Exploring The Trade-offs Between Strength And Sustainability For 3D Printing

Maxine Perroni-Scharf
MIT CSAIL
Cambridge, Massachusetts, USA
max1@mit.edu

Zhi Ray Wang
MIT CSAIL
Cambridge, Massachusetts, USA
zrwang@mit.edu

Mohammad Safazada
EPFL
Lausanne, Switzerland
safazada.mohammad@gmail.com

Muhammad Abdullah
Hasso Plattner Institute
Potsdam, Germany
muhammad.abdullah@hpi.de

Patrick Baudisch
Hasso Plattner Institute
Potsdam, Germany
patrick.baudisch@hpi.de

Stefanie Mueller
MIT CSAIL
Cambridge, Massachusetts, USA
stefmue@mit.edu

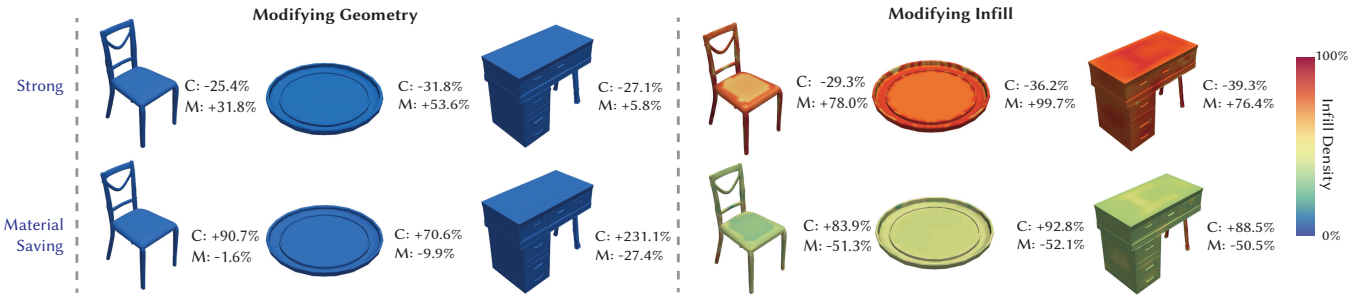

Figure 1: Two complementary controls for trading material usage against structural performance: modifying geometry (left) and modifying infill (right). We show examples of a strength-first setting (top) and a material-saving setting (bottom). C denotes the change in compliance (lower is stiffer) and M denotes the change in material usage, both relative to the baseline.

Abstract

The growing adoption of 3D printing presents significant environmental challenges, including material waste, high energy consumption, and the use of non-biodegradable polymers. To address this, we introduce a prototype tool for more sustainable fabrication. Our system features a user interface that reports key fabrication metrics, including material consumption and allows users to balance these against structural requirements such as mechanical strength. This is enabled by a multi-objective optimization framework driven by a differentiable simulator, which supports interactive refinement of designs toward improved resource efficiency. We evaluate the prototype in a simulation benchmark, demonstrating that it surfaces meaningful material–strength trade-offs across a range of functional shapes.

CCS Concepts

• **Applied computing** → **Engineering**; ***Computer-aided design***.

Keywords

Simulations, Optimization, Fused Deposition Modeling, 3D-Printing, Sustainability

ACM Reference Format:
Maxine Perroni-Scharf, Zhi Ray Wang, Mohammad Safazada, Muhammad Abdullah, Patrick Baudisch, and Stefanie Mueller. 2026. Exploring The Trade-offs Between Strength And Sustainability For 3D Printing. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3772363.3798396

1 Introduction

While 3D printing offers environmental benefits by enabling customized, on-demand fabrication, it also introduces unintended environmental consequences. For instance, an estimated 32% of all 3D printing material is discarded as waste [7], much of which is not biodegradable. To address this, the research community has proposed various solutions to make 3D printing processes more sustainable at different stages of the lifecycle.

These solutions, however, typically target a specific part of the sustainability life-cycle, ignoring the joint trade-offs of various competing objectives. For example, a tool that promotes eco-friendly materials might inadvertently increase print time or energy consumption. Crucially, sustainability goals must also be balanced against the object's functionality. We can save material by printing objects with less infill or thinner walls, but these changes could

CHI EA '26, Barcelona, Spain

ACM ISBN 979-8-4007-2281-3/26/04
https://doi.org/10.1145/3772363.3798396

compromise the structural integrity of the printed part. Effective integration of concurrent objectives calls for a holistic view of the print and simultaneous optimization across different sustainability metrics as well as performance metrics.

To this end, we present an interactive prototype system that helps users explore the trade-off between material usage and mechanical behavior. Given a text prompt and a simple load case, the system generates a small set of candidates spanning the trade-off space using a differentiable mechanics backend. Users can compare (1) geometry variants that trade stiffness for material usage and (2) graded infill profiles that redistribute material within a chosen geometry. The interface summarizes predicted material usage and mechanical metrics for each option, enabling users to select and export a configuration that matches their priorities.

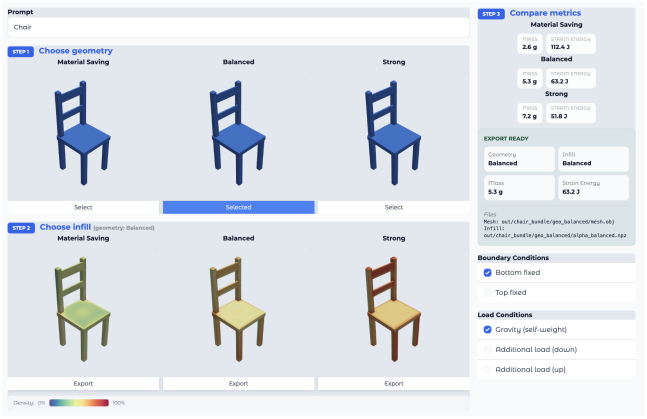

Figure 2: Prototype interface for exploring material–strength trade-offs. Users enter a prompt and specify a load case, then (Step 1) compare three geometry candidates, (Step 2) choose among three graded-infill profiles for the selected geometry, and (Step 3) review predicted metrics and export the chosen configuration.

2 Related work

2.1 Sustainability in fabrication

Prior HCI and fabrication research has examined sustainability across personal fabrication workflows, including reducing material use, energy consumption, and end-of-life waste. One common strategy is to reduce material during early prototyping [9]. Other systems focus on reuse [17]. Beyond reducing usage, researchers have explored alternative materials and end-of-life strategies [11, 19]. Other systems focus on deploying more sustainable filaments within functional prints [10].

In contrast to approaches that intervene primarily through low-fidelity previews, scrap insertion, or material substitution, our system focuses on design-time optimization. By jointly optimizing geometry and spatially varying infill, our approach reduces material usage while preserving structural performance, enabling users to explore these trade-offs during modeling rather than post hoc.

2.2 3D printing and its effect on the environment

Desktop FDM printing produces substantial waste due to iterative trial-and-error, failed prints, and support material. Empirical studies of fabrication workflows show that printing errors can lead to significant material and energy loss [14], and analyses of open workshop environments document high failure rates and recurring failure modes that compound waste across iterations [15]. Life-cycle assessment (LCA) studies further indicate that environmental impact is strongly correlated with part mass and print time, which determine both material usage and energy consumption [2, 5, 6].

These findings motivate tools that expose sustainability implications during design, when changes are inexpensive to make. Our system follows this direction by explicitly accounting for material usage alongside mechanical metrics, supporting informed trade-offs during design exploration.

2.3 Fabrication-aware design tools

A substantial body of HCI and graphics research integrates fabrication constraints and physical reasoning into interactive design tools. *Fab Forms* precomputes valid regions of parametric design spaces, enabling customization while maintaining manufacturability [12]. More recent work incorporates fabrication awareness into generative modeling: *Style2Fab* segments generated models based on functional requirements to support fabrication [4], and *MechStyle* integrates finite-element analysis into generative stylization to preserve structural integrity [3]. In computer graphics, structural optimization methods modify geometry or internal structure to satisfy mechanical constraints while reducing mass, including stress-based reinforcement [16], sparse internal frames [18], and optimized hollow or cellular interiors [8].

Our backend builds on recent differentiable, physics-driven optimization pipelines for generated geometry [21]. We extend this system by combining fabrication-aware structural optimization with sustainability-oriented objectives, including material usage. By co-optimizing geometry and graded infill and exposing the resulting trade-offs through an interactive interface, our approach supports designs that balance mechanical requirements with resource efficiency.

3 Method

Our method enables users to analyze and optimize the environmental and mechanical performance of a 3D object prior to fabrication. Our pipeline has three stages: (1) problem definition (design parameters and metrics), (2) differentiable simulation, and (3) optimization to expose trade-offs. We use TRELLIS + PhysiOpt as the backend physics-based optimization engine [20, 21]. PhysiOpt solves linear-elastic FEM on a hexahedral grid and returns per-element strain energies and element-wise sensitivities that can be backpropagated to TRELLIS latents for gradient-based updates [21]. Below we focus on our key contributions: (i) explicit total material-volume accounting with a volume-aware objective, (ii) spatially varying infill density computed from a material-benefit metric, and (iii) a UI that allows users to explore design–sustainability trade-offs.

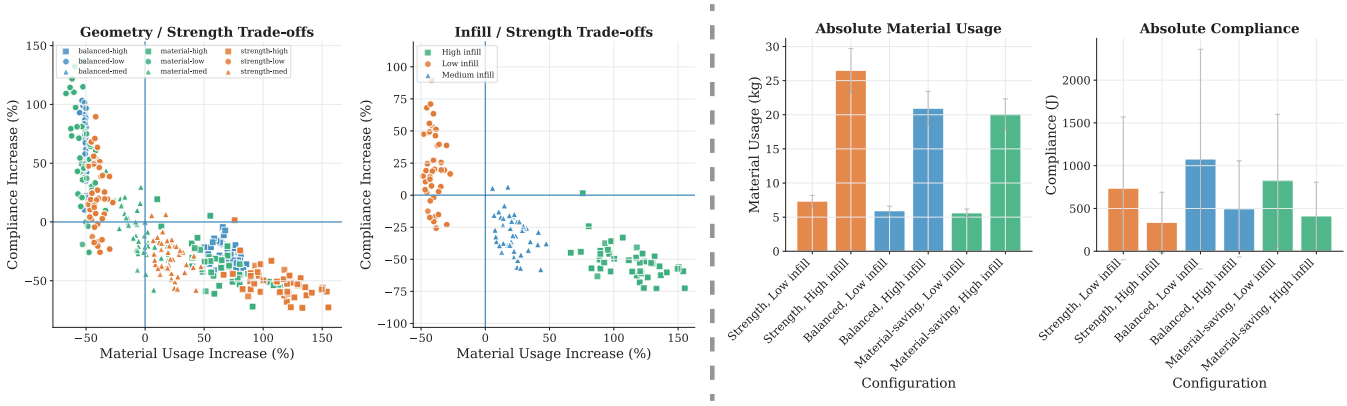

Figure 3: Prompt benchmark results over 40 prompts (simulation). Left: geometry-level trade-offs relative to the baseline. Middle: graded-infill trade-offs on the baseline geometry. Right: aggregate material usage and compliance for the different geometry–infill configurations across prompts.

3.1 Problem formulation

We define a 3D-printable design by (a) a geometry and (b) an infill field, evaluated under a user-specified load case (boundary conditions and external forces). We consider two performance axes: *sustainability* (material usage) and *mechanical viability* (stiffness and stress under load).

Users can explore the sustainability–strength trade-off in two ways: geometry optimization makes coarse, shape-level changes that reduce (or add) material, while graded infill fine-tunes where material is placed inside a chosen shape.

Design parameters

Geometry. Following PhysiOpt [21], geometry is parameterized by latent variables $\boldsymbol{\pi}$ decoded by TRELLIS into a sparse hexahedral grid with per-element densities $\rho_e \in [\rho_{\min}, 1]$ [20]. Let $\boldsymbol{\rho} = \{\rho_e\}_{e=1}^{N_e}$ denote the resulting density field over the N_e active elements.

Graded infill. We additionally assign a per-element infill multiplier $\alpha_e \in [\alpha_{\min}, 1]$, yielding an effective density

$$\tilde{\rho}_e \;=\; \rho_e \alpha_e. \tag{1}$$

We denote the effective-density field as $\tilde{\boldsymbol{\rho}} = \{\tilde{\rho}_e\}_{e=1}^{N_e}$. For our *balanced* baseline, we use a uniform infill setting $\alpha_e \equiv 0.5$.

Performance metrics

Material usage. Let v_e denote element volume. We define density-weighted material volume,

$$V_{\text{mat}}(\boldsymbol{\rho}) \;=\; \sum_{e=1}^{N_e} \rho_e \, v_e, \tag{2}$$

and under graded infill,

$$\widetilde{V}_{\text{mat}}(\tilde{\boldsymbol{\rho}}) \;=\; \sum_{e=1}^{N_e} \tilde{\rho}_e \, v_e. \tag{3}$$

(For our uniform hexahedral grid, v_e is constant across elements.)

Mechanical viability. We use PhysiOpt to evaluate mechanical response under load via linear-elastic FEM and report: (i) *compliance* $C(\boldsymbol{\rho})$ (lower implies a stiffer structure), and (ii) stress summaries (maximum von Mises stress max σ_{vM} [21]).

3.2 Differentiable simulation

Given $\boldsymbol{\pi}$ and a load case, TRELLIS decodes a hexahedral density grid $\boldsymbol{\rho}$ and PhysiOpt runs linear-elastic FEM to obtain displacements, stresses, and per-element strain energy. PhysiOpt also provides element-wise sensitivities that can be backpropagated to $\boldsymbol{\pi}$ for gradient-based updates [20, 21]. Our method uses this backend directly; the remaining sections describe what it adds.

3.3 Volume-aware objective for geometry optimization

We explicitly account for total material volume (as opposed to just minimizing *changes* in geometry) during geometry optimization. Using Eq. 2, we optimize the volume-aware objective

$$J(\boldsymbol{\rho}) \;=\; C(\boldsymbol{\rho}) \;+\; \lambda_{\text{stab}} \left| V_{\text{mat}}(\boldsymbol{\rho}) - V_0 \right| \;+\; \lambda_{\min} V_{\text{mat}}(\boldsymbol{\rho}), \tag{4}$$

where $C(\boldsymbol{\rho})$ is the compliance returned by PhysiOpt [21] and V_0 is the initial material volume. In our implementation, we apply the volume terms as additional element-wise gradients using $\partial V_{\text{mat}} / \partial \rho_e = v_e$: the stabilization term contributes $\lambda_{\text{stab}} \operatorname{sign}(V_{\text{mat}} - V_0)\, v_e$ and the minimization term contributes $\lambda_{\min} v_e$. These element-wise gradients are combined with the backend compliance sensitivities and backpropagated through TRELLIS to update $\boldsymbol{\pi}$ via the PhysiOpt optimization loop [21]. Varying $(\lambda_{\text{stab}}, \lambda_{\min})$ yields geometry variants that trade stiffness for material usage.

3.4 Sensitivity filtering

To reduce checkerboard artifacts when combining compliance and volume terms, we apply standard sensitivity filtering from topology optimization [1, 13]. We smooth element-wise sensitivities with a local, distance-weighted neighborhood average and backpropagate the filtered sensitivities through TRELLIS ($r_{\text{filt}} = 1.5$ grid units).

3.5 Graded infill density

Given a fixed geometry density field $\boldsymbol{\rho}$, we compute a spatially varying infill multiplier $\boldsymbol{\alpha} = \{\alpha_e\}$ from a single FEM evaluation. Let w_e denote the per-element strain energy returned by PhysiOpt [21]. We define a material-benefit score as an energy per unit material

"cost,"

$$s_e = \frac{w_e}{\rho_e v_e + \varepsilon}, \tag{5}$$

where $\varepsilon > 0$ avoids division by zero. We standardize s_e across elements to obtain $\hat{s}_e$ (z-score normalization) and map it to infill via a shifted sigmoid:

$$\alpha_e(\tau) = \alpha_{\min} + (1 - \alpha_{\min})\ \mathrm{sigmoid}\big(\beta(\hat{s}_e - \tau)\big), \tag{6}$$

where β controls the transition sharpness. We choose τ to match a target average infill $\bar{\alpha}$ over active elements which we solve by bisection (monotone in τ). In our implementation, the balanced setting uses $\bar{\alpha} = 0.5$. For manufacturability, we optionally smooth α on the voxel grid using local neighborhood averaging over occupied voxels; we account for this smoothing when solving this. Each candidate graded-infill design is evaluated by re-running the same PhysiOpt FEM solve with effective densities $\tilde{\rho}_e = \rho_e \alpha_e$ (Eq. 1), and we report $(\widetilde{V}_{\mathrm{mat}},\ C,\ \max \sigma_{\mathrm{vM}})$ to show the trade-off.

4 User interface and workflow

We provide a lightweight interface for exploring sustainability–strength trade-offs before fabrication. Users start from a text prompt and a simple load case, then make two decisions: (1) choose among geometry variants, and (2) choose among graded infill profiles for the selected geometry (Figure 2).

4.1 Layout

The UI is organized around a main exploration area and a metrics sidebar. A top input row accepts a prompt and triggers analysis under the selected boundary and load conditions. The main area then shows three geometry candidates side-by-side (Step 1). After a geometry is selected, the UI presents three graded-infill options for that geometry (Step 2), visualized as an infill-density overlay. The sidebar (Step 3) summarizes predicted metrics for each option, including material usage (volume) and compliance (lower is stiffer), and supports exporting the selected geometry and infill. We intentionally limit each stage to three candidates to reduce choice overload in a quick fabrication workflow while still offering a clear low–mid–high spectrum of alternatives. The workflow is as follows:

(1) **Provide an input.** User enters a prompt and specifies boundary and load conditions, and then runs the analysis.
(2) **Choose a geometry.** The system generates three geometry candidates spanning the material–stiffness trade-off and presents their previews and metrics for comparison.
(3) **Add graded infill and export.** For the selected geometry, the system computes three graded-infill profiles (low/medium/high average infill), re-evaluates predicted metrics, and allows exporting the chosen configuration.

5 Evaluation

We evaluate the prototype with a simulation benchmark that quantifies the material–stiffness–stress trade-offs exposed by our two controls (geometry variants and graded infill). All candidates are evaluated with the same TRELLIS and PhysiOpt backend (Section 3).

Protocol. For each prompt, we generate three geometry variants (material-saving, balanced, and strength-first). For each geometry, we compute three graded-infill profiles with target average infill $\bar{\alpha} \in \{0.2, 0.5, 0.7\}$ (low/medium/high). This yields 9 graded-infill candidates per prompt (3 geometries × 3 infills), plus geometry-only variants with uniform balanced infill ($\alpha_e \equiv 0.5$). We treat the unmodified baseline as the balanced geometry with uniform infill and report changes relative to this baseline. For each candidate we report material usage ($\tilde{V}_{\mathrm{mat}}$), compliance C (lower is stiffer), and maximum von Mises stress $\max \sigma_{\mathrm{vM}}$.

Results. Figure 3 shows that the geometry variants shift overall material usage and stiffness (left), while graded infill provides additional trade-offs for a fixed geometry (middle). To summarize the material–stiffness trade-off with simple statistics, we report the best material reduction achievable under a bounded stiffness loss, where stiffness loss is defined as an increase in compliance relative to the baseline. Under a +5% compliance limit, the best candidate reduces material by 3.3% ± 7.9% on average; relaxing the limit to +10% yields 5.8% ± 12.5%, and +20% yields 12.5% ± 16.3%.

6 Limitations

Our prototype focuses on material usage and mechanical proxies under simplified load cases using linear-elastic simulation. The reported print-time and energy quantities are estimates/proxies, and real prints may differ due to printer- and material-specific effects (e.g., anisotropy and layer adhesion). Extending the system to additional sustainability metrics and broader empirical validation are important directions for future work.

Our current evaluation assumes a single, simplified static load case and linear-elastic behavior; real objects may experience multi-directional loading, torsion, impacts, or cyclic fatigue, which can shift where reinforcement is needed. We also do not model buckling or contact/friction effects, so thin-walled designs that look acceptable in linear FEM may fail via instability in practice. Finally, while we note printer- and material-specific effects, our simulation does not explicitly capture print anisotropy, inter-layer adhesion strength, or voids from under-extrusion; users should treat predictions as comparative guidance rather than guarantees, and validate critical parts with conservative settings and real-world testing.

In future work, we plan to expand sustainability beyond material volume by incorporating print-time and energy estimates derived from printer profiles and toolpath (or slicer) predictions, and by adding end-of-life signals such as recyclability/recycled-content materials.

7 Conclusion

We presented an interactive prototype that helps users explore material–strength trade-offs prior to fabrication by comparing geometry variants and graded-infill profiles generated with a differentiable mechanics backend. In a simulation benchmark over 40 prompts, the candidate sets consistently contained multiple non-dominated options and enabled substantial material reductions under bounded stiffness loss. We see this as a step toward more resource-efficient personal fabrication workflows.

Acknowledgments

We would like to acknowledge the MIT Morningside Academy for Design (MIT MAD) for their support.

References

[1] Erik Andreassen, Anders Clausen, Mattias Schevenels, Boyan S Lazarov, and Ole Sigmund. 2011. Efficient topology optimization in MATLAB using 88 lines of code. *Structural and Multidisciplinary Optimization* 43, 1 (2011), 1–16.

[2] Jeremy Faludi, Cindy Bayley, Suraj Bhogal, and Myles Iribarne. 2015. Comparing environmental impacts of additive manufacturing vs traditional machining via life-cycle assessment. *Rapid Prototyping Journal* 21, 1 (2015), 14–33.

[3] Faraz Faruqi, Amira Abdel-Rahman, Leandra Tejedor, Martin Nisser, Jiaji Li, Vrushank Phadnis, Varun Jampani, Neil Gershenfeld, Megan Hofmann, and Stefanie Mueller. 2025. MechStyle: Augmenting Generative AI with Mechanical Simulation to Create Stylized and Structurally Viable 3D Models. In *Proceedings of the ACM Symposium on Computational Fabrication*. 1–15.

[4] Faraz Faruqi, Ahmed Katary, Tarik Hasic, Amira Abdel-Rahman, Nayeemur Rahman, Leandra Tejedor, Mackenzie Leake, Megan Hofmann, and Stefanie Mueller. 2023. Style2Fab: functionality-aware segmentation for fabricating personalized 3D models with generative AI. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–13.

[5] Simon Ford and Mélanie Despeisse. 2016. Additive manufacturing and sustainability: an exploratory study of the advantages and challenges. *Journal of cleaner Production* 137 (2016), 1573–1587.

[6] Malte Gebler, Anton JM Schoot Uiterkamp, and Cindy Visser. 2014. A global sustainability perspective on 3D printing technologies. *Energy policy* 74 (2014), 158–167.

[7] Mohammad Raquibul Hasan, Ian J Davies, Alokesh Pramanik, Michele John, and Wahidul K Biswas. 2024. Potential of recycled PLA in 3D printing: A review. *Sustainable manufacturing and service economics* 3 (2024), 100020.

[8] Lin Lu, Andrei Sharf, Haisen Zhao, Yuan Wei, Qingnan Fan, Xuelin Chen, Yann Savoye, Changhe Tu, Daniel Cohen-Or, and Baoquan Chen. 2014. Build-to-last: Strength to weight 3D printed objects. *ACM Transactions on Graphics (ToG)* 33, 4 (2014), 1–10.

[9] Stefanie Mueller, Sangha Im, Serafima Gurevich, Alexander Teibrich, Lisa Pfisterer, François Guimbretière, and Patrick Baudisch. 2014. WirePrint: 3D printed previews for fast prototyping. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. 273–280.

[10] Maxine Perroni-Scharf, Jennifer Xiao, Cole Paulin, Zhi Ray Wang, Ticha Sethapakdi, Muhammad Abdullah, Patrick Baudisch, and Stefanie Mueller. 2025. SustainaPrint: Making the Most of Eco-Friendly Filaments. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*. 1–13.

[11] Michael L Rivera, S Sandra Bae, and Scott E Hudson. 2023. Designing a sustainable material for 3D printing with spent coffee grounds. In *Proceedings of the 2023 ACM designing interactive systems conference*. 294–311.

[12] Maria Shugrina, Ariel Shamir, and Wojciech Matusik. 2015. Fab forms: Customizable objects for fabrication with validity and geometry caching. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–12.

[13] Ole Sigmund. 2007. Morphology-based black and white filters for topology optimization. *Structural and Multidisciplinary Optimization* 33, 4 (2007), 401–424.

[14] Ruoyu Song and Cassandra Telenko. 2017. Material and energy loss due to human and machine error in commercial FDM printers. *Journal of cleaner production* 148 (2017), 895–904.

[15] Ruoyu Song and Cassandra Telenko. 2019. Causes of desktop FDM fabrication failures in an open studio environment. *Procedia CIRP* 80 (2019), 494–499.

[16] Ondrej Stava, Juraj Vanek, Bedrich Benes, Nathan Carr, and Radomír Měch. 2012. Stress relief: improving structural strength of 3D printable objects. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–11.

[17] Ludwig Wilhelm Wall, Alec Jacobson, Daniel Vogel, and Oliver Schneider. 2021. Scrappy: Using scrap material as infill to make fabrication more sustainable. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–12.

[18] Weiming Wang, Tuanfeng Y Wang, Zhouwang Yang, Ligang Liu, Xin Tong, Weihua Tong, Jiansong Deng, Falai Chen, and Xiuping Liu. 2013. Cost-effective printing of 3D objects with skin-frame structures. *ACM Transactions on Graphics (ToG)* 32, 6 (2013), 1–10.

[19] Xin Wen, S Sandra Bae, and Michael L Rivera. 2025. Enabling recycling of multi-material 3D printed objects through computational design and disassembly by dissolution. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–21.

[20] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. 2025. Structured 3d latents for scalable and versatile 3d generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 21469–21480.

[21] Xiao Zhan, Clément Jambon, Evan Thompson, Kenney Ng, and Mina Konaković Luković. 2025. PhysiOpt: Physics-Driven Shape Optimization for 3D Generative Models. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers*. 1–11.