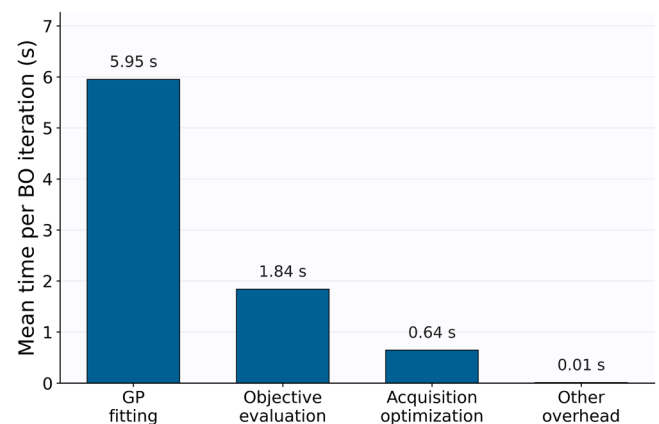
Overall, the optimization results demonstrate that, within a high-dimensional hybrid TPMS space evaluated on finite multi-cell specimens, stiffness-to-weight enhancement is achieved primarily through directional material redistribution and load-path alignment rather than through incremental geometric smoothing. The surrogate-assisted framework effectively uncovers this mechanically interpretable design trend, confirming that the observed performance gains originate from physically meaningful structural mechanisms rather than numerical artifacts or random search effects.

### 3.3. Capabilities, limitations, and future directions

The proposed framework demonstrates several methodological and computational capabilities that support its applicability for high-dimensional architected-material design.

First, the results show that surrogate-assisted optimization in an 18-dimensional implicit TPMS design space remains computationally tractable under finite-domain mechanical evaluation. All simulations were executed on a single CPU workstation without GPU acceleration, as described in Section 2.5. To quantify the complete computational cost of the Bayesian-optimization workflow, a wall-clock profiling analysis was performed across the full BO loop (20 initial samples + 800 BO iterations, seed = 42), with per-iteration timers placed around objective evaluation, GP model fitting, acquisition optimization, and residual overhead.

The resulting timing breakdown is summarized in Fig. 13, which reports the mean wall-clock time per BO iteration for each component of the workflow. Among the 800 BO iterations, GP model fitting accounted for 70.5% of total per-iteration wall time (mean 5.95 s/iteration, 4,762 s total), followed by objective evaluation, including geometry generation, feasibility screening, FEM evaluation, and post-processing, which accounted for 21.8% (mean 1.84 s/iteration, 1,471 s total). Acquisition



**Fig. 13.** Mean wall-clock time per BO iteration. Mean per-iteration wall-clock time is shown for BO iterations only, excluding the initial random evaluations for which GP fitting and acquisition optimization are not performed. GP model fitting dominates the BO-loop cost, followed by objective evaluation and acquisition optimization, while residual overhead is negligible. Checkpointing/I/O was not resolved as a separate timing component in this figure and is included, where present, in residual overhead.

optimization accounted for 7.6% (mean 0.64 s/iteration, 516 s total). Residual overhead was negligible at 0.1% (mean 0.01 s/iter). I/O and checkpointing were measured separately and contributed no resolved wall-clock time under the tested configuration. These results indicate that the dominant computational bottleneck of the BO loop is surrogate-model refitting as the training set grows, rather than the FEM-based objective evaluation itself.

Overall, the total wall-clock time for the 820-evaluation run was approximately 1.9 h. These results confirm that systematic exploration of hybrid TPMS morphologies in a moderately high-dimensional design space can be carried out with limited computational resources. They further suggest that future acceleration efforts should primarily target scalable GP training, approximate or sparse surrogate formulations, and more efficient acquisition optimization, rather than only the mechanical solver.

Second, the modular structure of the framework, separating geometry generation, voxel-based finite element evaluation, and surrogate-guided search, facilitates scalability and extensibility. Objective evaluations can be parallelized, and the mechanical solver can be replaced or enriched without altering the optimization logic. GPU acceleration or distributed computing would enable higher spatial resolution, larger multi-cell domains, or expanded evaluation budgets.

Third, the unified 18-dimensional TPMS representation enables continuous hybridization across canonical surface families within a single coherent parameter space. This structured yet expressive basis expands the reachable morphology spectrum beyond family-specific parameterizations while promoting geometric continuity and enabling connectivity control through screening. An optional feasibility-aware surrogate formulation can further improve robustness by reducing wasted evaluations and providing a framework that can be extended to incorporate geometric or process-related constraints through feasibility modeling.

In this context, the proposed 18-dimensional representation should be viewed as a deterministic and interpretable TPMS-derived design space that complements broader stochastic or spinodal descriptions. Stochastic and spinodal architectures offer important advantages, including statistical isotropy, smooth connectivity, and potentially favorable stress distribution under complex or multi-directional loading. The objective of the present work is not to replace these approaches, but to focus on a compact parameterization for stiffness-to-weight optimization under finite-domain uniaxial compression. This formulation provides straightforward reproducibility from an 18-coefficient vector, a low-dimensional continuous design space suitable for surrogate-assisted optimization, and interpretability through prescribed spatial modes associated with canonical TPMS families, harmonic terms, and mixed-frequency contributions.

However, some limitations remain. The mechanical simulations in the present study are restricted to small-strain, linear-elastic uniaxial compression. This setting is useful for isolating geometry-driven stiffness-to-weight efficiency and for comparing candidate architectures under controlled loading conditions, but it does not capture nonlinear deformation mechanisms such as elastic instability, post-buckling, plastic collapse, progressive damage, fatigue, tensile failure, or energy absorption under large deformation. Recent studies on 3D-printed cellular structures and architected metamaterials show that compressive response, post-yield deformation, damage tolerance, strength-toughness trade-offs, and energy absorption can depend strongly on material choice, fabrication process, local architecture, and deformation mode [37–39]. Extending the present framework to large-deformation kinematics, nonlinear constitutive laws, contact/self-contact, and experimentally validated collapse simulations would therefore be necessary for applications where strength, toughness, or energy absorption are primary objectives.

The geometric design space, although broader than conventional single-family TPMS parameterizations, remains limited to an 18-dimensional coefficient basis with fixed parameter bounds and a prescribed

target volume fraction. It should therefore not be interpreted as an exhaustive free-form topology space. Stochastic, spinodal, aperiodic, hierarchical, and topology-optimized architectures may access morphologies and multifunctional responses that are outside the present TPMS-derived representation. Future work could extend the framework through adaptive basis enrichment, generative geometric representations, topology optimization, or multi-scale parameterizations, enabling broader exploration of stiffness, robustness, acoustic response, and energy absorption within a unified optimization setting.

The optimization formulation is single-objective and focuses exclusively on effective stiffness. Other performance criteria, including strength, buckling resistance, anisotropy control, or energy absorption capacity, are not explicitly incorporated.

Despite these aspects, the present framework establishes a flexible and extensible platform for physics-informed, data-driven exploration of architected materials. Several future research routes naturally emerge. These include: (i) extension to nonlinear and instability-driven regimes to capture buckling and post-buckling behavior; (ii) integration of multi-objective Bayesian optimization to identify Pareto-optimal trade-offs among stiffness, strength, robustness, and energy absorption; (iii) incorporation of manufacturability-aware and process-constrained feasibility models; (iv) adaptive enrichment of the geometric basis to enable hierarchical or multi-scale architecture discovery; and (v) systematic investigation of size effects and boundary-condition sensitivity. Together, these directions position the methodology as a scalable foundation for next-generation computational design of hybrid architected materials.

#### 4. Conclusion

This work introduced and demonstrated a unified, mechanics-grounded framework for high-dimensional optimization of TPMS-based architected materials through an 18-dimensional coefficient representation integrated with surrogate-assisted global search. Unlike conventional family-specific parameterizations, the proposed formulation enables continuous hybridization across multiple TPMS types within a single coherent design space and evaluates performance on finite multi-cell specimens, thereby linking implicit geometry to physically consistent structural response.

The Bayesian optimization achieved a peak effective stiffness of  $E_{\text{eff}} = 216,536\text{MPa}$ , outperforming the 10-seed random-search baseline by 31.0% relative to its mean and exceeding all random-search runs. The optimized hybrid design exhibits a Gibson–Ashby scaling exponent of  $n = 1.082 \pm 0.015$  ( $R^2 = 0.999$ ), compared with  $n = 2.2\text{--}3.4$  for the canonical TPMS surfaces evaluated here, supporting a transition toward stretching-dominated load transfer. Axial stress maps further show that the optimized hybrid design has a 2.4-fold lower peak axial stress concentration than the canonical gyroid under identical loading, providing local field-level support for the proposed load-transfer mechanism. The optimization results indicate that stiffness enhancement arises primarily from directional redistribution of material and reorganization of load paths rather than from incremental parametric refinement. The optimized architecture develops vertically continuous load-bearing pathways that improve axial force transmission and reduce bending-dominated compliance. Under a fixed evaluation budget, the optimized hybrid designs achieved approximately 30% higher effective stiffness than the initial designs and approximately twofold improvement over representative canonical TPMS geometries, indicating that the gains reflect substantive structural reconfiguration.

The focus on stiffness maximization under uniaxial compression was intentionally chosen as a mechanically interpretable test case. Because the optimized morphology can be assessed against established engineering intuition, the emergence of vertically continuous load paths and stretching-dominated scaling provides confidence that the framework identifies physically meaningful design trends rather than numerical

artifacts. This creates a foundation for extending the approach to more complex nonlinear, multi-objective, and manufacturability-aware design problems.

The feasibility-aware formulation further revealed that peak performance lies near the connectivity boundary of the design space, highlighting a trade-off between axial stiffness maximization and geometric robustness. Overall, the proposed framework provides a computationally efficient and physically interpretable approach for exploring high-dimensional hybrid TPMS design spaces and for discovering stiffness-efficient architectures beyond the limitations of classical canonical surfaces.

### CRedit authorship contribution statement

**Luca Liu:** Conceptualization, Investigation, Data curation, Validation, Methodology, Writing – original draft. **Laura Maria Vergani:** Investigation, Methodology, Supervision, Funding acquisition, Writing – review & editing. **Federica Buccino:** Investigation, Validation, Methodology, Funding acquisition, Supervision, Writing – review & editing.

### Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Luca Liu reports financial support was provided by European Innovation Council. Federica Buccino reports financial support was provided by European Innovation Council. Laura Maria Vergani reports financial support was provided by European Innovation Council. Federica Buccino reports financial support was provided by Government of Italy. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgements

F. B and L. M. V. acknowledge the financial support from the European Union Horizon Europe research and innovation programme EIC Pathfinder Challenges under grant agreement No 101161602 (PANTAREI).

F.B. acknowledges financial support from the Italian Ministry of University and Research (MUR) through the Fondo Italiano per la Scienza (FIS 3) programme, under project IDEM (CUP D53C25002420001).

### Data availability

Data will be made available on request.

### References

- J.R. Greer, V.S. Deshpande, Three-dimensional architected materials and structures: Design, fabrication, and mechanical behavior, *MRS Bulletin* 44 (2019) 750–757, <https://doi.org/10.1557/mrs.2019.232>.
- M. Benedetti, A. du Plessis, R.O. Ritchie, M. Dallago, N. Razavi, F. Berto, Architected cellular materials: A review on their mechanical properties towards fatigue-tolerant design and fabrication, *Materials Science and Engineering: R: Reports* 144 (2021) 100606, <https://doi.org/10.1016/j.mser.2021.100606>.
- L. Liu, G. Raos, Assessment of reactive force fields for the failure of hydrocarbon chains: Insights from molecular dynamics and density functional theory, *Polymer* 348 (2026) 129732, <https://doi.org/10.1016/j.polymer.2026.129732>.
- A Review of Computational Methods in Materials Science: Examples from Shock-Wave and Polymer Physics n.d. <https://www.mdpi.com/1422-0067/10/12/5135/figure?id=f44-ijms-10-05135> (accessed May 7, 2026).
- M.J. Mirzaali, V. Mussi, P. Vena, F. Libonati, L. Vergani, M. Strano, Mimicking the loading adaptation of bone microstructure with aluminum foams, *Materials & Design* 126 (2017) 207–218, <https://doi.org/10.1016/j.matdes.2017.04.039>.
- M.J. Mirzaali, F. Libonati, D. Ferrario, L. Rinaudo, C. Messina, F.M. Olivieri, et al., Determinants of bone damage: An ex-vivo study on porcine vertebrae, *PLOS ONE* 13 (2018) e0202210, <https://doi.org/10.1371/journal.pone.0202210>.
- V. Tuninetti, S. Narayan, I. Ríos, B. Menacer, R. Valle, M. Al-lehaibi, et al., Biomimetic Lattice Structures Design and Manufacturing for High Stress, Deformation, and Energy Absorption Performance, *Biomimetics* (Basel) 10 (2025) 458, <https://doi.org/10.3390/biomimetics10070458>.
- A. du Plessis, N. Razavi, M. Benedetti, S. Murchio, M. Leary, M. Watson, et al., Properties and applications of additively manufactured metallic cellular materials: A review, *Progress in Materials Science* 125 (2022) 100918, <https://doi.org/10.1016/j.pmatsci.2021.100918>.
- L.J. Gibson, M.F. Ashby, *Cellular Solids: Structure and Properties*, 2nd ed., Cambridge University Press, 1997 <https://doi.org/10.1017/CBO9781139878326>.
- A.J. Lew, K. Jin, M.J. Buehler, Designing architected materials for mechanical compression via simulation, deep learning, and experimentation, *Npj Comput Mater* 9 (2023) 80, <https://doi.org/10.1038/s41524-023-01036-1>.
- Sokollu B, Gülcan O, Konukseven Eİ. Mechanical Properties Comparison of Strut-Based and Triply Periodic Minimal Surface Lattice Structures Produced by Electron Beam Melting 2022. <https://doi.org/10.2139/ssrn.4150322>.
- I. Maskery, L. Sturm, A.O. Aremu, A. Panesar, C.B. Williams, C.J. Tuck, et al., Insights into the mechanical properties of several triply periodic minimal surface lattice structures made by polymer additive manufacturing, *Polymer* 152 (2018) 62–71, <https://doi.org/10.1016/j.polymer.2017.11.049>.
- M.S. Hossain, M.M. Hossain, S. Nilufar, An Overview of Additive Manufacturing of Triply Periodic Minimal Surface (TPMS) Structures, *Polymers* 17 (2025) 3307, <https://doi.org/10.3390/polym17243307>.
- M. Zhao, X. Li, D.Z. Zhang, W. Zhai, TPMS-based interpenetrating lattice structures: Design, mechanical properties and multiscale optimization, *International Journal of Mechanical Sciences* 244 (2023) 108092, <https://doi.org/10.1016/j.ijmecsci.2022.108092>.
- Z. Li, X. Wang, Z. Wang, X. Li, X. Yu, S. Ramakrishna, et al., Emerging acousto-mechanical metamaterials: From physics-guided design to coupling-driven performance, *Materials Today* 89 (2025) 151–171, <https://doi.org/10.1016/j.matmod.2025.06.029>.
- Z. Li, K. Zeng, Z. Guo, Z. Wang, X. Yu, X. Li, et al., All-in-One: An Interwoven Dual-Phase Strategy for Acousto-Mechanical Multifunctionality in Microlattice Metamaterials, *Advanced Functional Materials* 35 (2025) 2420207, <https://doi.org/10.1002/adfm.202420207>.
- Z. Li, X. Wang, K. Zeng, Z. Guo, C. Li, X. Yu, et al., Unprecedented mechanical wave energy absorption observed in multifunctional bioinspired architected metamaterials, *NPG Asia Mater* 16 (2024) 45, <https://doi.org/10.1038/s41427-024-00565-5>.
- D.W. Abueidda, M. Bakir, R.K. Abu Al-Rub, J.S. Bergström, N.A. Sobh, I. Jasiuk, Mechanical properties of 3D printed polymeric cellular materials with triply periodic minimal surface architectures, *Materials & Design* 122 (2017) 255–267, <https://doi.org/10.1016/j.matdes.2017.03.018>.
- S.Z. Ahmad, M.H. Masood, M.U. Khattab, S.S. Ahmad, S.A.A. Zaidi, S.Z. Khan, Recent Developments in Novel TPMS Lattice Materials: Design Optimization, Performance Control, and Applications in Biomimetic Scaffolds, *Materials* (Basel) 18 (2025) 5209, <https://doi.org/10.3390/ma18225209>.
- W. Jiang, W. Liao, T. Liu, X. Shi, C. Wang, J. Qi, et al., A voxel-based method of multiscale mechanical property optimization for the design of graded TPMS structures, *Materials & Design* 204 (2021) 109655, <https://doi.org/10.1016/j.matdes.2021.109655>.
- F. Breish, C. Hamm, R. Kienzler, Beyond Global Mechanical Properties: Bioinspired Triply-Periodic Minimal Surface Cellular Solids for Efficient Mechanical Design and Optimization, *Advanced Engineering Materials* 27 (2025) 2402105, <https://doi.org/10.1002/adem.202402105>.
- S. Osher, J.A. Sethian, Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations, *Journal of Computational Physics* 79 (1988) 12–49, [https://doi.org/10.1016/0021-9991\(88\)90002-2](https://doi.org/10.1016/0021-9991(88)90002-2).
- B. Shahriari, K. Swersky, Z. Wang, R.P. Adams, N. de Freitas, Taking the Human Out of the Loop: A Review of Bayesian Optimization, *Proceedings of the IEEE* 104 (2016) 148–175, <https://doi.org/10.1109/JPROC.2015.2494218>.
- J. Mockus, *The Bayesian Approach to Local Optimization*, editor, in: J. Mockus (Ed.), *Bayesian Approach to Global Optimization: Theory and Applications*, Dordrecht, Springer Netherlands, 1989, pp. 125–156, [https://doi.org/10.1007/978-94-009-0909-0\\_7](https://doi.org/10.1007/978-94-009-0909-0_7).
- H.M. Sheikh, T. Meier, B. Blankenship, Z. Vangelatos, N. Zhao, P.S. Marcus, et al., Systematic design of Cauchy symmetric structures through Bayesian optimization, *International Journal of Mechanical Sciences* 236 (2022) 107741, <https://doi.org/10.1016/j.ijmecsci.2022.107741>.
- F. Buccino, A. Bunt, A. Lazell, L.M. Vergani, Mechanical Design Optimization of Prosthetic Hand's Fingers: Novel Solutions towards Weight Reduction, *Materials* 15 (2022) 2456, <https://doi.org/10.3390/ma15072456>.
- Wang Z, Hutter F, Zoghi M, Matheson D, Freitas N de. Bayesian Optimization in a Billion Dimensions via Random Embeddings 2016. <https://doi.org/10.48550/arXiv.1301.1942>.
- D. Eriksson, M. Pearce, J. Gardner, R.D. Turner, M. Poloczek, Scalable Global Optimization via Local Bayesian Optimization, in: *Advances in Neural Information Processing Systems*, 32, Curran Associates, Inc., 2019.
- Gelbart MA, Snoek J, Adams RP. Bayesian Optimization with Unknown Constraints 2014. <https://doi.org/10.48550/arXiv.1403.5607>.
- Hernández-Lobato JM, Gelbart MA, Adams RP, Hoffman MW, Ghahramani Z. A General Framework for Constrained Bayesian Optimization using Information-based Search 2016. <https://doi.org/10.48550/arXiv.1511.09422>.
- Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. arXivOrg 2019. <https://arxiv.org/abs/1912.01703v1> (accessed February 17, 2026).

- [32] Balandat M, Karrer B, Jiang DR, Daulton S, Letham B, Wilson AG, et al. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. arXivOrg 2019. <https://arxiv.org/abs/1910.06403v3> (accessed February 17, 2026).
- [33] Gardner J, Pleiss G, Bindel D, Weinberger KQ, Wilson A. GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration. ArXiv 2018.
- [34] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, et al., SciPy 1.0: fundamental algorithms for scientific computing in Python, Nat Methods 17 (2020) 261–272, <https://doi.org/10.1038/s41592-019-0686-2>.
- [35] C.R. Harris, K.J. Millman, S.J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, et al., Array programming with NumPy, Nature 585 (2020) 357–362, <https://doi.org/10.1038/s41586-020-2649-2>.
- [36] C.B. Sullivan, A.A. Kaszynski, PyVista: 3D plotting and mesh analysis through a streamlined interface for the Visualization Toolkit (VTK), Journal of Open Source Software 4 (2019) 1450, <https://doi.org/10.21105/joss.01450>.
- [37] X. Wang, X. Li, Z. Li, Z. Wang, W. Zhai, A ribbed strategy disrupts conventional metamaterial deformation mechanisms for superior energy absorption, Virtual and Physical Prototyping 19 (2024) e2337310, <https://doi.org/10.1080/17452759.2024.2337310>.
- [38] X. Wang, J. Deng, R. Zhao, L. Shen, Z. Wang, Compressive response and energy absorption of 3D-printed cellular structures: advances and challenges, Virtual and Physical Prototyping 21 (2026) e2638085, <https://doi.org/10.1080/17452759.2026.2638085>.
- [39] J. Deng, X. Wang, R. Zhao, Z. Li, K. Liu, K. Yuan, et al., Golden-Ratio-Guided Aperiodic Architected Metamaterials with Simultaneously Enhanced Strength and Toughness, Advanced Functional Materials 36 (2026) e16315, <https://doi.org/10.1002/adfm.202516315>.