Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Generative design for additive manufacturing through data-free optimization-derived neural weight subspaces

Don Pubudu Vishwana Joseph Jayakody

dxj219@student.bham.ac.uk

University of Birmingham https://orcid.org/0000-0002-9303-4351

Bailin Deng

Cardiff University https://orcid.org/0000-0002-0158-7670

Ajit Panesar

Imperial College London https://orcid.org/0000-0001-8801-8288

Hyunyoung Kim

University of Birmingham

Research Article

Keywords: Generative design for 3D printing, Diffusion-based shape generation, AI-driven design optimization, Neural latent space, Design space exploration

Posted Date: November 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11507263/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Generative design for additive manufacturing through data-free optimization-derived neural weight subspaces

Don Pubudu Vishwana Joseph Jayakody[1,2,*], Bailin Deng[4], Ajit Panesar[5], Hyunyoung Kim[3,*]

[1]Department of Mechanical Engineering, University of Birmingham, United Kingdom

[2] Topological Design Centre for Doctoral Training, University of Birmingham, United Kingdom

[3]School of Computer Science, University of Birmingham, United Kingdom

[4]School of Computer Science and Informatics, Cardiff University, United Kingdom

[5]Department of Aeronautics, Imperial College London, United Kingdom

*Corresponding Authors (DPVJ Jayakody: dxj219@student.bham.ac.uk, H Kim: h.kim.4@bham.ac.uk)

ABSTRACT

Generative design (GD) enables exploration of diverse shape variants that exhibit trade-offs among competing design objectives related to part manufacturability and performance, offering a transformational approach towards automated part design for advanced additive manufacturing (AM). However, existing GD methods often derive design diversity either from large datasets of existing designs or through topological modifications, limiting their applicability to data-scarce engineering problems, particularly for complex parts where manufacturability and preservation of topology and geometric features are critical.

In this work, we present a training-free generative design framework that leverages the optimisation trajectory of a single neural field-based shape optimisation to represent and explore neural field variants within a low-dimensional neural weight subspace. Our design space exploration is realised through a Langevin diffusion-inspired stochastic process with energy-conditioned guidance, enabling the generation of constraint-satisfying shape variants that provide

distinct objective trade-offs. The framework generalises across geometrically and topologically diverse complex parts, including freeform geometries with intricate surface features and non-trivial topology. We validate the method through simulations and physical additive manufacturing experiments, demonstrating support-free 3D printing of the generated shape variants and establishing a data-free route towards manufacturability-aware generative design for complex parts.

Key words: Generative design for 3D printing; Diffusion-based shape generation; AI-driven design optimization; Neural latent space; Design space exploration

INTRODUCTION

Additive manufacturing, also known as 3D printing, offers tremendous flexibility in the digital-to-physical conversion of complex geometries via a layer-by-layer fabrication method [1]. While recent advances in AM and widespread availability of low-cost, high-quality 3D printers have driven its adoption across numerous applications ranging from customised interactive devices to high-value aerospace part manufacturing. Realising the full potential of AM, however, requires equally transformative advances in computational design [2,3]. Recently, AM-oriented design optimisation [4] has attracted significant research interest and industrial demand, driven by its potential in maximising manufacturability and part performance-related objectives such as overhang angle compliance (towards support-free 3D printing) [5,6], surface reflectivity [7], strength [8] while reducing reliance on trial-and-error manual part design.

While AM offers a vast engineering design space, the increasing complexity of modern part manufacturing (inevitably) demands for methods that enable exploration of diverse shape variants as a pre-requisite for deciding the part optimality [9]. This poses a significant challenge for

traditional AI-driven design optimisation [5, 10] methods relying on the well-known inverse design principle, which performs iterative refinement of part geometry and converges to a single local optimum based on pre-defined objectives (i.e., a deterministic exploration of the design space).

To address this challenge, Generative design (GD) has emerged as a transformative design paradigm that realises the generation of constraint-satisfying shape variants, enabling engineers explore trade-offs between complex and often-competing objectives, rather than converging to a single sub-optimal solution. This core advantage has positioned GD as an essential component of early-stage engineering design [11], with applications spanning microstructure generation for AM, metamaterial design [12,13,14], structural and aerodynamic optimization [15,16], building design [17] and parameterised car body profile design [18]. In practice, generative design is mainly realised through two approaches: 1) As a by-product of topology optimisation, in which the material distribution is exploited for the "generation" of topology variants [19,20]; 2). As a variant of generative modelling, in which a data distribution is first learned, and subsequently sampled using stochastic methods [15].

In the context of TO-based GD approaches, a common approach is to train a model using a large dataset to learn latent representations or surrogate models, as means to enable shape variant generation. For example, Kim et al. [21] used a large dataset of topology optimised brackets for generating 2D depth image-based diffusion model to enable generative synthesis of new bracket shapes for mass manufacturing. Recently, Moon et al. [22] proposed using Generative Adversarial Networks (GANs) in latent space trained on a large dataset for generative synthesis of 3D topology optimized structures. Nobari et al. [23] showed that diffusion-based implicit neural fields can be used for topology optimization-based shape variant generation. It is important to understand that

pure geometric change cannot guarantee diversity in performance metrics [24]. Building on this idea, Kwon et al. [24] proposed a data-driven framework for performance-aware geometry augmentation to realise topology optimised shape generation. However, TO-based GD methods are often inapplicable for many design optimisation problems that involve parts with non-trivial topologies (e.g., Voronoi structures) and complex freeform geometries such as aircraft wing shapes [16], 3D car body shapes [25] and bicycle seat models [26], where original shape and aesthetic detail preservation remain critical.

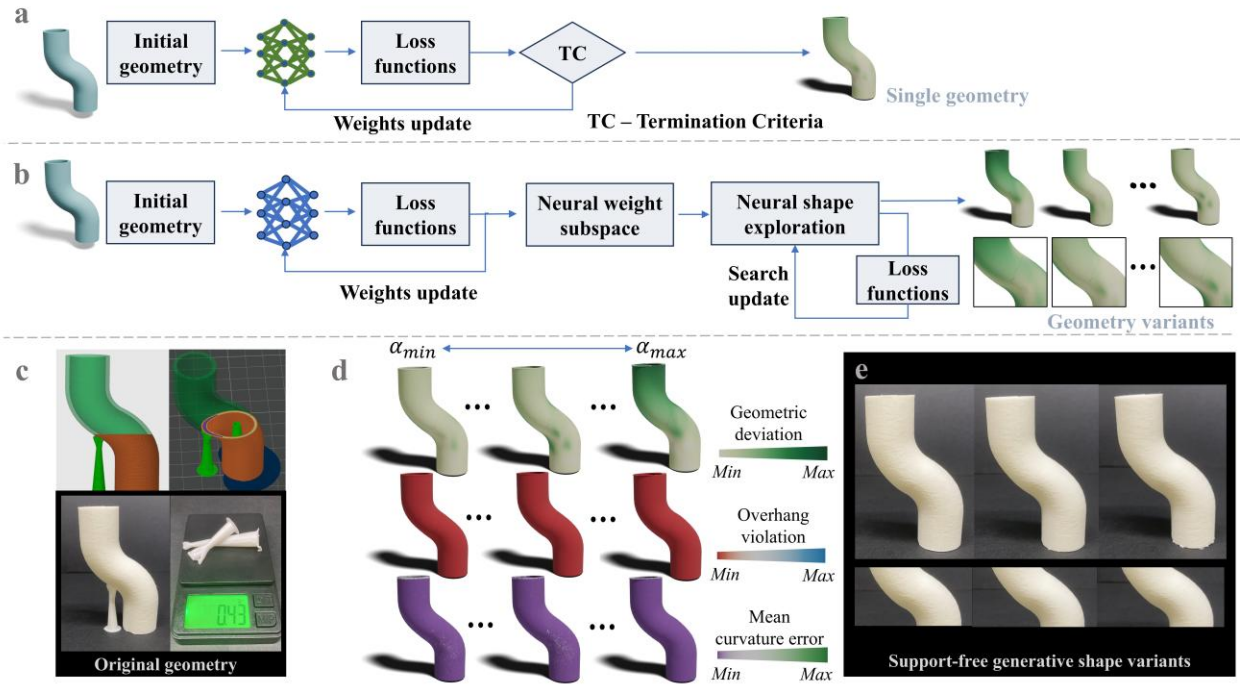


Fig 1: Overview of the neural field-based generative design concept (Representative geometry: S-pipe): (a) Conventional inverse design approach that finds a single optimal geometry; (b) The proposed generative design method which leverages neural field learning to design and explore a neural weight subspace for synthesis of performance-aware shape variants, offering unique objective trade-offs; (c) Original geometry that requires internal and external support structures; (d) Three selected geometry variants with unique geometric deviation (top), overhang (middle) and mean curvature error scalar maps and corresponding (e) support-free 3D printed geometry variants, eliminating material waste.

For freeform complex parts, shape variant generation mainly resorts to generative modelling, in which the idea is to learn a latent space for a given shape distribution and generate new samples by interpolation methods (e.g., diffusion model-based techniques). Danhaive et al. [15] showed that performance conditioned low dimensional latent space can be useful for targeted design space exploration of complex structures. Yang et al. [25] presented PhysGen framework, which learns a combined latent representation of geometry and physics-based properties for industrial design cases including 3D car body shapes with intricate details. Diffusion-based generation of implicit neural fields in latent space [27] has been employed for physics-based optimization of freeform shapes [8]. To minimise dependency on large datasets, Tewari et al. [28] proposed an integrated diffusion process with forward model for inverse design-based scene generation with minimal data (i.e., without direct supervision). Murali et al. [29] presented an interesting VAE-based approach that augments sparse shape data to create larger datasets to learn a latent space to generate human pose shapes. Similarly, diffusion-based sampling in the latent space has been proposed by several recent methods for microstructure synthesis for additive manufacturing [13,14,30].

In general, generative modelling-based methods depend on high-quality input data to learn a latent space, which can be challenging for data-scarce applications such as additive manufacturing-oriented part design. Acquiring and generation of performance data (for design optimization) are generally computationally expensive, making performance-aware shape variant generation challenging, while also diminishing manufacturability awareness within the generative process. Berzins et al. [31] proposed Geometry-Informed Neural Networks (GINNs), a data-free shape generation framework that trains neural fields based on design constraints and diversity constraints. While this method [31] well-suited for underdetermined design problems where large

geometric change is permissible (e.g., shape variant generation for a topology optimization problem), generating shape variants of an existing part while retaining its characteristic geometric features presents a different challenge. In fact, generative design of arbitrary complex geometric models characterised by non-trivial topology and intricate surface features for data-scarce applications has remained an unresolved problem. To the best of authors knowledge, generative design of manufacturable (e.g., 3D printable) and constraint-satisfying geometry variants of existing complex shapes in a data-free manner has been rarely (or never) studied in the existing literature.

Here we propose a generative design paradigm for advanced additive manufacturing, in which the engineering design space becomes the subject of shape representation and exploration, to enable data-free generation of manufacturability-optimized shape variants for a range of complex 3D freeform geometries (see Fig. 1). Given a neural field learned as an energy minimisation problem [5], and a "neural shape" (i.e., shape variant) can be created by imposing a neural vector field on the original shape, we hypothesise that the neural weight space paves the way for generation and exploration of a local shape space [32] without the need for training data. Our framework indeed shows how inverse-design based shape optimization can be converted into a generative design method for 3D printing. We introduce a low-dimensional neural weight subspace for effective design space exploration, for which we propose a Langevin diffusion-inspired stochastic approach coupled with energy term conditioning, to find low-energy variants within the optimisation landscape (see Fig. 2). This approach enables generation of neural shape variants, enabling generative design of geometric variants for arbitrary complex parts that offer unique design trade-offs while complying with manufacturability constraints. Therefore, the proposed data-free

generative neural shape method not only supports decision-making for initial design stage but also offers a fundamental shift in how generative design is performed for additive manufacturing.

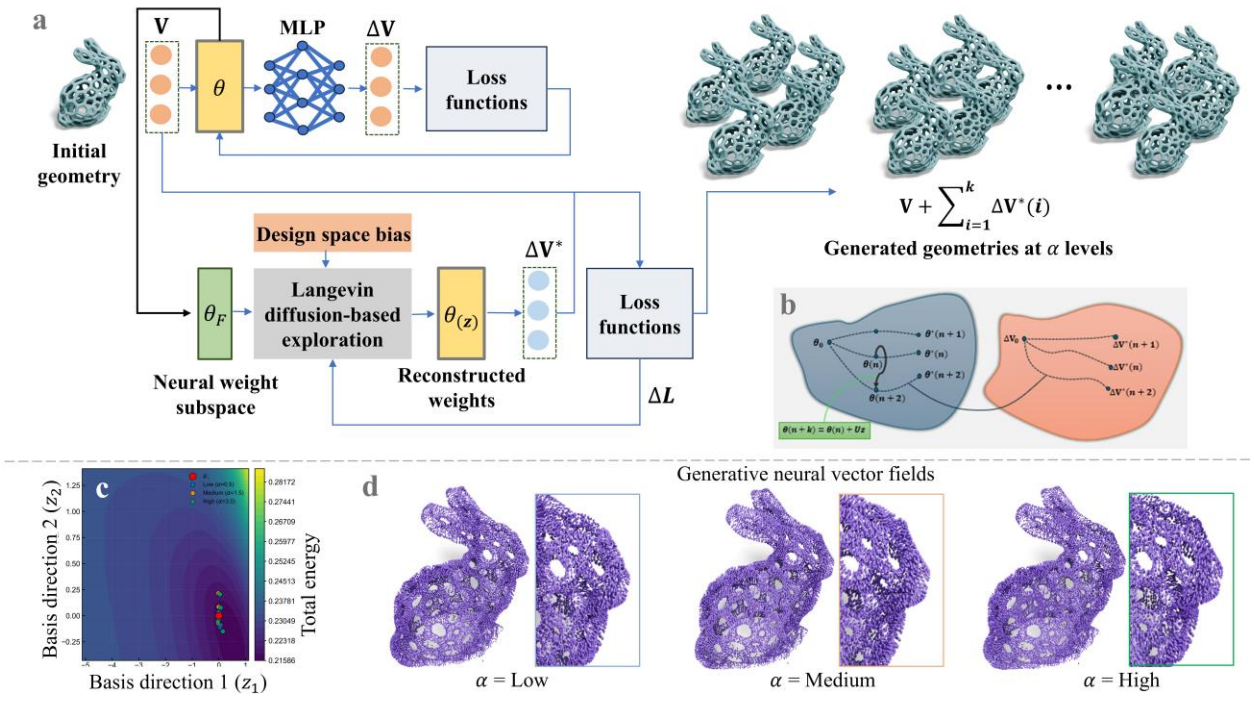


Fig 2: Generative design computational framework (wireframe bunny model): (a) The full computational pipeline which generates a low dimensional design space (neural weight subspace) to enable Langevin diffusion-based exploration of neural vector fields for data-free compliant shape generation of complex geometries including non-trivial topologies (right); (b) The correspondence between neural weight generation and neural field synthesis; (c) A representative example of generated shape variant samples on shape energy (2D) plot and (d) Representative neural vector field variant examples across three α (shape exploration) levels.

RESULTS

Low dimensional subspace of design variants

Our framework considers an input triangle mesh geometry $\mathcal{M}$ represented as a set of vertices $\mathbf{V} = \{\mathbf{v_i} \in \mathbb{R}^{n\times 3} | i = 1, \dots, n\}$ and faces $\mathbf{F} \in \mathbb{Z}^{m\times 3}$, and first learns an initial neural field $\Phi_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ [33,34] based on a shape energy loss function $\mathcal{L}$ to compute the optimal geometric change. The objectives of $\mathcal{L}$ (see Methods), which govern the optimality of the geometric change, are formulated to optimise overall part manufacturability (i.e., printability) such that the generatively designed shape variants are not only diverse in terms of geometry but also offer unique performance metrics, which can be crucial for advanced manufacturing applications. Specifically, our framework focuses on reducing overhang regions which increase print time and material waste due to support structure generation, while complying with other objectives (e.g., minimal deviation from original geometry) crucial for part quality. Conventional inverse design-based approaches use the learned neural vector field $\Phi_\theta(\mathbf{V})$ to extract the optimised geometry as $\mathbf{V}' = \mathbf{V} + \Phi_\theta(\mathbf{V})$ and iterate until the termination criteria are met (see Fig.1). In contrast, our goal is to employ the neural weight space defined by θ as a parameterised generator (i.e., sampler) of compliant neural field variants to realise data-free generation of manufacturability-improved shape variants (i.e., low-energy variants of $\mathcal{L}$). We hypothesise that this design problem can be solved by direct sampling of the updated (i.e., learned) neural weight parameter space according to a Boltzmann distribution as follows:

$$p(\theta) \propto \exp\left(-\frac{\mathcal{L}(\theta_c)}{T_s}\right) \quad (1)$$

where θ_c denotes the weight space of the initial neural field learning and T_s represents a temperature coefficient of the sampling process. However, due to its high-dimensionality, direct

exploration (e.g., via sampling) can be not only computationally prohibitive but also lead to uncontrollable geometric deformation. Existing frameworks that employ neural weight space for shape generation tasks [35,36] generally depend on a training on a large dataset, to either interpolate or apply diffusion-based sampling [35] to create new weights. In contrast, our method focuses on the objective landscape information [37] provided by the neural weight trajectory, which has been overlooked in existing generative design methods. We show that given a single neural field learned through an energy minimisation problem, a low-dimensional neural weight subspace can be generated without additional training (e.g., to learn a global prior), to realise effective exploration of local neural field variants (see Fig. 2). For this, neural weights generated at every n iteration during the initial neural field learning are recorded as a flattened vector as follows:

$$\boldsymbol{\theta^t} = \left[{\theta_1}^t, {\theta_2}^t \dots \dots, {\theta_P}^t\right] \in \mathbb{R}^P \tag{2}$$

in which P represents the number of neural weights (i.e., parameter space) that depends on the size of the number of hidden layers q and neurons r of the MLP. $t = 1, \dots, T$ represents the number of neural weight vector samples recorded during the initial neural field learning. Therefore, the full neural weight matrix $\mathbf{W} \in \mathbb{R}^{T \times P}$ can be represented as follows:

$$\mathbf{W} = \begin{bmatrix} {\theta_1}^{(1)} & {\theta_2}^{(1)} & \cdots & {\theta_P}^{(1)} \\ {\theta_1}^{(2)} & {\theta_2}^{(2)} & \dots & {\theta_P}^{(2)} \\ \vdots & \vdots & \ddots & \vdots \\ {\theta_1}^{(T)} & {\theta_2}^{(T)} & \dots & {\theta_P}^{(T)} \end{bmatrix} \tag{3}$$

$\mathbf{W}$ can therefore be considered a combined representation of the weight trajectory and the full parameter space, providing useful information of the shape energy landscape. However, its inherent high dimensionality can be problematic for subsequent processing, due to high

computational cost. To overcome this, we propose identifying the directions of T ($\ll P$) with highest variations (i.e., directions along which the shape energy has been permissive), which is significantly efficient compared to perturbing the full parameter space. To this end, we first compute the covariance matrix of $\mathbf{W}$ as follows:

$$\mathbf{C} = \mathbf{W_c}\mathbf{W_c}^{\boldsymbol{T}} \in \mathbb{R}^{T\times T} \tag{4}$$

in which $\mathbf{W_c} = \mathbf{W} - \bar{\theta}$ represents the mean centred weight matrix while $\bar{\theta} = \frac{1}{T}\sum_{t=1}^{T}\theta^t$ represents the mean weight vector. To formulate neural weight subspace $\mathbf{B}$ as a low dimensional basis of $\mathbf{W}$, we perform Singular Value Decomposition of $\mathbf{C}, \mathbf{SVD}\ (\mathbf{C}) = \mathbf{U\Sigma U}^{\top}$ and compute a k number of basis vectors (i.e., principal directions), as follows:

$$\mathbf{B} = \mathbf{W_c}^{\boldsymbol{T}}\mathbf{U_k} \in \mathbb{R}^{P\times k} \tag{5}$$

Note that $\mathbf{B}$ represents a design search space spanned by the dominant weight trajectory directions generated during the initial energy minimisation-based neural field learning. We propose this low dimensional neural weight subspace as the key to enabling data-free low-energy neural vector field variant generation. To better understand the influence of the weight trajectory, we first investigate how the weight recording frequency affects the subspace and generated shape variants by constructing the neural weight subspace using three different weight-recording frequencies. As shown in Fig. 3, our results reveal that selecting the weight recording frequency is a trade-off between the design space control and computational cost. More frequent recording (i.e., fine weight trajectory representation) improves the shape energy compliance (see Fig. 3(a)) of the generated neural vector fields, but also increases the weight matrix size, leading to increased computational cost (see Fig. 3(f)). For instance, neural weight subspace generated by weights

saved at every three or five iterations enable shape variant generation with consistent overhang reduction performance under lower deviation.

Data-free design space exploration via Langevin diffusion

The next step is to employ the neural weight subspace generated above to create diverse neural shape variants with distinct objective trade-offs. To this end, we introduce a latent vector $\mathbf{z} \in \mathbb{R}^k$, and formulate a stochastic neural weight reconstruction (i.e., mapping from weight subspace to full weight space) as follows:

$$\theta^*(\mathbf{z}) = \theta_c + \mathbf{B}\mathbf{z} \tag{6}$$

in which $\theta^*(\mathbf{z})$ represents the reconstructed (full) neural weight set. Note that (stochastic) perturbation of $\mathbf{z}$ leads to generation of unique combinations of basis directions, creating new neural weight sets that correspond to unique neural vector field variants without the need for additional data. To explore the low dimensional neural weight subspace in a controllable approach, we propose a Langevin diffusion (i.e., Langevin dynamics) theory [38]-based sampling algorithm, for generating variants of $\mathbf{z}$ as follows:

$$\mathbf{z}_{x+1} \leftarrow \mathbf{z}_x - \eta(\nabla_{\mathbf{z}} E + \beta \mathbf{v_c}) + \sqrt{2\eta T_s}\xi_x \tag{7}$$

where η, T_s, $\xi_x \sim \aleph(0, \mathbf{I})$, β and $\mathbf{v_c}$ represent the step size, temperature coefficient, Gaussian noise, control coefficient, and the directional control vector, respectively. $x = 0, \ldots, N_L$ represents the number of Langevin diffusion steps, while the temperature coefficient is defined as $T_s = \alpha\sigma_E$, in which $\sigma_E = 2$ denotes the energy scale parameter, and α represents the shape exploration (i.e., stochasticity) level parameter. It is important to understand that when the shape energy loss combines several objectives as in our framework, the proposed method induces low-energy neural

vector field variants that capture different trade-offs among the objectives. However, as the neural weight subspace is derived as a byproduct of the combined energy, explicit control over individual objectives (during shape generation) can be restricted. This can be a drawback for manufacturing-oriented applications, where the relative importance of design objectives can vary based on multiple factors including user requirements and process settings. To this end, we introduce an energy-conditioned design space exploration by identifying control directions for individual energy terms within the neural weight subspace. Let $\mathcal{L}_j$ denote the j^{th} energy term of the total shape energy $\mathcal{L}$ (see Methods), and the corresponding "control" direction can be estimated by projecting the gradient of $\mathcal{L}_j$ to the neural weight subspace, as follows:

$$\mathbf{g}_j = \mathbf{B}^\top \nabla_\theta \mathcal{L}_j \tag{8}$$

Therefore, the directional control vector of $\mathcal{L}_j$, $\mathbf{v}_{\mathrm{c(j)}} = \mathbf{g}_j \Big/ \|\mathbf{g}_j\|$ (i.e., the direction that mostly influences $\mathcal{L}_j$) in the neural weight subspace can be weighted by β (as shown in eq. (7)) to provide user control over the energy conditioning of neural field variant generation. For instance, a positive bias (e.g., +0.5) assigned to the overhang energy term control vector generates consistent overhang-minimised shape variants of the input model while providing unique distributions of mean curvature error, across multiple shape exploration α levels (see Fig. 3 (a-c)). In contrast, applying the same bias to the smoothness energy control direction enables the generation of shape variants with minimised mean curvature error, demonstrating the effect of the choice of energy term bias (see Fig. 4 (c-d)).

The sampling approach of eq. (7) realises an iterative refinement of the latent vector z (i.e., to find low-energy neural field variants), while shape exploration parameter, α, has a direct influence over the design space search region (see Fig. 4). Compared to assigning a fixed value, we observe that

multiple values within a pre-defined range for α (e.g., $\alpha_{low}, \alpha_{medium}, \alpha_{high}$) lead to a more systematic search, paving the way for distinct objective trade-offs among the generated neural shape variants. For example, Fig, 4(b) shows that the generated shape variants of Armadillo model offer unique trade-offs between overhang violation and deviation energy, demonstrating its advantage over single neural shape generated by the initial neural field [5] (i.e., inverse design approach as shown in Fig. 1 (a)). Thus, the design space exploration presented in this work provides a mechanism for energy-conditioned and controllable generative design for additive manufacturing.

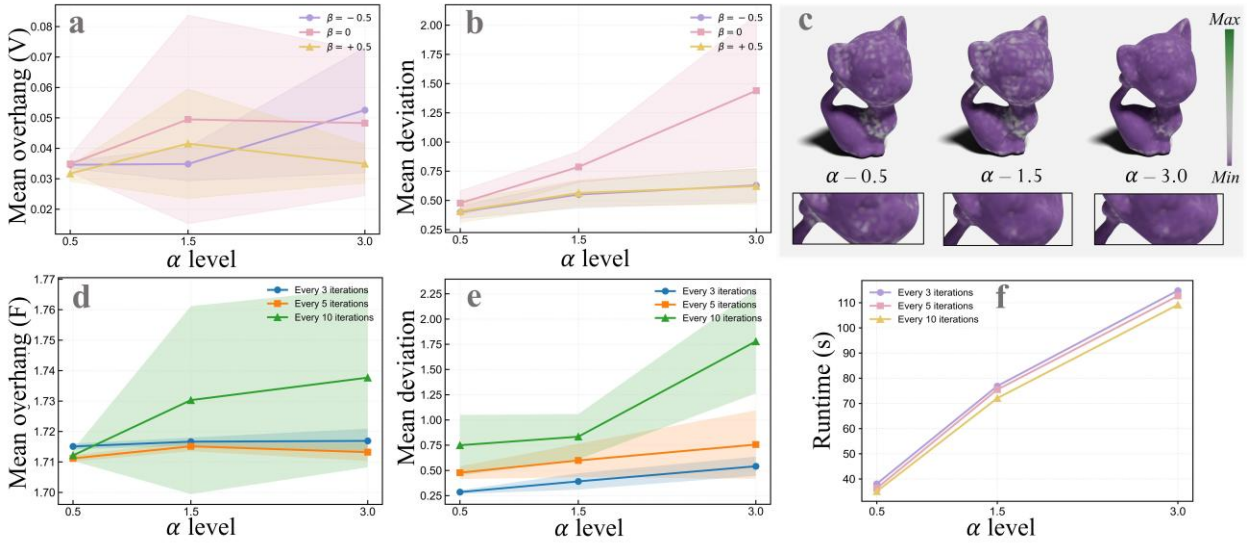


Fig. 3: Performance indicators of the generative design framework: (a) Mean overhang (vertex-level) and (b) mean deviation (from original geometry) variation across three different β values with identical (overhang) energy conditioned shape generation. (c) Mean curvature error scalar field maps for representative shape variants generated at different α values. (d) Mean overhang (face-level) and (e) mean deviation variation with different weight recording frequencies and (f) the effect of weight recording frequency on the computational runtime for generating neural shape variants.

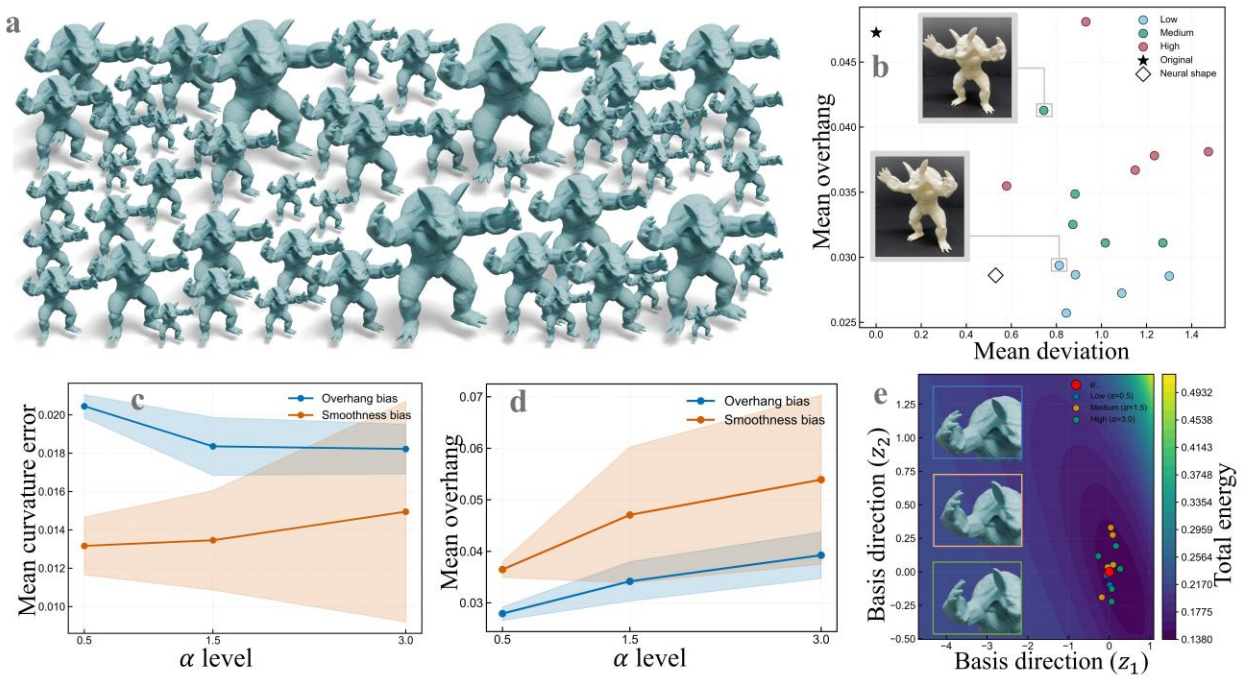


Fig. 4: Energy-conditioned generative design (Armadillo model): (a) An example 50-model dataset generated under 360 seconds (across five α levels, ten shape variants per α level); (b) Overhang angle violation vs. Mean deviation error trade-off plot for original, deterministic neural shape [5] and generated shape variants across three α levels; The influence of control (i.e., bias) energy term on (c) mean curvature error and (d) mean overhang violation of generated shape variants (across three α levels, five shape variants per α level); (e) Energy plot defined by two principal directions of the neural subspace with generated shape variants.

Table 01: Computational time of the neural field-based shape generation

Model	Vertices	Faces	Computational time				
			Initial neural field and weight subspace⋄	Shape generation⋆			
				$\alpha = 0.5$	$\alpha = 1.5$	$\alpha = 3.0$	Total time
Kitten	4103	8206	34.12	35.61	71.82	108.62	142.74
Wireframe Bunny	17279	35094	89.00	99.11	194.83	312.65	401.65
Yoda	54110	111300	267.33	280.51	576.21	880.00	1147.33
S-pipe	30609	61218	152.99	161.39	325.13	489.99	642.98
Krillin	45593	91280	222.80	248.07	490.45	720.97	943.77
Armadillo	5117	10230	39.55	41.46	85.85	130.72	170.26

⋆ $s = 5$ *at each alpha level (i.e., a total of 15 shape variants for each mesh)*

⋄ *Neural subspace computation time is between 0.04 -0.08 for the tested models*

All experiments were run with design space search bias set to overhang energy term

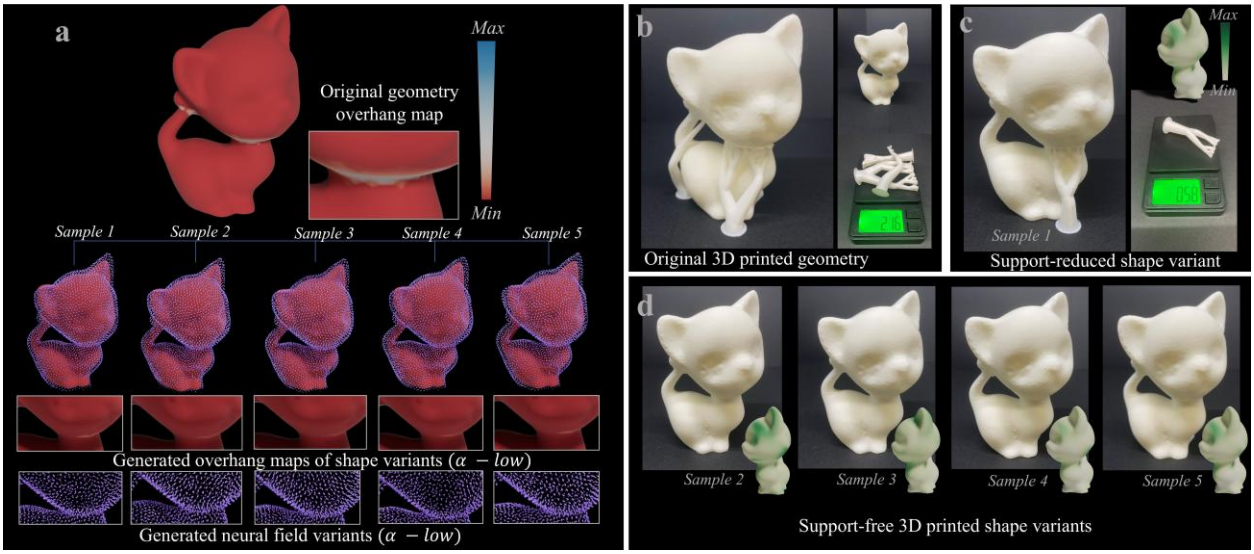


Fig. 5: Generative design for support-free 3D printing (Kitten model): (a) Overhang scalar field map of the original geometry (top) and overhang scalar field maps of generated shape variants (middle) with corresponding neural field variants (bottom); (b). 3D printed original geometry and the required support structures; (c) A 3D printed shape variant (α-Low) with reduced support structures (e) Support-free 3D printed shape variants (α-Low) generated by the proposed generative design framework

Generative neural shapes for manufacturability optimization

For most of the presented experiments, the user-defined number of shape variants (at each α level) $s = 1, \dots, N_S$ is set to five across three α exploration levels: $\alpha_{low} = 0.5, \alpha_{medium} = 1.5, \alpha_{high} = 3.0$. However, it can be easily changed depending on the nature of the application as shown in Fig. 4 (a), which depicts a 50-model dataset of Armadillo model generated with $N_S = 10$ across five different α levels. Therefore, the generative neural vector field variants at a given α level can be represented as:

$$N_\alpha = \{\Delta\mathbf{V}_1, \Delta\mathbf{V}_2, \dots, \Delta\mathbf{V}_{N_S}\} \tag{9}$$

Note that each (unique) $\Delta\mathbf{V}_s = \Phi_{\theta_s}(\mathbf{V})$ is computed using the reconstructed neural weights as $\theta_s = \theta_c + \mathbf{B}\mathbf{z}_s$. Consequently, the resultant geometric variants generated by imposing neural vector field variants on the input geometry (see Methods) are evaluated using the full shape energy loss function $\mathcal{L}$ to refine the latent vector $\mathbf{z}$ sampling process. The proposed framework enables efficient (see the computational time presented in Table 01) shape variant generation for a range of complex geometries that include intricate surface features and non-trivial topology.

More importantly, our method induces distinct, objective-conditioned distributions of scalar field maps (see Fig. 6(a-b)) while satisfying the required manufacturing constraints including fixed regions of the part geometry. This is particularly significant to additive manufacturing as it enables engineers to explore multiple manufacturability-aware shape variants and select designs according to application-specific trade-offs (see Table 02). For example, the S-pipe (see Fig. 1) is a complex mechanical part that originally requires support both externally and internally. Interestingly, the generated shape variants of S-pipe become overhang-free with distinct distributions of deviation and mean curvature error while ensuring the top and bottom surfaces remain fixed (i.e., no

displacement). The physical experiments presented in Fig. 1(e) further demonstrate how the generatively designed shape variants can be 3D printed without support structures (both externally and internally), showcasing the potential of the proposed method in minimising material wastage and print costs in additive manufacturing.

We further demonstrate the generality of our generative framework by performing simultaneous generation of manufacturability-improved shape variants for complex freeform geometries, compared to design optimization methods [5,6,7,8] that are limited by a single sub-optimal shape (see Fig. 4(b)). For example, among the five generated neural vector field variants of the Kitten model (at α_{low}), four completely eliminate overhang violations, while the remaining shape variant substantially reduces the overhang threshold-violating region, requiring minimal support compared to the original geometry. The proposed method is applicable to topologically non-trivial parts such as Wireframe bunny (see Supplementary Fig. S1) and Yoda models enabling the generative design of topology-preserving shape variants with distinct objective trade-offs, as shown in Fig. 6 (a-c). In addition, our energy-compliant neural field variant generation method can pave a new path for efficient dataset augmentation, as demonstrated by the example of Armadillo model dataset generation in under 360 seconds (see Fig. 4(e)). Supplementary Fig. S2 presents the physical validation of representative Armadillo shape variants, demonstrating the effectiveness of generative local shape changes induced by the proposed method for manufacturability improvement (i.e., energy compliance). Finally, we remark that this framework can also be useful as a post processing tool [39] for topology optimization (especially for neural field driven TO methods [40,41]) for finding efficient design trade-offs including manufacturability, under minimal deviation (i.e., minimal change) from the optimal topology layout, as shown in Supplementary Fig. S3.

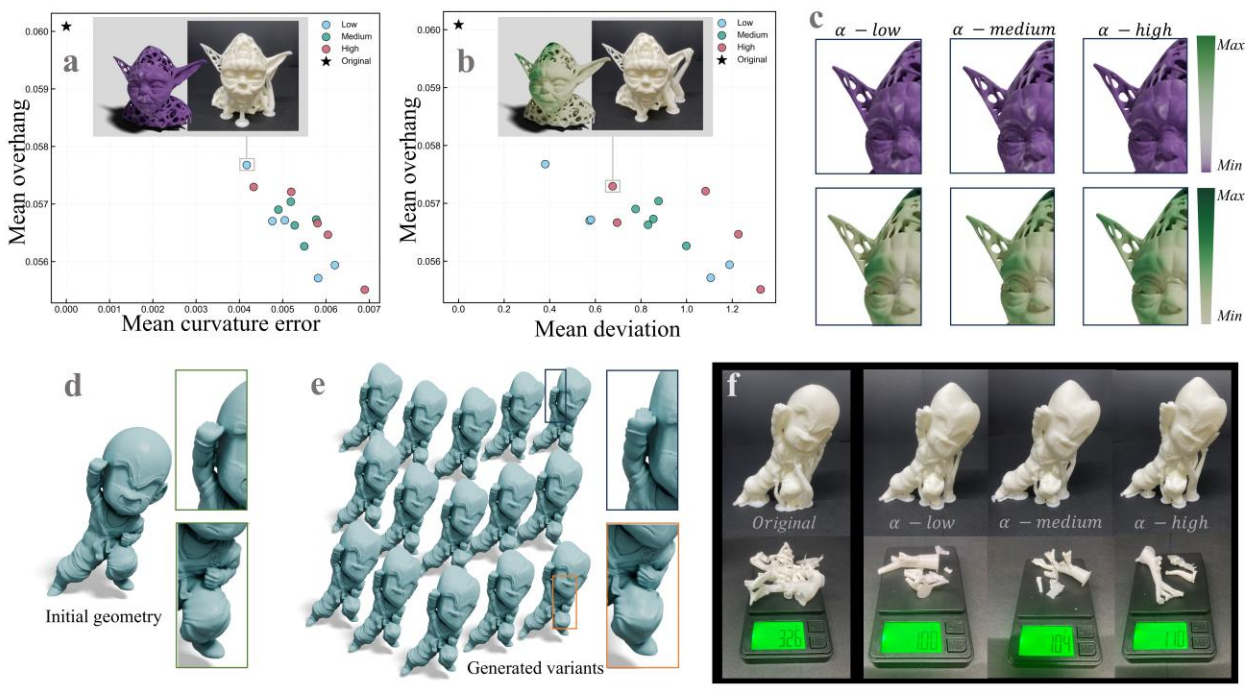


Fig. 6: Shape space exploration for objective trade-offs: (a) Overhang angle violation vs. Mean curvature error trade-off plot for original and generated shape variants across three α levels (Yoda model); (b) Overhang angle violation vs. Mean deviation error trade-off plot for original and generated shape variants across three α levels (Yoda model); (c) Representative shape variants of different α levels showing unique distributions of mean curvature and deviation (objective) values; (d) Initial geometry and generated shape variants (Krillin model); (e) 3D printed part comparison of original and representative shape variants (Krillin model)

Table 02: Computational statistics summary of the proposed generative framework

Model	Original geometry		Generated geometry samples				
	Overhang face ratio	Support structure print time (min)	Overhang face ratio (mean)	Support structure print time (mean)	Mean curvature error (mean)	Deviation (mean)	Reduced support print time (average)
S-pipe⋆	0.25	2	0.0	-	0.002	0.31	100%
Kitten⋆	1.04	16	0.14	0.2	0.009	0.35	98.75%
Yoda⋄	6.00	14	5.66	4.0	0.006	1.00	71%
Wireframe Bunny⋄	5.66	43	4.81	17.2	0.023	0.95	60%
Krillin⋆	5.72	10	0.90	3.2	0.005	0.52	68%

⋆ *Shape variants are generated with α- low setting,* ⋄ *Shape variants are generated with α- high setting*

Limitations and future work

Diffusion-based sampling on neural weight subspace

Currently, the neural weight subspace represents a local search space within the energy landscape, and as a result, the shape variants are indeed local geometric variants of the input geometry [32,42]. While this is desired for manufacturing-oriented design, future work can explore extending the current exploration towards finding more global shape variants [43]. While we employ a fixed value for the diffusion process iteration number, it is important to further investigate the influence of this on the quality and the nature of neural field generation, which we leave as future work. Although our framework is highly efficient for generative design of complex parts, the total computational cost can increase depending on the mesh resolution. For industrial applications where fine resolution of part geometry is required, future work can focus on integrating effective techniques such as coarse-to-fine neural field finetuning [5] or transitioning to Graphics Processing Unit (GPU)-based computing to maximise efficiency.

Search space control for generative neural fields

Our current framework achieves efficient, data-free generation of neural fields that produce geometric variants compliant with the shape energy and constraints. However, the proposed formulation of the shape energy loss inherently couples the individual energy terms, making independent control within the neural weight subspace challenging. For example, applying a bias to the overhang energy can simultaneously alter the other energy terms. Therefore, future work should focus on improving the control of energy conditioning [44,45], which can be crucial to enable a more unified neural field-based generative design framework with more objectives related to manufacturing processes and part functionality (e.g., part strength). For example, direct learning of PDE-based properties [46,47] on neural surfaces with minimal dependency on discretisation, can provide a meaningful extension for the proposed framework.

User-guided generative design for 3D printing

Our framework provides the foundation for user-controlled generative design of complex shapes, where the user input (e.g., shape exploration level, energy bias) has direct influence over the design space exploration. In future work, an integration of a user-friendly UI can establish a higher level of intuitiveness of the shape generation towards mass customization of complex parts [48]. In addition, an interesting extension would be to incorporate Large Language Model (LLM)-based user inputs into the presented shape generation [49, 50], enabling a more intent-aware generative design process for 3D printing that can better incorporate high-level user requirements.

DISCUSSION

The data-free generative design framework presented in this work demonstrates how a given neural field-based optimisation process can be converted into a source of design space exploration by constructing a low dimensional neural weight subspace that captures the optimisation trajectory. This enables the generation of shape variants across a range of complex geometries, offering distinct trade-offs between objectives related to part manufacturability and quality. In general, our framework establishes a fundamental shift from single-solution inverse design towards generative design, while enabling performance-aware shape generation for complex and data-scarce additive manufacturing applications without datasets.

As the neural weight subspace enables representation and exploration of shape variants, it significantly improves the generative process by preventing fragmented processing inherent to traditional data-driven generative design methods. The proposed energy-conditioned generation of neural field variants provides distinct trade-offs between complex and often-competing objectives based on design preferences, which demonstrates its significance towards advanced additive manufacturing applications, where the relative importance of design objectives often varies.

The physical 3D printing results further validate the effectiveness of the proposed method and more importantly, the manufacturability of generatively designed shape variants of complex geometries while minimising material waste and print time. We remark that our generative framework is process agnostic and can therefore be applicable or extended for other manufacturing methods (especially for different additive manufacturing methods), provided that the corresponding process-aware objectives and constraints guide the optimisation of the neural field.

While our work provides a transformational foundation for additive manufacturing oriented-generative design without training data, future work directions can focus on investigating the diversity of the neural field optimisation trajectory, LLM-based user intent integration and combining more part performance related objectives, which altogether can lead towards a more unified, human-in-the-loop [51] generative design process.

METHODS

Implementation

Our neural shape optimisation framework is implemented in Python on a laptop with an Intel(R) Core ™ i7 -12850HX CPU (2.1 GHz) + 32 GB RAM, running on Windows 10. We implement the neural field training process using the PyTorch library [52] with ReLU activation. ADAM optimiser [53] is used to optimise the neural weights using backpropagation. We use an MLP-based neural field with $q = 5$ hidden layers and $r = 192$ neurons per layer for all the experiments with a learning rate of $1e^{-4}$. In addition, the regularisation weights of the total loss function are fixed as follows: $\lambda_s = 0.1$, $\lambda_d = 0.05$ and $\lambda_o = 10$. Neural weight subspace is constructed with $k = 20$ basis directions while the iteration number $N_L = 80$ of the proposed Langevin diffusion-based algorithm is fixed for all experiments to manage the computational cost.

All our physical 3D printing experiments were conducted on a Bambu X1 Carbon 3D printer using Polylactic Acid (PLA) white filament material and Bambu Lab (Version 2.0.2.57) slicing software program was used for G-code generation. For all experiments, print speed, support generation overhang limit, layer height and infill percentage are set to 200 m/s (default speed), 60, 0.2mm and 15% (for fast printing), respectively.

Neural field learning

Our framework initiates by learning an initial neural vector field Φ_θ based on a shape energy loss function $\mathcal{L}$, to learn an intermediary optimal shape to govern the generative process. We refer the reader to the review by Xie et al. [54] for a detailed overview of neural field-based design. Note that Φ_θ can be evaluated at a given vertex $\mathbf{v_i}$ to compute the corresponding displacement vector (i.e., geometric change) as $\Delta\mathbf{v_i} = \Phi_\theta(\mathbf{v_i})$. We incorporate user-defined fixed regions (i.e., $\Delta\mathbf{v_i} = \mathbf{0}$) that prevent displacements in important (e.g., functional) regions of the part geometry, which is crucial for complex parts used in engineering applications. Then, by using the full neural vector field, $\Phi_\theta(\mathbf{V})$, the updated geometry can be represented as $\mathbf{V'} = \mathbf{V} + \Phi_\theta(\mathbf{V})$. We follow the conventional backpropagation rule to iteratively update the set of neural field weights θ, for given number q of iterations as follows:

$$\theta_c \leftarrow \underset{\theta}{argmin}\ \mathcal{L}\ (\Phi_\theta, \mathbf{V}, \mathbf{F}) \tag{10}$$

3D printability-oriented loss functions

Here we introduce the formulation of the energy E that governs the optimality of the neural deformation field in our 3D printing oriented-design optimisation framework. Our goal is to improve manufacturability through support structure reduction, while ensuring smooth geometric change (i.e., avoid sharp undesirable features) and minimal deviation from the original shape [5]. For this purpose, we introduce three loss functions as below:

Overhang loss:

To achieve support structure reduction in a geometry-agnostic manner (e.g., independent of tree-type or column-type supports), our method focuses on minimising overhang (i.e., unsupported) regions of $\mathcal{M}$, by introducing a vertex-level overhang loss as follows:

$$\mathcal{L}_o = \frac{1}{3n}\sum_{i=1}^{n} \sum_{f \in F_i} w_f \, \psi_f \tag{11}$$

In which $\psi_f = max(0, \cos \alpha_{max} - \mathbf{n}_\mathrm{f} \cdot \hat{\mathbf{d}})$ represents the per-face violation for a given face $f = (\mathbf{v_1}, \mathbf{v_2}, \mathbf{v_3}) \in \mathbf{F}$, where $|\cdot|$ denotes the dot product. $\mathbf{n}_\mathrm{f}$ and $\hat{\mathbf{d}}$ represent face normal vector and fixed print direction, respectively. We use face height-based weighting, $w_f = \frac{h_f}{\max_{f \in \mathbf{F}}(h_f)} \in [0,1]$, where $h_f = \frac{1}{3}(\mathbf{v_1} + \mathbf{v_2} + \mathbf{v_3})$ represents the (centroidal) height of f. Note that the overhang angle $\alpha_f = arccos(\mathbf{n}_\mathrm{f} \cdot \hat{\mathbf{z}})$ is used to determine overhang faces $\{F_{ov}\} \subseteq \mathbf{F}$ that exceed a user-defined overhang threshold, $\alpha_{max} = 135°$.

Smoothness loss:

Smoothness of the geometry induced by the optimal neural deformation vector field is important for manufacturing-oriented applications. While neural vector field-based shape deformation generally provides smooth displacements [5], discrete evaluation of Φ_θ at a given set of vertices can still be prone to local distortions (e.g., sharp dents). To minimise such distortions and improve the smoothness of the geometric change, we incorporate a direct smoothness loss using the cotangent Laplacian matrix $\mathbf{L}$ as follows:

$$\mathcal{L}_s = \frac{1}{n}\sum_{i=1}^{n} \|\mathbf{L}\mathbf{v_i}'\|^2 \tag{12}$$

Deviation loss:

Next, it is important that the deviation between the generated geometries and the original geometry remains minimal. For this purpose, we introduce the following loss function that penalises large deviations induced by the Φ_θ during the initial neural field learning:

$$\mathcal{L}_d = \frac{1}{n}\sum_{i=1}^{n}\|\mathbf{v_i}' - \mathbf{v_i}\|^2 \tag{13}$$

Total loss formulation and termination:

Our total loss function, which guides the initial neural vector field learning phase can be formulated by combining the above individual loss function as follows:

$$\mathcal{L} = \lambda_o\mathcal{L}_o + \lambda_s\mathcal{L}_s + \lambda_d\mathcal{L}_d \tag{14}$$

where λ_o, λ_s and λ_d represent the regularisation factors for the loss functions. Note that although the continuity of Φ_θ enables mesh-free learning of the optimal shape, termination criteria can be useful to guide the neural field learning effectively, depending on the nature of the application. In our case, we adopt an energy stagnation check as Jayakody et al. [5], and introduce a per-vertex geometric deviation threshold ϵ_{close} for the initial neural field through a soft constraint as follows:

$$^1/_{|V|}\left\|\mathbf{V} - \mathbf{V}_{(0)}\right\|_F^2 \leq \epsilon_{close} \tag{15}$$

Credit authorship contribution statement

Don Pubudu Vishwana Joseph Jayakody: Conceptualisation, Investigation, Methodology, Software, Validation, Visualisation, Writing – original draft, Writing – review and editing, **Bailin Deng:** Conceptualisation, Writing – review and editing, **Ajit Panesar**: Conceptualisation, Writing – review and editing, **Hyunyoung Kim:** Conceptualisation, Resources, Writing – review and editing, Supervision

Acknowledgements:

This work has been funded by the Engineering and Physical Sciences Research Council, UK Ph.D. Studentship EP/S02297X/1. We thank Mark Dennis (School of Physics and Astronomy, University of Birmingham) for helpful comments on this work and providing continuous support.

3D model courtesy: Kitten Model – AIM Shape Repository; Armadillo – Standford 3D scanning repository; Yoda, Wire Bunny– Cults3D model library; S-pipe – Thingiverse library

Competing interests

The authors declare no competing interests.

Data availability

No datasets were used for the development of the proposed framework. The 3D mesh models used for evaluation are available from the corresponding authors.

Additional Information

Correspondence and requests for materials should be addressed to Don Pubudu Vishwana Joseph Jayakody or Hyunyoung Kim.

REFERENCES

[1] MacDonald, E. & Wicker, R. Multiprocess 3D printing for increasing component functionality. *Science (80-.).* **353**, aaf2093 (2016).

[2] Margadji, C. & Pattinson, S. W. Hybrid reasoning for perception, explanation, and autonomous action in manufacturing. *Nat. Commun.* **17**, 6589 (2026).

[3] Margadji, C., Brion, D. A. J. & Pattinson, S. W. Iterative learning for efficient additive mass production. *Addit. Manuf.* **89**, 104271 (2024).

[4] Plocher, J. & Panesar, A. Review on design and structural optimisation in additive manufacturing: Towards next-generation lightweight structures. *Mater. Des.* **183**, 108164 (2019).

[5] Jayakody, D. P. V. J., Deng, B., Goonetilleke, R. S., Thomas-Seale, L. E. J. & Kim, H. Neural field-based shape optimization for manufacturability-aware support structure minimisation. *Addit. Manuf.* **122**, 105177 (2026).

[6] Jayakody, D. P. V. J., Deng, B., Goonetilleke, R. S., Thomas-Seale, L. E. J. & Kim, H. Part orientation fused shape optimization for minimisation of print time and material waste in extrusion-based 3D printing. *Comput. Des.* **192**, 104029 (2026).

[7] Tojo, K., Shamir, A., Bickel, B. & Umetani, N. Stealth Shaper: Reflectivity Optimization as Surface Stylization. in *ACM SIGGRAPH 2023 Conference Proceedings* (Association for Computing Machinery, 2023). doi:10.1145/3588432.3591542.

[8] Zhan, X., Jambon, C., Thompson, E., Ng, K. & Konaković Luković, M. PhysiOpt: Physics-Driven Shape Optimization for 3D Generative Models. *Proc. - SIGGRAPH Asia 2025 Conf. Pap. SA 2025* 0–11 (2025) doi:10.1145/3757377.3763884.

[9] Chen, H. *et al.* Generative Manufacturing: A requirements and resource-driven approach to part making. *Comput. Ind.* **168**, 104286 (2025).

[10] Wirth, M., Chapuis, J. N. & Shea, K. Inverse design of stochastic, voxelated thermo-viscoelastic digital materials. *Nat. Commun.* **16**, 11705 (2025).

[11] Tsumoto, R., Yaji, K., Nomaguchi, Y. & Fujita, K. Deep concept identification for generative design. *Adv. Eng. Informatics* **65**, 103354 (2025).

[12] Zheng, L., Karapiperis, K., Kumar, S. & Kochmann, D. M. Unifying the design space and optimizing linear and nonlinear truss metamaterials by generative modeling. *Nat. Commun.* **14**, 1–14 (2023).

[13] Yu, X., Peng, B. & Panesar, A. Smooth 3D transition cell generation based on latent space arithmetic. *Addit. Manuf.* **101**, 104714 (2025).

[14] Yu, L., Kok, Y. E., Parry, L., Özcan, E. & Maskery, I. Generating new cellular structures for additive manufacturing through an unconditional 3D latent diffusion model. *Addit. Manuf.* **101**, 104712 (2025).

[15] Danhaive, R. & Mueller, C. T. Design subspace learning: Structural design space exploration using performance-conditioned generative modeling. *Autom. Constr.* **127**, 103664 (2021).

[16] Yang, A., Wei, Z., Liem, R. P. & Fua, P. Physics-guided generative design with gradient-based source-space optimization on flow matching. *Struct. Multidiscip. Optim.* **69**, 136 (2026).

[17] Mokhtar, S., Danhaive, R. & Mueller, C. Neural implicit fields for performance-informed geometries in building design. *Adv. Archit. Geom. 2023* 397–410 (2023) doi:10.1515/9783111162683-030.

[18] Dogan, K. M., Suzuki, H., Gunpinar, E. & Kim, M.-S. A generative sampling system for profile designs with shape constraints and user evaluation. *Comput. Des.* **111**, 93–112 (2019).

[19] Briard, T., Segonds, F. & Zamariola, N. G - DfAM : a methodological proposal of generative design for additive manufacturing in the automotive industry. *Int. J. Interact. Des. Manuf.* **14**, 875–886 (2020).

[20] Oh, S., Jung, Y., Kim, S., Lee, I. & Kang, N. Deep Generative Design: Integration of Topology Optimization and Generative Models. J. Mech. Des. **141**, 111405 (2019).

[21] Kim, J., Kwon, Y. & Kang, N. Deep Generative Design for Mass Production. arXiv:2403.12098v1 (2024)

[22] Moon, J., Park, K., Park, S. & Rosen, D. W. A generative design framework for automated feasible design generation and decision-making in the early design process for additive manufacturing. *Adv. Eng. Informatics* **72**, 104526 (2026).

[23] Nobari, A. H., Regenwetter, L., Picard, C., Han, L. & Ahmed, F. Optimize Any Topology: A Foundation Model for Shape- and Resolution-Free Structural Topology Optimization. *arXiv:2510.23667* (2025).

[24] Kwon, Y. & Kang, N. Three-dimensional geometric augmentation guided by engineering uncertainty. *Eng. Appl. Artif. Intell.* **175**, 114648 (2026).

[25] You, Y., Zhao, C., Zhang, H., Xu, M. & Fua, P. PhysGen: Physically Grounded 3D Shape Generation for Industrial Design. *arXiv:2512.00422v2* (2025).

[26] Jiang, Z., Wen, H., Han, F., Tang, Y. & Xiong, Y. Data-driven generative design for mass customization : A case study. *Adv. Eng. Informatics* **54**, 101786 (2022).

[27] Zhang, B., Tang, J., Nießner, M. & Wonka, P. 3DShape2VecSet: A 3D Shape Representation for Neural Fields and Generative Diffusion Models. *ACM Trans. Graph.* **42**, (2023).

[28] Tewari, A. *et al.* Diffusion with Forward Models: Solving Stochastic Inverse Problems Without Direct Supervision. *Adv. Neural Inf. Process. Syst.* **36**, 1–14 (2023).

[29] Muralikrishnan, S. *et al.* GLASS: Geometric Latent Augmentation for Shape Spaces. *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.* **2022-June**, 18531–18540 (2022).

[30] Xue, T. *et al.* MIND: Microstructure INverse Design with Generative Hybrid Neural Representation. *Proc. - SIGGRAPH 2025 Conf. Pap.* 0–12 (2025) doi:10.1145/3721238.3730682

[31] Berzins, A. *et al.* Geometry-Informed Neural Networks. in *Proceedings of the 42nd International Conference on Machine Learning* (eds. Singh, A. et al.) vol. 267 3976–4004 (PMLR, 2025).

[32] Yang, Y. L., Yang, Y. J., Pottmann, H. & Mitra, N. J. Shape Space Exploration of Constrained Meshes. *ACM Trans. Graph.* **30**, 1–12 (2011).

[33] Yang, G., Belongie, S., Hariharan, B. & Koltun, V. Geometry Processing with Neural Fields. *in Thirty-Fifth Conference on Neural Information Processing Systems* (2021).

[34] Yang, X., Lin, G., Chen, Z. & Zhou, L. Neural Vector Fields: Implicit Representation by Explicit Learning. *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.* **2023-June**, 16727–16738 (2023).

[35] Erkoç, Z., Ma, F., Shan, Q., Nießner, M. & Dai, A. HyperDiffusion: Generating Implicit Neural Fields with Weight-Space Diffusion. *Proc. IEEE Int. Conf. Comput. Vis.* 14254–14264 (2023) doi:10.1109/ICCV51070.2023.01315.

[36] Holden, D., Komura, T. & Saito, J. Phase-functioned neural networks for character control. *ACM Trans. Graph.* **36**, (2017).

[37] Giannone, G., Srivastava, A., Winther, O. & Ahmed, F. Aligning Optimization Trajectories with Diffusion Models for Constrained Design Generation. in *Advances in Neural Information Processing Systems* (eds. Oh, A. et al.) vol. 36 51830–51861 (Curran Associates, Inc., 2023).

[38] Lemons, D. S. & Gythiel, A. Paul Langevin's 1908 paper '"On the Theory of Brownian Motion"' ['"Sur la the´orie du mouvement brownien,"' C. R. Acad. Sci. (Paris) 146, 530–533 (1908)]. *Am. J. Phys.* **65**, 1079–1081 (1997).

[39] Chen, H., Ha, D. Q., Carstensen, J. V. & Ahmed, F. TopoCtrl: Post-Optimization Topology Editing Toward Target Structural Characteristics. *arXiv:2603.26926* (2026).

[40] Kou, C. *et al.* Neural Topology Optimization Via Active Learning for Efficient Channel Design in Turbulent Mass Transfer. *Adv. Sci.* **12**, (2025).

[41] Chen, H., Joglekar, A. & Burak Kara, L. Topology Optimization Using Neural Networks With Conditioning Field Initialization for Improved Efficiency. *J. Mech. Des.* **146**, 61702 (2023).

[42] Schulz, A., Wang, H., Grinspun, E., Solomon, J. & Matusik, W. Interactive exploration of design trade-offs. *ACM Trans. Graph.* **37**, (2018).

[43] Schulz, A. *et al.* Interactive design space exploration and optimization for CAD models. *ACM Trans. Graph.* **36**, (2017).

[44] Li, C., Farkhoor, H., Liu, R. & Yosinski, J. Measuring the intrinsic dimension of objective landscapes. *6th Int. Conf. Learn. Represent. ICLR 2018 - Conf. Track Proc.* (2018).

[45] Du, Y. & Mordatch, I. Implicit generation and modeling with energy-based models. *Adv. Neural Inf. Process. Syst.* **32**, (2019).

[46] Aigerman, N. *et al.* Neural Jacobian Fields: Learning Intrinsic Mappings of Arbitrary Meshes. *ACM Trans. Graph.* **41**, (2022).

[47] Jain, P., Qu, Z., Chen, P. Y. & Stein, O. *Neural Monte Carlo Fluid Simulation*. *Proceedings - SIGGRAPH 2024 Conference Papers* vol. 1 (Association for Computing Machinery, 2024).

[48] Wang, X., Jiang, Z., Xiong, Y. & Liu, A. Human-LLM collaboration in generative design for customization. *J. Manuf. Syst.* **80**, 425–435 (2025)

[49] Yin, F. *et al.* ShapeGPT: 3D Shape Generation With a Unified Multi-Modal Language Model. *IEEE Trans. Multimed.* **27**, 4107–4120 (2025)

[50] Zheng, L., Kumar, S. & Kochmann, D. M. Algebraic language models for inverse design of metamaterials via diffusion transformers. *Nat. Mach. Intell.* **8**, 628–640 (2026).

[51] Jetton, C., Campbell, M. & Hoyle, C. Human-in-the-Loop Bayesian Optimization for Artificial Intelligence-Guided Design. *J. Mech. Des.* **148**, 31707 (2026).

[52] Paszke, A. *et al.* PyTorch: An Imperative Style, High-Performance Deep Learning Library. in *Advances in Neural Information Processing Systems* (eds. Wallach, H. et al.) vol. 32 (Curran Associates, Inc., 2019)

[53] Kingma, D. P. & Ba, J. L. Adam: A method for stochastic optimization. *3rd Int. Conf. Learn. Represent. ICLR 2015 - Conf. Track Proc*. 1–15 (2015)

[54] Xie, Y. *et al.* Neural Fields in Visual Computing and Beyond. *Comput. Graph. Forum* **41**, 641–676 (2022)