INT3RCEPT: PHYSICS-DRIVEN STRUCTURAL REDESIGN VIA 3D GENERATIVE INTERMEDIATE REPRESENTATIONS

Anonymous authors
Paper under double-blind review

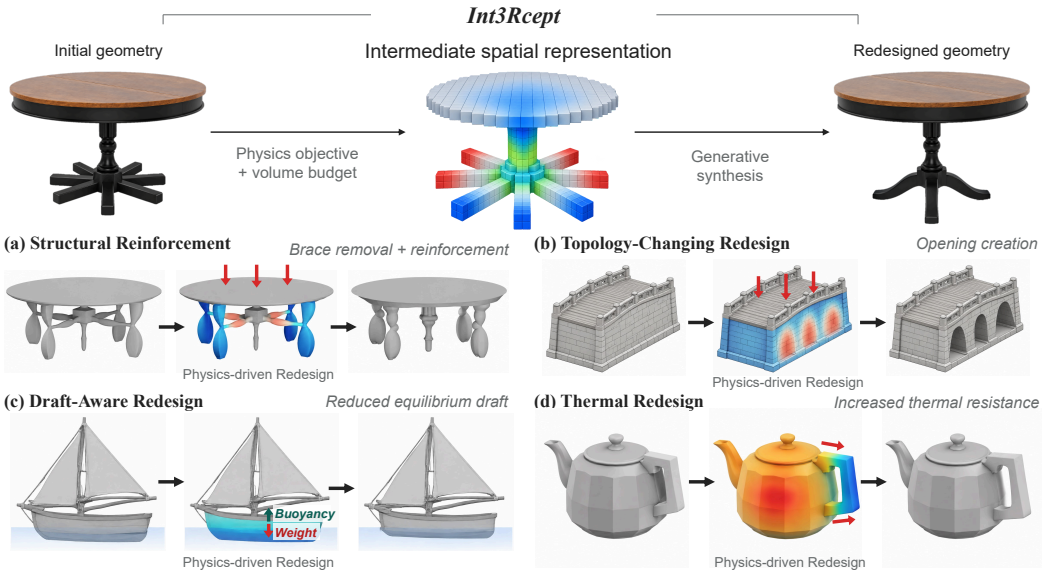

Figure 1: **Overview of Int3Rcept.** Int3Rcept intercepts an intermediate spatial representation of a pretrained 3D generator and performs physics-driven redesign. Examples include (a) structural reinforcement, (b) topology-changing redesign, (c) floating-body redesign, and (d) thermal redesign.

ABSTRACT

Modern 3D generative models synthesize visually compelling assets, but physics-driven redesign must support substantial structural changes without sacrificing visual quality. Existing approaches struggle to achieve both: learned-feature optimization preserves generative priors but provides only indirect control over structural changes, while explicit grid-based optimization offers direct spatial control but operates outside the native synthesis space of pretrained generators.

We introduce **Int3Rcept**, a physics-driven 3D redesign framework that intercepts and redesigns a pretrained generator's intermediate spatial representation under physical objectives. Our key idea is to use a spatially explicit representation as a direct design space while remaining compatible with downstream synthesis. Int3Rcept expands the intermediate structure into a volumetric design domain, optimizes a continuous density field under volume constraints, and converts it back through connectivity-aware discretization, followed by downstream synthesis.

Int3Rcept achieves stronger physical improvement than learned-feature optimization and higher visual quality than explicit grid-based optimization. It also shows strong controllability across target material volumes, enabling support pruning, reinforcement, and opening creation. We further demonstrate the interface across multiple physical objectives and pretrained 3D generators. Together, these results establish intermediate spatial representations as a practical bridge between physical optimization and 3D generative priors.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1 INTRODUCTION

Modern 3D generative models can synthesize visually compelling assets (Gao et al., 2022; Zhang et al., 2024; Xiang et al., 2025), yet an asset that looks plausible is not necessarily well suited for a physical task. Improving physical performance often requires substantial structural changes, such as thickening, removing or reallocating redundant structures, or creating new openings. A canonical example is compliance minimization, which seeks a stiffer structure under prescribed loads subject to a volume constraint (Bendsøe, 1989; Bendsøe & Sigmund, 2003; Sigmund, 2001). More broadly, tasks such as floating-body and thermal optimization likewise require reorganizing a prescribed material budget to satisfy different physical objectives. We refer to this broader class of problems as *physics-driven 3D redesign*: given an object, we seek to reorganize its 3D structure according to a physical objective while preserving its semantic identity and geometric characteristics.

Existing approaches to physics-driven optimization expose a trade-off. Explicit geometry optimization provides direct spatial control over shape or material distribution (Guo et al., 2024; Huang et al., 2024; Bendsøe, 1989; Bendsøe & Sigmund, 2003; Sigmund, 2001). However, these approaches either rely primarily on deformation of existing geometry, limiting structural reconfiguration, or operate outside pretrained generative representations and therefore do not directly benefit from learned geometric priors. Conversely, learned-feature optimizations preserve such priors by optimizing generative features or latent variables that encode geometry and spatial relationships (Zhan et al., 2025; Ijaz et al., 2025), but structural changes remain mediated through learned features rather than being directly controlled in space.

Our key observation is that several recent 3D generators expose *intermediate spatial representations* within their generative pipelines (Xiang et al., 2025; Ren et al., 2024; He et al., 2025). Such a representation offers a useful interface for physical redesign: it is spatially grounded enough to enable explicit material manipulation, yet it remains on the synthesis path of the pretrained generator. This suggests a third design space between latent optimization and external grid optimization. We instantiate this idea in **Int3Rcept**, a framework that performs physics-driven structural redesign directly in the intermediate spatial representation of a pretrained 3D generator.

Int3Rcept combines explicit spatial control for substantial structural changes with compatibility with downstream generative synthesis, thereby retaining high visual quality. Given a task-specific physical objective, we define task-specific semantic regions to control which parts of the object may be modified, preserved, or used to impose physical conditions. Int3Rcept then expands the intermediate structure into a volumetric design domain that permits material removal, addition, and redistribution beyond the initial geometry. Within this domain, a continuous density field is optimized under a prescribed material budget using gradients from differentiable physical objectives. Finally, connectivity-aware discretization converts the optimized density into a task-valid intermediate structure, from which frozen downstream stages synthesize the detailed final geometry.

We evaluate whether the choice of optimization space matters by comparing Int3Rcept with representative learned-feature and explicit-grid alternatives. Compared with PhysiOpt (Zhan et al., 2025), representing learned-feature optimization, Int3Rcept achieves substantially greater compliance reduction at comparable volume and exhibits a much broader response to changes in the target volume budget. Compared with explicit grid-based density optimization, Int3Rcept better preserves geometric quality and fine details. Beyond compliance minimization, we demonstrate floating-body redesign through equilibrium-draft optimization, thermally resistant handle redesign, and transfer across multiple pretrained 3D generators. These results support intermediate spatial representations as a practical interface between explicit physical redesign and learned 3D generative priors.

Our contributions are:

- We formulate *physics-driven 3D redesign* as structural reorganization under physical and volume constraints, and introduce **Int3Rcept**, which uses intermediate spatial representations of pretrained 3D generators as a design interface that combines explicit structural control and learned generative synthesis.
- We empirically validate the choice of optimization space through controlled comparisons with learned-feature and explicit-grid alternatives. Int3Rcept achieves stronger physical improvement than learned-feature optimization, while preserving higher geometric quality than explicit grid optimization.

- We demonstrate controllable redesign across material budgets, producing diverse structural changes including material reallocation, support pruning, reinforcement, and opening creation. We further show that the same design interface generalizes across multiple physical objectives and pretrained 3D generators.

2 Related Work

2.1 Physics-Driven Optimization for 3D Assets

Explicit geometry optimization. A long line of work optimizes explicit geometric representations under physical objectives. Differentiable simulation allows surface or volumetric geometry to be updated through gradients of mechanical objectives and constraints. PhysComp (Guo et al., 2024), for example, optimizes rest-shape geometry, material properties, and external forces under static equilibrium, while differentiable FEM frameworks such as DiffIPC (Huang et al., 2024) enable gradient-based optimization of meshes under physical objectives. These methods provide direct geometric control, but their design space is defined over explicit geometry rather than a learned generative representation. PhysComp optimizes the rest shape of an existing mesh, limiting structural reconfiguration beyond its topology. DiffIPC supports flexible shape optimization, but without a learned generative prior, the optimized geometry does not benefit from a learned shape prior.

Density-based topology optimization goes further by assigning spatially indexed material variables over an explicit design domain (Bendsøe, 1989; Bendsøe & Sigmund, 2003; Sigmund, 2001). This enables direct material redistribution and topology changes beyond deformation of an existing mesh. However, the optimization remains outside learned generative representations, so the resulting structure does not directly benefit from learned geometric priors.

Physics optimization with learned generative priors. Recent methods incorporate physical objectives or simulator feedback into pretrained 3D generative models. PhysiOpt (Zhan et al., 2025) optimizes learned features of pretrained 3D generative models using differentiable finite-element analysis under user-specified physical conditions. SATO (Ijaz et al., 2025) performs topology optimization over VAE latent variables that jointly encode part geometry and spatial relationships. DSO (Li et al., 2025) instead uses simulator-derived feedback to align a pretrained 3D generator toward physically stable outputs, operating at the generator-distribution level rather than through test-time redesign of individual instances.

In PhysiOpt, optimization is performed over latent variables or features of the underlying generative model, so structural changes remain mediated through these variables rather than directly specified through spatial design variables. In SATO, the latent variables capture part geometry and spatial relationships, but are not directly interpretable as volumetric design variables. DSO operates at a different level: its simulator-guided alignment improves the likelihood of physically stable generations but does not expose instance-specific spatial design variables for directly reconfiguring an object's structure. Across these approaches, physical improvement is mediated through learned features, latent variables, or model-level alignment rather than spatial design variables.

2.2 Intermediate Spatial Representations in 3D Generative Models

Several recent 3D generative models expose spatially structured representations within their generative pipelines. TRELLIS (Xiang et al., 2025) first generates a sparse spatial structure whose active locations provide a coarse but explicit representation of the object's spatial occupancy, and subsequently predicts Structured LATent (SLAT) features for detailed geometry and appearance. TRELLIS.2 (Xiang et al., 2026) similarly employs a sparse O-Voxel representation that explicitly organizes geometry and material information in 3D space, together with compact structured features for synthesis. OctFusion (Xiong et al., 2025) uses an octree-structured latent representation that combines learned features with an explicit hierarchical spatial structure across multiple scales.

These representations were introduced primarily to improve scalability and geometric fidelity in 3D generation, but their explicit spatial structure also makes them potential interfaces for structural modification. We instead explore this role for physics-driven optimization. Among possible intermediate representations, § 3 formalizes the properties required to serve as such an interface and contrasts this choice with learned-feature and explicit grid design spaces.

Table 1: Comparison of design-space properties relevant to physics-driven 3D redesign.

Design space	Spatial control	Direct volume control	Generative prior
Learned feature	Indirect	Indirect	✓
External grid	Explicit	Explicit	–
Intermediate spatial	**Explicit**	**Explicit**	✓

3 Limitations of Existing Design Spaces and Our Motivation

Given an initial asset $\mathcal{X}_0$ produced by a pretrained 3D generative model, we seek a redesigned asset $\mathcal{X}^*$ that improves a task-dependent physical objective under a prescribed material budget:

$$\min_{\mathcal{X}} \ \mathcal{J}_{\text{phys}}(\mathcal{X}) \qquad \text{s.t.} \qquad V(\mathcal{X}) = V_{\text{target}}. \tag{1}$$

The key question is where such optimization should act inside a 3D generator: the design space should expose explicit structural degrees of freedom while retaining the learned geometric prior.

3.1 A Trade-off Between Structural Control and Generative Prior

Existing approaches primarily expose generated 3D objects to physical optimization through either learned representations or explicit geometric domains.

Learned latent-space optimization. As discussed in § 2.1, prior methods optimize latent or feature variables while retaining generative priors. However, encoding spatial information does not imply exposing spatially indexed design variables: individual latent or feature coordinates generally do not correspond to specific 3D regions. As a result, geometric and structural changes remain mediated through these variables rather than being directly controlled in space.

Explicit grid optimization. As discussed in § 2.1, grid optimization provides direct spatial control over the density field. However, it operates outside the generator's native representation, so geometric details lost during optimization may not be recovered by subsequent generative reconstruction.

These two design spaces therefore expose a trade-off between explicit structural control and generative compatibility, as summarized in Table 1.

3.2 Motivating an Intermediate Spatial Interface

The above trade-off motivates a design interface that enables explicit spatial control while remaining compatible with generative synthesis. Consider a pretrained 3D generator decomposed at an intermediate representation $\mathcal{M}$ as

$$\mathcal{M} = G_{\text{pre}}(c), \qquad \mathcal{X} = G_{\text{post}}(\mathcal{M}), \tag{2}$$

where c denotes the conditioning input, G_{pre} denotes the upstream generative stages producing $\mathcal{M}$, and G_{post} denotes the remaining synthesis stages producing the output geometry $\mathcal{X}$. A pretrained generator may expose multiple such intermediate representations along its synthesis path, but not all are suitable as design interfaces for physical redesign.

We require two properties: ***(1) Explicit spatial controllability:*** the representation permits direct geometric and structural modification; and ***(2) Generative compatibility:*** the modified representation remains compatible with the frozen downstream stages of the pretrained generator.

Among these intermediate representations, we select a spatial representation satisfying both properties and denote its initial state by $\mathcal{M}_0$. We use $\mathcal{M}_0$ as the design interface for physics-driven redesign. The next section describes how we turn this interface into a physics-aware design space.

4 Physics-Driven Structural 3D Redesign

We introduce **Int3Rcept**, a physics-driven 3D redesign framework that optimizes an intermediate spatial representation of a pretrained 3D generator. Figure 2 illustrates the overall pipeline. Given a physical objective $\mathcal{J}_{\text{phys}}$ and task setup $\mathcal{S}$, we summarize the redesign as

$$\mathcal{M}_0 \xrightarrow{\mathcal{J}_{\text{phys}},\,\mathcal{S}} \mathcal{M}^*, \qquad \mathcal{X}^* = G_{\text{post}}(\mathcal{M}^*). \tag{3}$$

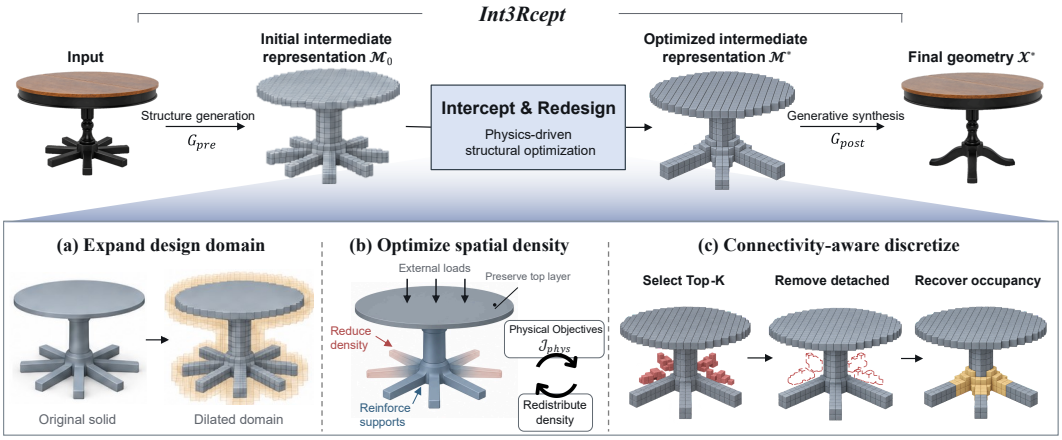

Figure 2: **Overview of Int3Rcept.** Int3Rcept intercepts an intermediate spatial representation of a pretrained 3D generator for physics-driven structural redesign, then resumes downstream generative synthesis from the optimized representation. The redesign consists of (a) design-domain expansion, (b) density optimization, and (c) connectivity-aware discretization.

4.1 Physical Objectives and Task Setup

We instantiate $\mathcal{J}_{\text{phys}}$ for load-bearing, floating bodies, and thermal tasks: compliance minimization, equilibrium draft minimization, and thermal resistance maximization. Throughout this section, $\boldsymbol{\rho}$ denotes the continuous density field representing spatial material distribution for physical evaluation, while $\mathcal{S}$ specifies semantic regions and physical boundary conditions for each task.

Compliance minimization. For load-bearing structures, we minimize structural compliance C, which measures deformation under the prescribed loading:

$$\mathbf{K}(\boldsymbol{\rho})\mathbf{u} = \mathbf{f}(\boldsymbol{\rho}), \qquad \mathcal{J}_{\text{phys}} = C(\boldsymbol{\rho}) = \mathbf{f}(\boldsymbol{\rho})^{\top}\mathbf{u}, \tag{4}$$

where $\mathbf{K}$ is the global stiffness matrix, $\mathbf{u}$ the displacement field, and $\mathbf{f}$ the total nodal force including self-weight and external loading.

Equilibrium-draft minimization. For floating bodies, the equilibrium draft d^*—the submerged depth below the waterline—is determined by the balance between weight and buoyancy:

$$\rho_{\text{w}} g V_{\text{sub}}(d^*, \boldsymbol{\rho}) = W(\boldsymbol{\rho}), \qquad \mathcal{J}_{\text{phys}} = d^*, \tag{5}$$

where ρ_{w} is the water density, V_{sub} the submerged volume, and W the object weight.

Thermal-resistance maximization. For thermal redesign, we solve steady-state heat conduction and maximize the thermal resistance R_{th} between the internal heat source and cooling region:

$$\mathbf{K}_{\text{th}}(\boldsymbol{\rho})\boldsymbol{\theta} = \mathbf{q}, \qquad R_{\text{th}} = \frac{\overline{T}_{\text{heat}} - T_{\text{cold}}}{P}, \qquad \mathcal{J}_{\text{phys}} = -R_{\text{th}}, \tag{6}$$

where $\mathbf{K}_{\text{th}}$ is the thermal conduction matrix, $\boldsymbol{\theta} = T - T_{\text{cold}}$ denotes the temperature rise relative to the fixed cooling temperature T_{cold}, $\mathbf{q}$ the heat-source vector, P the total heat generation, and $\overline{T}_{\text{heat}}$ the heat-generation-weighted mean temperature of the source region.

Detailed settings related to the physical objectives are provided in Appendix A.1.

Semantic control for region-aware redesign. We use semantic information to specify where redesign is allowed and how physical conditions apply, together defining the task setup $\mathcal{S}$. Preserve regions remain fixed while optimize regions are available for redesign; compliance additionally uses load and support regions, while thermal redesign uses heat-source and cooling regions. These regions are mapped to the intermediate spatial representation before constructing the design domain. Region construction and object-specific assignments are provided in § 5.1 and Appendix A.1.

4.2 Intermediate-Space Redesign and Synthesis

Given the task objective $\mathcal{J}_{\text{phys}}$ and setup $\mathcal{S}$ defined in § 4.1, Int3Rcept constructs a spatial design domain from the intermediate representation, optimizes its density distribution under the prescribed material budget, and converts the result back into a generator-compatible representation.

Design-domain expansion. Because the intermediate representation $\mathcal{M}_0$ consists of sparse spatial elements around the object surface rather than a filled solid volume, we first fill its interior and then expand the initial structure to permit material addition beyond the original geometry:

$$S_0 = \text{Fill}(\mathcal{M}_0), \qquad \Omega = \text{Dilate}(S_0, d), \tag{7}$$

where d controls the dilation extent in spatial elements. The expanded domain Ω defines the spatial extent within which redesign can occur.

Physics-driven density optimization. We exclude the preserve and task-specific fixed regions defined in § 4.1 from Ω, and denote the remaining editable elements by $\mathcal{A} \subseteq \Omega$. Each element $e \in \Omega$ is assigned a continuous design variable $x_e \in [\rho_{\min}, 1]$, collected in $\mathbf{x}$ and initialized from the occupancy of S_0; only variables in $\mathcal{A}$ are optimized, while the others remain fixed. We obtain the physical density field $\boldsymbol{\rho}$ by applying spatial filtering and a differentiable projection to $\mathbf{x}$:

$$\boldsymbol{\rho} = \Phi_{\beta,\eta}(\mathbf{x}), \tag{8}$$

where ρ_e denotes the resulting density of element e, β controls projection sharpness, and η the threshold. We then optimize

$$\mathbf{x}^*_{\mathcal{A}} = \arg\min_{\mathbf{x}_{\mathcal{A}}} \mathcal{J}_{\text{phys}}(\boldsymbol{\rho}; \mathcal{S}) \qquad \text{s.t.} \qquad V(\boldsymbol{\rho}) = \sum_{e\in\Omega} \rho_e = V_{\text{target}} = \alpha|S_0|, \tag{9}$$

where $|S_0|$ denotes the size of the initial solid and α specifies the target volume ratio. To favor changes that provide greater physical improvement per unit volume, we define

$$q_e = -\frac{\partial \mathcal{J}_{\text{phys}}/\partial x_e}{\max(\partial V/\partial x_e, \epsilon)}, \tag{10}$$

where $\epsilon > 0$ ensures numerical stability. The score q_e measures the physical benefit of changing element e relative to its local volume cost, guiding density toward more physically beneficial regions.

Connectivity-aware discretization. The optimized density field is continuous, whereas the downstream generator expects a discrete spatial structure. As illustrated in Fig. 2(c), we first select the highest-density elements according to the target volume. Physically less beneficial regions gradually lose density during optimization and may become detached after discretization. We remove components that violate task-dependent connectivity and recover the resulting volume deficit by progressively adding high-density admissible neighbors while preserving connectivity. The resulting discrete structure defines the optimized intermediate representation $\mathcal{M}^*$.

Generative synthesis. Finally, $\mathcal{M}^*$ is passed to the frozen downstream generator G_{post} to produce the geometry $\mathcal{X}^*$. Because optimization remains within the generator's native representation, the redesigned structure retains access to downstream synthesis for recovering detailed geometry.

Additional optimization, design-domain and discretization details are provided in Appendix A.

5 Experiments and Results

Fig. 3 presents qualitative examples of Int3Rcept, showing diverse redesigns including support reinforcement and removal, opening creation, and task-dependent reshaping while retaining high visual quality. We evaluate how these capabilities relate to the proposed choice of design space from three complementary perspectives: (i) ***Performance–Quality Trade-off***, comparing physical performance and visual quality against learned-feature and explicit-grid optimization; (ii) ***Versatility***, evaluating the same interface across different physical objectives and pretrained 3D generators; and (iii) ***Controllability***, evaluating how structural redesign adapts to different material budgets.



Figure 3: **Diverse physics-driven 3D redesigns with Int3Rcept.** (a)-1 **Compliance redesign at comparable volume.** Int3Rcept reallocates material through support reinforcement, seat/top thinning, and removal of low-benefit regions. (a)-2 **Compliance redesign under a 70% volume budget.** More substantial structural changes emerge, including material reallocation, support pruning, and the creation of openings or passages. (b) **Draft-aware redesign.** Floating bodies are reshaped to reduce equilibrium draft. (c) **Thermal redesign.** Heat-transfer paths are reconfigured to increase thermal resistance between internal heat sources and designated handle regions.

5.1 Experimental Setup

Data and models. We use TRELLIS (Xiang et al., 2025) as the primary 3D generator. Structural redesign uses segmented PartNet (Mo et al., 2019) renders as semantic cues, while floating-body and thermal redesign use FLUX.1-schnell (Black Forest Labs, 2024) generated images with SAM (Kirillov et al., 2023) masks. For thermal redesign, we designate an internal body region as the heat source and a handle region as the cooling surface. OmniPart (Yang et al., 2025) provides 3D part segmentation across these settings, from which we extract segmented occupancy to define task-specific regions. The resulting regions are manually inspected and filtered when necessary.

Baselines. We compare against PhysiOpt (Zhan et al., 2025) as learned-feature optimization and a grid-based SIMP baseline at 64^3 and 128^3 resolutions, hereafter referred to as Grid SIMP. For a controlled comparison with PhysiOpt, both methods start from the same geometry and are evaluated on the same paired subset of valid objects; we adapt PhysiOpt's volume objective to match our target-volume settings. For Grid SIMP, the extracted grid surfaces are smoothed before evaluation. We additionally report Int3Rcept w/o G_{post}, which extracts the optimized intermediate structure without downstream generative synthesis. Additional baseline details are provided in Appendix B.1.

Metrics. Physical metrics are evaluated on the synthesized geometry using the corresponding task-specific solver. For visual performance, we report CLIP-based aesthetic quality (Hentschel et al., 2022) and LIQE (Zhang et al., 2023) over 4 rendered views of each asset to assess perceptual quality.

Additional implementation and evaluation details are provided in Appendices A and B.

Table 2: **Main quantitative comparison across optimization spaces and physical objectives. Bold** and underlined values denote the best and second-best results, respectively, excluding the ablation. PhysiOpt is evaluated only for compliance following its available implementation. †OctFusion is evaluated on 33 separately generated objects due to its different conditioning interface; draft and thermal are omitted because the required categories are unavailable.

	Compliance			Draft			Thermal		
Method	$C_1/C_0 \downarrow$	Aesth.↑	LIQE↑	$H_1/H_0 \downarrow$	Aesth.↑	LIQE↑	$R_1/R_0 \uparrow$	Aesth.↑	LIQE↑
Intermediate-space opt.									
Int3Rcept (TRELLIS)	**0.483**	5.02	**4.31**	**0.625**	**4.71**	**4.59**	**2.688**	5.05	**4.63**
Int3Rcept (TRELLIS.2)	0.555	5.00	4.26	0.705	4.64	4.51	1.761	**5.11**	4.25
Int3Rcept (OctFusion)†	0.560	4.78	3.91	–	–	–	–	–	–
Learned-feature opt.									
PhysiOpt (TRELLIS)	0.716	**5.05**	**4.31**	–	–	–	–	–	–
Explicit grid opt.									
Grid SIMP-128	0.542	4.73	3.26	0.667	4.50	3.70	2.002	4.61	3.02
Grid SIMP-64	0.578	4.73	3.29	0.634	4.39	3.43	2.097	4.65	3.03
Ablation									
Int3Rcept w/o G_{post}	0.564	4.78	3.34	0.688	4.46	3.41	2.142	4.64	3.03

5.2 Comparison Across Optimization Spaces

Greater redesign freedom than in learned feature space. Table 2 shows that Int3Rcept achieves stronger physical improvement than PhysiOpt. For compliance minimization, Int3Rcept with TRELLIS reaches a normalized compliance of 0.483, compared with 0.716 for PhysiOpt. As motivated in § 3.2, we attribute this improvement to ***explicit spatial controllability***: rather than mediating structural changes through learned features, Int3Rcept directly modifies the intermediate spatial representation, providing freedom to search for redesigns. § 5.3 further shows that Int3Rcept achieves lower compliance than PhysiOpt at comparable volumes while enabling broader structural variation.

Higher aesthetic quality than explicit grid. Explicit grid optimization provides similarly direct spatial control but operates outside the pretrained generator. Grid SIMP remains competitive in physical performance, yet its Aesthetic and LIQE scores are consistently lower than those of Int3Rcept, as shown in Table 2. As motivated in § 3.2, this gap highlights the benefit of ***generative compatibility***: Int3Rcept remains compatible with downstream synthesis, enabling recovery of detailed geometry. Figure 4 further shows that direct grid optimization can remove fine geometric details and introduce surface artifacts.

Runtime and memory comparisons at multiple grid resolutions are provided in Appendix B.8.

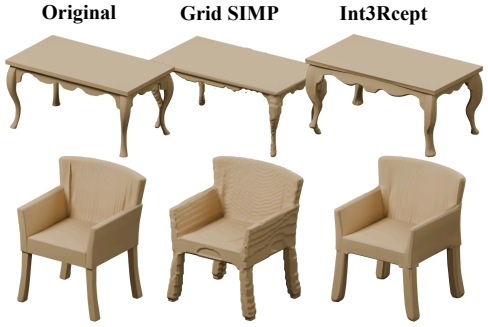

Figure 4: **Qualitative comparison with explicit grid optimization.** Int3Rcept better preserves detail and surface quality.

Versatility across physical objectives. Beyond compliance, Int3Rcept supports draft and thermal redesign, which are not supported by the available PhysiOpt implementation. Table 2 shows comparable draft performance to Grid SIMP with higher visual quality, while thermal redesign improves resistance over the initial design across generators despite generator-dependent variation after synthesis. Additional results are provided in Appendices B.5 and B.7.

Compatibility across generators. Table 2 shows that Int3Rcept improves compliance across TRELLIS, TRELLIS.2, and OctFusion. Additional qualitative results in Appendix B.4 show consistent structural redesigns across generators, supporting spatial intermediate representations as a reusable interface for physics optimization.

Physical consistency after synthesis. We verify that physical improvements in the intermediate representation persist after mesh synthesis. Mesh-space compliance follows the same decreasing trend as intermediate-space compliance despite differences from voxel discretization and mesh geometry; see Appendix B.3 for details.

5.3 Controllability Across Volume Settings

A common requirement in physics-driven design is to adapt structures under prescribed volume budgets. Int3Rcept's explicit spatial control allows the design to respond directly to changes in this constraint, which we evaluate by varying the target volume ratio α while holding all other settings fixed. Figure 5(a) reports the mean realized volume response across evaluated objects over the target volume settings. Because downstream synthesis can introduce a method-dependent but systematic offset between the optimized intermediate volume and the final mesh volume, we normalize each method by its output at $\alpha = 1.0$ to isolate responsiveness to changes in the target volume rather than absolute calibration. Int3Rcept exhibits a broad and monotonic response, ranging from 0.741 at $\alpha = 0.8$ to 1.357 at $\alpha = 1.3$, whereas PhysiOpt changes only from 0.969 to 1.133 over the same target-volume range. This compressed response is consistent with volume changes being mediated through learned features rather than directly controlled in the spatial design domain.

Figure 5(b) presents a representative example from the evaluation set, showing that Int3Rcept spans a broader region of the compliance–volume plane and achieves lower compliance than PhysiOpt at comparable realized volumes. Figure 5(c) visualizes the corresponding redesigns for the same example, where varying the target volume produces distinct structural changes including support removal and reinforcement.

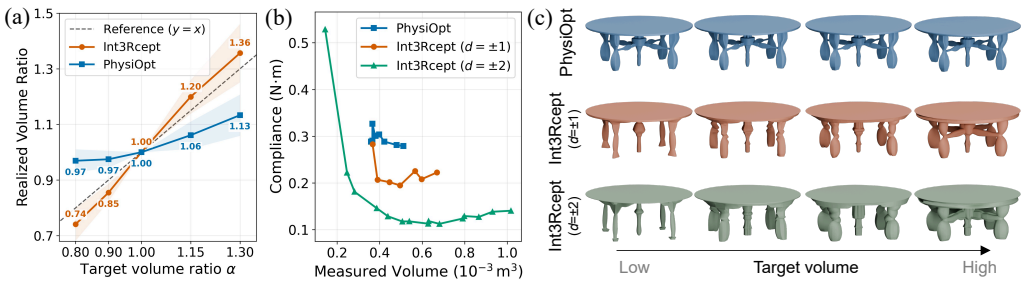

Figure 5: **Controllability across target volume settings.** (a) Realized volume response to varying target volume ratios across the evaluation set. (b) Compliance–volume sweep for a representative example. (c) Corresponding structural redesigns for the same example across target volumes.

6 Limitations and Conclusion

Int3Rcept combines explicit structural control with downstream generative synthesis through an intermediate spatial representation. This design enables substantial structural changes under physical objectives, but also introduces several limitations. Because physical constraints are enforced before synthesis, the final mesh volume may deviate from the optimized intermediate volume. In our TRELLIS implementation, the resulting volume shift is consistently around 7% across settings. In addition, large deviations of the optimized intermediate spatial representation from the original structure may make it harder for downstream synthesis to preserve object identity. Visual quality also remains generator-dependent, with OctFusion generally producing lower Aesthetic scores than TRELLIS. Future work could address these limitations by improving consistency between intermediate optimization and final synthesis, while broadening the range of achievable redesigns.

Overall, Int3Rcept provides a practical interface between explicit physical optimization and learned 3D generative priors, achieving a favorable ***performance–quality trade-off*** while preserving high visual quality under substantial structural changes. Its ***versatility*** enables the same interface to support multiple physical objectives and pretrained 3D generators, while its ***controllability*** allows redesign behavior to adapt systematically to different material budgets.

AI USE STATEMENT

In this work, we used generative AI tools to assist with translation, drafting and editing parts of the manuscript, implementation of the proposed methods, and data preprocessing, including data cleaning and reformatting. We did not use generative AI tools for other research tasks listed in the ICLR AI disclosure guidelines. All AI-assisted outputs were reviewed and, where applicable, modified and verified by the authors. In particular, AI-assisted implementation was inspected and tested for correctness, and AI-assisted data preprocessing was reviewed to ensure that it did not alter the intended content or properties of the data. AI-assisted translation, drafting, and editing were used to improve the presentation of the manuscript, and the authors verified the resulting text for accuracy and consistency with the underlying research. We take full responsibility for the final content of this work, including all text, claims, code, data processing, and other artifacts produced with the aid of generative AI.

REFERENCES

Martin Philip Bendsøe. Optimal shape design as a material distribution problem. *Structural Optimization*, 1:193–202, 1989. doi: 10.1007/BF01650949.

Martin Philip Bendsøe and Ole Sigmund. *Topology Optimization: Theory, Methods, and Applications*. Springer, 2003. doi: 10.1007/978-3-662-05086-6.

Black Forest Labs. FLUX. `https://github.com/black-forest-labs/flux`, 2024. Official inference repository for FLUX.1 models.

Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3d: A generative model of high quality 3d textured shapes learned from images. In *Advances in Neural Information Processing Systems*, volume 35, pp. 31841–31854, 2022.

Minghao Guo, Bohan Wang, Pingchuan Ma, Tianyuan Zhang, Crystal Elaine Owens, Chuang Gan, Joshua B. Tenenbaum, Kaiming He, and Wojciech Matusik. Physically compatible 3d object modeling from a single image. In *Advances in Neural Information Processing Systems*, 2024. URL `https://arxiv.org/abs/2405.20510`.

Xianglong He, Zi-Xin Zou, Chia-Hao Chen, Yuan-Chen Guo, Ding Liang, Chun Yuan, Wanli Ouyang, Yan-Pei Cao, and Yangguang Li. Sparseflex: High-resolution and arbitrary-topology 3d shape modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 14822–14833, 2025.

Simon Hentschel, Konstantin Kobs, and Andreas Hotho. Clip knows image aesthetics. *Frontiers in Artificial Intelligence*, 5:976235, 2022. doi: 10.3389/frai.2022.976235.

Yixin Hu, Teseo Schneider, Bolun Wang, Denis Zorin, and Daniele Panozzo. Fast tetrahedral meshing in the wild. *ACM Transactions on Graphics*, 39(4), July 2020. doi: 10.1145/3386569.3392385.

Zizhou Huang, Davi Colli Tozoni, Arvi Gjoka, Zachary Ferguson, Teseo Schneider, Daniele Panozzo, and Denis Zorin. Differentiable solver for time-dependent deformation problems with contact. *ACM Transactions on Graphics*, 43(3), 2024. doi: 10.1145/3657648.

Haroon Ijaz, Xuwei Wang, Wei Chen, Hai Lin, and Ming Li. Physically reliable 3d styled shape generation via structure-aware topology optimization in unified latent space. *Computer-Aided Design*, 183:103864, 2025. doi: 10.1016/j.cad.2025.103864.

Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4015–4026, 2023.

Ruining Li, Chuanxia Zheng, Christian Rupprecht, and Andrea Vedaldi. DSO: Aligning 3d generators with simulation feedback for physical soundness. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025. URL https://arxiv.org/abs/2503.22677.

Kaichun Mo, Shilin Zhu, Angel X. Chang, Li Yi, Subarna Tripathi, Leonidas J. Guibas, and Hao Su. Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 909–918, 2019.

Xuanchi Ren, Jiahui Huang, Xiaohui Zeng, Ken Museth, Sanja Fidler, and Francis Williams. Xcube: Large-scale 3d generative modeling using sparse voxel hierarchies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4209–4219, 2024.

Ole Sigmund. A 99 line topology optimization code written in MATLAB. *Structural and Multidisciplinary Optimization*, 21(2):120–127, 2001. doi: 10.1007/s001580050176.

Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21469–21480, 2025. doi: 10.1109/CVPR52734.2025.02000.

Jianfeng Xiang, Xiaoxue Chen, Sicheng Xu, Ruicheng Wang, Zelong Lv, Yu Deng, Hongyuan Zhu, Yue Dong, Hao Zhao, Nicholas Jing Yuan, and Jiaolong Yang. Native and compact structured latents for 3d generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14419–14429, 2026.

Bojun Xiong, Si-Tong Wei, Xin-Yang Zheng, Yan-Pei Cao, Zhouhui Lian, and Peng-Shuai Wang. Octfusion: Octree-based diffusion models for 3d shape generation. *Computer Graphics Forum*, 2025.

Yunhan Yang, Yufan Zhou, Yuan-Chen Guo, Zi-Xin Zou, Yukun Huang, Ying-Tian Liu, Hao Xu, Ding Liang, Yan-Pei Cao, and Xihui Liu. OmniPart: Part-aware 3D generation with semantic decoupling and structural cohesion. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers*, SA Conference Papers '25, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400721373. doi: 10.1145/3757377.3763872. URL https://doi.org/10.1145/3757377.3763872.

Xiao Zhan, Clément Jambon, Evan Thompson, Kenney Ng, and Mina Konaković Luković. PhysiOpt: Physics-driven shape optimization for 3d generative models. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers*. Association for Computing Machinery, 2025. doi: 10.1145/3757377.3763884. URL https://doi.org/10.1145/3757377.3763884.

Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan Xu, and Jingyi Yu. Clay: A controllable large-scale generative model for creating high-quality 3d assets. *ACM Transactions on Graphics*, 43(4):120:1–120:20, 2024. doi: 10.1145/3658146.

Weixia Zhang, Guangtao Zhai, Ying Wei, Xiaokang Yang, and Kede Ma. Blind image quality assessment via vision-language correspondence: A multitask learning perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14071–14081, 2023.

A Additional Implementation Details

A.1 Physical Objective Settings

Compliance. For load-bearing structures, we use Solid Isotropic Material with Penalization (SIMP) coupled with finite element analysis (FEM). The element stiffness is interpolated as

$$\mathbf{K}_e(\rho_e) = \rho_e^p \mathbf{K}_e^0, \tag{11}$$

where $\mathbf{K}_e^0$ is the stiffness matrix of a fully solid element and p is the SIMP penalization exponent. Static equilibrium is given by

$$\mathbf{K}(\boldsymbol{\rho})\mathbf{u} = \mathbf{f}(\boldsymbol{\rho}), \tag{12}$$

and compliance is evaluated as

$$C(\boldsymbol{\rho}) = \mathbf{f}(\boldsymbol{\rho})^\top \mathbf{u}. \tag{13}$$

The force includes density-dependent self-weight and a prescribed external load. For the main TRELLIS experiments, we interpret the normalized domain as a $1\,\mathrm{m}$ cube and use Young's modulus $E = 10^9\,\mathrm{Pa}$, Poisson ratio $\nu = 0.3$, material density $650\,\mathrm{kg/m^3}$, gravitational acceleration $g = 9.81\,\mathrm{m/s^2}$, and $p = 3$. A total downward load of $800\,\mathrm{N}$ is distributed uniformly over the designated load-region nodes.

Equilibrium draft. For floating-body redesign, we estimate the equilibrium waterline using a differentiable approximation of submerged volume. For a candidate waterline height h, the submerged fraction is approximated as

$$R(h, \boldsymbol{\rho}) = \frac{\sum_i \rho_i \sigma((h - z_i)/T)}{\sum_i \rho_i}, \tag{14}$$

where z_i is the vertical position of element i, σ is a smooth approximation of the below-water indicator, and T controls its smoothness. The equilibrium waterline is determined by balancing weight and buoyancy, and the resulting draft is minimized. We set the object density to 0.3 times the water density.

Thermal resistance. For thermal redesign, we solve steady-state heat conduction,

$$\mathbf{K}_{\mathrm{th}}(\boldsymbol{\rho})\boldsymbol{\theta} = \mathbf{q}, \qquad \boldsymbol{\theta} = T - T_{\mathrm{cold}}, \tag{15}$$

and maximize the thermal resistance between the internal heat source and the handle cooling region,

$$R_{\mathrm{th}} = \frac{\overline{T}_{\mathrm{heat}} - T_{\mathrm{cold}}}{P} = \frac{\mathbf{q}^\top \boldsymbol{\theta}}{P^2}. \tag{16}$$

We distribute a total heat generation of $P = 1\,\mathrm{W}$ uniformly over the designated internal heat-source region and fix the cooling region at $25^\circ\mathrm{C}$. All remaining exterior surfaces are insulated. The solid material is modeled as homogeneous and isotropic with $k_{\mathrm{solid}} = 200\,\mathrm{W/(m\,K)}$, with density-dependent conductivity

$$k(\rho) = k_{\min} + (k_{\mathrm{solid}} - k_{\min})\rho^3. \tag{17}$$

Objects are scaled to a common initial height of $10\,\mathrm{cm}$ for thermal evaluation.

A.2 Optimization Details

Table 3: Optimization settings for the main TRELLIS experiments.

Setting	Compliance	Draft	Thermal
Resolution	64^3	64^3	64^3
$\rho_{\min}$	10^{-3}	10^{-3}	10^{-3}
Filter radius	2.5 voxels	2.5 voxels	2.5 voxels
Projection β	$1 \to 32$	$1 \to 32$	$1 \to 16$
Iterations	60	60	300
Optimizer	Projected gradient	Projected gradient	Projected gradient
Move limit	0.2	0.12	0.15
Default α	0.7	1.0	1.0
Default dilation d	1	2	$2/1^\dagger$

†For thermal optimization, the editable band extends two voxels inward and outward on the body and one voxel inward and outward on the main handles. Preserved regions remain fixed.

Density parameterization. For the main TRELLIS experiments, optimization is performed on a 64^3 intermediate spatial representation. Each raw design field $\mathbf{x}$ is spatially filtered and differentiably projected to obtain the physical density field $\boldsymbol{\rho}$ used by the corresponding physical solver. We

use $\rho_{\min} = 10^{-3}$ and a filter radius of 2.5 voxels. The projection sharpness β is progressively increased from 1 to 32, while the projection threshold η is adjusted to maintain the prescribed volume constraint.

Sensitivity-based update. We use the sensitivity score q_e defined in Eq. (10) to update the editable design variables. A move-limited projected-gradient update redistributes density toward elements with larger physical benefit while enforcing the prescribed volume constraint. The projection sharpness is progressively increased during optimization; task-specific move limits and continuation settings are summarized below.

Task-specific settings. Table 3 summarizes the optimization settings used for the main TRELLIS compliance, draft and thermal experiments.

A.3 Semantic Region Assignments

We define object-specific semantic regions according to the task. For chairs, semantically distinctive parts such as armrests are typically preserved while the supporting structure is optimized. For tables, we preserve the top surface layer, apply external loads to the designated load region, and optimize the supporting structure below. For floating bodies, upper structures such as sails and deck superstructures are preserved while the hull is optimized. For thermal redesign, an internal body region is designated as the heat source and the handle as the cooling region, while semantically distinctive parts outside the target redesign region are preserved when necessary. All semantic assignments are manually inspected and refined before optimization.

A.4 Design-Domain Sensitivity

We examine the effect of the dilation extent d, which controls how far the design domain expands beyond the initial structure. As shown in Fig. 6, increasing d provides greater spatial freedom for material redistribution. Under the same volume budget, the optimized structure can progressively move material farther from its initial configuration, as illustrated by the downward redistribution of the seat structure. This highlights the trade-off between preserving the initial geometry and allowing broader structural redesign.

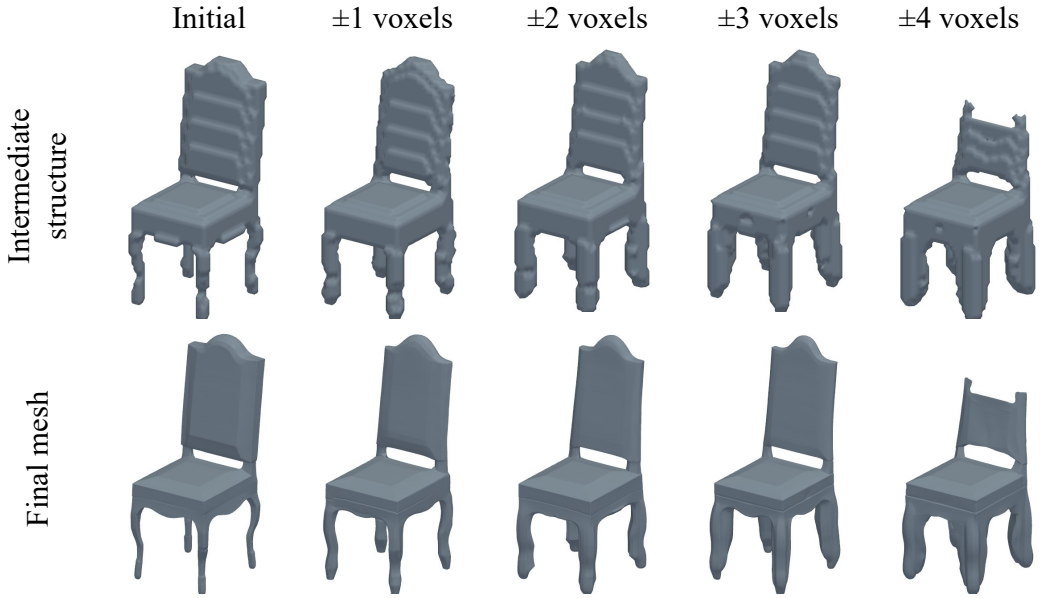

Figure 6: **Effect of design-domain expansion.** Increasing the dilation extent enlarges the accessible redesign space, enabling broader material redistribution under a fixed volume budget.

A.5 Connectivity-Aware Discretization

Algorithm 1 summarizes the discretization procedure used to convert the optimized continuous density field into a task-valid intermediate representation. Naive top-K selection can preserve the target volume while producing detached components. Retaining only the valid connected component restores connectivity but may reduce the realized volume. Our full correction therefore recovers the resulting volume deficit while preserving task-dependent connectivity. As shown in Fig. 7, our full correction recovers this volume deficit while preserving task-dependent connectivity.

Algorithm 1 Connectivity-Aware Discretization

Require: Optimized density field $\boldsymbol{\rho}^*$, editable region $\mathcal{A}$, target volume V_{target}, task setup $\mathcal{S}$

1: Select the highest-density elements in $\mathcal{A}$ according to the target volume.
2: Identify connected components and remove those that violate the task-specific connectivity criterion defined by $\mathcal{S}$.
3: Recover the resulting volume deficit by progressively adding high-density admissible neighbors while preserving task validity.
4: Combine the redesigned and fixed regions to obtain the discrete optimized structure.
5: Map the resulting structure to the optimized intermediate representation $\mathcal{M}^*$ for downstream generative synthesis.
6: **return** $\mathcal{M}^*$



Figure 7: **Connectivity-aware discretization.** Top-K selection can create detached components; component filtering restores connectivity, and occupancy recovery restores the target volume.

B Evaluation Details

B.1 PhysiOpt Comparison

Shared initialization. For a controlled comparison, Int3Rcept and PhysiOpt start from the same initial geometry. We use the solidified spatial coordinates employed by Int3Rcept and sample the corresponding TRELLIS SLat features once from the input image. Decoding this shared representation yields the common baseline mesh used for both methods. PhysiOpt optimizes the learned features associated with this representation, while Int3Rcept performs density-based redesign on the corresponding spatial structure.

PhysiOpt setup. PhysiOpt is optimized for 60 steps from the shared initialization with a target volume ratio of 1.0. We use a learning rate of 0.01, a compliance weight of 2.0, and automatically balanced volume regularization. Inpainting is applied periodically during optimization and once at the final step.

Mesh processing and FEM validity. Final PhysiOpt meshes are processed using the TRELLIS cavity-filling procedure. Outputs from both methods are tetrahedralized using fTetWild (Hu et al., 2020) and evaluated under the same FEM parameters, loads, and supports. After tetrahedralization,

we retain the largest connected material component and exclude a case if this removes more than 1% of the tetrahedralized material volume. Three PhysiOpt outputs fail this criterion, while all Int3Rcept outputs remain valid, resulting in a paired evaluation set of 61 objects.

Table 4: Cases excluded from the paired PhysiOpt comparison by the volumetric FEM validity criterion.

Object	Removed material volume
Chair/2894	47.8%
Lamp/14565	2.38%
Lamp/14615	48.6%

B.2 Volume-Sweep Evaluation

We vary the target volume ratio from $\alpha = 0.8$ to 1.3. For PhysiOpt, we adapt its original volume-preservation objective to target the same prescribed volume ratios as Int3Rcept. Because fixed semantic regions constrain the minimum achievable volume and the finite design domain limits the maximum, not every object can realize the full sweep. We therefore retain only objects for which all target settings are feasible and both methods produce valid outputs throughout the sweep, yielding a fixed paired subset of 24 objects. In the main paper, final mesh volumes are normalized by each method's output at $\alpha = 1.0$ to isolate the relative response to changes in the target volume.

B.3 Intermediate-to-Final Physical Correspondence

We evaluate whether physical improvements obtained in the intermediate representation are preserved after downstream generative synthesis. At each optimization checkpoint, compliance is measured on the intermediate density representation and independently evaluated on the corresponding synthesized and tetrahedralized mesh. Each trajectory is normalized by its own initial compliance before aggregation to account for representation-dependent differences in absolute values.

We use a fixed cohort with valid FEM evaluations at every checkpoint. Among the 64 PartNet objects, one chair is excluded because its intermediate meshes fail the tetrahedralization validity criterion at intermediate checkpoints, yielding 63 objects for the trajectory analysis. This exclusion applies only to the trajectory evaluation; its final synthesized mesh remains valid.

Figure 8 shows the mean relative compliance trajectories. The synthesized-mesh trajectory follows the same decreasing trend as the intermediate-space trajectory throughout optimization, indicating that the physical improvements obtained during density optimization persist after generative synthesis. Absolute differences arise from discretization and geometric differences between the intermediate representation and synthesized mesh.

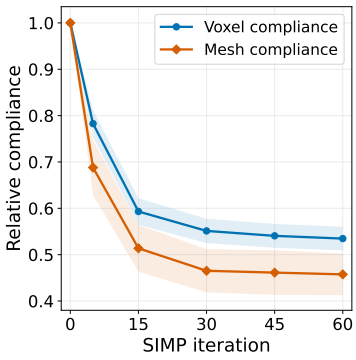

Figure 8: **Physical correspondence after generative synthesis.** Mean relative compliance trajectories over 63 objects; shaded regions indicate 95% confidence intervals.

B.4 Cross-Generator Evaluation

We evaluate whether the intermediate-space redesign interface transfers beyond the primary TRELLIS model by applying Int3Rcept to TRELLIS.2 and OctFusion. The same physics-driven redesign pipeline is used once a compatible intermediate spatial representation is exposed, while generator-specific synthesis stages remain frozen.

TRELLIS.2 is evaluated using the same physical redesign formulation as the main TRELLIS experiments. OctFusion is evaluated separately on 33 generated objects because its conditioning interface does not support the same image-conditioned evaluation protocol used for TRELLIS. We therefore use independently generated OctFusion objects and evaluate compliance redesign on this set. Draft and thermal evaluations are omitted for OctFusion because the required floating-body and cookware categories are unavailable in the evaluated generation setting.

These experiments assess transfer of the design interface across pretrained 3D generators rather than direct paired comparison between generators. Quantitative results are reported in Table 2, while Fig. 9 shows representative qualitative redesigns.

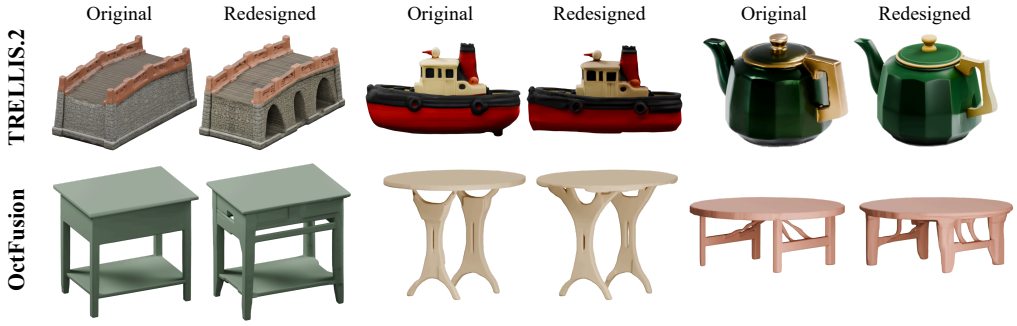

Figure 9: **Cross-generator redesigns with Int3Rcept** using TRELLIS.2 and OctFusion.

B.5 Draft Evaluation

For floating-body redesign, we evaluate the equilibrium draft on the final synthesized mesh under the same buoyancy conditions used during optimization. Draft is reported as the submerged depth below the equilibrium waterline, normalized by the mesh height. Table 5 reports per-example measurements before and after redesign.

Table 5: Per-example equilibrium-draft evaluation on six floating-body examples. Draft is reported as a percentage of mesh height at relative density 0.3.

Example	D_0	D_1	Reduction ↑
1	22.70	15.17	33.2%
2	24.80	14.55	41.3%
3	22.51	13.86	38.4%
4	25.53	14.56	43.0%
5	13.22	9.67	26.8%
6	10.59	6.13	42.1%
Mean	–	–	37.5%

B.6 Thermal Evaluation

For thermal redesign, we evaluate steady-state thermal resistance on the final synthesized mesh using the same physical conditions as during optimization. The designated internal source region receives

a total heat input of $1\,\mathrm{W}$, while the cooling region on the handle is fixed at $25^{\circ}\mathrm{C}$; all remaining exterior surfaces are insulated. The final mesh is evaluated using solid FEM with homogeneous isotropic conductivity $k = 200\,\mathrm{W}/(\mathrm{m\,K})$ after scaling the initial object height to $10\,\mathrm{cm}$. Higher thermal resistance indicates stronger conductive isolation between the heated body and the handle.

Table 6 reports the per-example thermal resistance before and after redesign. Int3Rcept increases thermal resistance across all nine examples, with an average ratio of $R_1/R_0 = 2.688$.

Table 6: Per-example thermal-resistance evaluation for Int3Rcept with TRELLIS. Higher values indicate better thermal isolation.

Example	R_0 (K/W)	R_1 (K/W)	$R_1/R_0 \uparrow$	Increase ↑
K2	1.0466	2.3123	2.209	120.9%
K3	1.2273	2.6082	2.125	112.5%
K4	1.0927	2.1003	1.922	92.2%
M2	0.9802	1.5388	1.570	57.0%
M3	0.8440	1.2106	1.434	43.4%
M4	0.8113	1.5131	1.865	86.5%
P2	0.4576	0.9255	2.022	102.2%
P3	0.6887	5.2609	7.638	663.8%
P4	0.5799	1.9771	3.410	241.0%
Mean	**0.8587**	**2.1607**	**2.688**	**168.8%**

B.7 Physical Field Visualizations

We provide additional qualitative visualizations of the physical fields before and after redesign. For both compliance and thermal evaluation, the final synthesized meshes are tetrahedralized with fTetWild (Hu et al., 2020) and evaluated using FEM. For compliance minimization, the redesigned structures exhibit reduced normalized displacement under the prescribed loading. For thermal-resistance maximization, the redesigned geometries alter the normalized temperature distribution to increase thermal isolation between the heated body and the designated handle region.

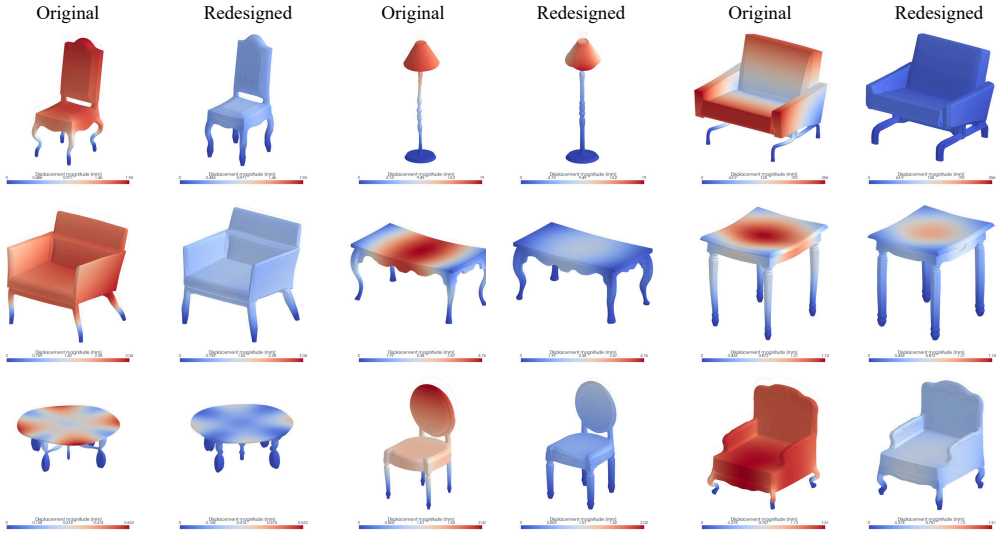

Figure 10: **Displacement fields before and after compliance redesign.** Redesigned structures exhibit lower normalized displacement under the prescribed loading.

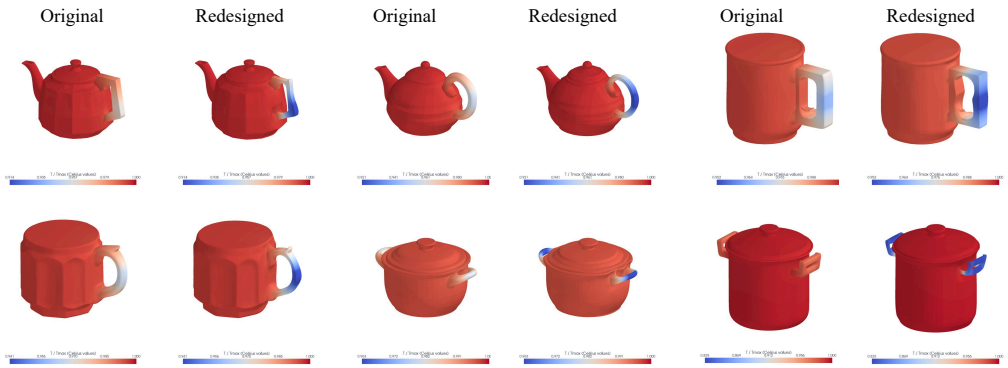

Figure 11: **Temperature fields before and after thermal redesign.** Redesigned geometries increase thermal isolation between the heated body and the designated handle region.

B.8 Computational Efficiency

We evaluate computational efficiency on 64 PartNet objects. Each method is executed three times under the same evaluation setting, and we report the average of the per-object median measurements. Runtime is measured as synchronized wall-clock time from prepared inputs to the exported mesh, excluding model loading and Python/CUDA initialization. Peak GPU and CPU memory are measured per process.

Table 7: Computational efficiency over 64 PartNet objects. Runtime and memory are averaged over per-object medians of three runs.

Method	Resolution	Runtime ↓ (s)	GPU Mem. ↓ (GiB)	CPU Mem. ↓ (GiB)
Grid SIMP-128	128^3	875.57	9.76	8.91
Grid SIMP-64	64^3	103.67	**1.69**	**1.97**
Int3Rcept + TRELLIS	64^3	**84.68**	15.12	6.71

Int3Rcept + TRELLIS achieves lower end-to-end runtime than Grid SIMP at both evaluated resolutions, with $10.34\times$ and $1.22\times$ speedups over Grid SIMP-128 and Grid SIMP-64, respectively. The pretrained generative model incurs higher GPU memory usage, revealing a runtime–memory trade-off.