TopoCtrl: Post-Optimization Topology Editing Toward Target Structural Characteristics

Hongrui Chen[1] **Dat Ha** [2] **Josephine V. Carstensen**[2] **Faez Ahmed**[1] *
[1]Department of Mechanical Engineering
[2]Department of Civil and Environmental Engineering
Massachusetts Institute of Technology
Cambridge, MA, 02139 USA

Topology optimization can generate high-performance structures, but designers often need to revise the resulting topology in ways that reflect fabrication preferences, structural intuition, or downstream design constraints. In particular, they may wish to explicitly control interpretable structural characteristics such as member thickness, characteristic member length, the number of joints, or the number of members connected to a joint. These quantities are often discrete, non-smooth, or only available through a forward evaluation procedure, making them difficult to impose within conventional optimization pipelines. We present TopoCtrl, a post-optimization control framework that repurposes the latent space of a pre-trained topology foundation model for explicit characteristic-guided editing. Given an optimized topology, TopoCtrl encodes it into the latent space of a latent diffusion model, applies partial noising to preserve instance similarity while creating room for modification, and then performs regression-guided denoising toward a prescribed target characteristic. The core concept is to train a lightweight regression model on latent representations annotated with evaluated structural characteristics, and to use its gradient as a differentiable guidance signal during reverse diffusion. This avoids the need for characteristic-specific reformulations, hand-derived sensitivities, or iterative optimization. Because the method operates through partial noising of an existing topology latent, it preserves overall structural similarity while still enabling characteristic controls. Across representative control tasks involving both continuous and discrete structural characteristics, TopoCtrl produces target-aligned topology modifications while better preserving structural coherence and design intent than indirect parameter tuning or naive geometric post-processing. These results highlight latent diffusion with learned regression guidance as a practical route for explicit and instance-preserving control in topology optimization.

1 INTRODUCTION

Topology optimization has become a powerful paradigm for engineering design, enabling the synthesis of lightweight, high-performance structures through the optimal distribution of material within a prescribed domain. Yet in practical workflows, the optimized result is rarely the end of the design process. Engineers and designers often want to investigate, trust, and revise the solution in ways that reflect downstream constraints, fabrication preferences, and structural intuition. In particular, they may wish to adjust interpretable characteristics of the design, such as member thickness, the number of joints, characteristic member length, or the number of members connected to a joint, while still preserving the overall structural similarity and mechanical performance of the topology. Such requests arise naturally in practice, but they remain difficult to realize within conventional topology optimization pipelines, especially when the desired quantity is discrete, non-smooth, or only available through a forward evaluation procedure. Despite their practical importance, there is still a lack of methods that allow fast, explicit, and instance-preserving control over these challenging structural characteristics.

Existing alternatives partially address this need. One option is to reformulate the optimization problem by introducing additional constraints, penalties, or parameterized controls tailored to the target characteristic [1,2,3,4,5,6]. While effective for certain quantities, these approaches typically require problem-specific derivations, careful tuning, and repeated optimization solves. The difficulty is even greater for characteristics such as the number of joints or the number of members incident to a joint, which are discrete and not naturally amenable to gradient-based optimization [7,8,9,10,11]. Even varying standard optimization parameters, such as filter size, can introduce some design variability, but this variability is indirect and limited, and it does not amount to explicit control over most of the characteristics that practitioners care about [12, 13, 14]. Even varying standard optimization parameters, such as filter size, can introduce some design vari-

*Address all correspondences to hongrui@mit.edu

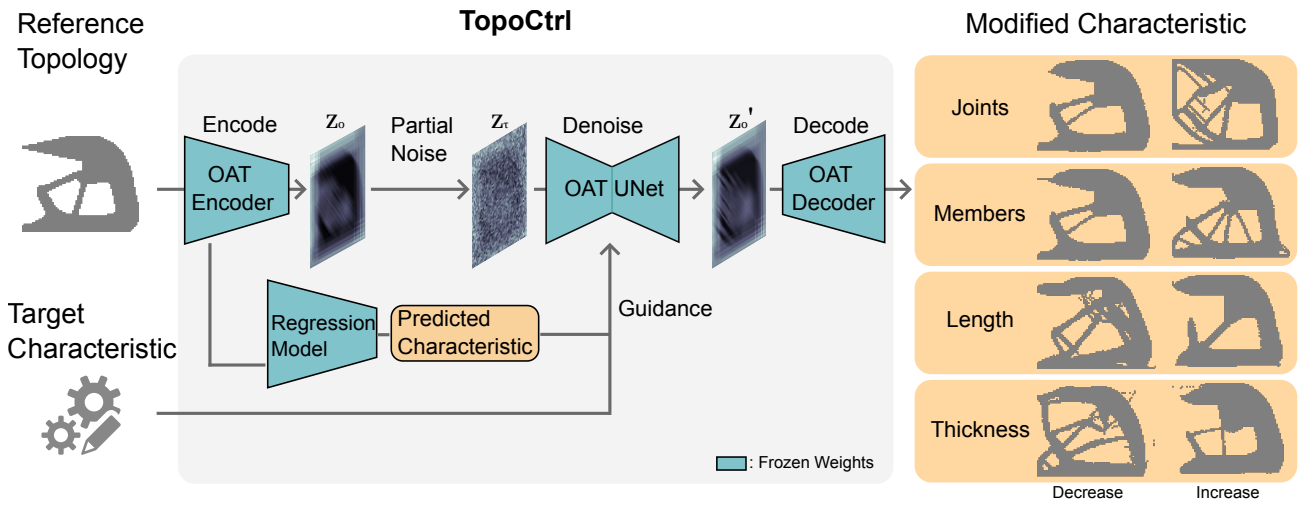


Fig. 1: TopoCtrl uses a pretrained OAT with frozen weights. Given a reference topology and a user-specified target characteristic (number of joints, maximum members connected to a joint, maximum member length, and member thickness), TopoCtrl encodes the input into OAT's structured latent, applies partial noising. A small regression model is trained on a small dataset of latent-characteristic pairs to predict the target characteristic. Starting from the noisy latent, we perform diffusion denoising with regression-model guidance toward the target characteristic to recover an edited clean latent that preserves structural similarity. Finally, the edited latent is decoded to yield the edited topology with characteristics matching the target.

ability, but this variability is indirect and limited, and it does not amount to explicit control over most of the characteristics that practitioners care about. As a result, designers are often left choosing between slow re-optimization, weak indirect control, or ad hoc geometric edits that can easily compromise structural behavior.

Data-driven generative models have recently expanded the design space of topology optimization by learning to predict optimized structures across diverse loading conditions, boundary conditions, and geometric settings [15, 16, 17, 18]. These methods demonstrate rapid topology generation. However, their primary role is still forward synthesis: given a specification, they generate a candidate topology. They do not directly address the problem of editing an existing optimized design toward a user-specified structural characteristic while preserving its similarity. Furthermore, forward evaluation is relatively easy for many characteristics. Given a topology, one can measure its thickness, count its joints, or extract a characteristic length. The challenge lies in the inverse problem of integrating such evaluations into a controllable design pipeline when only the forward path is readily available. To our knowledge, there is still no general tool that bridges this gap and enables explicit topology control in this setting.

In this work, we address this problem through a regression-guided latent diffusion pipeline, which we call *TopoCtrl*. Figure 1 shows an overview of the proposed method. Our approach builds upon a pre-trained topology foundation model called Optimized Any Topology (OAT) in latent diffusion form and repurposes its structured latent space for controllable post-optimization editing [18, 19]. Using a relatively small dataset of topologies annotated with evaluated structural characteristics, we train a lightweight regression model that maps the latent representation of the diffusion model to the target characteristic. Once trained, this regressor provides a differentiable guidance signal during denoising, allowing the latent trajectory to be steered toward a desired characteristic value. This approach therefore avoids the need to derive problem-specific sensitivities or reformulate the original topology optimization problem. Instead, we leverage the learned structural knowledge of the diffusion model together with a learned forward predictor of the characteristic to perform explicit control directly in latent space.

Our method relies on partial noising. Rather than sampling a completely new design, we begin from the latent representation of an existing optimized topology and inject only a limited amount of noise before running the reverse diffusion process with regression guidance [20, 21, 22, 23, 24, 25]. Partial noising preserves a portion of the original topology information, which helps maintain geometry consistency, while still creating enough freedom for the latent to be manipulated toward the desired target. The subsequent denoising process then allows the foundation model to reconstruct a topology by drawing on the structural knowledge encoded during pre-training [18].

The resulting framework is particularly appealing for characteristics that are difficult to differentiate and cumbersome to impose through classical optimization. By decoupling characteristic evaluation from optimization, TopoCtrl opens a practical route for explicit control in settings where only a forward evaluator is available. At the same time, be-

cause the method operates on top of a pre-trained foundation model, it inherits broad geometric and physical information from that model and can be applied across structures of varying scales, aspect ratios, and characteristic types.
Our main contributions are:

- TopoCtrl: We introduce a characteristic-control pipeline built on a pre-trained latent diffusion topology foundation model, where a lightweight regression model is trained in latent space and used to guide denoising toward prescribed structural characteristics.
- Explicit control of difficult characteristics: We show that the proposed framework enables explicit control over challenging characteristics, including quantities that are discrete, hard to differentiate, or otherwise inconvenient to integrate into conventional topology optimization formulations.
- Instance-preserving control through partial noising: We demonstrate that partial noising of an existing topology latent provides an effective mechanism for preserving structural similarity and design intent while still allowing guided modification toward the target characteristic.
- General applicability across structures and tasks: We show that the latent-guided framework is naturally compatible with structures of different scales and aspect ratios, and can be adapted to different characteristic-control tasks using the same underlying pipeline.

The code for this work can be found at: `https://github.com/HongRayChen/TopoCtrl`

2 RELATED WORK

Our review focuses on data-driven topology optimization, characteristic control in topology optimization, and diffusion-based latent editing and guidance.

Data-driven topology optimization Topology optimization distributes material within a design domain to satisfy structural objectives under constraints such as volume fraction. Classical methods include homogenization-based design, shape and topology viewpoints beyond homogenization, material interpolation through Solid Isotropic Material with Penalization (SIMP), and level-set formulations for boundary evolution [26, 27, 28, 29, 30]. However, their iterative solve process can be computationally expensive, which motivates learning-based surrogates for rapid generation and exploration.

Data-driven topology optimization addresses this bottleneck by learning a direct mapping from problem specification to a near-optimal structure. Prior work spans conditional GAN and transfer-learning approaches [31, 32], supervised iteration-free predictors [33, 34], hybrid learning-and-optimization pipelines [35], and architectures designed to improve generalization, implicit representation, multiresolution behavior, and three-dimensional generation [36, 37, 38, 39, 40]. These methods demonstrate that a dataset of optimized topologies can be learned and used to accelerate topology generation. However, their main role is still the forward synthesis of designs from given boundary conditions and other problem descriptors; they predict an optimized design. They are generally not designed for instance-preserving edits, where the goal is not to generate a new topology from scratch, but to modify one existing design while retaining its structural similarity.

Generative modeling for topology optimization has progressed from GAN-based methods conditioned on physical fields [15], to resolution-free implicit generators [41], and more recently to diffusion-based design generation [42, 16, 17]. This progression culminates in OAT, a latent diffusion foundation model trained on a large and diverse corpus of optimized structures across varying boundary conditions, loads, shapes, and resolutions [18]. Compared with earlier data-driven pipelines, such models better capture the broader distribution of optimized designs and provide a structured latent space in which diverse topologies can be represented with a uniform latent tensor. For the present work, this is especially important as explicit characteristic control often pushes a design away from its original optimum, so a strong generative model is needed to reconstruct a topology without significant compromise in compliance.

Editing and characteristic control in topology optimization Although optimized topologies can satisfy performance objectives, they often do not directly satisfy downstream preferences regarding manufacturability, interpretability, or structural style. As a result, a substantial body of work introduces user-facing controls into topology optimization. One class of methods focuses on feature size, thickness, and related geometric descriptors, including interactive feature-size control, graded porosity, and explicit thickness or feature control under density and level-set formulations [13, 14, 43, 6, 44]. Another class introduces appearance-oriented guidance through user-drawn patterns, pattern libraries, patterned shells, multi-pattern constraints, and other designer-in-the-loop mechanisms [45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56]. These works show the value of exposing interpretable controls to designers, but they typically do so by reformulating and re-solving the optimization problem itself.

For the characteristics studied in this paper, the literature is more fragmented. Thickness and length scale are among the most extensively studied geometric controls. In density-based and SIMP settings, minimum and maximum member size or length scale are commonly imposed through projection, morphology, filtering, or explicit geometric constraints [57, 58, 1, 12, 2, 3, 4, 59, 5]. Related works also pursue direct or explicit thickness and length-scale control, including explicit SIMP-based formulations and discrete-object approaches [60, 61, 62]. In level-set settings, explicit feature control, skeleton- or boundary-distance-based constraints, uniform thickness formulations, and derivable skeleton representations provide more direct geometric handles [6, 44, 63, 64, 65, 66, 67, 68]. Closely related complexity-control formulations further regulate tunnels, cavities, thinning, and global structural complexity, often with the goal of simplifying or regularizing the topology [11, 69, 70, 71, 72,

73, 74, 75, 76, 77, 78].

In contrast, controls related to the number of joints or the number of members connected to a joint are much less developed in continuum topology optimization. Related works appear more naturally in truss and layout optimization, where one can directly restrict the number of nodes, simplify connectivity, or impose local nodal stability conditions [7, 8, 79, 80, 9, 10]. Some continuum-based topology-control works also indirectly affect structural complexity and connectivity by limiting holes, tunnels, or other topological events [71, 72, 73, 74, 75]. Nevertheless, explicit post-hoc control over joint count or local joint valence remains difficult. These quantities are often discrete, non-smooth, and only available through a forward evaluation pipeline, making them awkward to embed into conventional gradient-based topology optimization.

These studies highlight both the promise of classical control mechanisms. Many of the differentiable characteristics with explicit controls can already deliver accurate and structurally optimized solutions. If a desired characteristic admits a carefully designed differentiable formulation, one may incorporate it into the optimization problem, but doing so typically requires substantial problem-specific derivation, careful tuning, and repeated iterative solves. Even when the intended edit is local, the optimization update is still global and iterative, and the resulting design may drift from the original instance. While we are not directly competing with existing works on characteristic control, specifically for post-optimization control, we see a limitation: designers may want to preserve the overall similarity of a topology while adjusting one difficult characteristic, rather than solving a new optimization problem from scratch.

Diffusion models, latent editing, and guidance Diffusion models learn a data distribution by reversing a gradual noising process [81, 82, 83]. Practical sampling and control are enabled by accelerated samplers and guidance mechanisms such as Denoising Diffusion Implicit Models (DDIM) and classifier-free guidance [20, 84], while strong empirical performance has made diffusion models a dominant generative framework [85]. Latent diffusion improves efficiency by carrying out this process in a compressed latent space rather than directly in pixel space [19], and few-step generation has been further improved through distillation and consistency-based formulations [86, 87, 88].

Diffusion models also demonstrate their success in context aware edits. Prior work demonstrates text- and semantic-guided image editing through attention control, inversion, and instruction-based manipulation [21, 89, 90, 22, 23, 91]. Other works show that sparse interactive constraints, such as point dragging, can be propagated through diffusion features to produce coherent geometric modifications while maintaining instance similarity [92, 93]. These prior works demonstrated that editing in a learned latent or feature space is often more structurally aware than direct image-space warping, because the model can reconstruct changes using the semantic and geometric information captured during training.

Recent work also shows that diffusion models can be steered by auxiliary predictive models or learned readouts rather than only by text prompts or classifier scores. Plug-and-play guidance, universal guidance, diffusion-feature readouts, and related guidance-by-prediction strategies provide practical mechanisms for controlling the denoising trajectory with external objectives [94, 95, 96, 21]. This is especially attractive for topology control. In many structural settings, the target characteristic is easy to evaluate in the forward direction once a topology is given, but difficult to differentiate through a conventional optimization pipeline. A learned regression model in latent space offers an alternative that can serve as a lightweight differentiable proxy for the characteristic and provide guidance during denoising without requiring a full reformulation of the original topology optimization problem.

Motivated by these developments, we investigate whether a latent diffusion foundation model for topology optimization can support instance-preserving characteristic control through regression-guided denoising. Compared to direct geometric edits or post-processing, latent-space guidance offers access to the broader structural knowledge learned by the generative model. Relative to reformulating topology optimization with handcrafted characteristic constraints, it avoids characteristic-specific derivations and repeated global solves. This makes it a promising framework for explicit control of challenging structural characteristics, particularly those that are discrete, non-smooth, or otherwise inconvenient to impose in classical optimization.

3 Proposed Method

TopoCtrl performs characteristic control by guiding the denoising trajectory of a pre-trained latent diffusion model with a lightweight regression network. Given an optimized topology T, we first encode it into the latent space of OAT and apply partial noising so that sampling remains tied to the original structural instance while still permitting controlled variation. We then predict the target structural characteristic from the estimated clean latent and use the resulting gradient to steer reverse diffusion toward a prescribed value. The final latent is decoded back to the topology space by the OAT autoencoder.

3.1 Diffusion model

We build TopoCtrl on top of the pre-trained OAT latent diffusion model [18]. OAT is well-suited for controllable topology editing because it maps topologies of varying sizes and aspect ratios into a fixed latent tensor, while preserving a neural rendering path that reconstructs the topology at arbitrary resolution. This fixed latent representation makes it convenient to train a single lightweight control model and to perform guidance directly in latent space.

Neural-field autoencoder. Let $T \in [0,1]^{H\times W}$ denote an optimized density field defined on a rectangular design domain of height H and width W. OAT first encodes T into a

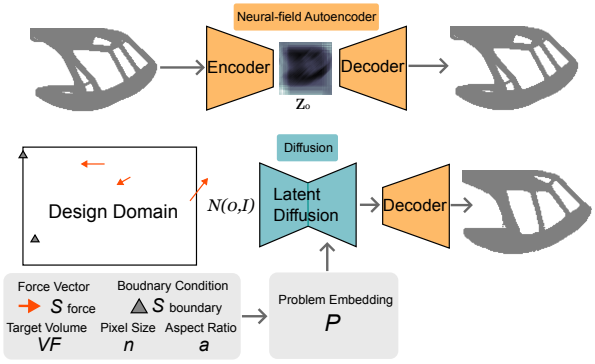


Fig. 2: An overview of the OAT framework. First, an autoencoder is trained which the encoder encodes problems of arbitrary shapes into a structured fixed resolution latent z_0 and the decoder decodes it back to the original topology. Given the problem configuration in the design domain, the problem embedding P is fed into the diffusion model as a conditioning vector. The output from the latent diffusion model is a fixed-resolution latent, which the decoder decodes into a topology.

latent tensor z through a neural-field autoencoder,

$$z = E(T), \qquad \phi = D(z), \qquad \tilde{T} = R(\phi, c, n), \tag{1}$$

where $E(\cdot)$ is the encoder, $D(\cdot)$ is the decoder, and $R(\cdot)$ is a neural renderer. The renderer reconstructs the topology from decoder features ϕ using spatial coordinates c and cell sizes n, which allows the decoded density field to be queried at arbitrary resolution. The autoencoder is trained with the reconstruction objective

$$\mathscr{L}_{\mathrm{AE}} = \left\| \tilde{T} - T_{\mathrm{GT}} \right\|_1, \tag{2}$$

where T_{GT} denotes the ground-truth topology.

Problem representation. Each topology optimization problem is conditioned on the design-domain geometry, boundary conditions, loads, and target volume fraction. Following OAT, we denote the raw problem specification as

$$\hat{P} = \left(S_{\mathrm{boundary}}, S_{\mathrm{force}}, \mathrm{VF}, n, a \right), \tag{3}$$

where S_{boundary} is the point set describing support locations and directional constraints, S_{force} is the point set describing load locations and force vectors, $\mathrm{VF} \in (0,1)$ is the target volume fraction, n encodes the cell size, and $a \in \mathbb{R}^2$ encodes the aspect ratio of the domain. These quantities are embedded into a fixed conditioning vector

$$P = \mathrm{concat}\Big(\mathrm{BPOM}_b(S_{\mathrm{boundary}}), \mathrm{BPOM}_f(S_{\mathrm{force}}), \mathrm{MLP}_{\mathrm{vf}}(\mathrm{VF}), \mathrm{MLP}_{\mathrm{cell}}(n), \mathrm{MLP}_{\mathrm{ratio}}(a) \Big), \tag{4}$$

which is supplied to the diffusion model during training and inference.

Latent diffusion. Let $z_0 = E(T)$ denote the clean latent of the input topology. OAT defines a forward-noising process

$$z_t = \sqrt{\bar{\alpha}_t}\, z_0 + \sqrt{1-\bar{\alpha}_t}\, \varepsilon, \qquad \varepsilon \sim \mathscr{N}(0, I), \tag{5}$$

where $t \in \{0, \ldots, T\}$ is the diffusion timestep and $\bar{\alpha}_t$ is the cumulative noise schedule. Rather than predicting ε directly, OAT uses the velocity parameterization

$$v = \sqrt{\bar{\alpha}_t}\, \varepsilon - \sqrt{1-\bar{\alpha}_t}\, z_0, \tag{6}$$

and trains a conditional UNet v_θ with the objective

$$\mathscr{L}_{\mathrm{LDM}} = \mathbb{E}_{z_0, t, \varepsilon} \left[\left\| v - v_\theta(z_t, t \mid \hat{P}) \right\|_2^2 \right]. \tag{7}$$

Partial noising and editable sampling. TopoCtrl uses partial noising. Instead of starting the reverse process from pure noise, we begin from an intermediate timestep $\tau \ll T$. Given an input topology, we compute its clean latent z_0 and form a partially noised latent

$$z_\tau = \sqrt{\bar{\alpha}_\tau}\, z_0 + \sqrt{1-\bar{\alpha}_\tau}\, \varepsilon. \tag{8}$$

We then run the reverse diffusion process from τ back to 0 under regression guidance. Since $\bar{\alpha}_\tau$ remains non-negligible, z_τ still preserves the dominant semantic content of the original design, including its coarse load path and global layout, while the injected noise creates sufficient flexibility for controlled modification and stochastic sampling of multiple plausible candidates. In contrast, starting from full noise would largely discard instance similarity and turn the process into unconditional regeneration.

3.2 Regression model

Medial-axis characteristic evaluation. We compute four structural characteristics from each topology using a medial-axis-based post-processing pipeline. Medial-axis skeletonization is a morphological technique used to extract the skeleton of an image. The skeleton consists of all points that are equidistant from the image's boundary. Starting from a binarized topology, we compute the medial axis and its associated distance field, which together provide a compact representation of the structural centerlines and local thickness. Because raw medial axes are sensitive to small boundary protrusions, we prune short terminal branches with a nominal threshold of 10 pixels. We then identify skeleton junctions and merge nearby thick-region junctions so that visually continuous multi-member joints are counted as a single joint; specifically, candidate joints are merged when their local thickness exceeds 6 pixels and their distance is

below 8 pixels. The number of merged joint representatives defines the joint-count characteristic, while the maximum number of incident branches at any merged joint defines the maximum-members-per-joint characteristic. To compute member length, we remove the merged joint regions, split the remaining skeleton into individual members, and take the average length of the two longest members. Because the OAT latent has a fixed spatial size and does not by itself encode the original topology scale, this length is normalized by the longer side of the topology domain. For the thickness characteristic, we use the distance field along the medial axis, convert it to local diameter, and take the 90th percentile value; this quantity is likewise normalized by the longer domain dimension.

Regression model architecture and training. Let y denote one of the four target characteristics and let $\tilde{y}$ denote its normalized version. We train four separate regression models, one for each characteristic: number of joints, maximum members going into a joint, average top-two member length, and 90th-percentile thickness. Each regressor takes the latent tensor z as input and outputs a single scalar prediction,

$$\hat{y} = f(z), \tag{9}$$

where f is a lightweight convolutional network. In our implementation, f_ψ consists of four convolutional blocks with channel sizes 32, 64, 128, and 256, each followed by batch normalization and SiLU activation, with spatial downsampling performed by strided convolutions. The resulting feature tensor is globally averaged and passed through a two-layer multilayer perceptron to produce a scalar output. Each regressor is trained independently with mean-squared error,

$$\mathscr{L}_{\text{reg}} = \mathbb{E}_{(z,\tilde{y})}\left[\|f(z) - \tilde{y}\|_2^2\right]. \tag{10}$$

3.3 Guidance and denoising

At inference time, the trained regressor is used to guide the denoising toward a prescribed target characteristic. Let $y^\star$ denote the desired characteristic value and let $\tilde{y}^\star$ be its normalized target. Starting from the partially noised latent z_τ, we apply guidance at each reverse step $t = \tau, \tau-1, \ldots, 1$.

Regression guidance on the estimated clean latent. Given the current latent z_t, we first query the diffusion model to estimate the corresponding clean latent,

$$\hat{z}_0(z_t, t) = \sqrt{\bar{\alpha}_t}\, z_t - \sqrt{1-\bar{\alpha}_t}\, v_\theta(z_t, t \mid \hat{P}). \tag{11}$$

We then evaluate the regressor on this estimated clean latent and define the control loss

$$\mathscr{L}_{\text{ctrl}}(z_t) = \left\| f_\psi\big(\hat{z}_0(z_t, t)\big) - \tilde{y}^\star \right\|_2^2. \tag{12}$$

Parameter	Value
Regression model training epochs	15
Regression training learning rate	0.0001
DDIM Total Steps (t_{total})	100
DDIM Partial Steps (τ)	30
Guidance strength (λ)	10

Table 1: Hyperparameter configuration for the topology control tasks.

The latent is updated by a gradient step

$$\tilde{z}_t = z_t - \nabla_{z_t}\Big(\lambda_{\text{ctrl}}\, \mathscr{L}_{\text{ctrl}}(z_t)\Big), \tag{13}$$

where λ_{ctrl} controls the strength of characteristic guidance.

DDIM-style reverse update. After the gradient update, we recompute the diffusion prediction on the guided latent $\tilde{z}_t$ and perform a deterministic DDIM reverse step. Specifically, we define

$$\tilde{v}_\theta = v_\theta(\tilde{z}_t, t \mid \hat{P}), \tag{14}$$

$$\hat{z}_0(\tilde{z}_t, t) = \sqrt{\bar{\alpha}_t}\, \tilde{z}_t - \sqrt{1-\bar{\alpha}_t}\, \tilde{v}_\theta, \tag{15}$$

and the corresponding noise estimate

$$\hat{\varepsilon}(\tilde{z}_t, t) = \sqrt{1-\bar{\alpha}_t}\, \tilde{z}_t + \sqrt{\bar{\alpha}_t}\, \tilde{v}_\theta. \tag{16}$$

The reverse update is then

$$z_{t-1} = \sqrt{\bar{\alpha}_{t-1}}\hat{z}_0(\tilde{z}_t, t) + \sqrt{1-\bar{\alpha}_{t-1}}\hat{\varepsilon}(\tilde{z}_t, t). \tag{17}$$

Repeating this process from $t = \tau$ to $t = 0$ yields the final controlled latent, which is decoded by the OAT neural-field autoencoder to produce the edited topology. Because the process begins from a partially noised version of the original instance rather than from pure noise, the resulting design tends to preserve the similarity of the input while moving toward the prescribed characteristic target.

We summarize the algorithm in Algorithm 1. Table 1 includes the hyperparameters used for training the regression model and the configurations for partial denoising with characteristics guidance. We use a partial noising to total step ratio similar to existing latent diffusion-based edit and guidance methods [92].

Algorithm 1 Regression-guided latent-based topology control

Encode latent: $z_0 \leftarrow E(T)$
Add partial noise: sample $\varepsilon \sim \mathcal{N}(0, I)$ and set $z_\tau \leftarrow q_\tau(z_0, \varepsilon)$
Normalize target characteristic: $\tilde{y}^\star \leftarrow \mathrm{Norm}(y^\star)$
Initialize $z_t \leftarrow z_\tau$
Denoising: for $t = \tau, \ldots, 1$ do
    Query the diffusion model: $v_\theta(z_t, t \mid \hat{P})$
    Estimate the corresponding clean latent:
        $\hat{z}_0(z_t, t) \leftarrow \sqrt{\bar{\alpha}_t}\, z_t - \sqrt{1-\bar{\alpha}_t}\, v_\theta(z_t, t \mid \hat{P})$
    Evaluate the control loss:
        $\mathscr{L}_{\mathrm{ctrl}}(z_t) \leftarrow \left\| f_\psi\big(\hat{z}_0(z_t, t)\big) - \tilde{y}^\star \right\|_2^2$
    Apply regression guidance to the current latent:
        $\tilde{z}_t \leftarrow z_t - \nabla_{z_t}(\lambda_{\mathrm{ctrl}}\, \mathscr{L}_{\mathrm{ctrl}}(z_t))$
    Recompute the diffusion prediction on the guided latent:
        $\tilde{v}_\theta \leftarrow v_\theta(\tilde{z}_t, t \mid \hat{P})$
    Recover the guided clean latent and noise estimate:
        $\hat{z}_0(\tilde{z}_t, t) \leftarrow \sqrt{\bar{\alpha}_t}\, \tilde{z}_t - \sqrt{1-\bar{\alpha}_t}\, \tilde{v}_\theta$
        $\hat{\varepsilon}(\tilde{z}_t, t) \leftarrow \sqrt{1-\bar{\alpha}_t}\, \tilde{z}_t + \sqrt{\bar{\alpha}_t}\, \tilde{v}_\theta$
    Take one DDIM reverse step:
        $z_{t-1} \leftarrow \sqrt{\bar{\alpha}_{t-1}}\, \hat{z}_0(\tilde{z}_t, t) + \sqrt{1-\bar{\alpha}_{t-1}}\, \hat{\varepsilon}(\tilde{z}_t, t)$
Denote the final denoised latent by z_0'
Decode structure: $\phi' \leftarrow D(z_0')$; $T' \leftarrow R(\phi', c, s)$

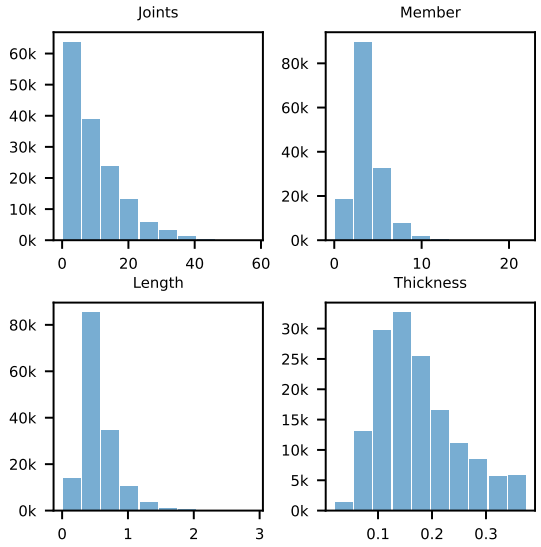


Fig. 3: The characteristic distributions of the filtered 150K-sample training dataset used for regression model training. We show the four quantities used in TopoCtrl: the number of joints, the maximum number of members incident to a joint, the normalized average of the top two longest members, and the normalized 90th-percentile thickness. The dataset is moderately imbalanced, with joint counts concentrated toward the lower end, most joint complexities lying between 3 and 5, characteristic length concentrated around intermediate values, and thickness spread more broadly across the range.

4 RESULTS

In this section, we set up experiments to demonstrate the characteristic control of the number of joints, the maximum number of members going into a joint, the top two average member lengths, and the 90th percentile thickness. These four characteristics are abbreviated as joints, member, length, and thickness in this section. For the topologies, we select only from the testing dataset of OAT. We conduct the experiment on a PC equipped with an RTX Pro 6000 graphics card and 96GB of VRAM. We generated 64 latent edit results for each of the characteristic control operations for the best-of-N analysis. On average, the inference time on a single topology is around 0.5 seconds per sample.

For the regression model, we use the pretraining set of the OAT dataset and filter it to a subset of 150K samples for characteristic evaluation and training. This setup is intentionally challenging. The OAT dataset is generated with a fixed filter radius, and for a given boundary condition and volume fraction, the topology optimization solver produces a single deterministic solution. As a result, the training data does not explicitly provide multiple valid solutions with systematically varied characteristics for the same problem instance. Training the regression model for one characteristic takes around 3 minutes.

For performance evaluation, we use only the test set of OAT. This means that neither OAT nor the regression model saw the test set during training. We select 100 topologies and run characteristic control on each of them. Across the four control tasks, we evaluate 16 target settings in total. For every topology-target pair, we generate 64 stochastic samples from the diffusion model and report both qualitative examples and best-of-64 quantitative results.

4.1 Regression model training

We begin by examining the characteristic distribution of the filtered 150K-sample training set. Figure 3 shows the dataset statistics for the four controlled quantities. We observe that the number of joints is concentrated toward the lower end of the range, indicating that highly articulated structures are comparatively rare in the data. The maximum number of members incident to a joint is most often between 3 and 5, which is consistent with the dominant junction patterns observed in the OAT topologies. The normalized average of the two longest members is concentrated around 0.5 to 0.75, while the normalized 90th-percentile member thickness is more broadly distributed. These distributions confirm that the regression task is moderately imbalanced, especially for the more extreme structural cases.

Figure 4 shows the prediction performance of the latent-space regression models on the held-out test examples. The regressors capture the characteristic trends well enough to serve as useful guidance signals during denoising. The continuous quantities, namely characteristic length and thickness, exhibit smoother prediction behavior, while the discrete quantities, such as joint count and maximum joint complexity, are more challenging near the tails of the distribution. Nevertheless, the results indicate that a relatively small regression model can extract characteristic information from

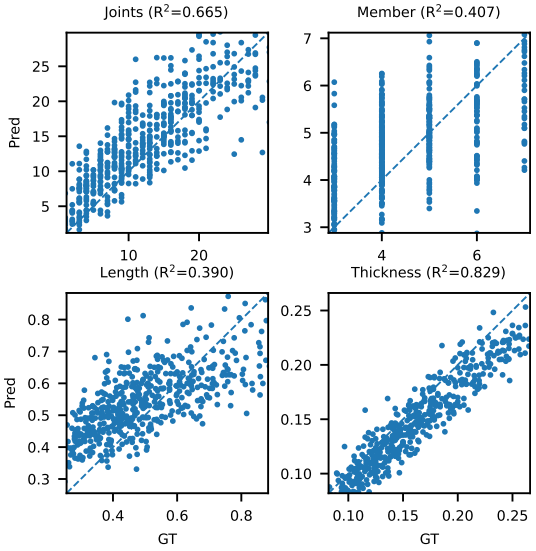


Fig. 4: Prediction performance of the latent-space regression models on the testing set. Each plot compares the ground-truth characteristic computed from the topology against the value predicted from the latent representation. The regressors recover the main trend of each characteristic. The continuous quantities, namely characteristic length and thickness, exhibit smoother prediction behavior, while the discrete quantities, such as joint count and maximum joint complexity, are more challenging near the tails of the distribution. The length prediction showed slightly higher error around the lower and higher ends of the characteristic values. This also caused the subsequent slight increase in error with the length control evaluations.

the latent representation of the pre-trained diffusion model.

4.2 Characteristic control

We next evaluate whether the trained regressors can guide the reverse diffusion process toward prescribed target characteristics. For each of the 100 test topologies, we generate 64 samples per target setting. The final characteristic is evaluated using the same geometric analysis used for dataset annotation. In particular, the structural quantities are recomputed from the decoded topology using medial-axis skeletonization and the associated post-processing pipeline, ensuring that training and evaluation are consistent.

Figure 5 summarizes the best-of-64 characteristic accuracy across the four tasks. For the discrete-valued controls, namely the number of joints and the maximum number of members incident to a joint, we use swarm plots. For the continuous controls, including normalized characteristic length and normalized 90th-percentile thickness, we use violin plots. In all cases, the best sample is selected from the 64 stochastic generations according to the smallest absolute deviation from the target value under the same medial-axis skeletonization with identical post-processing settings. TopoCtrl achieves good target accuracy, although performance degrades somewhat at the most extreme target values, such as very high joint counts, the largest target length, or highly extreme thickness values. This behavior is expected, since those cases are both rarer in the data and farther from the nominal optimum represented by the original topology.

As a benchmark, we also generate comparison topologies by rerunning the testing set of OAT with varying filter radii. For each topology in the test set, we varied the filter radius from 1.5 to 12 with 7 reruns to obtain a variety of solutions (Figre 6), similar to previous papers in interactive edits [14]. This provides a useful baseline for the amount of characteristic variation that can be obtained through a simple modification of the conventional topology optimization pipeline. Varying the filter radius directly changes feature scale and can indirectly influence topology complexity. However, it remains an indirect mechanism and does not provide explicit instance-level control over most of the evaluated characteristics.

To quantify the reachable target span, we report the coverage metric in Table 2. Let the prescribed target interval for a given characteristic be $[c_{\min}, c_{\max}]$, and let the interval actually reached by a method over the evaluated results be $[\hat{c}_{\min}, \hat{c}_{\max}]$. We define coverage as

$$\text{Coverage} = \frac{\max(0, \min(\hat{c}_{\max}, c_{\max}) - max(\hat{c}_{\min}, c_{\min}))}{c_{\max} - c_{\min}}. \tag{18}$$

For example, if the target number of joints spans $[2, 40]$ and a method only reaches $[10, 20]$, the resulting coverage is $10/38$. Using this metric, we observe that TopoCtrl achieves substantially higher coverage than the baseline obtained by changing the filter radius alone. This is particularly important for the discrete controls, where indirect solver modifications provide only limited variation.

In addition to characteristic accuracy, Table 2 reports compliance, Intersection over Union (IoU), volume-fraction error, and failure rate. Following the original OAT evaluation protocol, we count a result as a failure if its compliance exceeds two times that of the ground-truth topology under the same loading and boundary conditions. TopoCtrl exhibits a somewhat higher failure rate than the original OAT, which is expected because characteristic-controlled topologies are intentionally pushed away from the optimal solution for stiffness. Even so, the IoU remains high across tasks, indicating that the generated outputs remain strongly tied to the original topology and function effectively as edits.

4.2.1 Joint control

We first consider control over the number of joints. Figure 7 visualizes the latent trajectory for representative examples: the ground-truth topology, the encoded latent z_0, the partially noised latent z_τ, the denoised guided latent z_0', and the decoded controlled topologies corresponding to lower and higher target values. Figure 8 then shows six qualitative examples selected from the result set for target joint counts of 5, 10, 20, and 40.

The number of joints can be effectively steered by the regression-guided denoising process. This effect is visible across topologies of varying initial complexity. When fewer

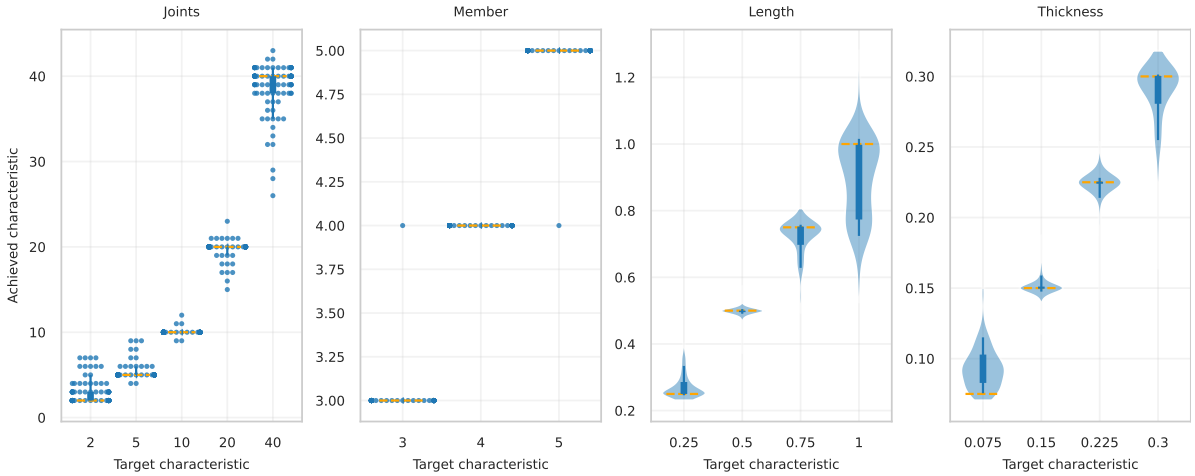


Fig. 5: Characteristic control accuracy of TopoCtrl over the 100 test topologies using best-of-64 sampling. For the discrete-valued characteristics: the number of joints and the maximum number of members incident to a joint, we use swarm plots. For the continuous-valued characteristics: the normalized characteristic length and normalized 90th-percentile thickness, we use violin plots. In each case, the selected result is the sample whose evaluated characteristic is closest to the requested target under the same medial-axis-based analysis used for dataset annotation. TopoCtrl achieves good target accuracy, with somewhat reduced precision at more extreme target values.

	Joints				Member				Length				Thickness			
Best of	Coverage%	VFE%	IoU%	Failure Rate%	Coverage%	VFE%	IoU%	Failure Rate%	Coverage%	VFE%	IoU%	Failure Rate%	Coverage%	VFE%	IoU%	Failure Rate%
2	66.63	0.75	76.95	45.40	64.00	0.68	80.09	38.67	48.69	0.74	77.83	43.50	68.81	1.27	74.25	49.00
4	76.21	0.44	78.82	34.20	82.50	0.40	81.68	27.67	58.62	0.42	79.42	32.00	74.05	0.78	75.74	38.00
8	83.50	0.24	80.18	21.60	91.00	0.23	82.85	17.67	66.31	0.23	80.41	21.50	77.42	0.51	76.89	29.00
16	87.89	0.12	81.07	15.80	98.00	0.11	83.66	10.33	72.60	0.12	81.15	15.00	80.13	0.34	77.71	22.00
32	91.76	0.06	81.81	11.20	99.00	0.06	84.25	6.00	78.05	0.06	81.92	9.75	83.05	0.18	78.56	19.00
64	94.82	0.03	82.51	7.40	99.50	0.03	84.93	4.67	83.84	0.03	82.55	7.00	85.02	0.10	79.21	16.00
BM	23.89	0.21	76.45	N/A	56.00	0.21	76.45	N/A	43.80	0.21	76.45	N/A	16.94	0.21	76.45	N/A

Table 2: Best-of-N quantitative results for TopoCtrl across the four characteristic-control tasks on the testing set. We report coverage, volume-fraction error (VFE), IoU, and failure rate for N=2,4,8,16,32,64, together with the benchmark method obtained by rerunning topology optimization with varying filter radius. Coverage measures how much of the prescribed target interval is actually reached by the method. The failure rate is defined as the percentage of edited topologies whose compliance exceeds two times that of the corresponding ground-truth topology, following the evaluation criterion used in OAT. TopoCtrl achieves higher target coverage than the benchmark while maintaining good IoU, indicating that the generated results remain close to the original instance and function effectively as characteristic-preserving edits.



Fig. 6: Sample benchmark result from the testing dataset re-run with varying filter size (FS). With an increase in filter size, the members become generally thicker, demonstrating some effect in characteristic control.

joints are requested, the generated topologies generally become simpler, with reduced branching. When more joints are requested, the structures typically become more articulated and develop more complex internal connectivity. Reaching very high targets, such as 40 joints, is more challenging for initially simple structures, but the trend remains clear.

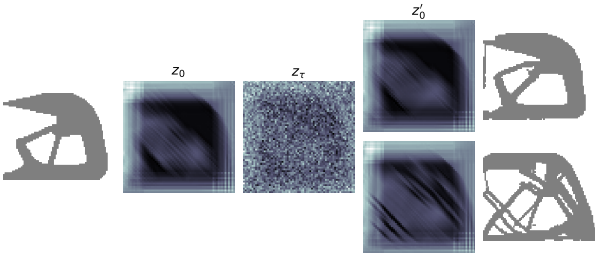


Fig. 7: TopoCtrl latents for controlling the number of joints. The topology is first converted into latent z_0, then partial noise is added to obtain z_τ, after which regression-guided denoising is applied to steer the latent toward a lower or higher target joint count, producing the edited latent z_0' . Decoding z_0' gives the controlled topology. The visualization shows how the latent trajectory preserves the overall structural similarity while still allowing the diffusion model to reorganize the connectivity toward the requested level of joint complexity.

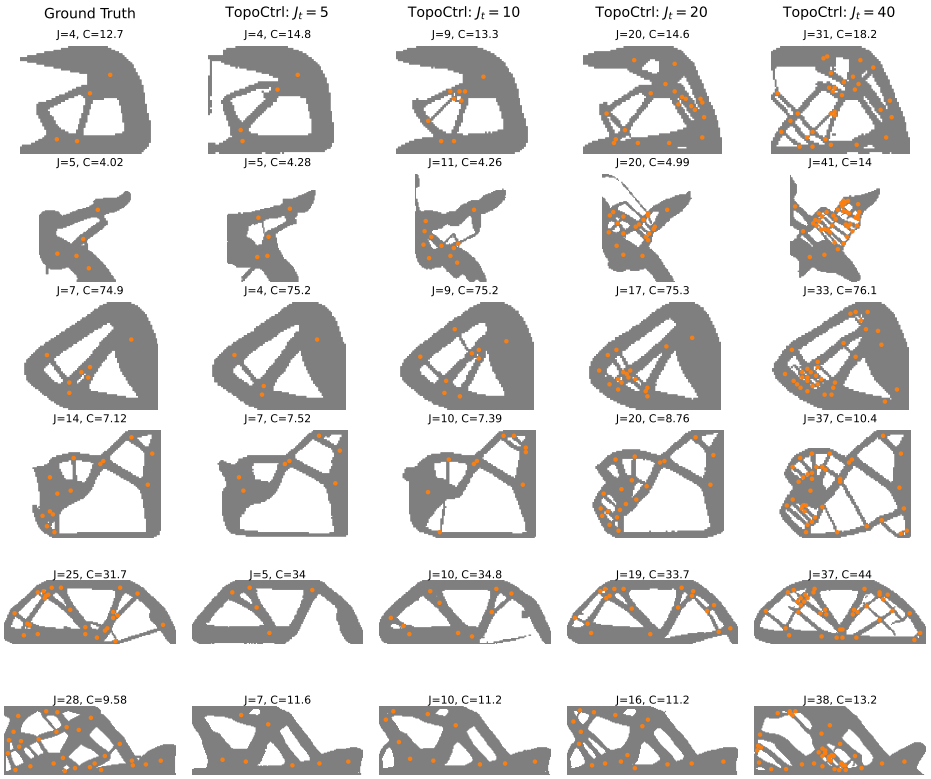


Fig. 8: Selected examples of joint-count control using TopoCtrl. We define a joint where the medial axis skeletonization branches out. We show six reference topologies together with edited results for target joint counts of 5, 10, 20, and 40. The method successfully steers the topology toward simpler or more articulated structures depending on the target. With lower target values, the generated topologies generally contain fewer branches and connection points. With higher target values, the structures become more complex and develop richer internal connectivity, although reaching very high joint counts can be challenging for initially simple designs. The actual structure generated from TopoCtrl may contain more or fewer joints than the desired target, but in general, they center around the target joints.

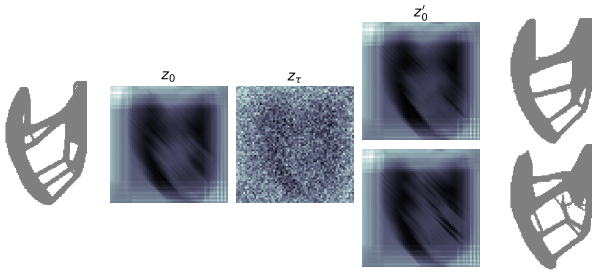


Fig. 9: TopoCtrl latents for controlling the maximum number of members going into a joint.

4.2.2 Member joint complexity control

We next control the maximum number of members going into a joint. Figure 9 shows representative latent trajectories, and Figure 10 visualizes six examples with target values of 3, 4, and 5.

The behavior is qualitatively similar to joint-count control, but now the modification is more localized to the structure of the junctions. When the target is 3, the resulting topologies tend to favor simpler Y-like or chain-like connectivity with fewer highly connected nodes. When the target is increased to 5, the model produces thicker and more complex junction regions with more members meeting at the same node. This control is noteworthy because it targets a discrete and difficult-to-differentiate structural quantity that is cumbersome to impose directly in conventional topology optimization.

4.2.3 Length control

We next study control of normalized characteristic length, defined here as the normalized average of the two longest members extracted from the skeletonized topology. Figure 11 visualizes the corresponding latent trajectories, and Figure 12 shows six examples with target values of 0.25, 0.5, 0.74, and 1.0.

Despite the fact that explicit feature-length variation is not part of the original OAT dataset, the regression-guided denoising process still succeeds in steering the topology toward shorter or longer dominant members. Lower target values lead to structures with more fragmented or interrupted long paths, while higher target values encourage the formation of longer continuous loops or members. In particular, for the largest target value of 1.0, we observe generated topologies with visibly elongated structural loops that increase the characteristic member length. This behavior is especially interesting because classical feature-scale control in topol-

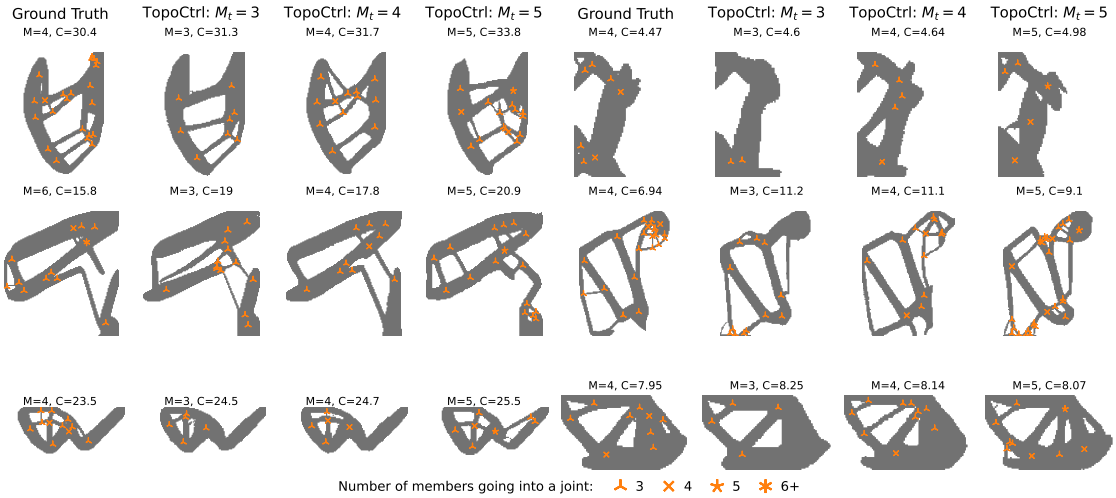


Fig. 10: Selected examples of the maximum number of members going into a joint control. We show six topologies edited toward target values of 3, 4, and 5 for the maximum number of members connected to a joint. Similar to joint-count control, lower target values encourage simpler connectivity, while higher target values lead to thicker and more complex junction with more members meeting at the same node.

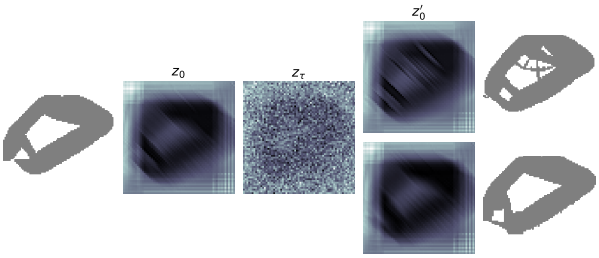


Fig. 11: TopoCtrl latents for controlling normalized characteristic length, defined as the normalized average of the two longest members extracted from the topology skeleton.

ogy optimization is typically framed as limiting the largest admissible member size, whereas TopoCtrl can also move the design in the opposite direction and deliberately increase characteristic member length.

4.2.4 Thickness control

Finally, we consider control of normalized member thickness, measured as the normalized 90th-percentile thickness of the topology. Since the characteristic is evaluated at the 90th-percentile, we can still get some members with thicker and thinner than the threshold. Figure 13 shows representative latent trajectories, and Figure 14 presents six examples with target values of 0.075, 0.15, 0.225, and 0.3.

Although the OAT training data is produced with a fixed filter size, TopoCtrl is still able to vary the thickness in both directions. When the target thickness is reduced, the generated topologies tend to contain thinner members and more branching. When the target thickness is increased, the structures become visibly thicker. This result shows that the proposed guidance mechanism can extrapolate beyond the nominal feature scale implied by the original dataset generation process.

4.3 Cantilever beam

Despite all the previous topologies experimented with being from the testing set, which neither OAT nor the regression model saw during training, generally, the topologies are still within the overall randomized boundary condition style. To verify TopoCtrl's performance on a common, out-of-distribution topology-optimization test case, we use a cantilever-beam example and control its characteristics (Figure 15). We observe successful characteristic control across all four characteristics.

These experiments show that TopoCtrl can transform a pre-trained topology foundation model from a forward generator into a control tool. By combining partial noising with regression-guided denoising, the method enables explicit manipulation of structural characteristics that are discrete, non-smooth, or otherwise difficult to integrate into conventional optimization formulations.

5 LIMITATIONS

This work has several limitations that warrant further investigation. First, we have not exhaustively explored all possible structural characteristics, and the current characteristics rely on manually chosen post-processing values. In addition, we did not obtain manual labeling for these characteristics, which may further provide more intuitive labeling. For some characteristics, the regression models also show room for improvement in predictive accuracy. Furthermore, due to the stochastic nature of diffusion models, we relied on up to 64 samples for a single control instance to obtain reliable estimates; with fewer samples, the accuracy remains insufficient, limiting the method's usefulness as a highly responsive interactive tool for practitioners. Finally, although we demonstrate the approach on both the testing set and a cantilever beam example, it has not yet been extensively evaluated across the full range of common topology optimization

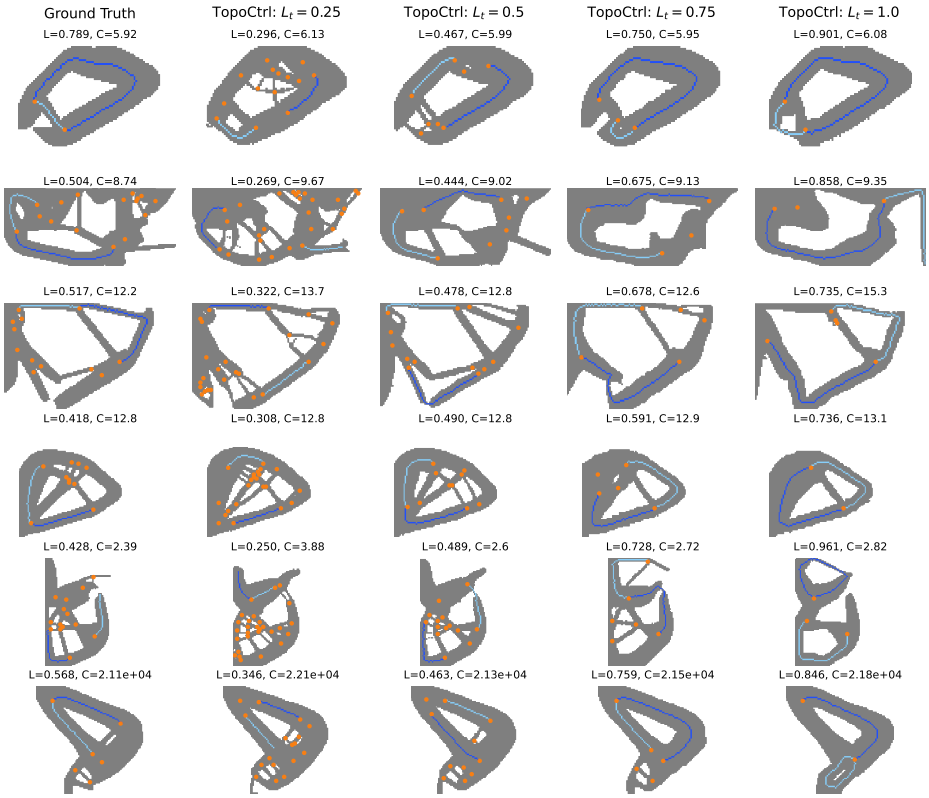


Fig. 12: Selected examples of characteristic-length control using TopoCtrl. We show six topologies edited toward target values of 0.25, 0.5, 0.75, and 1.0 for the normalized average length of the two longest members. We highlight the longest member in dark blue and the second longest member in light blue. Lower target values generally produce structures with shorter members, while higher target values encourage the formation of longer continuous loops or members.

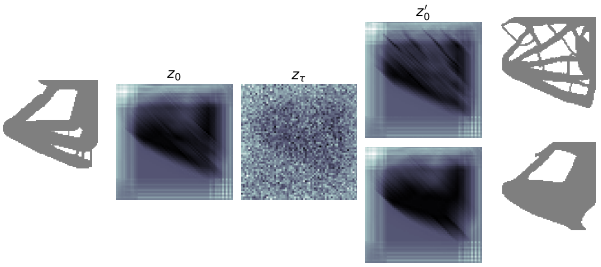


Fig. 13: TopoCtrl latents for controlling normalized member thickness, measured here as the normalized 90th-percentile thickness of the topology.

problems. As a result, the generalizability of the proposed approach remains to be thoroughly explored.

6 CONCLUSION AND FUTURE WORK

In this work, we introduced TopoCtrl, a regression-guided latent diffusion framework for post-optimization control of structural characteristics in topology optimization. Across the experiments, we demonstrated control over four representative characteristics. These results show that TopoCtrl can operate not only on continuous geometric characteristics but also on discrete and difficult-to-differentiate characteristics that are cumbersome to impose through conventional topology optimization pipelines. The method achieved good target-reaching behavior, with broader reachable characteristic coverage than the benchmark based on varying the filter radius, while still maintaining good similarity to the original topology. The results highlight the main advantage of TopoCtrl as it provides a practical route to explicit characteristic control in settings where only a forward evaluation procedure is readily available.

Several directions can extend the present work. First, the four characteristics studied here represent only a small initial set of continuous and discrete quantities chosen to demonstrate the effectiveness of the framework. The same underlying pipeline can be applied to many other structural characteristics, with the practical limitation being the ability of the regression model to reliably learn the desired evaluator from latent representations. Second, future work can explore more diverse and application-driven characteristics. For example, quantities such as the number of holes, the minimum angle between members at joints, or other measures of geometric complexity, manufacturability should also be compatible with the TopoCtrl framework as long as they can be evaluated and labeled for training. Extending the method in this direction would further increase its usefulness as a practical design interface for engineers, especially in cases where the target quantity is easy to measure but difficult to optimize for directly. Finally, an important long-term direction is ex-

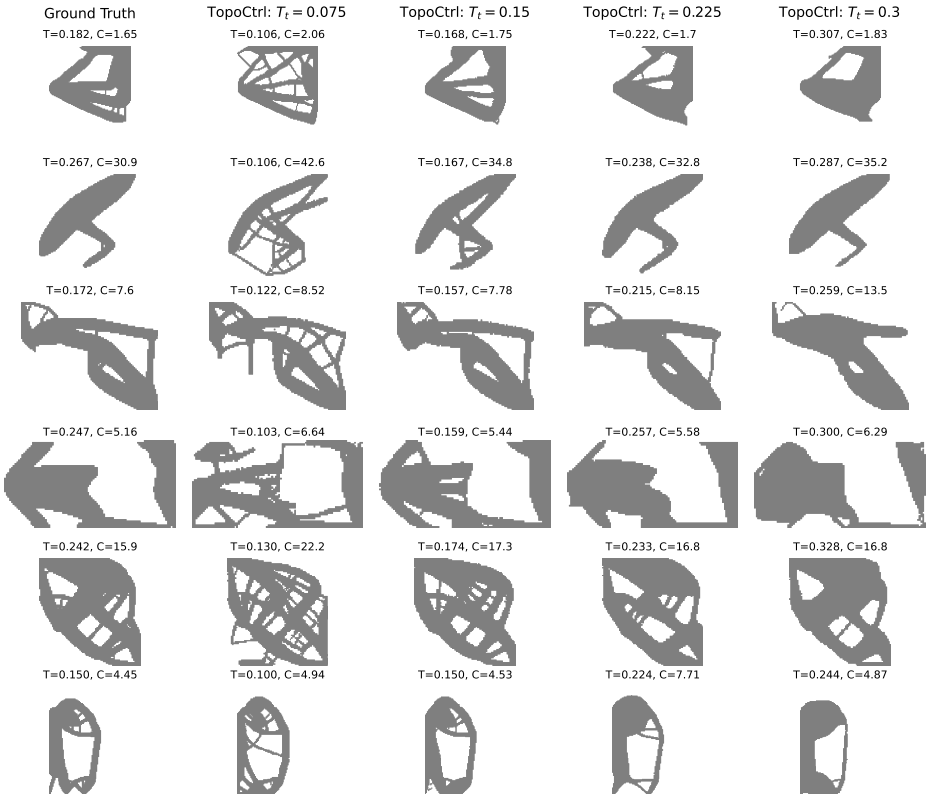


Fig. 14: Selected examples of thickness control using TopoCtrl. We show six topologies edited toward normalized thickness values of 0.075, 0.15, 0.225, and 0.3. When the target thickness is reduced, the resulting topologies tend to contain thinner members and more branching. When the target thickness is increased, the dominant load-bearing members become thicker.

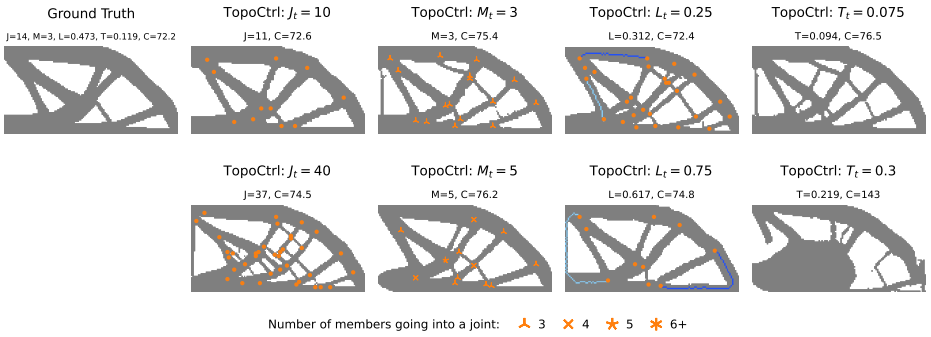


Fig. 15: Characteristic control of a cantilever beam. The characteristics and compliance of the ground truth topology are evaluated. Lower and higher bound control is then established for the four characteristics. The lower range characteristics are shown in the top row, and the higher range characteristics in the bottom row. Successful control of all four characteristics is observed in the cantilever beam example.

tension to three-dimensional topology optimization. Once a 3D topology foundation model becomes available, the same regression-guided latent control pipeline can be generalized to volumetric designs, enabling post-optimization manipulation of 3D structural characteristics under complex boundary conditions. This would open the door to characteristic control for practical 3D engineering components and broaden the applicability of TopoCtrl beyond planar topologies.

ACKNOWLEDGEMENT

The authors would like to thank the support of Tata Steel and MIT Generative AI Impact Consortium.

References

[1] Guest, J. K., Prévost, J. H., and Belytschko, T., 2004. "Achieving minimum length scale in topology optimization using nodal design variables and projection functions". *International Journal for Numerical Methods in Engineering,* ***61***(2), pp. 238–254.

[2] Guest, J. K., 2009. "Imposing maximum length scale in topology optimization". *Structural and Multidisciplinary Optimization,* ***37***(5), pp. 463–473.

[3] Lazarov, B. S., and Wang, F., 2017. "Maximum length scale in density based topology optimization". *Computer Methods in Applied Mechanics and Engineering,* ***318***, pp. 826–844.

[4] Fernández, E. F., Yang, K.-k., Koppen, S., Alarcón, P., Bauduin, S., and Duysinx, P., 2020. "Imposing minimum and maximum member size, minimum cavity size, and minimum separation distance between solid members in topology optimization". *Computer Methods in Applied Mechanics and Engineering,* ***368***, p. 113157.

[5] Zobaer, T., and Sutradhar, A., 2023. "Maximum thickness control in topology optimization using an inflection-point-based geometric constraint". *Computer Methods in Applied Mechanics and Engineering,* ***414***, p. 116171.

[6] Allaire, G., Jouve, F., and Michailidis, G., 2016. "Thickness control in structural optimization via a level set method". *Structural and Multidisciplinary Optimization,* ***53***(6), pp. 1349–1382.

[7] Prager, W., 1978. "Optimal layout of trusses with finite numbers of joints". *Journal of the Mechanics and Physics of Solids,* ***26***(4), pp. 241–250.

[8] Kanno, Y., and Fujita, S., 2018. "Alternating direction method of multipliers for truss topology optimization with limited number of nodes: A cardinality-constrained second-order cone programming approach". *Optimization and Engineering,* ***19***, pp. 327–358.

[9] Ohsaki, M., and Katoh, N., 2005. "Topology optimization of trusses with stress and local constraints on nodal stability and member intersection". *Structural and Multidisciplinary Optimization,* ***29***, pp. 190–197.

[10] Cai, Q., Feng, R., and Zhang, Z., 2022. "Topology optimization of truss structure considering nodal stability and local buckling stability". *Structures,* ***40***, pp. 64–73.

[11] Han, H., Guo, Y., Chen, S., and Liu, Z., 2021. "Topological constraints in 2d structural topology optimization". *Structural and Multidisciplinary Optimization,* ***63***, pp. 39–58.

[12] Sigmund, O., 2007. "Morphology-based black and white filters for topology optimization". *Structural and Multidisciplinary Optimization,* ***33***(4–5), pp. 401–424.

[13] Ha, D. Q., and Carstensen, J. V., 2023. "Human-informed topology optimization: interactive application of feature size controls". *Structural and Multidisciplinary Optimization,* ***66***, p. 59.

[14] Schiffer, G., Ha, D. Q., and Carstensen, J. V., 2023. "Hitop 2.0: combining topology optimisation with multiple feature size controls and human preferences". *Virtual and Physical Prototyping,* ***18***(1).

[15] Nie, Z., Lin, T., Jiang, H., and Kara, L. B., 2021. "TopologyGAN: Topology optimization using generative adversarial networks based on physical fields over the initial domain". *Journal of Mechanical Design,* ***143***(3), p. 031715.

[16] Giannone, G., and Ahmed, F., 2023. "Diffusing the optimal topology: A generative optimization approach". *arXiv preprint arXiv:2303.09760*.

[17] Giannone, G., Srivastava, A., Winther, O., and Ahmed, F., 2023. "Aligning optimization trajectories with diffusion models for constrained design generation". In Advances in Neural Information Processing Systems, Vol. 36, pp. 51830–51861.

[18] Heyrani Nobari, A., Regenwetter, L., Picard, C., Han, L., and Ahmed, F., 2025. "Optimize any topology: A foundation model for shape- and resolution-free structural topology optimization". *arXiv preprint arXiv:2510.23667*.

[19] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B., 2022. "High-resolution image synthesis with latent diffusion models". In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

[20] Song, J., Meng, C., and Ermon, S., 2021. "Denoising diffusion implicit models". In International Conference on Learning Representations.

[21] Kim, G., Kwon, T., and Ye, J. C., 2022. "Diffusionclip: Text-guided diffusion models for robust image manipulation". In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

[22] Kawar, B., Zada, S., Lang, O., Tov, O., Chang, H., Dekel, T., Mosseri, I., and Irani, M., 2023. "Imagic: Text-based real image editing with diffusion models". In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

[23] Mokady, R., Hertz, A., Aberman, K., Pritch, Y., and Cohen-Or, D., 2023. "Null-text inversion for editing real images using guided diffusion models". In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

[24] Shi, Y., Xue, C., Liew, J. H., Pan, J., Yan, H., Zhang, W., Tan, V. Y. F., and Bai, S., 2024. "Dragdiffusion: Harnessing diffusion models for interactive point-based image editing". In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

[25] Mou, C., Wang, X., Song, J., Shan, Y., and Zhang, J., 2024. "Dragondiffusion: Enabling drag-style manipulation on diffusion models". In International Conference on Learning Representations.

[26] Bendsøe, M. P., and Kikuchi, N., 1988. "Generating optimal topologies in structural design using a homogenization method". *Computer Methods in Applied Mechanics and Engineering,* ***71***(2), pp. 197–224.

[27] Rozvany, G. I. N., Zhou, M., and Birker, T., 1992. "Generalized shape optimization without homogenization". *Structural Optimization,* ***4***, pp. 250–252.

[28] Bendsøe, M. P., and Sigmund, O., 1999. "Material interpolation schemes in topology optimization". *Archive of Applied Mechanics,* ***69***, pp. 635–654.

[29] Sethian, J. A., and Wiegmann, A., 2000. "Structural

boundary design via level set and immersed interface methods". *Journal of Computational Physics,* ***163***(2), pp. 489–528.

[30] Allaire, G., Jouve, F., and Toader, A.-M., 2004. "Structural optimization using sensitivity analysis and a level-set method". *Journal of Computational Physics,* ***194***(1), pp. 363–393.

[31] Behzadi, M. M., and Ilieş, H. T., 2022. "GANTL: Toward practical and real-time topology optimization with conditional generative adversarial networks and transfer learning". *Journal of Mechanical Design,* ***144***(2), p. 021711.

[32] Sharpe, C., and Seepersad, C. C., 2019. "Topology design with conditional generative adversarial networks". In Proceedings of the ASME 2019 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE), p. V02AT03A062.

[33] Li, B., Huang, C., Li, X., Zheng, S., and Hong, J., 2019. "Non-iterative structural topology optimization using deep learning". *Computer-Aided Design,* ***115***, pp. 172–180.

[34] Yu, Y., Hur, T., Jung, J., and Jang, I. G., 2019. "Deep learning for determining a near-optimal topological design without any iteration". *Structural and Multidisciplinary Optimization,* ***59***, pp. 787–799.

[35] Zhang, Z., Li, Y., Zhou, W., Chen, X., Yao, W., and Zhao, Y., 2021. "TONR: An exploration for a novel way combining neural network with topology optimization". *Computer Methods in Applied Mechanics and Engineering,* ***386***, p. 114083.

[36] Wang, D., Xiang, C., Pan, Y., Chen, A., Zhou, X., and Zhang, Y., 2022. "A deep convolutional neural network for topology optimization with perceptible generalization ability". *Engineering Optimization,* ***54***(6), pp. 973–988.

[37] Hu, J., He, Y., Xu, B., Wang, S., Lei, N., and Luo, Z., 2024. "IF-TONIR: Iteration-free topology optimization based on implicit neural representations". *Computer-Aided Design,* ***167***, p. 103639.

[38] Rawat, S., and Shen, M.-H. H., 2019. "A novel topology optimization approach using conditional deep learning". *arXiv preprint arXiv:1901.04859*.

[39] Keshavarzzadeh, V., Alirezaei, M., Tasdizen, T., and Kirby, R. M., 2021. "Image-based multiresolution topology optimization using deep disjunctive normal shape model". *Computer-Aided Design,* ***130***, p. 102947.

[40] Zheng, S., He, Z., and Liu, H., 2021. "Generating three-dimensional structural topologies via a U-Net convolutional neural network". *Thin-Walled Structures,* ***159***, p. 107263.

[41] Heyrani Nobari, A., Giannone, G., Regenwetter, L., and Ahmed, F., 2024. "NITO: Neural implicit fields for resolution-free topology optimization". *arXiv preprint arXiv:2402.05073*.

[42] Mazé, F., and Ahmed, F., 2023. "Diffusion models beat GANs on topology optimization". In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 37, pp. 9108–9116.

[43] Schmidt, M.-P., Pedersen, C. B. W., and Gout, C., 2019. "On structural topology optimization using graded porosity control". *Structural and Multidisciplinary Optimization,* ***60***, pp. 1437–1453.

[44] Guo, X., Zhang, W., and Zhong, W., 2014. "Explicit feature control in structural topology optimization via level set method". *Computer Methods in Applied Mechanics and Engineering,* ***272***, pp. 354–378.

[45] Schiffer, G., Schmidt, M.-P., Pedersen, C. B. W., and Carstensen, J. V., 2024. "Interactive infill topology optimisation guided by user drawn patterns". *Virtual and Physical Prototyping,* ***19***(1), p. e2361864.

[46] Navez, T., Schmidt, M.-P., Sigmund, O., and Pedersen, C. B. W., 2022. "Topology optimization guided by a geometrical pattern library". *Structural and Multidisciplinary Optimization,* ***65***, p. 108.

[47] Zhu, L., Zuo, T., Wang, C., Wang, Q., Yu, Z., and Liu, Z., 2024. "Hole appearance constraint method in 2d structural topology optimization". *Mathematics,* ***12***(17), p. 2645.

[48] Meng, X., Zhang, L.-Y., Zhao, Z.-L., and Xie, Y. M., 2023. "A direct approach to achieving efficient free-form shells with embedded geometrical patterns". *Thin-Walled Structures,* ***185***, p. 110559.

[49] Li, Y., Zhang, Z., Luo, J., Peng, W., Zhou, W., and Yao, W., 2024. "Concurrent topology optimization of shells with pattern-guided infills for intuitive design and additive manufacturing". *Computer Methods in Applied Mechanics and Engineering,* ***418***, p. 116485.

[50] Zhang, Z., Jiang, L., Yarlagadda, T., Zheng, Y., and Usmani, A., 2024. "A novel multi-pattern control for topology optimization to balance form and performance needs". *Engineering Structures,* ***303***, p. 117581.

[51] Li, Z., Lee, T.-U., and Xie, Y. M., 2023. "Interactive structural topology optimization with subjective scoring and drawing systems". *Computer-Aided Design,* ***160***, p. 103532.

[52] Li, Z., Lee, T.-U., and Xie, Y. M., 2025. "Interactive 3d structural design in virtual reality using preference-based topology optimization". *Computer-Aided Design,* ***180***, p. 103826.

[53] Zhu, S., Hu, J., Qi, J., Wang, L., Guo, J., Ma, J., and Zhu, G., 2025. "Sketch-guided topology optimization with enhanced diversity for innovative structural design". *Applied Sciences,* ***15***(5), p. 2753.

[54] Zhang, W., Wang, Y., Youn, S.-K., and Guo, X., 2024. "Machine learning powered sketch aided design via topology optimization". *Computer Methods in Applied Mechanics and Engineering,* ***419***, p. 116651.

[55] Mueller, C. T., and Ochsendorf, J., 2015. "Combining structural performance and designer preferences in evolutionary design space exploration". *Automation in Construction,* ***52***, pp. 70–82.

[56] Zhang, W., Zhuang, X., Guo, X., and Youn, S.-K., 2026. "A novel approach for topology optimiza-

tion mirroring human intention: Introducing a pattern-embedded filter built by machine learning". *Computer Methods in Applied Mechanics and Engineering,* ***455***, p. 118882.

[57] Poulsen, T. A., 2003. "A new scheme for imposing a minimum length scale in topology optimization". *International Journal for Numerical Methods in Engineering,* ***57***, pp. 741–760.

[58] Zhou, M., Shyy, Y. K., and Thomas, H. L., 2001. "Checkerboard and minimum member size control in topology optimization". *Structural and Multidisciplinary Optimization,* ***21***, pp. 152–158.

[59] Song, L., Gao, T., Wang, J., and Zhang, W., 2024. "Directional maximum length scale control in density-based topology optimization". *Computers & Structures,* ***292***, p. 107236.

[60] Zhang, W., Zhong, W., and Guo, X., 2014. "An explicit length scale control approach in SIMP-based topology optimization". *Computer Methods in Applied Mechanics and Engineering,* ***282***, pp. 71–86.

[61] Zuo, T., Han, H., Wang, Q., Zhao, Q., and Liu, Z., 2025. "An explicit topology and thickness control approach in SIMP-based topology optimization". *Computers & Structures,* ***307***, p. 107631.

[62] Carroll, J. D., and Guest, J. K., 2022. "Topology optimization of uniform thickness structures using discrete object projection". *Structural and Multidisciplinary Optimization,* ***65***, p. 271.

[63] Xia, Q., and Shi, T., 2015. "Constraints of distance from boundary to skeleton: For the control of length scale in level set based structural topology optimization". *Computer Methods in Applied Mechanics and Engineering,* ***295***, pp. 525–542.

[64] Wang, Y., Zhang, L., and Wang, M. Y., 2016. "Length scale control for structural optimization by level sets". *Computer Methods in Applied Mechanics and Engineering,* ***305***, pp. 891–909.

[65] Liu, J., Li, L., and Ma, Y., 2018. "Uniform thickness control without pre-specifying the length scale target under the level set topology optimization framework". *Advances in Engineering Software,* ***115***, pp. 204–216.

[66] Tran, Q.-D., Kambampati, S., Kim, H. A., and Jang, G.-W., 2022. "Level set-based structural optimization with uniform wall thickness for wire-fed metal additive manufacturing". *Journal of Mechanical Design,* ***144***(7), p. 071702.

[67] Huang, J., and Liu, J., 2024. "Derivable skeletons in topology optimization for length scale control". *Computer Methods in Applied Mechanics and Engineering,* ***421***, p. 116778.

[68] Wang, Y., Hu, T., Zou, Z., Wang, Y., and Zhang, X., 2025. "Constraint-free length scale control for topology optimization using the velocity field level set method". *Structural and Multidisciplinary Optimization,* ***68***, p. 46.

[69] Wang, Q., Han, H., Wang, C., and Liu, Z., 2022. "Topological control for 2d minimum compliance topology optimization using SIMP method". *Structural and Multidisciplinary Optimization,* ***65***, p. 38.

[70] Zuo, T., Wang, C., Han, H., Wang, Q., and Liu, Z., 2022. "Explicit 2d topological control using SIMP and MMA in structural topology optimization". *Structural and Multidisciplinary Optimization,* ***65***, p. 293.

[71] Zuo, T., Han, H., and Liu, Z., 2023. "Explicit tunnels and cavities control using SIMP and MMA in structural topology optimization". *Computer-Aided Design,* ***158***, p. 103482.

[72] Liang, Y., Yan, X., and Cheng, G., 2022. "Explicit control of 2d and 3d structural complexity by discrete variable topology optimization method". *Computer Methods in Applied Mechanics and Engineering,* ***389***, p. 114302.

[73] Zhao, Z.-L., Zhou, S., Cai, K., and Xie, Y. M., 2020. "A direct approach to controlling the topology in structural optimization". *Computers & Structures,* ***227***, p. 106141.

[74] He, Y., Zhao, Z.-L., Cai, K., Kirby, J., Xiong, Y., and Xie, Y. M., 2022. "A thinning algorithm based approach to controlling structural complexity in topology optimization". *Finite Elements in Analysis and Design,* ***207***, p. 103779.

[75] He, Y., Zhao, Z.-L., Lin, X., and Xie, Y. M., 2023. "A hole-filling based approach to controlling structural complexity in topology optimization". *Computer Methods in Applied Mechanics and Engineering,* ***416***, p. 116391.

[76] Zhang, W., Liu, Y., Wei, P., Zhu, Y., and Guo, X., 2017. "Explicit control of structural complexity in topology optimization". *Computer Methods in Applied Mechanics and Engineering,* ***324***, pp. 149–169.

[77] Zhang, W., Zhou, J., Zhu, Y., and Guo, X., 2017. "Structural complexity control in topology optimization via moving morphable component (MMC) approach". *Structural and Multidisciplinary Optimization,* ***56***, pp. 535–552.

[78] Zhang, W., Li, D., Zhang, J., and Guo, X., 2016. "Minimum length scale control in structural topology optimization based on the moving morphable components (MMC) approach". *Computer Methods in Applied Mechanics and Engineering,* ***311***, pp. 327–355.

[79] Fairclough, H., and Gilbert, M., 2020. "Layout optimization of simplified trusses using mixed integer linear programming with runtime generation of constraints". *Structural and Multidisciplinary Optimization,* ***61***, pp. 1977–1999.

[80] Lu, H., and Xie, Y. M., 2023. "Reducing the number of different members in truss layout optimization". *Structural and Multidisciplinary Optimization,* ***66***, p. 52.

[81] Sohl-Dickstein, J., Weiss, E. A., Maheswaranathan, N., and Ganguli, S., 2015. "Deep unsupervised learning using nonequilibrium thermodynamics". In Proceedings of the 32nd International Conference on Machine Learning, Vol. 37 of *Proceedings of Machine Learning Research*, pp. 2256–2265.

[82] Ho, J., Jain, A., and Abbeel, P., 2020. "Denoising diffusion probabilistic models". In Advances in Neural

Information Processing Systems.

[83] Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B., 2021. "Score-based generative modeling through stochastic differential equations". In International Conference on Learning Representations.

[84] Ho, J., and Salimans, T., 2022. Classifier-free diffusion guidance.

[85] Dhariwal, P., and Nichol, A., 2021. "Diffusion models beat GANs on image synthesis". In Advances in Neural Information Processing Systems.

[86] Salimans, T., and Ho, J., 2022. "Progressive distillation for fast sampling of diffusion models". In International Conference on Learning Representations.

[87] Song, Y., Dhariwal, P., Chen, M., and Sutskever, I., 2023. "Consistency models". In Proceedings of the 40th International Conference on Machine Learning.

[88] Luo, S., Tan, Y., Huang, L., Li, J., and Zhao, H., 2023. Latent consistency models: Synthesizing high-resolution images with few-step inference.

[89] Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., and Cohen-Or, D., 2022. Prompt-to-prompt image editing with cross attention control.

[90] Brooks, T., Holynski, A., and Efros, A. A., 2023. "Instructpix2pix: Learning to follow image editing instructions". In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

[91] Cao, M., Wang, X., Qi, Z., Shan, Y., Qie, X., and Zheng, Y., 2023. "Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing". In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV).

[92] Zhao, X., Guan, J., Fan, C., Xu, D., Lin, Y., Pan, H., and Feng, P., 2024. "Fastdrag: Manipulate anything in one step". In Advances in Neural Information Processing Systems.

[93] Liu, H., Xu, C., Yang, Y., Zeng, L., and He, S., 2024. "Drag your noise: Interactive point-based editing via diffusion semantic propagation". In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

[94] Go, H., Lee, Y., Kim, J.-Y., Lee, S., Jeong, M., Lee, H. S., and Choi, S., 2023. "Towards practical plug-and-play diffusion models". In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1962–1971.

[95] Luo, G., Darrell, T., Wang, O., Goldman, D. B., and Holynski, A., 2024. "Readout guidance: Learning control from diffusion features". In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 8217–8227.

[96] Bansal, A., Chu, H.-M., Schwarzschild, A., Sengupta, S., Goldblum, M., Geiping, J., and Goldstein, T., 2024. "Universal guidance for diffusion models". In International Conference on Learning Representations.