

Research paper

Generating new cellular structures for additive manufacturing through an unconditional 3D latent diffusion model

Leijian Yu ^a, Yong En Kok ^b, Luke Parry ^a, Ender Özcan ^c, Ian Maskery ^a,*

^a Centre for Additive Manufacturing, Faculty of Engineering, University of Nottingham, Nottingham NG8 1BB, UK

^b Computer Vision Laboratory, School of Computer Science, University of Nottingham, Nottingham, NG8 1BB, UK

^c Computational Optimisation and Learning (COL) Lab, School of Computer Science, University of Nottingham, Nottingham, NG8 1BB, UK

ARTICLE INFO

Keywords:

Additive manufacturing
Unconditional 3D latent diffusion model
New cellular structure design
Triply periodic continuous surfaces
Mechanical testing

ABSTRACT

Advances in additive manufacturing (AM) have facilitated the fabrication of cellular structures inspired by those in the natural world. But the design of complex, tessellating cellular structures remains a challenge for human designers, and only a small number of geometries, defined either by connected walls or struts, or by surface equations, have been investigated. This study introduces generative deep learning to the problem, with the aim of synthesising novel cellular geometries producible by AM. Our unconditional 3D latent diffusion model (U3LDM) explores the design space from a new class of training data comprising 10,650 unit cells. A critical task involved developing a varied set of cell geometries based on random permutations of trigonometric surface equations. This was coupled with a stringent set of pass/fail tests to ensure the generated structures possessed structural connectivity and could tessellate in 3D. The new cellular structures were analysed numerically using finite element analysis, fabricated by polymer AM, and subjected to compression tests to verify their manufacturability and mechanical properties. Results indicate that the U3LDM is capable of generating new 'unseen' cellular structures with geometries and mechanical properties consistent with those of the training specimens. This method also demonstrates the potential universal technique for creating nature-inspired and AM-manufacturable structures beyond the currently limited set of human-derived geometries.

1. Introduction

Additive manufacturing (AM) mitigates some traditional manufacturing constraints (CNC tool access, for example) and enables the production of geometrically complex parts, making it effective in fabricating cellular structures [1]. As illustrated in Fig. 1, these are widely observed in biological systems, including the stems of plants, butterfly wings, and animal bones. Such structures have attracted significant attention due to their high specific strength and surface area [2], their potential use in orthopaedic implants [3], and their energy absorption characteristics under axial compression and complex loads [4]. Thanks to the capabilities of AM, cellular structures have been applied in engineering domains including thermal management [5], biological tissue engineering [6], vibration sensing [7] and various other fields.

The speed and quality of geometric modelling of cellular structures remain crucial factors at the design stage [8]. Various commercial computer-aided-design (CAD) packages, such as Autodesk Netfabb [9], Rhinoceros 3D [10] and nTopology [11], along with non-commercial options like MSLattice [12] and FLatt Pack [13], have been developed

for designing cellular structures. However, modelling cellular structures with complex geometries remains challenging with existing CAD software [14]. This challenge persists because it requires experienced designers and extensive trial-and-error efforts to explore the design space [15], which anyway is still constrained by designers' knowledge and past experience.

To tackle the above challenge, some studies have used optimisation algorithms to design and adjust cellular structures. Alkhatib et al. [16] designed structures with the objective of mechanical isotropy through topology optimisation (TO). A multi-patch isogeometric topology optimisation (MP-ITO) method was proposed by Gao et al. [17] for the design of periodic or graded cellular structures, and Zhang et al. [18] introduced a Kriging-assisted multiscale TO method to achieve maximum natural frequencies in heterogeneous cellular structures. However, optimising cellular lattice structures remains challenging due to the need for systematic and accurate characterisation, as well as knowledge of the design and manufacturing constraints [19].

* Corresponding author.

E-mail address: ian.maskery@nottingham.ac.uk (I. Maskery).

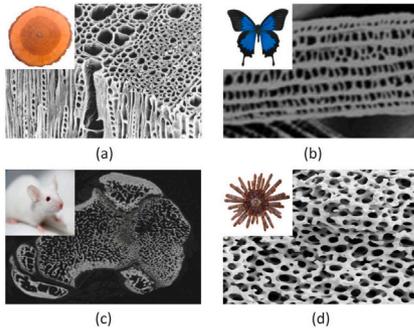


Fig. 1. Biological cellular structures from nature. (a) English oak hardwood [20], (b) butterfly wings [21], (c) mice epiphysis [22], (d) transverse of sea urchin [23].

Machine learning has also been employed to enhance design capabilities beyond brute force computation [24]. Machine learning models, such as artificial neural networks (ANNs), are widely implemented to predict and optimise structural properties such as the stiffness matrix [25], Poisson's ratio [26] and buckling load [27]. These are powerful tools for designing structures with target mechanical properties. However, human intervention in data labelling leads to models focusing on the connections between labels and samples, limiting their capability to create truly new structures [28].

Unconditional generative methods, such as variational autoencoders (VAEs) and generative adversarial networks (GANs), explore the design space without human-specified constraints and are capable of generating a vast number of new designs. A VAE model has been utilised by Zhang et al. to reconstruct 3D digital rocks [29]. However, VAEs are ineffective at capturing complex data distributions [30], and the detail representation and quality of VAE-generated images cannot compete with that of GANs [31]. ScaffoldGAN was previously proposed for the synthesis of bone scaffolds [32], and a combination of a two-dimensional GAN and a similarity-based stacking approach was used to generate 3D structures based on X-ray CT images of porous media [33]. Shen et al. developed a GAN to generate nature-inspired architected materials [28]. However, GANs can suffer from training instability and mode collapse [34]. The latter refers to the case where, despite the model's stochastic nature and different initial conditions, the GAN consistently produces very similar designs [35]. Compared to GAN, diffusion model (DM) approaches are more resistant to model collapse and training instability [36].

In the study from Jadhav et al. [25], a 3D DM was employed for generating cellular structures, either unconditionally or conditionally. It directly operated on the solution fields of trigonometric surface equations, which subsequently limited the spatial resolution of the generated designs, thereby reducing the accuracy of structural fabrication and analysis [37]. This yields structural details with high spatial frequency, which require extensive training to model accurately [38]. The recently developed latent diffusion model (LDM) [38] conducts diffusion and denoising processes in the latent space, thereby reducing computational costs compared to a traditional DM. Herron et al. [35] explored the use of a conditional 3D LDM to generate structures with strain energy similar to those originating from solid isotropic material with penalisation (SIMP) optimisation. However, the 3D LDM of Herron et al. is guided by labels during the generation process, and its design space is influenced by conditional labels. A limitation of this method is that it only generates binary 3D voxel grids for the designed structures, which are unable to capture smooth gradients and geometrical transitions at high spatial frequencies. In addition, generative models developed in computer science do not always generate high-quality, high-fidelity images or structures in specific research domains, such as the generation of kirigami metamaterials [39].

Consequently, there is currently a lack of documented investigation into the unconditional 3D latent diffusion model (U3LDM) for AM

design within a broad design space. Our study aims to be the first to develop a U3LDM for a new context — designing cellular structures with new geometries for AM. Uniquely, the training data comprised a comprehensively diverse set of cell geometries generated by a triply periodic continuous surface (TPCS) approach. Unlike recent studies, which are based on unconditional GAN and conditional LDM that employ binarised voxel arrays [28,32,33,35], our approach uses U arrays that result from the trigonometric surface equations and preserve the geometry and smooth contours of the underlying 3D mathematical surfaces. The U3LDM learns from the provided dataset in the latent space and is capable of sampling from the learned design space to create new and AM-manufacturable cellular structures. We demonstrate the capabilities of this method and quantitatively assess its effectiveness. The newly generated cells are here investigated using numerical analysis and physical testing. This new design approach will not only facilitate the creation of manufacturable cellular structures with never before seen geometries, but also lays the foundation for developing new designs with performance that outstrips that which is currently accessible through conventional design methods.

Following this introduction, the methodologies for cell geometry generation, dataset construction, and the U3LDM are introduced in Section 2. The results of geometry comparison, numerical analysis, AM fabrication and compression tests of the generated designs are presented in Section 3. Section 4 summarises the work and discusses future directions.

2. Methodology

2.1. Surface-based cellular structure design

We first generate a dataset comprising AM-manufacturable cellular structures to represent the design space. Inspired by triply periodic minimal surfaces (TPMS), we present the TPCS approach to create a vast array of cellular structures with diverse geometries and properties. Before introducing the surface equations, some shorthand terms related to these equations are presented. k_i define cellular function periodicities:

$$k_i = 2\pi n_i \quad (1)$$

where $i = x, y, z$, indicating that the design space spans the x , y and z directions. n_i is the number of cell repetitions along each direction. The following sine and cosine functions are then formulated:

$$S_i = \sin(k_i \frac{i}{L_i}) \quad (2)$$

$$S_{2i} = \sin(2k_i \frac{i}{L_i}) \quad (3)$$

and

$$C_i = \cos(k_i \frac{i}{L_i}) \quad (4)$$

$$C_{2i} = \cos(2k_i \frac{i}{L_i}) \quad (5)$$

where L_i represents the sizes of the cellular structure along each axis. Sets of trigonometric functions in each direction are represented as follows:

$$SCX \in \{S_x, S_{2x}, C_x, C_{2x}, 1\} \quad (6)$$

$$SCY \in \{S_y, S_{2y}, C_y, C_{2y}, 1\} \quad (7)$$

$$SCZ \in \{S_z, S_{2z}, C_z, C_{2z}, 1\} \quad (8)$$

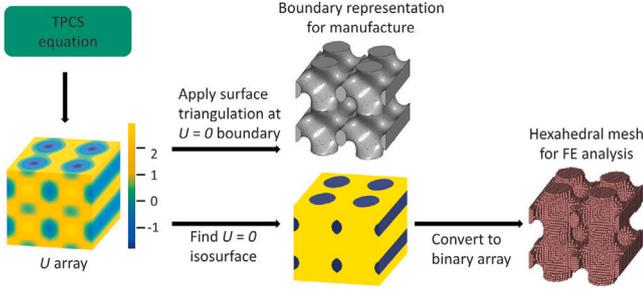


Fig. 2. The main methodology of TPCS cellular structure design. This example shows the creation of a cellular structure with a $2 \times 2 \times 2$ array of cells. The 3D array or surface, U , is defined by summations of trigonometric terms (see Eqs. (1)–(9)).

Then the TPCS approximation of the cellular structure is given by summations of the form:

$$U_G = \begin{cases} a_1 \left(\sum_{j=1}^3 SCX_j SCY_j SCZ_j \right)^m - t^m, & G = 1 \\ \left(a_1 \sum_{j=1}^3 SCX_j SCY_j SCZ_j - \right. \\ \left. a_2 \sum_{j=4}^6 SCX_j SCY_j SCZ_j \right)^m - t^m, & G = 2 \end{cases} \quad (9)$$

where U_G represents the U array, which is a 3D matrix prescribed by the trigonometric functions spanning positive and negative values. The boundary between positive and negative values (which is the $U = 0$ surface) can be interpreted as a solid–void boundary, meaning U can be treated as representing a 3D geometry. a_1 , a_2 , m and G are parameters that define the complexity of the generated geometries, with G being the number of summative terms, 1 or 2. SCX_j , SCY_j and SCZ_j represent randomly selected elements from the sets SCX , SCY and SCZ , respectively, which are used as the j th sub-term of Eq. (9). The variable t controls the positions of the $U = 0$ surface, and is used to control the solid volume fraction of the resulting cellular structure.

As illustrated in Fig. 2, the $U = 0$ boundary separates solid and void regions, with the solid regions forming the cellular structure suitable for finite element (FE) analysis and AM. A large number of candidate cellular structures with diverse geometries can be generated by the combinational exploration of randomly generated values of G , a_1 , a_2 , SCX , SCY , SCZ , m , and t based on Eq. (9). For computational efficiency, predefined ranges are utilised for these variables. Drawing inspiration from existing TPMS equations in the literature [13], a_1 and a_2 are assigned integer values from -4 to 4 . The parameter m takes integer values from $\{1, 2, 3\}$, and t ranges as a float from -3 to 3 . SCX , etc. are selected randomly from the sets given in Eqs. (6)–(8), and G is either 1 or 2. We generate 3D U arrays with a resolution of $128 \times 128 \times 128$ voxels per lattice cell. U arrays with minimum values ranging from -12.5 to -2.4 , maximum values from 2.4 to 12.5 and a maximum standard deviation of 5 are chosen as candidates for dataset construction. Following this, geometrical tests designed to ensure manufacturability are applied to further refine the dataset.

2.2. Manufacturability requirements

U arrays are converted to binarised voxel data based on the $U = 0$ isosurface, and five constraints are applied to verify the manufacturability of the cell geometries.

- m1. The volume fraction, ρ^* , representing the proportion of the cell composed of solid material, is checked. Unit cells with ρ^* ranging from 0.1 to 0.8 pass this test.
- m2. The number of solid regions in the unit cell is determined by evaluating the 3D connectivity of each voxel. Only cells containing one solid region are acceptable.
- m3. The presence of solid elements at the cell faces is checked. Only cells which possess some solid elements at all faces pass this test.
- m4. The cell faces are examined to determine if each one contains more than 5% solid voxels. This is to prevent thin neck regions in the resulting cellular structures, which would possibly be unmanufacturable, and would anyway possess limited strength.
- m5. Opposing cell face pairs are analysed for similarity by calculating the ratio of overlapping solid elements to the total number of overlapping and distinct solid elements. Ratios falling between 0.95 and 1.05 indicate that the unit cell exhibits approximate periodicity, and therefore can be tessellated to form a larger continuous cellular structure.

U arrays meeting all these criteria undergo min–max normalisation to scale them to the range of 0 to 1 and are then utilised for training the U3LDM. The process is summarised in Fig. 3.

Examples of normalised U arrays with various ρ^* values along with their corresponding boundary representations are shown in Fig. 4. The dataset comprises normalised U arrays spanning ρ^* values from 0.1 to 0.8 , incremented by 0.01 . 150 samples are preserved for each volume fraction, resulting in a total of 10,650 normalised U arrays using the described methodology. The dataset is randomly split into training and testing subsets, with 90% and 10% of the samples allocated to each set, respectively.

2.3. Unconditional 3D latent diffusion model

Fig. 5 depicts the architecture of the U3LDM developed here. It consists of two main components: an external variational autoencoder-generative adversarial network (VAE-GAN) and an internal DM. The VAE-GAN encodes U arrays into latent space, while DM operates on these latent representations. After training, DM can generate new latent representations, which are then decoded to U arrays using the VAE decoder. Notably, geometry tests (see Section 2.2) are embedded in the U3LDM to ensure that the generated designs are tessellable and manufacturable. These processes are described more completely in the following sections.

2.3.1. Variational autoencoder-generative adversarial network

The VAE consists of an encoder, $p_\theta(z_u|U)$, which maps the input U to the latent variable z_u , and a decoder, $p_\theta(\hat{U}|z_u)$, which reconstructs \hat{U} from the latent representation. Both the encoder and decoder are neural networks composed of 3D convolutional layers. Subsequently, the decoder acts as the generator within the GAN framework, while a discriminator distinguishes between real data and that generated by the decoder. We incorporated a GAN discriminator in our VAE in order to employ an adversarial training strategy. This has been shown to improve VAE compression and reconstruction capabilities [40–42].

The VAE is optimised using three loss terms. The first is the reconstruction loss, \mathcal{L}_{rec} , which comprises both a per-voxel loss function and a perceptual loss function, \mathcal{L}_{perc} , [43] that relies on high-level features.

$$\mathcal{L}_{rec}(U, \hat{U}) = \|U - \hat{U}\|_1 + \mathcal{L}_{perc}(U, \hat{U}) \quad (10)$$

The second loss term regularises the latent variables sampled from the encoder-mapped distribution towards a Gaussian prior distribution. This is achieved by minimising the Kullback–Leibler (KL) divergence between the prior distribution $p(z_u)$ and the encoder-mapped distribution $p_{\theta_1}(z_u|U)$.

$$\mathcal{L}_{KL}(U) = KL(p_{\theta_1}(z_u|U)||p(z_u)) \quad (11)$$

The third loss term is derived from the manufacturability constraints on binarised U arrays, U_b . This guides the model in maintaining the regularities of the training samples and is defined as:

$$\mathcal{L}_{phy}(U_b, \hat{U}_b) = \lambda_1 |\rho_{U_b}^* - \rho_{\hat{U}_b}^*| + \lambda_2 (N_s - 1) + \lambda_3 \sum_{i=1}^6 (\|F_i(U_b) - F_i(\hat{U}_b)\|_1) \quad (12)$$

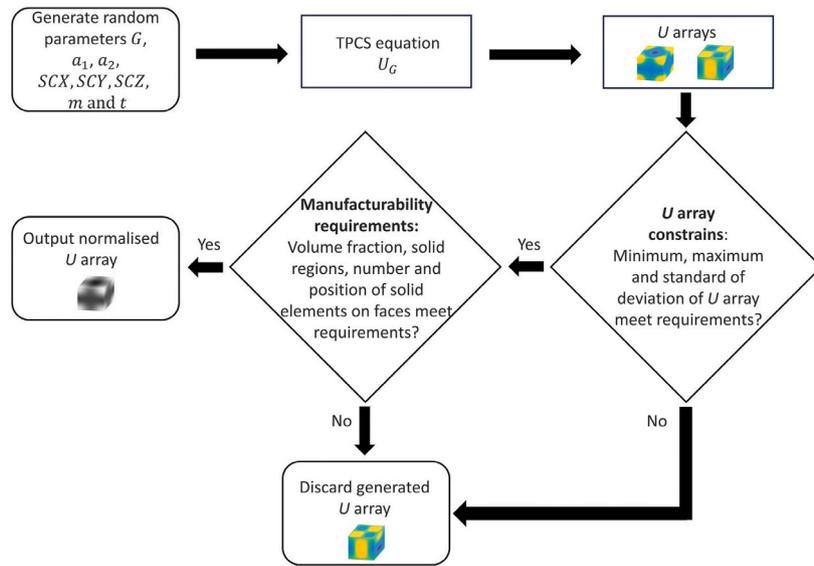


Fig. 3. TPCS-based U array dataset generation, featuring randomised parameters in the underlying 3D trigonometric functions, and the application of pass/fail geometry tests to ensure the resulting generated cells can form larger manufacturable cellular structures.

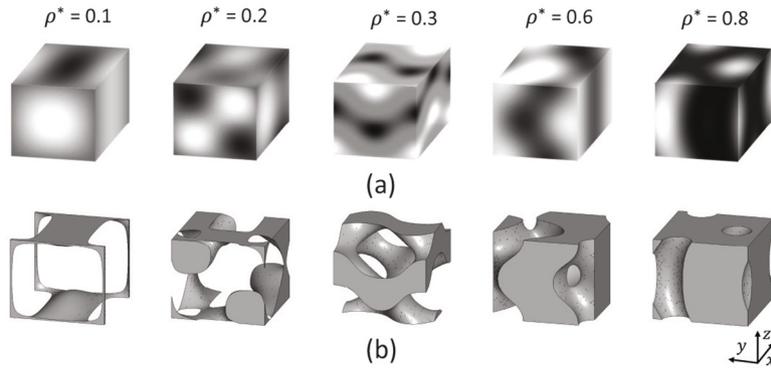


Fig. 4. TPCS generated unit cells. (a) Normalised U arrays, (b) boundary representations.

where the three sub-terms are derived from manufacturing constraints (m) in Section 2.2 as follows: the first sub-term is based on m1, the second on m2, and the third on a combination of m3, m4 and m5. These three constraints respectively guide the reconstructed U arrays to preserve the volume fraction, retain the number of solid regions, and match the cell faces with those of the input. Specifically, $\rho_{U_b}^*$ and $\rho_{\hat{U}_b}^*$ are the volume fractions of the input and reconstructed U_b , respectively. N_s is the number of solid regions produced by \hat{U} . $F(U_b)$ and $F(\hat{U}_b)$ denote the cell faces of the input and reconstructed U_b , respectively. λ_1 , λ_2 and λ_3 control the weights of the different loss components.

Finally, the VAE loss is represented as follows:

$$\mathcal{L}_{VAE} = \lambda_4 \mathcal{L}_{rec}(U, \hat{U}) + \lambda_5 \mathcal{L}_{KL}(U) + \mathcal{L}_{phy}(U_b, \hat{U}_b) \quad (13)$$

where λ_4 and λ_5 are hyperparameters used to control the weighted contribution of different high-level loss terms.

Regarding the GAN loss, the Hinge loss [44] is adopted:

$$\mathcal{L}_{GAN} = -\mathbb{E}[D_{GAN}(G_{GAN}(z_u))] \quad (14)$$

$$\mathcal{L}_{D_{GAN}} = \mathbb{E}[\max(0, 1 - D_{GAN}(U))] + \mathbb{E}[\max(0, 1 + D_{GAN}(G_{GAN}(z_u)))] \quad (15)$$

where $G_{GAN}(z_u)$ and $D_{GAN}(U)$ are the outputs of the GAN generator and discriminator, respectively.

The VAE is trained initially to improve the stability of the training procedure, and the GAN is initiated after s_{GAN} steps. So the final loss

functions for the VAE-GAN are:

$$\mathcal{L}_{FG} = \begin{cases} \mathcal{L}_{VAE}, & s_{tra} < s_{GAN} \\ \mathcal{L}_{VAE} + \lambda_6 \lambda_{adap} \mathcal{L}_{GAN}, & s_{tra} \geq s_{GAN} \end{cases} \quad (16)$$

$$\mathcal{L}_{FD} = \begin{cases} 0, & s_{tra} < s_{GAN} \\ \mathcal{L}_{D_{GAN}}, & s_{tra} \geq s_{GAN} \end{cases} \quad (17)$$

where \mathcal{L}_{FG} and \mathcal{L}_{FD} are generator and discriminator losses, respectively. λ_{adap} adjusts the weight of the adversarial loss autonomously based on the gradient norms of the losses [38], while λ_6 controls the effect of the adaptive adversarial loss. s_{tra} is the number of training steps.

2.3.2. Denoising diffusion probabilistic model

The 3D denoising diffusion probabilistic model (DDPM) for generating U arrays is derived from 2D DDPM [45] by replacing 2D convolution with 3D convolution. This model includes both forward diffusion and reverse denoising processes. The noising procedure is a Markov process, and it can be defined as:

$$q(z_u^{(t)} | z_u^{(t-1)}) \sim \mathcal{N}(z_u^{(t)}; \sqrt{1 - \beta_t} z_u^{(t-1)}, \beta_t \mathbf{I}) \quad (18)$$

where $z_u^{(t)}$ is the latent representation after t diffusion steps. $\beta_t \in (0, 1)$ controls the noising schedule, and \mathbf{I} is the identity matrix.

Once the reverse distribution $q(z_u^{(t-1)} | z_u^{(t)})$ can be learned, the noise can gradually be removed in the backward steps, and this process can be approximated by a neural network with parameter θ . The diffusion

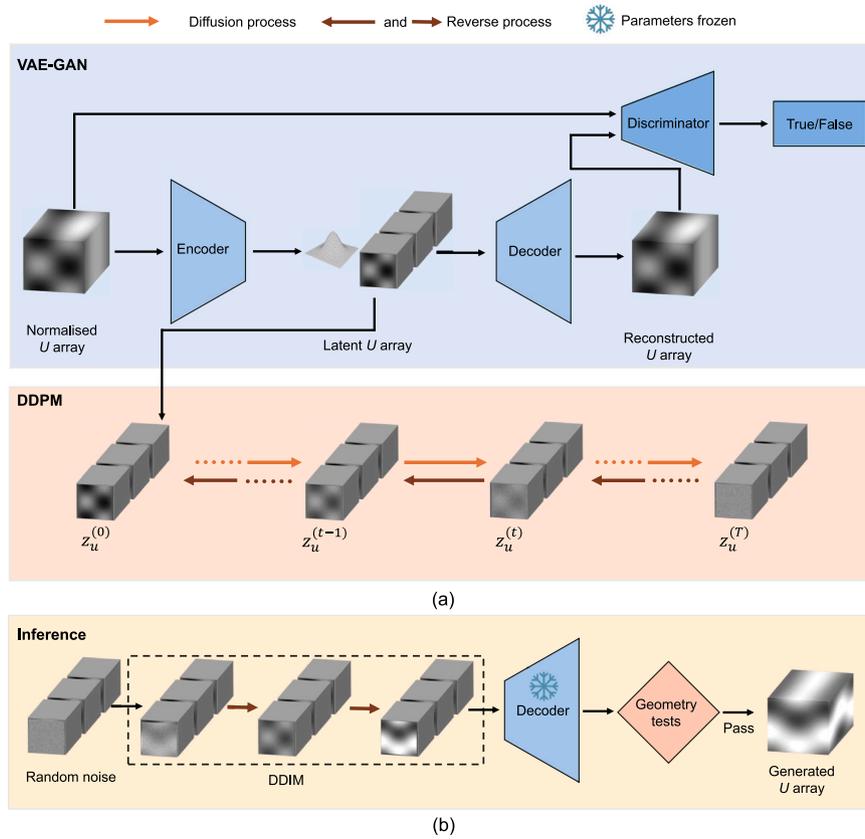


Fig. 5. The principle of the U3LDM. (a) Training stage. U arrays are mapped into the latent space and reconstructed by the VAE-GAN. A DDPM learns the latent representations through forward diffusion and reverse denoising. (b) Inference process. The DDIM generates latent representations from random noise, which the pretrained decoder reconstructs back to U arrays. Geometry tests (see Section 2.2) ensure that the generated designs are tessellable and manufacturable.

neural network is trained by minimising the distance between the predicted and added noises, which is defined as:

$$\mathcal{L}_{MAE} = \mathbb{E}_{z_u^{(0)}, t, \epsilon^{(t)}} [\|\epsilon^{(t)} - \epsilon_{\theta}^{(t)}(z_u^{(t)})\|_1] \quad (19)$$

where $\epsilon^{(t)}$ and $\epsilon_{\theta}^{(t)}$ are the added noise and predicted noise at step t .

Nichol et al. [46] demonstrated that the model's performance can be improved through the optimisation of variational lower-bound (VLB), so we adopted \mathcal{L}_{VLB} from Ho et al. [45]. Finally, the learning objective for the DDPM is defined as follows:

$$\mathcal{L}_{DDPM} = \mathcal{L}_{MAE} + \mathcal{L}_{VLB} \quad (20)$$

2.3.3. Denoising diffusion implicit model

The main drawback of DDPM is that each step in the denoising process involves some randomness, and it requires a large number of steps to generate a new design. To accelerate the sampling process with the pre-trained DDPM, the denoising diffusion implicit model (DDIM) [47] is proposed. In DDIM the reverse diffusion process becomes deterministic, thereby enhancing the consistency and stability of the generation process. Consequently, DDIM enables latent variable generation with fewer reverse steps, reducing computational costs and accelerating the sampling process.

2.4. Model architecture

The model architecture draws inspiration from Rombach et al. [38] and incorporates 3D convolution. The DDPM [45] leverages a U-Net [48] for predicting noise due to its excellent capability to capture and fuse multi-scale features. Table 1 summarises the VAE and U-Net architectures. The VAE model operates with 4 feature map resolutions spanning from $128 \times 128 \times 128$ down to $16 \times 16 \times 16$ and

Table 1
Parameters in the VAE and U-Net model architectures.

	VAE	U-Net
Input channels	1	16
Base channels	32	256
Channel multipliers	[1, 2, 4, 4]	[1, 2, 4, 4]
Global attention layer	0	1
Attention layers on spatial resolution	0	[16, 8, 4]
Residual blocks z channels	1	1
Output channels	1	16

incorporates 1 convolutional residual block per resolution level. The input normalised U arrays, initially sized at $128 \times 128 \times 128 \times 1$, are encoded into a latent representation of $16 \times 16 \times 16 \times 16$ by the encoder, and subsequently decode back to the initial size by the decoder. The U-Net architecture includes 4 feature map resolutions ranging from $16 \times 16 \times 16$ to $2 \times 2 \times 2$, with 1 convolutional residual block per resolution level. From the resolution level of $16 \times 16 \times 16$ down to $4 \times 4 \times 4$, a global attention layer with a single head is employed between convolution blocks. The GAN discriminator is developed from the PatchGAN discriminator [49] and extended to process 3D data.

2.5. Training procedure

Both the VAE-GAN and DDPM were implemented using PyTorch 1.13.0 and trained on an Nvidia Tesla A100 40 GB GPU. Following Rombach et al. [38], they were optimised using Adam and AdamW, respectively. For the VAE-GAN, the learning rate was set to 4.5×10^{-6}

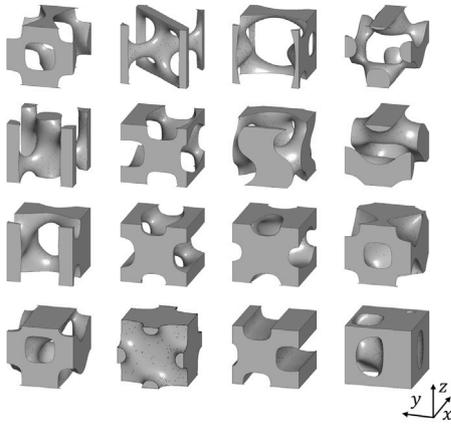


Fig. 6. U3LDM-generated designs, arranged from left to right and top to bottom, with ρ^* values ranging from 0.12 to 0.72 in increments of 0.04.

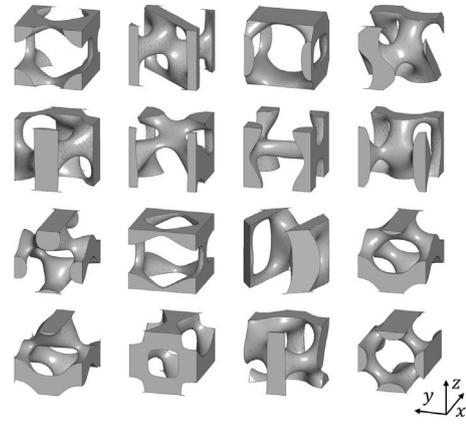


Fig. 7. U3LDM-generated designs with $\rho^* = 0.26$.

with an accumulated grad batch size of 24. λ_5 and λ_6 were set as 1×10^{-6} and 0.5, respectively [38]. The GAN discriminator was invoked after 30,000 steps. To maintain consistency in the scale of various loss terms, we assigned λ_1 , λ_2 , λ_3 and λ_4 the values 1, 0.1, 1×10^{-5} and 1×10^{-6} , respectively. The DDPM was trained with a learning rate of 5×10^{-5} , and the accumulated grad batch size was 96. The diffusion process involved 1000 time steps with a linear noise schedule. Both the VAE-GAN and DDPM were trained until their respective losses converged. The results presented in Section 3 were derived from the VAE-GAN model trained for 124 epochs and the DDPM model trained for 343 epochs.

3. Results

3.1. Cellular structure generation

Each initial U3LDM starting condition produces a unique cell design. These were evaluated qualitatively to assess the performance of the developed U3LDM. For this evaluation, 1000 normalised U arrays were generated with a batch size of 1, and each denoising process contained 200 steps. The average GPU usage was 7948 MB, with an average generation time of 35.51 s per U array meeting the geometrical requirements outlined in Section 2.2.

Examples of generated designs, with ρ^* ranging from 0.12 to 0.72 in increments of 0.04, are presented in Fig. 6. These results indicate that the U3LDM has successfully learned the principles required for designing cells similar to TPCS structures.

Fig. 7 depicts different cell types selected with the same ρ^* , illustrating the U3LDM's ability to generate multiple geometries with the same ρ^* . Therefore, the trained U3LDM is capable of generating designs across a wide range of ρ^* , as well as creating various new cell geometries not contained in the training dataset. This suggests that U3LDM is effective in preventing instances of mode collapse.

A qualitative comparison between structures designed by the U3LDM and similar structures from the training dataset is shown in Fig. 8. The similarity of these structures was evaluated by calculating the similarity of their U_b arrays.

In Fig. 8, the structure shown in the top left was created by the U3LDM with a ρ^* value of 0.16. The other two structures at the top were the most geometrically similar structures identified from the training dataset with the same ρ^* . The bottom three structures were the three most similar ones with any ρ^* in the training dataset. It can be observed that the structures designed by the U3LDM share a combination of geometrical features present in the training dataset, however are not identical. Notably, the generated structures do not share the same ρ^* as the three most similar structures. This demonstrates the model's ability

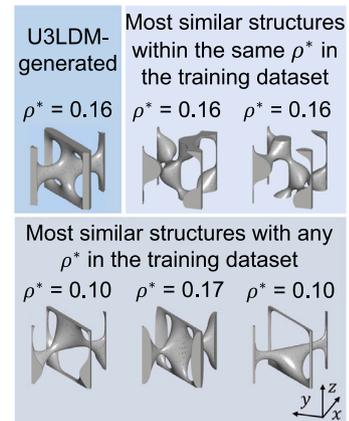


Fig. 8. Comparison of the generated design with similar structures in the training dataset. The top left structure represents the U3LDM-generated design. To its right are the most similar structures that have the same ρ^* value from the training dataset. On the bottom row are the three most similar structures of any ρ^* from the training dataset.

to combine different parts from various structures while maintaining smooth connectivity in the generated designs.

To analyse how distinct U3LDM-generated cells are from the training dataset, 190 cells that pass our manufacturability tests were generated in a single run of U3LDM. We categorised latent representations of generated and training cells using 3D t-distributed Stochastic Neighbour Embedding (t-SNE) [50] with default parameters from Pedregosa et al. [51]. The results are shown in Fig. 9, which indicates that U3LDM-generated designs are distinct from training samples.

The Chamfer distance, as previously used by Jadhav et al. [25], was used to quantify the similarity of the generated and training cells. It describes the average distance between pairs of nearest neighbours of two sets of point clouds. The cellular structures were converted into point clouds using PyMeshLab [52], with each point cloud spanning the coordinate range from 0 to 10 along the x , y , and z axes. For each point cloud of generated designs, the most similar point cloud in the training dataset was determined by the minimum pairwise Chamfer distance. The average Chamfer distance between these generated cellular structures and their closest training samples was 0.48. Given this is > 0 , it further confirms that the U3LDM learns the characteristic geometrical feature sets associated with TPCS structures and preserves key features while generating similar designs that are not explicitly present in the training dataset.

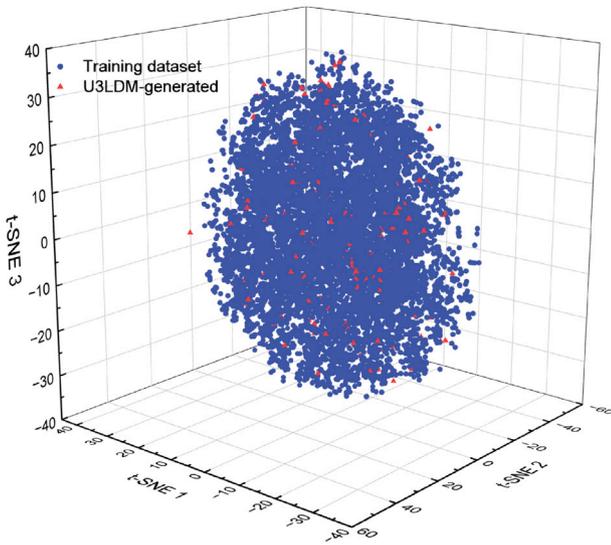


Fig. 9. 3D t-SNE visualisation of the latent representations of U3LDM-generated and training cellular structures.

3.2. Finite element analysis

To further evaluate the efficacy of the developed U3LDM for creating cells similar to TPCS geometries, the elastic moduli of the generated structures were investigated. Gibson and Ashby [53] provided comprehensive insights into the physical properties of cellular solids. Among their findings, the relationship between the relative elastic modulus E^* and volume fraction ρ^* is one of the most significant factors for cellular structural design. Hexahedral FE meshes were created based on the $U = 0$ isosurface, where each solid element in the cellular structure corresponds to a hexahedral element in the FE mesh. This is the same methodology previously employed by Maskery et al. [13] to programmatically create FE meshes for surface-based cellular structures and export them as Abaqus FE analysis input files.

The generated cells were subjected to compressive displacements at the nodes of their top planes, equivalent to 1% of each model's height. Elements on the bottom surface of each model were constrained by the Abaqus ZSYMM boundary condition, which restricted these elements to translation only in the x - y plane and rotation only around the z axis. Further details regarding the configuration and methodology for the FE analysis can be found in [54].

In this study, three unit cells generated by the U3LDM were selected for each ρ^* ranging from 0.15 to 0.75 in increments of 0.05, giving a total of 65 models for FE analysis. To validate their performance compared to randomly generated TPCS structures, we chose the three most similar designs with different volume fractions and the two most similar designs with the same volume fraction in the training dataset.

$E^*(\rho^*)$ for each unit cell under compressive loading in the z -direction is presented in Fig. 10. There is no discernible elastic modulus deviation for the U3LDM-generated models compared to those from the training dataset. We can conclude that the U3LDM not only produces cellular structures with similar geometries to the training dataset, but also with broadly comparable mechanical properties to those of the TPCS structures.

3.3. Fabrication and compression tests

This section validates the manufacturability of U3LDM-generated cell designs, and provides some experimental verification of their mechanical properties to compare with TPCS cellular structures.

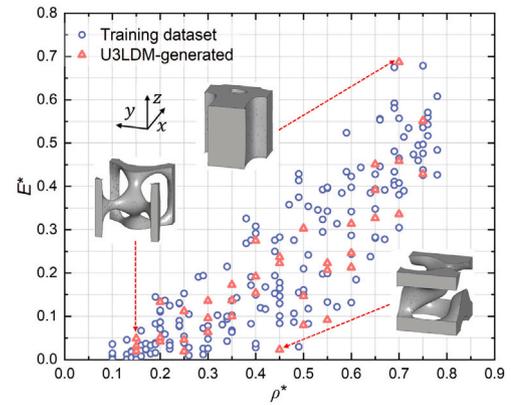


Fig. 10. Comparison of E^* between U3LDM-generated designs and similar structures in the training dataset.

Specimen #1 Specimen #2 Specimen #3

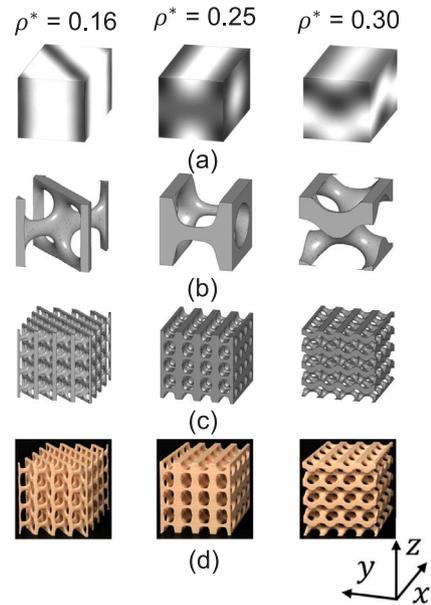


Fig. 11. U3LDM generated and fabricated structures. (a) Normalised U arrays of unit cells; (b) boundary representations of unit cells; (c) cellular structures comprising $4 \times 4 \times 4$ cells; (d) fabricated specimens comprising $4 \times 4 \times 4$ cells.

As depicted in Fig. 11, three designs with ρ^* values of 0.16, 0.25, and 0.30 were randomly selected from the U3LDM-generated models. Structures comprising $4 \times 4 \times 4$ cells were fabricated using a Bambu Lab X1 Carbon 3D printer [55] from polylactic acid (Bambu Lab PLA Basic). Each specimen had dimensions of $50 \times 50 \times 50$ mm and was fabricated according to the parameters detailed in Table 2. ρ^* of the manufactured cellular specimens was calculated as the ratios of their mass densities to the density of a solid specimen made with the same printer settings (infill pattern, etc.). The mass densities were determined via electronic balance and calliper to determine mass and volume, respectively. The fact that all specimens were built without failure, and the average difference in ρ^* between the manufactured specimens and the corresponding STL models was only 1.2%, confirms the manufacturability and fidelity of the U3LDM-generated designs.

The fabricated specimens were subjected to mechanical compression using an Instron 5966 universal testing machine equipped with a 50 kN load cell. The deformation rate was set to 0.1 mm/min. The resulting stress-strain curves are shown in Fig. 12. Because of the

Table 2

Parameters used by the Bambu Lab X1-Carbon Combo 3D Printer for specimens fabrication.

Parameter	Value
Material	PLA basic
Nozzle temperature	220 °C
Nozzle diameter	0.4 mm
Layer thickness	0.16 mm
Print speed	500 mm/s
Heated bed temperature	55 °C

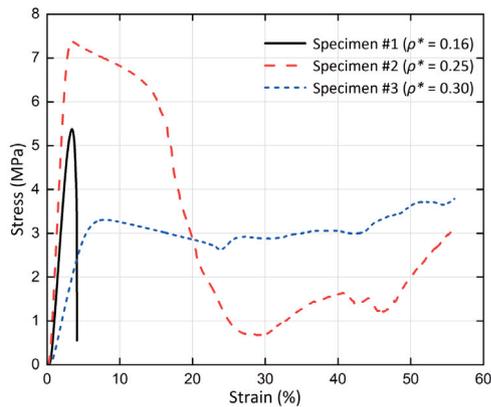


Fig. 12. Stress–strain curves of the U3LDM-generated cellular structures under compression.

differences in cell geometry, they exhibit distinct behaviour under compression. Specimen #1 exhibits brittle failure and collapses at around 5% strain. Specimen #2 showed an initial elastic response followed by plastic deformation, eventually leading to structural collapse at approximately 15% strain. Specimen #3 demonstrated a prolonged plastic plateau, akin to those of bending-dominated cellular structures. These behaviour profiles are generally consistent with those observed in conventional TPMS and strut-based cellular structures made by AM [54,56].

To determine the E^* of these three structures, the compressive modulus of elasticity of Bambu Lab PLA Basic was obtained from the literature [57], with a value used in this study of 2.7 GPa. The E^* values for U3LDM-generated cells obtained from FE analysis and via compression of the fabricated specimens are summarised in Table 3. E^* generally shows an increasing trend with ρ^* , but not for every increment in ρ^* . Specimen #3, despite having a larger ρ^* than specimens #1 and #2, exhibits the lowest E^* among them. Nonetheless, this remains entirely consistent with the structural characteristics derived from TPCS equations; i.e., specimen #3 has anisotropic mechanical properties owing to its geometry; the stiffness in other axial loading directions is much higher. Specifically, E^* in the x and y loading directions was found by FE analysis to be 101.6×10^{-3} and 101.0×10^{-3} , respectively. Table 3 demonstrates that the elastic modulus of fabricated specimens is well described by the FE model, with an average over-estimation, of just 6%. Unit cells generated by the U3LDM can therefore be reliably tessellated to form repeating cellular structures, which offers confirmation that the manufacturability requirements described in Section 2.2 are effective.

A final interesting observation from the U3LDM-generated cellular structures is illustrated in Fig. 13. Unlike conventional TPMS or strut-based geometries, which are always rotationally symmetric, specimen #1 exhibits very different connectivity and symmetry in different planes. It possesses a body-centred-cubic-like (BCC) geometry in the x - y plane and a gyroid-like structure in the y - z plane. Fig. 8 proves that this type of cellular structure is not present in the training dataset. Hence, it can be regarded as a hybrid structure effectively combining

Table 3

Comparison of E^* of structures comprising $4 \times 4 \times 4$ U3LDM-generated cells obtained from FE analysis and compression tests.

	ρ^* (%)	E^* from FEA $\times 10^{-3}$	E^* from compression tests $\times 10^{-3}$
Specimen #1	16	89.9	84.8
Specimen #2	25	135.0	122.6
Specimen #3	30	28.0	27.4

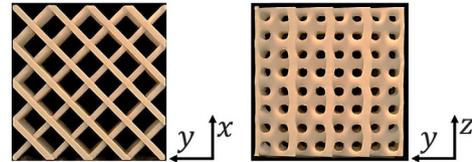


Fig. 13. Different views of the specimen #1. In the x - y plane, its geometry appears identical to the commonly studied body-centred-cubic-like (BCC) cell type, while in the y - z plane, it appears like a gyroid or other surface-based cell type.

the characteristics of geometries from the training dataset. This further demonstrates the potential of the U3LDM to create highly novel, manufacturable cellular designs, and implies that the same technique could be applied to other 3D structures (e.g., natural, biological structures) to synthesise useful materials for a wide range of applications. Since data-driven inverse design methods for AM rely heavily on the quality and variety of datasets [25], they will also benefit from the U3LDM's ability to enrich and expand dataset diversity.

3.4. Discussion

To the best of our knowledge, the U3LDM presented here is the first to utilise the unconditional LDM for generating 3D surfaced-based cellular structures. Most GAN-based [32,58], DM-based [59], and the latest LDM-based [35] 3D structure generation approaches do not guarantee structures that can be tessellated and are limited to a single representation of the structure.

While a DM-based approach named GLU3D [25] for generating 3D surface-based cellular structures does exist, it is trained by a set of 22 families of known TPMS structures and their combinations. It employs a combination of unconditional generation and masked conditioning to enable conditional generation during the training stage, which allows for a broad exploration of the design space. Despite this, predefined labels may still limit the extent of the explored design space, preventing it from fully capturing the variety of the entire dataset. Since GLU3D directly operates at the solution fields of trigonometric equations with a resolution of $64 \times 64 \times 64$, it generates meshes which cannot accurately represent complex volumes of the structure, such as thin and highly tortuous walls or struts [25]. Furthermore, GLU3D does not address the presence of unconnected regions of the cell, which can result in convergence failures during FEA and make the designed structures unmanufacturable with some AM processes.

Our U3LDM goes further by being trained on a new family of diverse TPCS geometries not previously available. It generates smoothly-varying U arrays with a resolution of $128 \times 128 \times 128$ from the latent space, and the inference process requires less than 8 GB of GPU memory. The high-resolution U arrays generated by our U3LDM, which pass a stringent set of pass/fail geometry tests, offer adaptability for repeating cellular that are fabricable by AM.

However, there are two limitations observed in our U3LDM's ability to generate structures. First, only 19% of the generated cell designs are currently able to pass the geometry tests for tessellation and manufacture. To better understand this, the generated results are further individually evaluated through the manufacturability requirements m1 to m5 (see Section 2.2), and success rates are 98.3%, 45.9%, 95.6%,

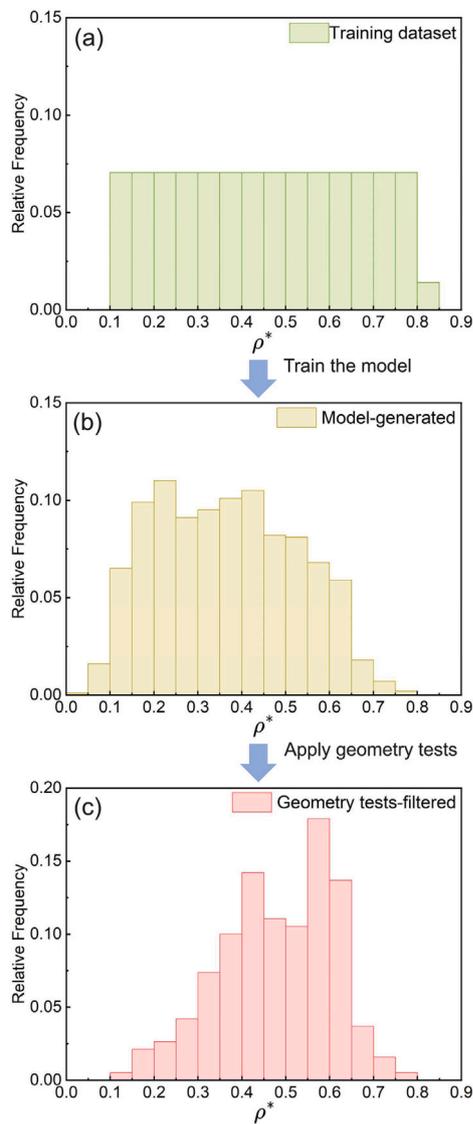


Fig. 14. Distribution of ρ^* across the training dataset, created by the unconditional latent diffusion model, and the U3LDM-generated results that pass geometry tests.

83.8%, and 32.3%, respectively. The low pass rates for the manufacturability requirements m2 and m5 are because we do not allow any disconnected solids and only permit 5% differences on the opposite faces to ensure tessellation, which are structural features the LDM does not learn well. The success rate of U3LDM could potentially be improved by imposing the geometry constraints during the diffusion and denoising processes rather than only at the compression and reconstruction stage, as we did here. By contrast, if we follow the evaluation matrix employed in GLU3D [25], where any structure with $\rho^* > 0.05$ is considered a successful generation, the comparable success rate for our model would be 99.9%.

Another limitation is that, although our U3LDM is trained with cells with ρ^* uniformly distributed between 0.1 to 0.8, as shown in Fig. 14, high and low volume fractions are under-represented in the U3LDM output. We understand this discrepancy arises because the U3LDM learns from samples with ρ^* values both below and above a specific value to generate a design with this ρ^* value. A large number of LDM-generated low ρ^* samples fail the given geometry tests. This is due to the thin connections in low ρ^* structures, which the LDM struggles to generate effectively. Additionally, the low ρ^* structures contain small solid face regions where small discrepancies between opposing faces are

magnified in significance. As a result, the final output of our U3LDM, shown in Fig. 14(c), spans approximately the same ρ^* range as the training dataset, but with fewer generated designs at low and high volume fractions. This could be mitigated by implementing optimal network architectures and training hyperparameters.

4. Conclusions and future work

This study demonstrates how a U3LDM can be used to generate entirely new cellular structures. A dataset comprising 10,650 random surface-based cell designs was constructed, the VAE-GAN and DDPM models were trained, and new ‘unseen’ cellular structures were created by DDIM and the VAE’s decoder. The U3LDM is capable of creating similar geometries to those in the training dataset, as well as novel structures which combine several distinct geometrical properties previously only seen in separate cell types, such as the body-centred-cubic and gyroid.

Evidence from numerical modelling, fabrication by additive manufacture, and compression testing indicate that the U3LDM can effectively identify and replicate features of complex cellular geometries, which is potentially useful for the creation of intricate nature-inspired structures for AM. Additionally, the properties of the U3LDM-generated structures could be manipulated through refinement of the training dataset.

A significant improvement here over previous efforts in this area is that we employed smoothly-varying arrays of a wide variety of TPCS geometries instead of binarised voxel models or limited cell types in the training dataset. This allowed the generated design to preserve gradients and transitions to create smooth surfaces, especially in curved and angled areas, while also presenting a wide range of geometries and properties. Another advancement of the U3LDM is its ability to generate $128 \times 128 \times 128$ arrays from the latent space that pass the geometry tests, ensuring they can be tessellated and manufactured by AM processes while preserving more structural detail than existing methods.

Future research directions include quantifying geometric similarities between the training dataset and U3LDM-generated designs, as well as optimising the U3LDM configurations, such as adjusting the number of layers and weights for different loss terms to improve the efficacy of design generation.

CRedit authorship contribution statement

Leijian Yu: Writing – review & editing, Writing – original draft, Validation, Methodology, Formal analysis, Data curation, Conceptualization. **Yong En Kok:** Methodology, Conceptualization. **Luke Parry:** Writing – review & editing, Supervision, Methodology, Funding acquisition, Conceptualization. **Ender Özcan:** Writing – review & editing, Supervision, Methodology, Funding acquisition, Conceptualization. **Ian Maskery:** Writing – review & editing, Supervision, Project administration, Methodology, Funding acquisition, Data curation, Conceptualization.

Declaration of competing interest

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

Acknowledgements

This work is supported by the Engineering and Physical Sciences Research Council New Horizons project [grant number EP/X016854/1]. Thank Mark East, Adam Whitbread and Richard Homer for helping with this work. The authors acknowledge the access to GPU resources provided by the University of Nottingham’s Ada High Performance Computing (HPC) service and HPC Midlands Plus [grant number EP/T022108/1].

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.addma.2025.104712>. STL files for the U3LDM-generated structures shown in the Results section are available in the Supplementary material associated with this article.

Data availability

STL files for the U3LDM-generated structures shown in the Results section are available in the Supplementary material associated with this article. Other data is available on request.

References

- [1] T. Liu, S. Guessasma, J. Zhu, W. Zhang, Designing cellular structures for additive manufacturing using Voronoi–Monte Carlo approach, *Polymers* 11 (7) (2019) 1158.
- [2] S. Ataollahi, A review on additive manufacturing of lattice structures in tissue engineering, *Bioprinting* (2023) e00304.
- [3] P. Nogueira, J.P. Magrinho, M.B. Silva, A.M. de Deus, M.F. Vaz, Compression properties of cellular iron lattice structures used to mimic bone characteristics, *Proc. Inst. Mech. Eng. Part L: J. Materials: Des. Appl.* (2024) 14644207241241799.
- [4] H. Yin, W. Zhang, L. Zhu, F. Meng, J. Liu, G. Wen, Review on lattice structures for energy absorption properties, *Compos. Struct.* 304 (2023) 116397.
- [5] K. Shimoyama, A. Komiya, Multi-objective Bayesian topology optimization of a lattice-structured heat sink in natural convection, *Struct. Multidiscip. Optim.* 65 (1) (2022) 1.
- [6] S. Vijayavenkataraman, L. Zhang, S. Zhang, J.Y. Hsi Fuh, W.F. Lu, Triply periodic minimal surfaces sheet scaffolds for tissue engineering applications: An optimization approach toward biomimetic scaffold design, *ACS Appl. Bio Mater.* 1 (2) (2018) 259–269.
- [7] J. Xu, Z. Wang, H.-Y. Nie, Y. Wei, Y. Liu, Self-powered smart vibration absorber in situ sensing and energy harvesting, *Adv. Intell. Syst.* (2024) 2300792.
- [8] N. Letov, Y.F. Zhao, Geometric modeling of advanced cellular structures with skeletal graphs, *Int. J. Mech. Sci.* (2024) 109087.
- [9] Fusion 360 with Netfabb, <https://www.autodesk.ca/en/products/netfabb/overview>.
- [10] Rhino – rhinoceros 3D <https://www.rhino3d.com/>.
- [11] Next-Generation Engineering Design Software, <https://design.ntopology.com/demo/>.
- [12] O. Al-Ketan, R.K. Abu Al-Rub, MSLattice: A free software for generating uniform and graded lattices based on triply periodic minimal surfaces, *Mater. Des. Process. Commun.* 3 (6) (2021) e205.
- [13] I. Maskery, L.A. Parry, D. Padrão, R.J. Hague, I.A. Ashcroft, FLatt Pack: A research-focussed lattice design program, *Addit. Manuf.* 49 (2022) 102510.
- [14] N. Letov, Y.F. Zhao, A geometric modelling framework to support the design of heterogeneous lattice structures with non-linearly varying geometry, *J. Comput. Des. Eng.* 9 (5) (2022) 1565–1584.
- [15] J. Hu, M. Li, X. Yang, S. Gao, Cellular structure design based on free material optimization under connectivity control, *Computer-Aided Des.* 127 (2020) 102854.
- [16] S.E. Alkhatib, A. Karrech, T.B. Sercombe, Isotropic energy absorption of topology optimized lattice structure, *Thin-Walled Struct.* 182 (2023) 110220.
- [17] J. Gao, X. Wu, M. Xiao, V.P. Nguyen, L. Gao, T. Rabczuk, Multi-patch isogeometric topology optimization for cellular structures with flexible designs using Nitsche's method, *Comput. Methods Appl. Mech. Engrg.* 410 (2023) 116036.
- [18] Y. Zhang, L. Gao, M. Xiao, Maximizing natural frequencies of inhomogeneous cellular structures by Kriging-assisted multiscale topology optimization, *Comput. Struct.* 230 (2020) 106197.
- [19] Z. Jihong, Z. Han, W. Chuang, Z. Lu, Y. Shangqin, W. Zhang, A review of topology optimization for additive manufacturing: Status and challenges, *Chin. J. Aeronaut.* 34 (1) (2021) 91–110.
- [20] M.P. Ansell, Wood microstructure–A cellular composite, in: *Wood Composites*, Elsevier, 2015, pp. 3–26.
- [21] A. Elbaz, B. Gao, Z. He, Z. Gu, Hepatocyte aggregate formation on chitin-based anisotropic microstructures of butterfly wings, *Biomimetics* 3 (1) (2018) 2.
- [22] P. Salmon, Application of bone morphometry and densitometry by X-ray micro-CT to bone disease models and phenotypes, *Micro-Comput. Tomogr. (Micro-CT) Med. Eng.* (2020) 49–75.
- [23] T. Yang, Z. Jia, Z. Wu, H. Chen, Z. Deng, L. Chen, Y. Zhu, L. Li, High strength and damage-tolerance in echinoderm stereom as a natural bicontinuous ceramic cellular solid, *Nat. Commun.* 13 (1) (2022) 6083.
- [24] K. Guo, Z. Yang, C.-H. Yu, M.J. Buehler, Artificial intelligence and machine learning in design of mechanical materials, *Mater. Horizons* 8 (4) (2021) 1153–1172.
- [25] Y. Jadhav, J. Berthel, C. Hu, R. Panat, J. Beuth, A. Barati Farimani, Generative lattice units with 3D diffusion for inverse design: GLU3D, *Adv. Funct. Mater.* (2024) 2404165.
- [26] S.V. Patil, K. Balakrishna Rao, G. Nayak, Prediction of recycled coarse aggregate concrete mechanical properties using multiple linear regression and artificial neural network, *J. Eng. Des. Technol.* 21 (6) (2023) 1690–1709.
- [27] F. Liu, T. Niu, J.-G. Gong, H. Chen, F.-Z. Xuan, Experimental and numerical investigations on buckling behaviours of axially compressed cylindrical–conical–cylindrical shells at elevated temperature, *Thin-Walled Struct.* 184 (2023) 110549.
- [28] S.C.-y. Shen, M.J. Buehler, Nature-inspired architected materials using unsupervised deep learning, *Commun. Eng.* 1 (1) (2022) 37.
- [29] T. Zhang, Y. Yang, A. Zhang, 3D reconstruction of porous media using a batch normalized variational auto-encoder, *Comput. Geosci.* 26 (5) (2022) 1261–1278.
- [30] Y. Zhang, P. Seibert, A. Otto, A. Raßloff, M. Ambati, M. Kästner, DA-VEGAN: Differentiably augmenting VAE-GAN for microstructure reconstruction from extremely small data sets, *Comput. Mater. Sci.* 232 (2024) 112661.
- [31] P. Chi, J. Sun, X. Luo, R. Cui, H. Dong, Reconstruction of 3D digital rocks with controllable porosity using CVAE-GAN, *Geoenergy Sci. Eng.* 230 (2023) 212264.
- [32] H. Zhang, L. Yang, C. Li, B. Wu, W. Wang, Scaffoldgan: Synthesis of scaffold materials based on generative adversarial networks, *Computer-Aided Des.* 138 (2021) 103041.
- [33] P. Siegkas, Generating 3D porous structures using machine learning and additive manufacturing, *Mater. Des.* 220 (2022) 110858.
- [34] M. Shi, F. Xie, J. Yang, J. Zhao, X. Liu, F. Wang, Cutout with patch-loss augmentation for improving generative adversarial networks against instability, *Comput. Vis. Image Underst.* 234 (2023) 103761.
- [35] E. Herron, J. Rade, A. Jignasu, B. Ganapathysubramanian, A. Balu, S. Sarkar, A. Krishnamurthy, Latent diffusion models for structural component design, *Computer-Aided Des.* (2024) 103707.
- [36] J. Wang, Z. Zhang, C. Han, Exploring latent fingerprint synthesis with diffusion probabilistic models, in: *Proceedings of the Future Technologies Conference*, Springer, 2024, pp. 78–89.
- [37] G. Dong, Y. Tang, Y.F. Zhao, A 149 line homogenization code for three-dimensional cellular materials written in matlab, *J. Eng. Mater. Technol.* 141 (1) (2019) 011005.
- [38] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10684–10695.
- [39] G. Felsch, V. Slesarenko, Generative models struggle with kirigami metamaterials, *Sci. Rep.* 14 (1) (2024) 19397.
- [40] A.B.L. Larsen, S.K. Sønderby, H. Larochelle, O. Winther, Autoencoding beyond pixels using a learned similarity metric, in: *International Conference on Machine Learning*, PMLR, 2016, pp. 1558–1566.
- [41] R. VanRullen, L. Reddy, Reconstructing faces from fMRI patterns using deep generative neural networks, *Commun. Biology* 2 (1) (2019) 193.
- [42] K. Uchida, T. Shibuya, Y. Takida, N. Murata, S. Takahashi, Y. Mitsufuji, MoLA: Motion generation and editing with latent diffusion enhanced by adversarial training, 2024, arXiv preprint arXiv:2406.01867.
- [43] J. Johnson, A. Alahi, L. Fei-Fei, Perceptual losses for real-time style transfer and super-resolution, in: *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, the Netherlands, October 11–14, 2016, Proceedings, Part II 14*, Springer, 2016, pp. 694–711.
- [44] J.H. Lim, J.C. Ye, Geometric gan, 2017, arXiv preprint arXiv:1705.02894.
- [45] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, *Adv. Neural Inf. Process. Syst.* 33 (2020) 6840–6851.
- [46] A.Q. Nichol, P. Dhariwal, Improved denoising diffusion probabilistic models, in: *International Conference on Machine Learning*, PMLR, 2021, pp. 8162–8171.
- [47] J. Song, C. Meng, S. Ermon, Denoising diffusion implicit models, 2020, arXiv preprint arXiv:2010.02502.
- [48] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, Springer, 2015, pp. 234–241.
- [49] P. Isola, J.-Y. Zhu, T. Zhou, A.A. Efros, Image-to-image translation with conditional adversarial networks, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1125–1134.
- [50] L. Van der Maaten, G. Hinton, Visualizing data using t-SNE, *J. Mach. Learn. Res.* 9 (11) (2008).
- [51] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al., Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* 12 (2011) 2825–2830.
- [52] A. Muntoni, P. Cignoni, PyMeshLab, Zenodo, 2021, <http://dx.doi.org/10.5281/zenodo.4438750>.
- [53] L.J. Gibson, ASHBY. MF cellular solids structure and properties, 1997, Edition: Cambridge Solid State Science Series.
- [54] I. Maskery, L. Sturm, A.O. Aremu, A. Panesar, C.B. Williams, C.J. Tuck, R.D. Wildman, I.A. Ashcroft, R.J. Hague, Insights into the mechanical properties of several triply periodic minimal surface lattice structures made by polymer additive manufacturing, *Polymer* 152 (2018) 62–71.

- [55] Bambu Lab X1-Carbon 3D Printer — uk.store.bambulab.com, 2024, <https://uk.store.bambulab.com/collections/3d-printer/products/x1-carbon-3d-printer>. (Accessed 10 June 2024).
- [56] D. De Aquino, I. Maskery, G. Longhitano, A. Jardini, E. Del Conte, Investigation of load direction on the compressive strength of additively manufactured triply periodic minimal surface scaffolds, *Int. J. Adv. Manuf. Technol.* 109 (2020) 771–779.
- [57] A. Brancaccio, S. Palladino, Contactless X-band detection of steel bars in concrete: A preliminary numerical and experimental analysis, 2024, Preprints.
- [58] X. Zheng, T.-T. Chen, X. Jiang, M. Naito, I. Watanabe, Deep-learning-based inverse design of three-dimensional architected cellular materials with the target porosity and stiffness using voxelized Voronoi lattices, *Sci. Technol. Adv. Mater.* 24 (1) (2023) 2157682.
- [59] X. Lyu, X. Ren, Microstructure reconstruction of 2D/3D random materials via diffusion-based deep generative models, *Sci. Rep.* 14 (1) (2024) 5041.