Proceedings of the ASME 2026
International Design Engineering Technical Conferences and Computers and Information in Engineering Conference
IDETC/CIE2026
August 23–26, 2026, Houston, TX

IDETC2026-XXXX

SEMANTIC INVERSE DESIGN OF CELLULAR STRUCTURES IN IMPLICIT FIELD SPACE

Yunsheng Wang[1], Zu Yashengjiang[1], Bojing Yu[1], Shaoliang Yang[1], Jun Wang[1]

[1]Santa Clara University, Santa Clara, CA

ABSTRACT

*In cellular structures design, engineers often seek candidates that are "lightweight yet stiff". Since the parameter space is high-dimensional and difficult to navigate using intuition alone, mapping such intent to geometry remains challenging, and existing inverse design methods largely rely on numerical optimization or generative models. To bridge this gap, we introduce a semantic inverse design framework for cellular structures that maps engineering descriptors to quantitative targets in property space. Rather than directly generating geometry, we use a retrieval-and-verification strategy: it parses descriptor queries (*e.g.*, low stiffness, medium density), retrieves candidate geometries from implicit functions, and ranks them using a surrogate model trained on voxel-FEM homogenization labels. Across the evaluation prompt suite, using a fixed top-K retrieval setting with $K = 5$, the proposed framework achieves a joint Success@5 of 92.5%, defined as the fraction of queries for which at least one returned candidate satisfies both the stiffness target within tolerance and all hard constraints, together with a stiffness mean relative error (MRE) of 3.4% for the top-ranked result. A Gibson–Ashby scaling check further shows that candidates remain physically consistent over the explored density range.*

1. INTRODUCTION

Cellular structures have become an important class of engineered materials for achieving mechanical performance that are difficult to obtain with bulk solids alone, such as high stiffness-to-weight ratios, independently tunable effective properties, and enhanced energy absorption capability. By distributing material through cellular geometries, these structures can combine low mass with high stiffness, exhibit tailored compliance, and provide controlled energy absorption [1–3]. Recent developments in additive manufacturing have increased interest in these materials by facilitating the production of complex unit-cell geometries with consistent quality [4–6].

While analysis tools for cellular structures, such as voxel-based numerical homogenization, finite element analysis (FEA) [7], and analytical scaling laws (*e.g.*, Gibson–Ashby relations [8]), are well established and used to predict properties, the corresponding design problem remains far less mature. Effective properties can be predicted with high accuracy for a given unit-cell geometry using numerical homogenization techniques. In contrast, the inverse problem of identifying geometries that achieve prescribed targets, such as density or stiffness, remains fundamentally challenging. From an engineering design perspective, inverse problems are often ill-posed and admit multiple feasible solutions, which complicates systematic synthesis even when accurate forward models, such as finite element analysis (FEA) and homogenization-based property prediction, are available [9, 10]. The relationship between geometry and effective response is highly nonlinear, and the associated high-dimensional design space makes exploration computationally prohibitive [11].

This difficulty is further compounded by a disconnect between functionality requirement specification and parameterized geometric representations. In many CAD-based design frameworks, geometry is parameterized using low-level variables (*e.g.*, unit-cell thickness parameters, spline control points, level-set coefficients, or implicit function weights), which are computationally convenient but only loosely connected to high-level engineering intent (*e.g.*, target stiffness, mass constraints, natural frequency requirements, or energy absorption goals). As a result, designers often rely on iterative trial-and-error workflows to obtain satisfactory structures [12]. Related efforts, including our prior work, have explored alternative representations and design-space constructions to improve controllability and expressiveness [13, 14]. However, these approaches still require designers to manually tune geometric parameters (*e.g.*, implicit function coefficients or unit-cell parameters) rather than directly specifying desired performance targets such as stiffness or density.

Existing approaches to inverse design of cellular structures have largely followed two complementary directions. Topology optimization provides physics-based solutions and can generate mechanically efficient micro-structures, but it is commonly computationally intensive and often produces designs specialized to a particular loading or boundary condition. Data-driven approaches, including deep generative models, enable rapid candidate generation but can yield geometries that are difficult to interpret and may violate fundamental engineering or manufacturing constraints such as connectivity or minimum feature size. These limitations reduce their effectiveness as tools for guided design exploration [15, 16].

In this work, we propose a framework for semantic inverse design of cellular structures that emphasizes engineering validity, interpretability, and controlled exploration. Rather than directly generating geometry from unconstrained language input, the method maps qualitative engineering descriptors to quantitative property targets and constraints, then retrieves candidate unit cells from an analytic implicit-function library. The retrieved candidates are subsequently verified and ranked using a surrogate model trained on homogenization data [17]. By formulating inverse design as a retrieval-and-verification problem, the framework provides a more transparent and controllable alternative to unconstrained generative approaches, while still allowing users to specify design intent in an intuitive descriptor-based form.

The main contributions of this work are:

- We propose a semantic inverse design framework for cellular structures that grounds qualitative engineering intent into explicit quantitative targets and hard constraints.
- We introduce a retrieval-and-verification pipeline in implicit field space that retrieves candidate unit-cell geometries and ranks them using surrogate-assisted property prediction.
- We demonstrate that the proposed framework achieves strong target-matching performance and physically consistent behavior, highlighting its potential as an alternative to unconstrained text-to-geometry generation.

2. RELATED WORK

Our research draws upon developments in computational inverse design, geometric representation, and semantic interfaces for engineering applications. This section reviews representative work in these areas to position our retrieval-based framework relative to existing approaches.

2.1. Inverse Design of Cellular Structures

Designing cellular structures with custom mechanical properties involves searching through an enormous range of possibilities. Inverse design tackles this by reversing the standard workflow: instead of creating a structure to see how it performs, we start with the performance we want and let an algorithm determine the internal geometry that achieves it. To solve this layout problem, researchers typically rely on two main approaches: Topology Optimization and Data-Driven Generative Models.

2.1.1. Topology Optimization The inverse problem of identifying material distributions that optimize performance objectives subject to constraints has traditionally been addressed through topology optimization. Density-based formulations remain the dominant paradigm for generating mechanically efficient structures, owing to their strong theoretical grounding and general applicability [11, 18]. However, topology optimization typically requires repeated solutions of governing partial differential equations and relies on rigid, parameter-specific objective functions. This computational cost and reliance on numerical targets limit its flexibility for conceptual design exploration.

2.1.2. Data-Driven Generative Models To alleviate these computational burdens, recent research has increasingly explored data-driven alternatives. Machine learning models and deep generative approaches accelerate the synthesis of structures by learning mappings between design variables and precise performance metrics [19]. Generative models enable exploration within continuous latent spaces and can rapidly produce candidate designs conditioned on target numerical responses [20].

Even though these generative and optimization methods are highly efficient, they still rely heavily on exact numerical inputs to guide the design process. This creates a bottleneck during the early stages of design, where engineers usually think in terms of qualitative goals rather than precise mechanical properties. To bridge this gap, our work introduces a framework that moves away from strictly parameter-driven generation. Instead, we propose a semantic inverse design approach that supports descriptor-driven retrieval of structures.

2.2. Implicit Representations for Cellular Structures

Implicit geometric representations have become essential for designing complex cellular structures. Analytic implicit functions are particularly well-suited for modeling repeating microstructures, because they naturally define smooth surfaces and inherent continuity [21].

Recently, data-driven methods have applied neural networks to learn implicit representations directly from structural data, which allows for highly flexible reconstructions of complex cellular geometries [22, 23]. Despite their ability to capture fine architectural details, it remains challenging to control these learned neural implicit models to meet the exact engineering constraints required for cellular materials, such as densities, stiffness, or strict periodic boundary conditions.

To address this lack of control, our work relies on analytic implicit functions whose parameters admit direct physical interpretation. By grounding the geometry in analytic formulations, we enable a more controlled exploration of the cellular design space. This structured representation is critical for our objective of mapping semantic descriptions directly to cellular structures.

2.3. Semantic Interfaces for Computer-Aided Design

The growing complexity of computational design tools has motivated research into more intuitive human-computer interfaces, including those driven by natural language. Early efforts focused on translating text commands into predefined sequences of CAD operations. More recent advances in language modeling have enabled text-conditioned generation of three-dimensional geometry [24, 25].

While such systems can rapidly generate shapes that resemble engineered structures, they are typically trained on visual data and do not explicitly encode physical reasoning. As a result, generated geometries may satisfy aesthetic criteria while failing to meet mechanical performance requirements.

Our work departs from purely visual paradigms by treating language as a mechanism for querying physical behavior rather than appearance. Instead of generating geometry directly from text, semantic descriptions are mapped to quantitative performance constraints, and candidate designs are retrieved based on their simulated mechanical response.

3. METHODOLOGY

We propose a semantic inverse design paradigm in this work, shown in Figure 1. Given a descriptor-based design request, the framework first parses the semantic intent into structured targets and constraints, then retrieves multiple candidate unit cells from a verified dataset of implicit fields, and finally ranks them using fast property prediction validated against homogenization.

3.1. The Analytic Implicit Dataset

We construct a dataset as a rigorous foundation for retrieval. Unlike voxel approaches that operate in a discrete space, we define our unit cells using continuous implicit trigonometric functions [26], as shown in Table 1. Some representative samples of the generated geometries are shown in Figure 2. By varying the isovalue C (See Sec. 3.1.2) and applying periodicalization transforms (See Sec. 3.1.3), we generate a diverse range of configurations. Table 2 summarizes the key configuration parameters used to generate the dataset.

3.1.1. Mathematical Representation. Let $\Omega \subset \mathbb{R}^3$ be the domain of the unit cell, typically defined as $[-1, 1]^3$. The geometry is represented by a scalar field $f : \Omega \to \mathbb{R}$, where the solid phase is defined by the inequality $f(\mathbf{x}) \leq 0$. The zero level set $\partial S = \{x \in \Omega \mid f(x) = 0\}$ represents the solid–void interface. This implicit representation ensures that the surface is watertight and differentiable almost everywhere, which is critical for extracting high-quality triangular meshes via isosurface extraction algorithms [27]:

$$\mathcal{S} \triangleq \{\mathbf{x} \in \Omega \mid f(\mathbf{x}; \boldsymbol{\phi}) \leq 0\}, \partial\mathcal{S} = \{\mathbf{x} \in \Omega \mid f(\mathbf{x}; \boldsymbol{\phi}) = 0\}, \quad (1)$$

where $\boldsymbol{\phi}$ denotes the parameters of the implicit function family.

3.1.2. Function Families and Connectivity. The dataset is built upon classic Triply Periodic Minimal Surfaces (TPMS) and generalized trigonometric families [28, 29]. A generalized parametric form for these families involves a weighted sum of frequency components modulated by crystallographic vectors, along with an isovalue threshold C. The parameter C serves as a control knob for the relative density $\bar{\rho}$. For the analytic families considered here, restricting the isovalue range empirically preserves open and connected topologies for the retained samples, reducing the occurrence of isolated clusters that often appear in unconstrained generative outputs. Table 1 details the specific analytic expressions used in our library.

3.1.3. Coordinate Distortion for Geometric Variation. To expand the geometric diversity of the dataset and simulate functionally graded architectures, we introduce a deterministic spatial distortion parameterized by a randomly sampled exponent r. We define a spatial gradient function $p(x, y, z) = x + y + z$ along the principal diagonal of the unit cell and normalize it to $\bar{p} \in [0, 1]$ over the sampling grid. We then construct a non-linear scaling field:

$$\mathrm{PER}(x, y, z) = \tfrac{1}{2} + \tfrac{1}{2}\left(1 - \bar{p}(x, y, z)\right)^r, r \sim \mathcal{U}(r_{\min}, r_{\max}), \quad (2)$$

and apply the coordinate transformation $(x', y', z') = (x, y, z)/\mathrm{PER}(x, y, z)$. While this directional warping significantly enriches the morphological variety of the implicit fields, it can inherently disrupt the exact periodicity at the cell boundaries. To mitigate this, we compute and record a boundary consistency score (defined as the scalar difference between opposing faces of the SDF volume) for each generated sample. This score serves as a vital metadata metric for downstream constraint-based filtering (Sec. 3.3), ensuring only geometrically valid and physically tileable candidates are retrieved.

3.1.4. Metadata and Normalization For each generated sample, we compute and store a metadata vector containing the relative density and a boundary consistency score that measures the difference between opposing faces of the SDF volume. The analytic trigonometric functions define the zero-level set geometry, but they do not naturally constitute true Signed Distance Fields (SDFs) because their gradient magnitudes vary significantly across the domain. If fed directly into a neural network, these varying gradients would obscure the true geometric thickness of the structural features. To ensure consistent feature representation for the 3D CNN surrogate model, all implicit fields are mathematically normalized [30]. Specifically, we re-scale the fields such that the gradient magnitude is approximately unity ($||\nabla f|| \approx 1$) near the solid-void interface. This guarantees that the scalar values in the discretized volume strictly correspond to the true Euclidean distance to the nearest surface, providing a standardized geometric input for reliable retrieval and training.

3.2. Semantic Interface

3.2.1. Request schema. A user query is mapped to a structured request containing target properties such as effective stiffness tensor components, a set of constraints, and the number of requested candidates.

3.2.2. Parsing and Grounding. We employ a rule-based parsing module to map qualitative engineering descriptors to numeric targets and constraint bins derived from dataset statistics and the evaluation prompt suite. In this work, we ground stiffness targets via a normalized effective stiffness component E_x^* (defined as the ratio of the effective stiffness to the solid base material stiffness). We also enforce hard constraints on the relative density ρ and periodicity score s_{per} (Table 3). Extensions to multi-component stiffness targets and anisotropy constraints are left for future work.

We implement semantic grounding by discretizing linguistic modifiers (*e.g.*, "low/medium/high density" and "weak/moderate/strong periodicity") into numeric constraint bins or thresholds, while mapping stiffness descriptors (*e.g.*, "low/high stiffness") to representative normalized target values derived from the dataset distribution (Table 3). Each query specifies a normalized effective stiffness target $E_x^\star$ together with hard constraints on relative density ρ and periodicity score s_{per}, where lower s_{per} values indicate better boundary consistency and therefore stronger periodicity. For each query, the system returns the top-K candidates ($K = 5$ in our evaluation), and retrieval success is evaluated under a relative tolerance τ on the stiffness target while enforcing all hard constraints. The density intervals

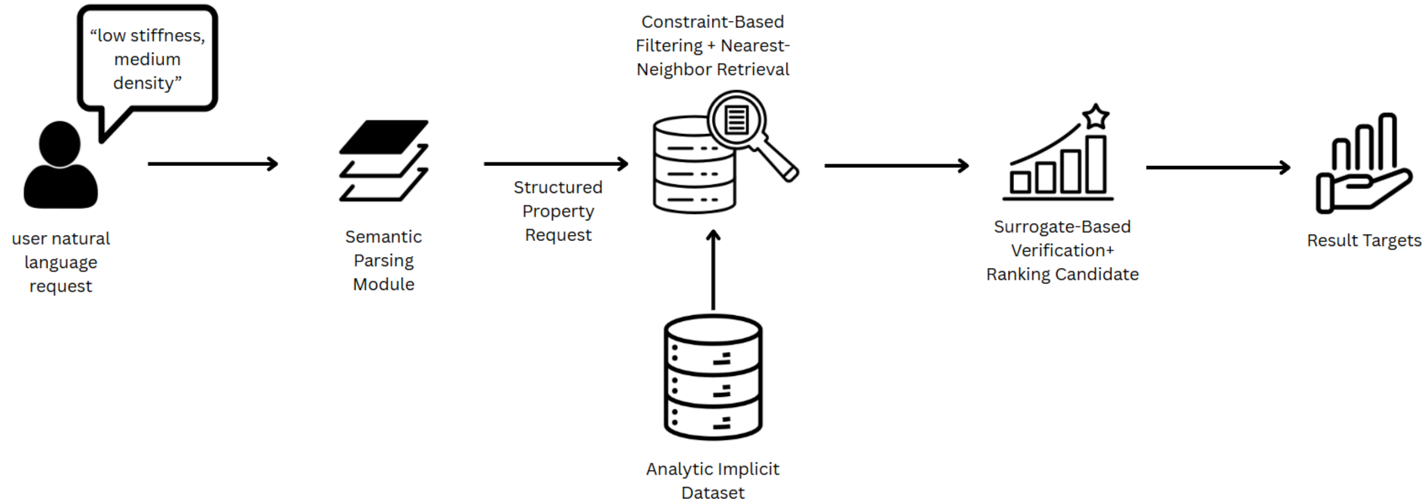

FIGURE 1: Flowchart of the semantic design framework. A descriptor-based user request is first translated into property targets and physical constraints via a semantic parsing module. Candidate unit-cell geometries are then retrieved from an implicit dataset using constraint-based filtering and nearest-neighbor search, followed by surrogate-based verification and ranking to yield the final matched designs.

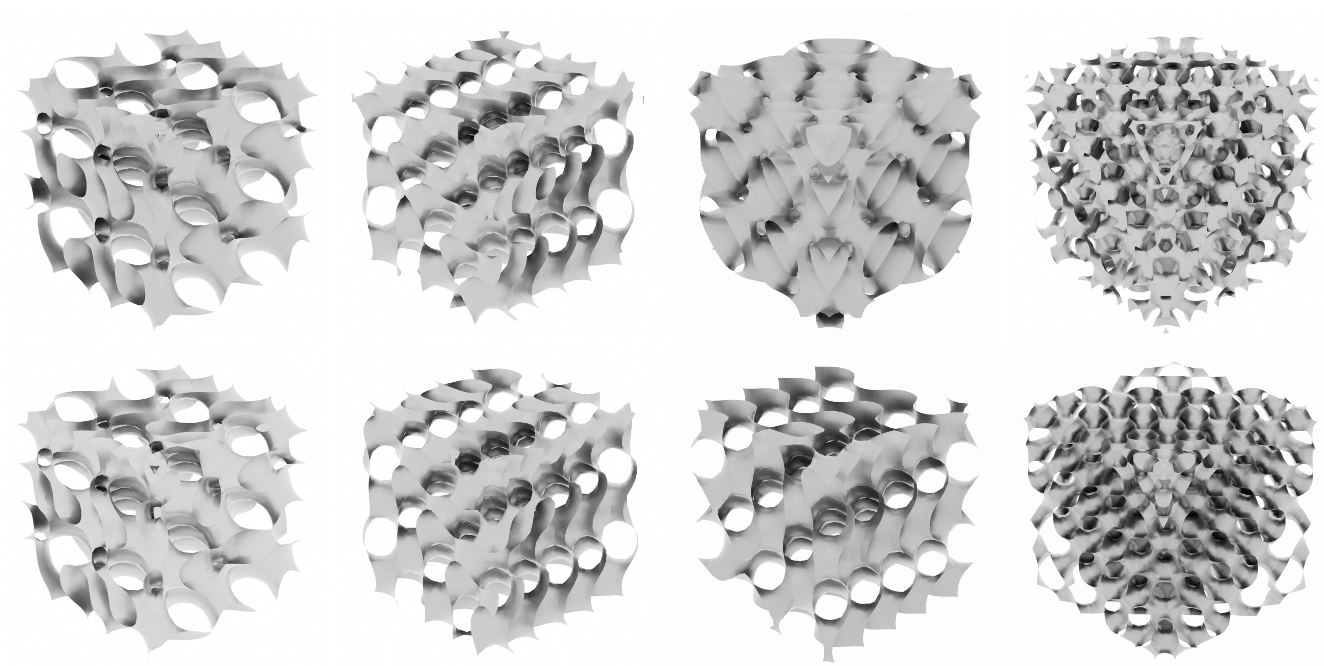

FIGURE 2: Representative samples from our dataset. These geometries demonstrate the diversity of the generated database, ranging from simple unit cells to complex periodic lattices derived from the functions in Table 1.

TABLE 1: Library of analytic implicit functions used for dataset generation. These equations define the implicit level-set surface. The parameter C is the isovalue threshold that controls the relative density, where the solid phase is enclosed by $f(X, Y, Z) - C \leq 0$. To ensure periodicity, we define the arguments as $X = 2\pi x$, $Y = 2\pi y$, and $Z = 2\pi z$.

ID / Name	Implicit Field Equation $f(X,Y,Z) - C$
Trig_f1	$\cos(X) + \cos(Y) + \cos(Z) - C$
Trig_f33	$\cos(X+Y-Z) + \cos(X-Y+Z) + \cos(X-Y-Z) - C$
Trig_f23I	$\cos(X+Y) + \cos(X-Y) + \cos(X+Z) - C$
Trig_f23II	$\cos(X+Y) + \cos(X+Z) + \cos(Y+Z) - C$
Trig_f24II	$\cos(X+Y) + \cos(X-Y) + \cos(X+Z) - \cos(X-Z) - C$
Trig_f25II	$\cos(X-Y) + \cos(X+Z) + \cos(X-Z) + \cos(Y+Z) - \cos(Y-Z) - C$
Trig_f25IV	$\sin(X-Y) + \sin(X+Z) + \sin(X-Z) + \sin(Y+Z) - \sin(Y-Z) - C$
Trig_f34II	$\cos(X+Y+Z) + \cos(X+Y-Z) + \cos(X-Y+Z) - \cos(X-Y-Z) - C$
Trig_f1_f33	$0.7 \cdot f_{\text{Trig\_f1}} + 0.3 \cdot f_{\text{Trig\_f33}} - C$
Trig_f34I_f86I	$0.6\left(\sum \cos(X \pm Y \pm Z)\right) + 0.4\left(\cos(2X)\cos(2Y) + \cos(2Y)\cos(2Z) + \cos(2Z)\cos(2X)\right) - C$

TABLE 2: Configuration of the generated implicit-field dataset. The parameters are derived directly from the generation script settings.

Parameter	Value / Setting
Total Samples (N)	20,000
Grid Resolution	$32 \times 32 \times 32$
Split Ratio	70% Train / 15% Val / 15% Test
Function Family	Original Trigonometric Set (10 types)
Periodicalization	Enabled (Exponent $r \in [0.3, 0.9]$)
Isovalue (C)	Randomized $\in [0.01, 0.5]$
Domain	$[-1.0, 1.0]^3$

TABLE 3: Semantic grounding of qualitative descriptors into quantitative targets and constraints used in the prompt suite. Stiffness targets are mapped to normalized effective values derived from the dataset distribution.

Descriptor	Property	Quantitative grounding
low stiffness	target E_x	$E_x^* = 0.1980$ (Example)
high stiffness	target E_x	$E_x^* = 0.4372$ (Example)
low density	constraint ρ	$\rho \in [0.00, 0.25]$
medium density	constraint ρ	$\rho \in [0.25, 0.60]$
high density	constraint ρ	$\rho \in [0.60, 1.00]$
weak periodicity	constraint s_{per}	$s_{per} > 0.0017$
moderate periodicity	constraint s_{per}	$0.0010 < s_{per} \leq 0.0017$
strong periodicity	constraint s_{per}	$s_{per} \leq 0.0010$

are chosen empirically to provide a practical semantic partition of the sampled design space and to remain consistent with the representative query targets used in the qualitative case studies.

3.3. Reliable Retrieval and Verification

Our core contribution is to replace stochastic generation with a robust retrieval-based pipeline, as detailed in Algorithm 1. Let each candidate i be represented by an analytic implicit field $f(\cdot; \phi_i)$ and a (normalized) property vector $\tilde{\mathbf{p}}_i$ stored as metadata (*e.g.*, ρ and E_x). Given a target $\tilde{\mathbf{p}}^\star$ and hard constraints $C(\cdot)$, retrieval is defined as

$$\hat{i} = \arg\min_{i \in \mathcal{F}} \left\|\tilde{\mathbf{p}}_i - \tilde{\mathbf{p}}^\star\right\|_2, \quad \mathcal{F} = \{i \mid C(f(\cdot; \phi_i)) = \text{true}\}. \tag{3}$$

3.3.1. Filtering by Constraints. The search space is first reduced to a feasible set by applying hard constraints. We enforce geometric validity through explicit constraints, thereby filtering out disconnected or boundary-inconsistent candidates from the dataset before retrieval.

3.3.2. Nearest Neighbor Retrieval. Within the feasible set $\mathcal{F}$, we perform nearest-neighbor search in the normalized property space using Euclidean distance. Candidates are scanned in ascending distance to the target (Line 7, Algorithm 1) and accepted if they satisfy the surrogate-based target tolerance (Lines 9–17). This yields multiple valid options when different analytic families match the same intent.

3.4. Surrogate-Based Verification

Evaluating the properties of retrieved candidates with the Finite Element Method (FEM) is computationally prohibitive. Therefore, we integrate a physics-informed surrogate model into the retrieval loop [31] to enable real-time interaction.

Algorithm 1 Retrieval and Verification Pipeline

Require: Dataset $\mathcal{D} = \{(f(\cdot;\boldsymbol{\phi}_i),\ \mathbf{p}_i)\}_{i=1}^{N}$, target $\mathbf{p}^\star$, constraints $C(\cdot)$, desired count K
Ensure: Candidate set $\mathcal{R}$

$\mathcal{F} \leftarrow \emptyset$
for $i = 1$ to N **do**
  if $C(f(\cdot;\boldsymbol{\phi}_i)) = \texttt{true}$ **then**
    $\mathcal{F} \leftarrow \mathcal{F} \cup \{i\}$
  end if
end for
Sort $\mathcal{F}$ by $d_i = \|\mathbf{p}_i - \mathbf{p}^\star\|_2$ ascending
$\mathcal{R} \leftarrow \emptyset$
for each i in sorted $\mathcal{F}$ **do**
  if $|\mathcal{R}| = K$ **then**
    break
  end if
  $\tilde{E}_{x,i} \leftarrow \textsc{SurrogatePredict}(f(\cdot;\phi_i))$
  if $\left|\tilde{E}_{x,i} - E_x^\star\right| / E_x^\star \le \tau$ **then**
    $\mathcal{R} \leftarrow \mathcal{R} \cup \{i\}$
  end if
end for
return $\mathcal{R}$

3.4.1. Architecture. We employ a 3D Convolutional Neural Network (CNN) designed for volumetric regression. The network takes the $32 \times 32 \times 32$ discretized SDF volume as a single-channel input tensor. Feature extraction is performed through four sequential 3D convolutional blocks. Each block applies a 3D convolution (kernel size $3 \times 3 \times 3$, stride 1, padding 1) with an increasing number of feature channels (16, 32, 64, and 128, respectively), followed by Batch Normalization and a ReLU activation function. To progressively downsample the spatial dimensions, $2 \times 2 \times 2$ Max Pooling layers (stride 2) are inserted immediately after the second and fourth convolutional blocks. The resulting $8 \times 8 \times 8 \times 128$ spatial feature map is then reduced to a 128-dimensional latent vector via a Global Average Pooling (GAP) layer. Finally, a Multi-Layer Perceptron (MLP) consisting of two hidden layers (64 and 32 neurons, respectively, with ReLU activations) processes this vector to linearly regress the normalized effective stiffness E_x. The detailed layer configurations are summarized in Table 4.

3.4.2. Physics-Informed Training. The network is trained on a subset of the analytic dataset labeled with high-fidelity homogenization. To ensure the model respects physical laws, we augment the standard Mean Squared Error loss with a penalty term inspired by Physics-Informed Neural Networks (PINNs) [32]. The composite loss function $\mathcal{L}_{\text{total}}$ is defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{MSE}} + \lambda \mathcal{L}_{\text{phys}}, \tag{4}$$

where $\mathcal{L}_{\text{MSE}}$ is the mean squared error on the predicted stiffness E_x^*. The penalty term $\mathcal{L}_{\text{phys}}$ explicitly encodes the Gibson–Ashby scaling laws for cellular solids. By defining $E_{\text{low}}(\rho)$ and $E_{\text{up}}(\rho)$ as the theoretical bending-dominated and stretching-dominated bounds for a given relative density ρ, the physics loss is formulated as:

$$\begin{aligned}\mathcal{L}_{\text{phys}} = \frac{1}{N}\sum_{i=1}^{N}\Big(&\max\big(0, E_{\text{low}}(\rho_i) - \tilde{E}_{x,i}^{\text{pred}}\big)^2 \\ &+ \max\big(0, \tilde{E}_{x,i}^{\text{pred}} - E_{\text{up}}(\rho_i)\big)^2\Big).\end{aligned} \tag{5}$$

This formulation discourages the network from predicting physically impossible stiffness values that fall outside the theoretical envelope, thereby helping the verification module remain reliable over unseen or sparse regions of the design space.

4. EXPERIMENTS

In this section, we evaluate the framework's ability to satisfy users' semantic intent.

4.1. Evaluation Protocol

To evaluate whether the proposed semantic inverse design pipeline can return candidates that satisfy user intent, we construct a prompt suite based on combinations of qualitative descriptors for stiffness and density. Each prompt is grounded into a structured request containing (i) a target axial stiffness value $E_x^\star$, (ii) a density constraint ρ, and (iii) a periodicity constraint s_{per}, according to the semantic bins defined in Table 3. For every query, the system retrieves and ranks the top-K candidates, with $K = 5$ throughout this study.

4.2. Experimental Setup

4.2.1. Voxel-Based FEM Homogenization. We generate ground-truth labels for effective elastic properties using voxel-based computational homogenization, following the procedure outlined in Algorithm 2. The normalized SDF, $\Phi(\mathbf{x}) = f(\mathbf{X}) - C$, is thresholded at the zero-level set to create a binary voxel model $\chi(\mathbf{x}) \in \{0, 1\}$, where $\chi(\mathbf{x}) = 1$ if $\Phi(\mathbf{x}) \le 0$ and 0 otherwise. Periodic boundary conditions (PBCs) are applied to the displacement field to ensure representative volume element (RVE) consistency.

4.2.2. Evaluation Metrics. We use a template-based prompt suite covering diverse combinations of stiffness, density, and structural requirements. Three evaluation metrics are used to assess the efficacy of the proposed framework:

- **Success Rate:** the percentage of queries for which at least one returned candidate among the top-K results satisfies the requested stiffness target within a relative tolerance τ while also satisfying all hard constraints;
- **Feasibility Rate:** the percentage of returned candidates that satisfy all hard constraints specified by the query;
- **Physical Consistency:** the degree to which retrieved candidates remain broadly consistent with Gibson–Ashby[1] style stiffness–density trends over the explored density range.

TABLE 4: Detailed architecture and hyperparameters of the 3D CNN surrogate model.

Layer Type	Hyperparameters	Activation	Output Shape
Input	Single-channel SDF	-	$32 \times 32 \times 32 \times 1$
Conv3D Block 1	$3 \times 3 \times 3$, 16 channels, stride 1	ReLU	$32 \times 32 \times 32 \times 16$
Conv3D Block 2	$3 \times 3 \times 3$, 32 channels, stride 1	ReLU	$32 \times 32 \times 32 \times 32$
MaxPool3D 1	$2 \times 2 \times 2$, stride 2	-	$16 \times 16 \times 16 \times 32$
Conv3D Block 3	$3 \times 3 \times 3$, 64 channels, stride 1	ReLU	$16 \times 16 \times 16 \times 64$
Conv3D Block 4	$3 \times 3 \times 3$, 128 channels, stride 1	ReLU	$16 \times 16 \times 16 \times 128$
MaxPool3D 2	$2 \times 2 \times 2$, stride 2	-	$8 \times 8 \times 8 \times 128$
Global Avg Pool	-	-	$1 \times 1 \times 1 \times 128$
Flatten	-	-	128
MLP Dense 1	64 neurons	ReLU	64
MLP Dense 2	32 neurons	ReLU	32
Output	1 neuron	Linear	1 (Predicted E_x)

Algorithm 2 Voxel FEM Homogenization (Periodic Unit Cell)

Require: Binary voxel grid $\chi(\mathbf{x}) \in \{0, 1\}$, base material $\mathbb{C}_0$, set of macro-strain load cases $\{\bar{\boldsymbol{\varepsilon}}^{(m)}\}_{m=1}^{M}$

Ensure: Effective stiffness $\mathbb{C}^*$

Build voxel mesh with DOFs at grid nodes; assign $\mathbb{C}(\mathbf{x}) = \chi(\mathbf{x})\mathbb{C}_0$

Impose periodic boundary conditions on opposite faces (displacement fluctuations periodic; tractions anti-periodic)

for $m = 1$ to M **do**

  Apply macro strain $\bar{\boldsymbol{\varepsilon}}^{(m)}$ via affine displacement $u^{(m)}(\mathbf{x}) = \bar{\boldsymbol{\varepsilon}}^{(m)}\mathbf{x} + \tilde{u}^{(m)}(\mathbf{x})$

  Assemble linear system $\mathbf{K}\mathbf{u}^{(m)} = \mathbf{f}^{(m)}$

  Solve for $\mathbf{u}^{(m)}$ using Conjugate Gradient with tolerance ϵ

  Compute voxel stresses $\boldsymbol{\sigma}^{(m)}(\mathbf{x}) = \mathbb{C}(\mathbf{x}) : \boldsymbol{\varepsilon}^{(m)}(\mathbf{x})$

  Volume-average $\bar{\boldsymbol{\sigma}}^{(m)} = \langle \boldsymbol{\sigma}^{(m)}(\mathbf{x}) \rangle_\Omega$

end for

Fit $\mathbb{C}^*$ from $\bar{\boldsymbol{\sigma}}^{(m)} = \mathbb{C}^* : \bar{\boldsymbol{\varepsilon}}^{(m)}$ for all m

return $\mathbb{C}^*$

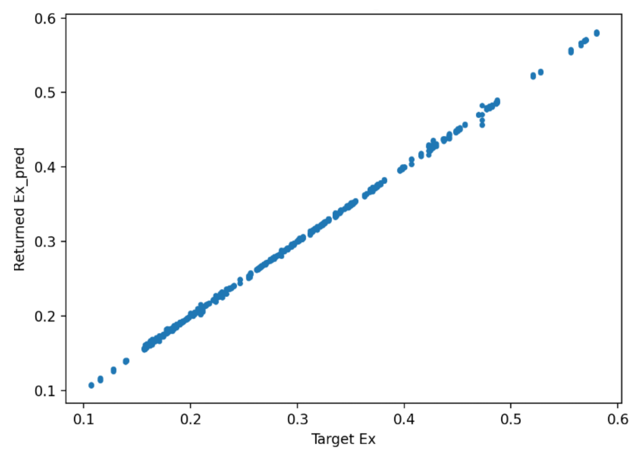

FIGURE 3: Global quantitative results for target matching in stiffness. Each marker corresponds to one test query, plotting the requested target stiffness E_x against the stiffness of the returned design. The near-diagonal trend indicates strong agreement between requested and retrieved stiffness values, with goodness of fit quantified by the coefficient of determination R^2.

4.3. Main Quantitative Results

4.3.1. Target Matching Success. Figure 3 summarizes target-matching performance for stiffness, with the goodness of fit quantified by the coefficient of determination R^2. Each marker corresponds to one evaluation query, where the requested target stiffness E_x is plotted against the stiffness of the returned design predicted by the surrogate model used for verification and ranking. The near-diagonal alignment indicates that the retrieval-and-verification pipeline returns candidates whose predicted stiffness remains close to the requested target. Under a relative tolerance of $\tau = 0.1$, a query is counted as successful if at least one returned candidate among the top-K results satisfies

$$\frac{\left|E^{\text{ret}}_{x,i,k} - E^{\text{tar}}_{x,i}\right|}{E^{\text{tar}}_{x,i}} \le \tau, \qquad \tau = 0.1, \tag{6}$$

while also satisfying all hard constraints, where $E^{\text{tar}}_{x,i}$ denotes the target stiffness for query i and $E^{\text{ret}}_{x,i,k}$ denotes the stiffness of the kth returned candidate for that query. Accordingly, the Success Rate is defined as

$$\text{Success Rate} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\Big(\exists k \in \{1, \ldots, K\} : \frac{|E^{\text{ret}}_{x,i,k} - E^{\text{tar}}_{x,i}|}{E^{\text{tar}}_{x,i}} \le \tau, \; C_{i,k} = 1\Big) \times 100\%. \tag{7}$$

[1] Gibson–Ashby scaling refers to the classical power-law relationship between the effective mechanical properties of a cellular material and its relative density, and is widely used for evaluating structural plausibility.

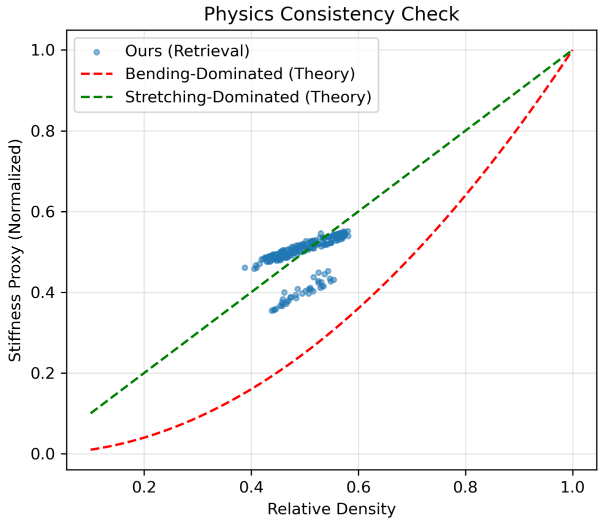

FIGURE 4: Physical consistency check against Gibson–Ashby scaling. Retrieved candidates (blue circles) do not exactly follow either theoretical trend, but instead form an inclined band between the bending-dominated and stretching-dominated regimes (red and green dashed lines), indicating broadly plausible stiffness–density behavior over the explored density range.

where $\mathbf{1}(\cdot)$ is the indicator function. The stiffness mean relative error (MRE) is defined as

$$\mathrm{MRE} = \frac{1}{N}\sum_{i=1}^{N} \min_{k\in\{1,\ldots,K\}} \frac{\left|E^{\mathrm{ret}}_{x,i,k} - E^{\mathrm{tar}}_{x,i}\right|}{E^{\mathrm{tar}}_{x,i}} \times 100\%. \quad (8)$$

where the MRE is reported for the top-1 returned candidate. Under these definitions, the framework achieves a Success Rate of 92.5% and a stiffness MRE of 3.4% across the evaluation suite. These results suggest that the proposed semantic grounding and retrieval strategy is effective for descriptor-driven exploration in property space.

4.3.2. Physical Consistency Check. A key validation step in this work is the physical consistency check shown in Figure 4. We plot normalized effective stiffness against relative density for a representative set of retrieved candidates and compare their distribution against Gibson–Ashby-style scaling trends for cellular structures, where the linear upper trend represents stretching-dominated behavior and the quadratic lower trend represents bending-dominated behavior [33]. As shown in Figure 4, the retrieved candidates do not lie exactly on either theoretical curve, but instead form an inclined band between these two idealized regimes over the explored density range. We interpret this as a physics sanity check, indicating that the returned designs remain broadly consistent with established stiffness–density behavior, rather than as a proof of optimality or a rediscovery of the governing law. The observed spread at fixed density further suggests that the framework can return geometrically distinct candidates with different structural efficiencies, thereby offering multiple feasible design options to the user.

4.4. Qualitative Case Studies

Table 5 illustrates representative outputs of the system for descriptor-based design queries. For example, in response to a query requesting high stiffness under a prescribed medium-density range, the system retrieves multiple distinct topologies with relative densities around 0.57 and normalized stiffness around 0.33. Despite satisfying similar performance targets, these structures exhibit significant geometric diversity, where some resemble truss-like networks and others resemble shell-based TPMS geometries. Because all candidates are drawn from an analytic implicit-function library, the retrieved structures are well-defined and more suitable for downstream meshing and fabrication workflows than unconstrained voxel-native outputs.

5. DISCUSSION

The results presented above demonstrate the efficacy of the semantic inverse design framework. In this section, we discuss the implications of these findings, the computational advantages over traditional methods, and current limitations.

5.1. Computational Efficiency vs. Accuracy

A primary motivation of this work is to reduce the time required to reach a valid candidate design. Classical topology optimization for a 3D unit cell often requires many repeated FEM evaluations, leading to substantial online computational cost [34]. In contrast, the proposed framework shifts most of the cost to the offline stage, where the dataset is generated and the surrogate model is trained. Once this preprocessing is complete, the online retrieval phase is efficient, with an average inference latency of approximately 1.02 seconds per query. This makes the method attractive for interactive design exploration, where rapid feedback is more important than exact solver-in-the-loop optimality at every step.

5.2. Physical Consistency and Manufacturability

The consistency of the retrieved candidates with Gibson–Ashby-style bounds provides additional evidence that the framework remains physically grounded over the explored density range. Compared with unconstrained voxel or point-cloud generation, restricting the search space to analytic implicit functions introduces a strong geometric prior that biases the returned structures toward smooth and connected forms. In addition, the implicit representation is convenient for downstream meshing and fabrication-oriented processing, reducing the need for extensive geometric repair after retrieval.

5.3. Limitations and Future Work

Despite the strong target-matching performance, the current framework has several limitations. First, retrieval quality depends on the coverage of the precomputed analytic library; when a requested property combination lies far from the sampled distribution, the returned candidates may be less optimal or less diverse. Second, the current interface relies on structured descriptor grounding rather than fully open-ended language understanding, and no large language model (LLM) is currently used to interpret more complex or ambiguous natural-language intent. Third, the present study focuses primarily on normalized stiffness, density,

TABLE 5: Dialogue-style qualitative results. Each case shows one dialogue row followed by the top-K visual outputs (K=5).

Prompt 0: “low stiffness, medium density”. **Targets:** $E_x^\star$=0.1980, $\rho^\star$=0.5809, $s_{\mathrm{per}}^\star$=0.00160. **Top-1:** id=11403, E_x=0.1981, ρ=0.5811, s_{per}=0.0013.

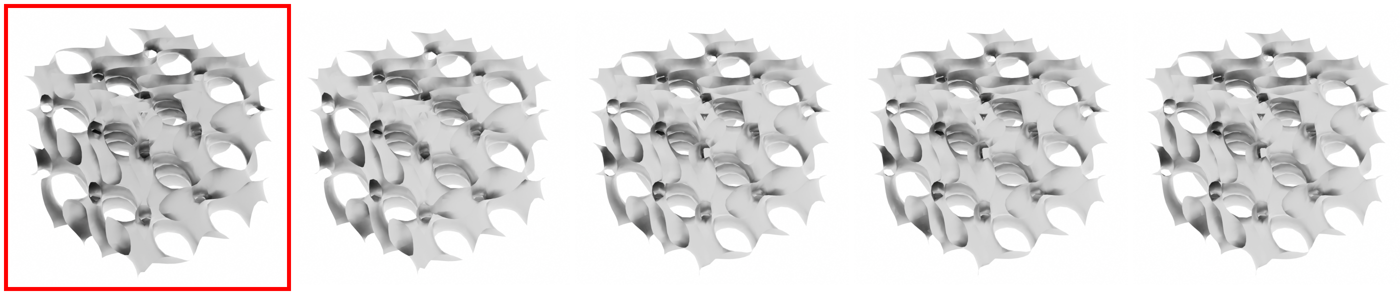

Prompt 1: “high stiffness, high density”. **Targets:** $E_x^\star$=0.4372, $\rho^\star$=0.6642, $s_{\mathrm{per}}^\star$=0.00094. **Top-1:** id=9977, E_x=0.4370, ρ=0.6631, s_{per}=0.0006.

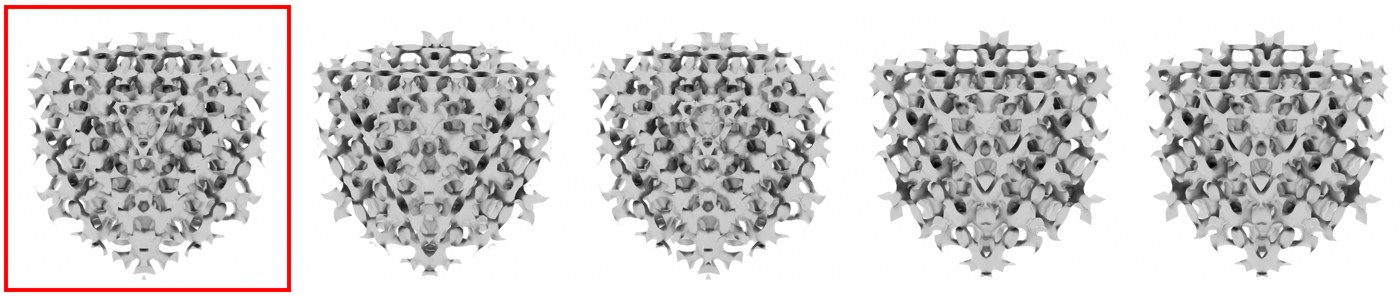

Prompt 2: “high stiffness, medium density”. **Targets:** $E_x^\star$=0.3364, $\rho^\star$=0.5736, $s_{\mathrm{per}}^\star$=0.00192. **Top-1:** id=9654, E_x=0.3368, ρ=0.5745, s_{per}=0.0019.

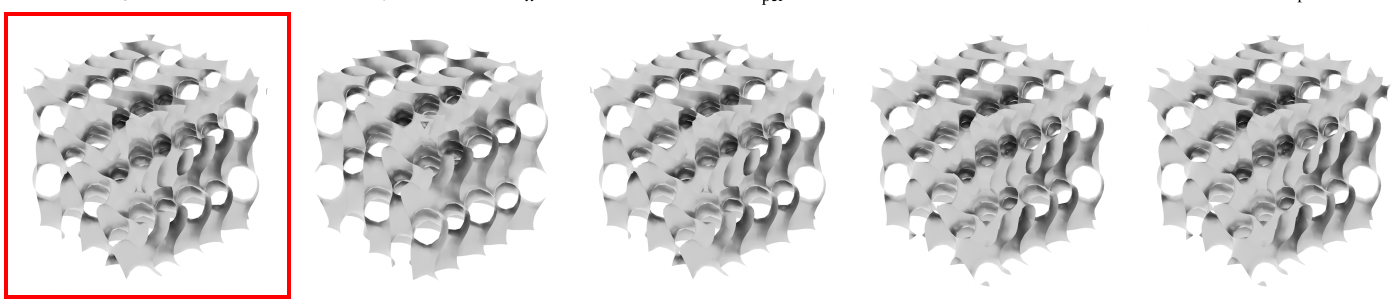

and basic structural descriptors, and does not yet address broader multi-property or multi-physics objectives.

Future work will expand the analytic library, incorporate local refinement after retrieval, and support richer interactive workflows for iterative engineering design. In particular, LLM-based conversational interfaces may help translate user intent into design constraints, especially when combined with multimodal interaction and richer downstream property models. The conceptual interface evolution is shown in Figure 5.

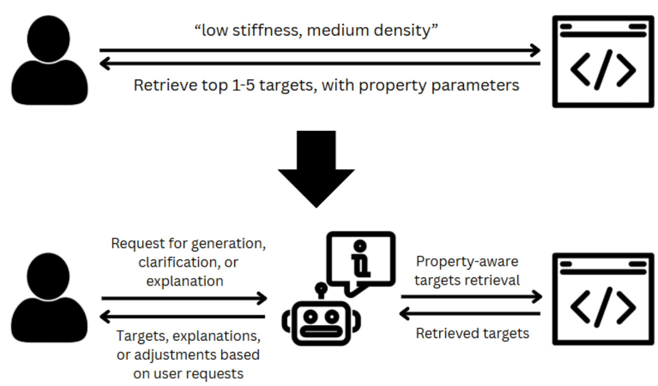

FIGURE 5: Conceptual interface evolution. The current system performs single-turn retrieval from semantic queries (top), while a future extension may support interactive dialogue-based design refinement using LLM (bottom).

6. CONCLUSION

We introduced a semantic inverse design framework for cellular structures that supports descriptor-driven exploration in property space. Rather than directly generating geometry from language input, the proposed method maps qualitative descriptors to targets and constraints, retrieves candidate unit cells from an implicit-function library, and verifies their performance using a surrogate model trained on homogenization data. This retrieval-and-verification formulation provides a more controllable alternative to unconstrained generative design workflows.

Across the evaluation suite, the framework demonstrates strong target-matching performance, including a 92.5% query-level success rate and a 3.4% stiffness mean relative error, while remaining consistent with Gibson–Ashby-style scaling behavior over the explored density range. These results suggest that the proposed framework is a promising step toward practical inverse design of cellular structures. Future work will focus on richer property objectives, expanded analytic libraries, and tighter integration of retrieval with local refinement and interactive design workflows.

REFERENCES

[1] Yeo, Seon Ju, Oh, Min Jun and Yoo, Pil J. "Structurally controlled cellular architectures for high-performance ultra-lightweight materials." *Advanced Materials* Vol. 31 No. 34 (2019): p. 1803670.

[2] Costanza, Girolamo, Solaiyappan, Dinesh and Tata, Maria Elisa. "Properties, applications and recent developments of cellular solid materials: A review." *Materials* Vol. 16 No. 22 (2023): p. 7076.

[3] Iqbal, Safdar and Kamiński, Marcin. "Review study on mechanical properties of cellular materials." *Materials* Vol. 17 No. 11 (2024): p. 2682.

[4] Nazir, Aamer, Abate, Kalayu Mekonen, Kumar, Ajeet and Jeng, Jeng-Ywan. "A state-of-the-art review on types, design, optimization, and additive manufacturing of cellular structures." *The International Journal of Advanced Manufacturing Technology* Vol. 104 No. 9 (2019): pp. 3489–3510.

[5] Opgenoord, Max MJ and Willcox, Karen E. "Design for additive manufacturing: cellular structures in early-stage aerospace design." *Structural and Multidisciplinary Optimization* Vol. 60 No. 2 (2019): pp. 411–428.

[6] Diani, Andrea, Nonino, Carlo and Rossetto, Luisa. "Melting of phase change materials inside periodic cellular structures fabricated by additive manufacturing: experimental results and numerical simulations." *Applied Thermal Engineering* Vol. 215 (2022): p. 118969.

[7] Schneider, Matti. "Voxel-based finite elements with hourglass control in fast Fourier transform-based computational homogenization." *International Journal for Numerical Methods in Engineering* Vol. 123 No. 24 (2022): pp. 6286–6313.

[8] Zhong, Haozhang, Song, Tingting, Li, Chuanwei, Das, Raj, Gu, Jianfeng and Qian, Ma. "The Gibson-Ashby model for additively manufactured metal lattice materials: Its theoretical basis, limitations and new insights from remedies." *Current Opinion in Solid State and Materials Science* Vol. 27 No. 3 (2023): p. 101081.

[9] Rosen, David W. "Computer-aided design for additive manufacturing of cellular structures." *Computer-Aided Design and Applications* Vol. 4 No. 5 (2007): pp. 585–594.

[10] Gallet, Adrien, Rigby, Samuel, Tallman, TN, Kong, Xiangxiong, Hajirasouliha, Iman, Liew, Andrew, Liu, Dong, Chen, Liang, Hauptmann, Andreas and Smyl, Danny. "Structural engineering from an inverse problems perspective." *Proceedings of the Royal Society A* Vol. 478 No. 2257 (2022): p. 20210526.

[11] Sigmund, Ole and Maute, Kurt. "Topology optimization approaches: A comparative review." *Structural and multidisciplinary optimization* Vol. 48 No. 6 (2013): pp. 1031–1055.

[12] Kyratzi, Sofia and Azariadis, Philip. "A Constraint-based Framework to Recognize Design Intent during Sketching in Parametric Environments." *Computer-Aided Design & Applications* Vol. 18 No. 3 (2021).

[13] Wang, Yunsheng, Chen, Songhao and Jin, Kevin. "Prometheus Chatbot: Knowledge Graph Collaborative Large Language Model for Computer Components Recommendation." *arXiv preprint arXiv:2407.19643* (2024).

[14] Habibi, Milad, Wang, Jun and Fuge, Mark. "When is it actually worth learning inverse design?" *Journal of Mechanical Design* Vol. 148 No. 6 (2026): p. 061704.

[15] Van't Sant, Sikko, Thakolkaran, Prakash, Martínez, Jonàs and Kumar, Siddhant. "Inverse-designed growth-based cellular metamaterials." *Mechanics of Materials* Vol. 182 (2023): p. 104668.

[16] Nakarmi, Sushan, Daphalapurkar, Nitin P, Lee, Kwan-Soo, Kim, Jihyeon, Leiding, Jeffery A, Dattelbaum, Dana M and Luscher, Darby J. "Inverse design of cellular structures with the targeted nonlinear mechanical response." *Scientific Reports* (2025).

[17] Wang, Jun and Rai, Rahul. "Generative design of conformal cubic periodic cellular structures using a surrogate model-based optimisation scheme." *International Journal of Production Research* Vol. 60 No. 5 (2022): pp. 1458–1477.

[18] Shin, Seungyeon, Shin, Dongju and Kang, Namwoo. "Topology optimization via machine learning and deep learning: a review." *Journal of Computational Design and Engineering* Vol. 10 No. 4 (2023): pp. 1736–1766.

[19] Lu, Yifu and Tong, Liyong. "Concurrent topology optimization of cellular structures and anisotropic materials." *Computers & Structures* Vol. 255 (2021): p. 106624.

[20] Yang, Shaoliang, Wang, Jun and Wang, Kang. "NURBS-OT: An Advanced Model for Generative Curve Modeling." *Journal of Mechanical Design* Vol. 147 No. 3 (2025): p. 031703.

[21] Jia, Jiao, Da, Daicong, Hu, Jianxing and Yin, Sha. "Crashworthiness design of periodic cellular structures using topology optimization." *Composite Structures* Vol. 271 (2021): p. 114164.

[22] Park, Jeong Joon, Florence, Peter, Straub, Julian, Newcombe, Richard and Lovegrove, Steven. "Deepsdf: Learning continuous signed distance functions for shape representation." *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*: pp. 165–174. 2019.

[23] Mescheder, Lars, Oechsle, Michael, Niemeyer, Michael, Nowozin, Sebastian and Geiger, Andreas. "Occupancy networks: Learning 3d reconstruction in function space." *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*: pp. 4460–4470. 2019.

[24] Li, Xueyang, Song, Yu, Lou, Yunzhong and Zhou, Xiangdong. "Cad translator: An effective drive for text to

3d parametric computer-aided design generative modeling." *Proceedings of the 32nd ACM International Conference on Multimedia*: pp. 8461–8470. 2024.

[25] Chen, Zilong, Wang, Feng, Wang, Yikai and Liu, Huaping. "Text-to-3d using gaussian splatting." *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*: pp. 21401–21412. 2024.

[26] Turk, Greg and O'brien, James F. "Modelling with implicit surfaces that interpolate." *ACM Transactions on Graphics (TOG)* Vol. 21 No. 4 (2002): pp. 855–873.

[27] Schaefer, Scott and Warren, Joe. "Dual marching cubes: Primal contouring of dual grids." *12th Pacific Conference on Computer Graphics and Applications, 2004. PG 2004. Proceedings.*: pp. 70–76. 2004. IEEE.

[28] Jones, Alistair, Leary, Martin, Bateman, Stuart and Easton, Mark. "TPMS designer: A tool for generating and analyzing triply periodic minimal surfaces." *Software Impacts* Vol. 10 (2021): p. 100167.

[29] Xu, Yonglai, Pan, Hao, Wang, Ruonan, Du, Qiang and Lu, Lin. "New families of triply periodic minimal surface-like shell lattices." *Additive Manufacturing* Vol. 77 (2023): p. 103779.

[30] Oleynikova, Helen, Millane, Alexander, Taylor, Zachary, Galceran, Enric, Nieto, Juan and Siegwart, Roland. "Signed distance fields: A natural representation for both mapping and planning." *RSS 2016 workshop: geometry and beyond-representations, physics, and scene understanding for robotics*. 2016. University of Michigan.

[31] Wu, Jiajun, Chen, Qiang, Jiang, Jindong, Chatzigeorgiou, George and Meraghni, Fodil. "Adaptive deep homogenization theory for periodic heterogeneous materials." *Composite Structures* Vol. 340 (2024): p. 118171.

[32] Lawal, Zaharaddeen Karami, Yassin, Hayati, Lai, Daphne Teck Ching and Che Idris, Azam. "Physics-informed neural network (PINN) evolution and beyond: A systematic literature review and bibliometric analysis." *Big Data and Cognitive Computing* Vol. 6 No. 4 (2022): p. 140.

[33] Zhong, Haozhang, Das, Raj, Gu, Jianfeng and Qian, Ma. "Low-density, high-strength metal mechanical metamaterials beyond the Gibson-Ashby model." *Materials Today* Vol. 68 (2023): pp. 96–107.

[34] Yago, Daniel, Cante, Juan, Lloberas-Valls, Oriol and Oliver, Javier. "Topology optimization methods for 3D structural problems: a comparative study." *Archives of Computational Methods in Engineering* Vol. 29 No. 3 (2022): pp. 1525–1567.