

# VisiPrint: Previewing 3D-Print Appearance from Real Material Samples

Maxine Perroni-Scharf  
MIT CSAIL

Cambridge, Massachusetts, USA  
max1@mit.edu

Faraz Faruqi  
MIT CSAIL

Cambridge, Massachusetts, USA  
ffaruqi@mit.edu

SooYeon Ahn

Department of AI Convergence  
Gwangju Institute of Science and  
Technology  
Gwangju, Republic of Korea  
syahn1160@gm.gist.ac.kr

Raul Hernandez  
MIT CSAIL

Cambridge, Massachusetts, USA  
rhern@mit.edu


Szymon Rusinkiewicz  
Department of Computer Science  
Princeton University  
Princeton, New Jersey, USA  
smr@princeton.edu

William Freeman  
MIT

Cambridge, Massachusetts, USA  
billf@mit.edu

Stefanie Mueller  
MIT CSAIL

Cambridge, Massachusetts, USA  
stefanie.mueller@mit.edu



**Figure 1: The *VisiPrint* workflow.** The user slices a 3D model [16] using their preferred software (e.g., Cura), then uploads a screenshot of the sliced view and a photo of the intended filament or printed exemplar. *VisiPrint* generates an appearance-faithful preview that reflects both slicing patterns and material characteristics.

## Abstract

We present *VisiPrint*, a tool for appearance-first previews of 3D-printed objects. Existing print preview slicers focus on toolpaths, not appearance, while pure rendering software is complex and cannot automatically reproduce slicing patterns. Prior work highlights persistent gaps between digital previews and printed results, such as color shifts, gloss/translucency changes, and layer-line highlights, motivating the creation of *VisiPrint*, an appearance-focused support tool. The *VisiPrint* algorithm combines slicer screenshots with

filament photos via a custom diffusion-based synthesis pipeline. We present both a standalone user interface for *VisiPrint* compatible with any slicer and an Ultimaker Cura Plugin. We evaluate *VisiPrint* through a user study showing it is significantly faster, easier to use, and more faithful than alternatives: within a time-limit, participants completed 100% of preview tasks with *VisiPrint*, versus 63% with Cura and 13% with Blender. *VisiPrint* narrows the gap between design intent and printed appearance, complementing settings-centric tools with appearance-driven decision support.



This work is licensed under a Creative Commons Attribution 4.0 International License.  
CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2278-3/26/04  
<https://doi.org/10.1145/3772318.3790851>

## CCS Concepts

• Applied computing → Engineering: Computer-aided design.

## Keywords

3-D Printing, Previewing Tools, Stylized Fabrication

**ACM Reference Format:**

Maxine Perroni-Scharf, Faraz Faruqi, SooYeon Ahn, Raul Hernandez, Szymon Rusinkiewicz, William Freeman, and Stefanie Mueller. 2026. *VisiPrint: Previewing 3D-Print Appearance from Real Material Samples*. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3772318.3790851>

**1 Introduction**


3D printing has transformed prototyping and manufacturing, empowering users to fabricate custom objects with a wide range of materials and printer settings. However, there is a critical gap that remains in the 3D printing workflow: the ability to accurately preview what a printed object will *look* like before fabrication. For Fused Deposition Modeling (FDM), the most popular type of 3D printing [35], most slicers and software offer print previews, but these primarily highlight layer-by-layer toolpaths rather than the material and visual characteristics of the final object [71]. As a result, users are often left to guess how prints will turn out, which can lead to wasted time, filament, and effort when outcomes do not meet expectations. Prior works document persistent mismatches between print-previews and actual appearance arising from color-difference errors [52], interactions between gloss and perceived color [15], and FDM-specific surface cues such as anisotropic layer-line highlights [25]. From an HCI perspective, these mismatches are not just rendering errors—they directly influence whether users judge a print as acceptable for display, prototyping, or everyday use. Prior perceptual work shows that observers are sensitive to small changes in color, gloss, and surface finish even when geometry is identical [31, 36, 83]. Because prior work already links both slicing patterns and material properties to perceived quality [25, 26], a preview that reflects *both* is especially valuable for HCI.

On the other hand, modeling and CAD tools like Blender and Fusion 360 can generate high-fidelity renders, but they do not account for the printing process. A render may depict a smooth, glossy object, while the printed result, shaped by layer height, support removal, artifacts, and filament texture, looks dramatically different [44]. Neither these tools nor typical slicers capture the interplay between slicing patterns and material appearance that defines FDM’s distinctive look [49]. Recently, mainstream slicers have released dedicated features to mitigate visible artifacts in 3D prints such as Prusa’s seam painting [3], and Cura’s recent *Z-Seam* controls which enable fine-tuning of surface details [4], demonstrating that appearance is a practical concern in everyday workflows. Community discussions frequently raise 3D-printing issues due to appearance mismatches, such as color banding [1], needing to print personal color swatches to see how filament will actually look, and transparency falling short of “clear” claims [2], underscoring that appearance is a practical pain point for makers. Studies of fabrication platforms similarly show that missing appearance cues creates uncertainty, can increase the number of trial prints, and limits reuse of designs [7, 24]. Thus, print appearance can shape user decisions about whether to print an object, which filament to use, and whether the result meets a user’s standards [26, 63].

Meanwhile, recent advances in computer vision techniques have brought fast, expressive visual synthesis to other domains (e.g., SDXL [61]), while diffusion-based material transfer continues to



**Figure 2: The Benchy boat being previewed in silver PLA across different slicing patterns and software programs.**



**Figure 3: The Benchy boat previewed in home material samples (the blue and clear spheres) and samples found online, including challenging translucent and rainbow filaments.**

progress (e.g., *ZeST* [13]). Physically based pipelines estimate SVBRDFs and re-render the object [5, 17, 27, 48, 79], but they rely on controlled capture and calibrated alignment that everyday makers rarely have. Exemplar-driven methods such as *MatCap* [54] and recent lightweight material-transfer models [13, 50, 76, 84] are easy to use but ignore FDM-specific cues like layer lines and seams. Yet no existing system leverages computer vision to preview the appearance of 3D prints from a user’s specific sliced input and chosen filament. As we demonstrate in our evaluation, out-of-the-box models struggle with domain-specific challenges such as slicing-induced microgeometry and limited training data. For example, we find that off-the-shelf SDXL cannot reliably synthesize images of 3D-printed objects, as this remains an underrepresented domain in its training set. To our knowledge, this is the first paper to directly target these domain-specific challenges for 3D-printed appearance preview.

At the same time, AI-assisted functionality has become standard in creative software, from photo editors (e.g., Photoshop) [77] to illustration tools and animation pipelines [14]. Bringing such capabilities to fabrication previews is a natural next step. But unlike purely generative image tasks, outputs must correspond to objects that will physically exist, not just be plausible renderings. This requires preserving key features such as geometry and slicing patterns. While 3D printing software has begun adding more appearance-oriented features (e.g., easier color/material selection and themed visual presets), reflecting demand for faster, more informative appearance decisions [74], to date no available tool employs AI to preview the 3D-printed look directly from a user’s own sliced input and materials on hand. Our work builds on HCI systems such as *Aldeation* [80], which show that workflow-aware, transparent

AI can support creative exploration while keeping designers in control. We focus on an appearance-first preview that helps designers decide whether and how to print. A preview that reflects the user's actual toolpaths and filament reduces visual uncertainty [9, 40], and because of this, we treat appearance preview as a post-slicing step for setting expectations about the print.

To address this gap, we present *VisiPrint* (shown in Fig. 1), a post-slicing preview tool that operates alongside a user's preferred slicer and printer. *VisiPrint* accepts a slicer screenshot and a single material photo as inputs, and synthesizes previews that reflect both slicing patterns and filament properties. As shown in Fig. 2 and Fig. 3, it works with any slicer software and with any material example, requiring only image inputs from the slicer of choice's UI and material. *VisiPrint* aims to predict how prints are likely to look, rather than to predict print success. It complements, not replaces, slicer checks (e.g., supports, adhesion, print settings).

## 2 Related Work

Research in digital fabrication spans both *appearance-focused* tools, which address how prints look, and *feasibility-focused* tools, which ensure they succeed. From an HCI perspective, appearance is critical: it shapes perception, acceptance, and interaction with fabricated objects, yet it remains less explored than reproducibility. For example, Alcock et al. [7] found that users on Thingiverse struggled with missing cues about how designs would actually print, underscoring the importance of effective previewing in reducing uncertainty during the 3D-printing process.

To situate our contribution, we first review work on appearance in 3D printing, then contrast it with feasibility, interaction, and vision-assisted fabrication, highlighting how each frames the user experience of making and evaluating artifacts.

### 2.1 Appearance-Focused 3D Printing Tools and Studies

The gap between how prints are previewed and how they actually appear is well documented, and motivates the need for appearance-first tools. Yuan et al. [83] emphasize the absence of standards for reproducing appearance on curved, glossy, or translucent prints. Galati et al. [26] propose evaluation frameworks where aesthetic quality is judged by users and treated as a primary workflow concern. Together, these perspectives underscore that appearance is not incidental but central, shaping perception, satisfaction, and adoption. This motivates *VisiPrint*, which aims to preview aesthetic properties more faithfully than existing systems. By grounding previews in exemplars of real 3D-printed materials and matching color profiles, *VisiPrint* addresses these gaps, allowing users to anticipate printed appearance.

Color discrepancies are commonly quantified with the CIEDE2000 formula [52]. In dentistry, for example, full-color 3D-printed casts deviate unacceptably from design intent [8, 68], and similar mismatches occur in other medical models [29]. For education and communication, lifelike visual properties add value beyond geometry alone. In prosthetics, appearance is tied directly to acceptance: studies have found that more natural-looking 3D-printed cosmetic covers improved aesthetics to user satisfaction [21, 63]. These examples reinforce the case for appearance as a first-order design

concern in HCI, one that *VisiPrint* elevates from afterthought to previewable design variable.

Beyond color, other material properties significantly influence appearance. Condor et al. [15] show that surface gloss alters perceived color and propose gloss-aware correction workflows. For translucent parts, appearance varies with thickness and subsurface scattering; Brunton et al. [11] estimate optical properties from samples and apply inverse rendering to reproduce target translucency. Orientation also plays a role: Filip and Vitek [25] find that observers consistently notice anisotropic highlights and layer-line banding. Elkhuizen et al. [23] demonstrate the reproduction of paintings by jointly capturing gloss, color, and surface relief, making appearance, not geometry, the primary target. In manufacturing, Zhang et al. [88] model orientation tradeoffs that explicitly include appearance, while Hartcher-O'Brien et al. [31] show users distinguish surface finishes tied to FDM parameters beyond roughness metrics. Shape and lighting matter as well: Jiang et al. [36] show perceived color shifts with geometry and illumination, while Hlayhel et al. [32] find that "naturalness" in color prints depends on elevation and roughness. To accommodate these effects, *VisiPrint* previews not only material color but also gloss, translucency, and slicing artifacts, enabling users to anticipate banding direction and surface finish.


Mainstream slicers also reflect demand for better visible outcomes, with features like *PrusaSlicer*'s seam-hiding [3] and Cura's *Z-Seam* controls [4]. While such features and prior work reduce discrepancies between intent and outcome, *VisiPrint* instead foregrounds expectation management. By integrating color, gloss, translucency, and artifacts, it reframes previews as tools for aligning user expectations with reality.

### 2.2 Interactive Fabrication Tools

Feasibility and reproducibility remain central to digital fabrication research. Prior work often treats slicer parameters and machine settings as first-class artifacts for sharing and repeatability: *SliceHub* [24] embeds slicing previews and metadata into repositories, while Subbaraman et al. [75] introduce workflows for recording and reusing parameters. Other systems address variability and failure, including 3DPFIX for diagnosing failed prints [43] and uncertainty frameworks [40].

*VisiPrint* addresses a complementary question: What will it look like? While not fundamental to success, visible outcome influences choices such as filament, orientation, slicing strategy, or whether to print at all. Appearance-first previews can complement settings-first and feasibility-first tools. By positioning the preview itself as an HCI intervention, *VisiPrint* reframes fabrication workflows around expectations management.

Vision and interactivity have also been used to provide feedback before and during fabrication. *Don't Mesh Around* uses real-time scanning for clay printing [58]. *SensiCut* applies speckle imaging and deep learning to identify materials for safer laser cutting [19]. *Patching Physical Objects* employs depth sensing to repair prints [78], while *MultiFab* monitors multi-material prints via 3D scanning [70]. Other systems embed information directly into appearance: G-ID encodes slicer-parameter textures [20], while *AnisoTag* introduces anisotropic micro-patterns [53]. Interactive fabrication has connected users more directly through direct manipulation



**Figure 4: The *VisiPrint* interface. (a) The user uploads a sliced model screenshot and selects or uploads a material. (b) Advanced settings control color and luminance transfer, slicing pattern influence, custom prompts, and fine-tuned model usage. (c) The user generates a preview of the 3D printed object.**

[82], AR-integrated robotic printing (*RoMA*) [60], mid-process reshaping (*FormFab*) [59], mechanism remixing (*Grafter*) [67], and shape-changing interface prototyping (*MorpheusPlug*) [39]. Spatial preview tools further situate models in context, such as *Mix&Match* [73], *pARam* [72], and *TinkerXR* [10]. These works emphasize feasibility or interaction, while *VisiPrint* emphasizes pre-fabrication appearance, enabling informed decisions before committing material.

### 2.3 Image-Based Appearance Transfer and Diffusion Models

Applying imaging and vision techniques to computational fabrication is still an emerging area of research, though appearance transfer and image generation are well studied. *VisiPrint* builds on material appearance transfer techniques, primarily *ZeST* [13], which applies pre-trained image-to-image translation with depth and illumination conditioning. Related methods include *Cross-image Attention* [6], which uses diffusion with semantic correspondences to transfer appearance without retraining, *ReStyle3D* [90], which aligns correspondences across multiple views, and Deschaintre et al. [18], which adapts deep appearance models via targeted fine-tuning for large-scale transfer. Complementary to these transfer pipelines, a rich line of physically based work estimates material parameters for predictive rendering, including single-image SVBRDF recovery [17, 48, 85], two-shot acquisition [5], high-resolution inverse rendering [27], and physically based editing from a single photo [79], with semi-procedural material priors [89] and procedural approaches to material synthesis [46].

Neural style transfer [28, 37] and diffusion approaches such as *Alchemist* [69] enable parametric control over gloss and roughness, while *Material Palette* [51] and Kocsis et al. [42] extract material sets and perform probabilistic decomposition from single views. Image-space appearance editing and swapping spans both classical and recent techniques [50, 76, 84]. Like *VisiPrint*, *Anisotropic MatCap* [54] demonstrates that spherical material samples can capture how a material responds under a fixed illumination. Building on these advances, *VisiPrint* integrates FDM slicing patterns into *ControlNet*-guided Stable Diffusion [66, 86]. Conditioning on slicer-derived Canny edges and depth preserves geometry, lamination, and seam placement, while the exemplar transfers hue and sheen—explicitly coupling *material* and *process* for previews. To our

knowledge, *VisiPrint* is the first system to combine exemplar-based transfer with slicer-derived structure in this way; SVBRDF/inverse-rendering, *MatCap*, and generic editing pipelines do not simultaneously reconstruct both real filament appearance and FDM slicing patterns. In contrast to these pipelines, which focus on material properties or generic image edits, *VisiPrint* specifically targets FDM fabrication: it fuses exemplar-based transfer with slicer-encoded depth and lamination structure so that previews reflect both real filament appearance and actual slicing patterns.

## 3 User Interface

Figure 4 shows the *VisiPrint* interface, designed for both novice and expert users. It supports a simple end-to-end workflow from model preparation to preview generation, while offering advanced controls for fine-grained customization.

### 3.1 Core Workflow


As illustrated in Figure 4a, the workflow consists of three steps:

- (1) **Upload Sliced Geometry:** The user provides a screenshot of the sliced model (e.g., from Cura or Bambu Studio), which encodes object shape, slicing pattern, and view direction.
- (2) **Select or Upload Material:** A reference image of a printed object may be uploaded, or a preset chosen from the curated *Material Library*. Uploaded materials can also be saved for reuse.
- (3) **Generate Preview:** Clicking *Generate* produces a photorealistic preview in under one minute, displayed alongside the input slicing image and material exemplar.

### 3.2 Advanced Settings

The expandable *Advanced Settings* panel (Figure 4b) provides expert controls: **Color Influence**, for the degree of color transfer from the exemplar, **Luminance Influence**, for the impact of exemplar brightness and shading, **Slicing Influence**, for the visibility of slicing contours in the output, **Text Prompt Override**, for custom prompts and adjustable influence weights, **Inference Steps**, for the trade-off between speed and quality, and **Fine-tuned Model**, an option to use a Stable Diffusion variant trained on 3D print imagery.

By combining a streamlined workflow with advanced customization, *VisiPrint* enables both quick previews and detailed control over



**Figure 5: VisiPrint architecture:** The system takes as input a material exemplar and a 3D geometry. The material is encoded using an Image Prompt Adapter (IP Adapter), while the geometry is sliced and rendered. Slicing, geometry, and latent illumination cues guide image generation via a multi-ControlNet pipeline. The output is refined using diffusion inpainting (SDXL) and finalized through color correction, producing a photorealistic preview of the 3D-printed object.

visualization. The “Slicing Influence” slider in Fig. 4b provides a user-facing control for adjusting the relative contribution of Canny edges and depth cues in our *ControlNet* conditioning (with “Slicing Influence” chosen as a more user-friendly term); this corresponds directly to the edge/depth weight settings evaluated in the ablation study in Section 7.4.

The interface also includes a small “Preview Limitations” information panel that clearly explains that outputs are simulated, notes that appearance depends on the exemplar’s lighting, and reminds users that *VisiPrint* is for appearance-only decisions; a screenshot is provided in the Supplemental Material.

### 3.3 Cura Plugin

To allow users to interact with the preview functionality within their current workflows, we implemented a dedicated Ultimaker Cura plugin that lets users access *VisiPrint* directly in their slicing software (Figure 4c). Once a model is sliced, the plugin automatically captures a screenshot of the 3D viewport as the *target* geometry. Users can select or manage a local library of material exemplars directly in Cura, then generate previews via the Extensions / *VisiPrint* menu. The captured geometry and chosen material are processed by the *VisiPrint* pipeline, and the resulting preview is displayed in a zoomable viewer embedded within Cura. This integration eliminates the need for manual screenshot handling, providing a seamless loop between slicing and appearance evaluation.

## 4 Method

*VisiPrint* generates photorealistic previews of 3D-printed objects by combining material transfer with geometry and slicing-aware synthesis (Fig. 5). Users see both overall shape and how slicing and

material affect appearance. At its core is a generative model conditioned on material exemplars and structural cues. We describe the core components of our method below, and quantitatively evaluate them later in Section 7.4.

### 4.1 Input Normalization and Capture

*Slicer screenshots.* The inputs to *VisiPrint* are slicer agnostic, and can be taken with any preview color or slicing pattern/angle. However, we recommend choosing a preview color that contrasts with the build plate color in the slicer preview and disabling overlays (such as travel moves). Once a screenshot is uploaded, preprocessing removes the background and computes edges and depth only inside the silhouette. Exemplars are single photos of printed samples (e.g., spheres) captured in a controlled tabletop setup. We used a soft lightbox positioned at a 45° overhead angle with white interior walls, providing diffuse global illumination while still producing a single soft specular highlight visible in the exemplar images. A ceiling light contributed weak ambient lighting. The same lightbox and camera position were used for both exemplar spheres and all evaluation photographs, ensuring consistent illumination and viewing geometry across materials.

### 4.2 Material Transfer

Our pipeline extends the material transfer pipeline proposed in *ZeST* [13] by incorporating additional slicing cues and accommodations for the 3D printed appearance domain. Our pipeline consists of three stages: (i) extract material features, (ii) compute geometry and structure guidance, (iii) synthesize previews conditioned on both.

Material features are encoded with CLIP [62]. Depth is estimated with DPT [64], while grayscale masking preserves illumination. From slicer screenshots we extract depth and Canny edges [12], encoding both shape and lamination.

Previews are generated with Stable Diffusion XL (SDXL) [61] and multi-ControlNet [86]. Depth and edges guide shape and layer lines; CLIP embeddings and silhouettes constrain region and appearance. The result approximates printed look for the chosen filament and slicing parameters.


### 4.3 Incorporating Slicing Patterns

FDM prints exhibit anisotropic appearance from extrusion, so previews must retain slicing detail. From screenshots, we compute a Canny edge map highlighting lamination, to complement the depth map conditioning proposed in *ZeST* [13]. Because slicer screenshots vary in brightness and contrast, fixed Canny thresholds can miss layer lines, especially on dark views. We therefore use a mask-aware, adaptive Canny (Fig. 6). We first obtain an object mask from the alpha matte and replace background pixels with the median interior intensity to avoid a silhouette edge. After a Gaussian blur, we choose the high threshold by searching a gradient percentile until interior edge density falls within a target band (about 12–25% of masked pixels for FDM layer lines). Finally, edges are restricted to an eroded mask and small components are removed to reduce speckle.


As can be seen in Fig. 7, depth maps ensure shading and global shape, while Canny edges preserve slicing and internal contours. Depth alone may blur hollow regions; edges alone lose shading. Users can balance the two via a *Slicing Influence* slider: higher values emphasize lamination, lower values emphasize shading. Defaults balance both.

### 4.4 Color Correction and Artifact Reduction

Canny edges aid structure but can shift tone. To correct this, we apply statistical color transfer [65] within the object silhouette (eroded to reduce boundary noise). Exemplar RGB mean/variance is matched, with contrast weight 0.5 and blend factor 0.8 to avoid



**Figure 6: Visualization of a challenging, dark input with and without adaptive Canny map generation. Adaptive Canny more effectively picks up the slicing details of the input.**



**Figure 7: We use both Canny edges and depth maps for conditioning the generated images: depth-only smooths shading but loses slicing detail, while edges preserve lamination but weaken shape. Combining both balances these effects. Red-circled regions highlight depth-only errors. In hollow regions (e.g., Benchy vent), edges preserve contours while depth retains global shape. Without edges, spurious holes may appear.**

overcorrection. This improves tonal coherence while retaining generative detail (Fig. 8). To prevent unwanted structure from material photos (e.g., a donut hole or logo) from leaking into previews, we crop the exemplar to its foreground only. We also restrict edits to the object mask from the target view so the background does not influence results (Fig. 10).

### 4.5 Material Exemplar Geometry Ablation

While users may provide any image of a printed material as an exemplar to *VisiPrint*, how well texture and lighting are captured is affected by the geometry of the exemplar. To identify which shapes are most effective as material example geometries, we conducted an ablation study across 12 geometries from our dataset. Each geometry was used as a material exemplar, and we evaluated the results using CLIP similarity [62], a metric that quantifies semantic similarity between generated previews and real photographs of printed objects. CLIP was chosen because it is more robust to variations in orientation and lighting than other metrics, making it suitable for evaluating our model’s output images against the corresponding ground-truth photographs.

We found that spheres yielded the highest similarity scores (mean CLIP = .818), likely due to their even curvature and ability to capture material qualities under varied lighting, and cylinders performed worst (mean CLIP = .795). This preference for spherical exemplars mirrors *MatCap*-based workflows [54], where a single photographed sphere is used to encode how a material responds to a fixed illumination environment and is then mapped onto arbitrary geometries. In our case, the printed exemplar’s appearance is injected as a conditioning signal into the diffusion model rather than used as a direct normal-lookup texture as in classical *MatCap* rendering. This allows *VisiPrint* to preserve FDM-specific surface cues such as lamination edges and layer-line highlights via slicer-derived depth and edge maps. While the performance gap between the best and worst geometries was modest (.023), we recommend using spheres for printing a new material sample for optimal results. This finding is specific to our pipeline and default hyperparameters, and performance may vary in other settings.




Figure 8: Color correction improves correspondence between exemplar and output.




Figure 9: Silver *VisiPrint* boat, without a custom prompt and with the custom prompt “A very shiny 3D printed boat”. Users can adjust previews with optional prompting.




Figure 10: Cropping the material exemplar removes background pixels and reduces artifact transfer, such as the clear spurious hole on the front of the boat.

#### 4.6 Optional Text Prompt Conditioning

Users may provide text prompts for creative control, though this is not required or recommended for realism. By default, a BLIP caption [47] extracts the main noun and appends it to “A 3D printed” (e.g., “A 3D printed vase”). Users can override with custom prompts that can modify the output’s material appearance (Fig. 9).

#### 4.7 Optional Fine-tuning

While SDXL is a powerful generative model, our early experiments show that it does not handle generations of 3D printed objects well (Fig. 11). To better align SDXL with the appearance of real



Figure 11: Top row: SDXL outputs without fine-tuning. Bottom row: SDXL outputs with fine-tuning. Fine-tuning improves generated outputs of SDXL when prompted to visualize a 3D printed cup.

prints, we optionally apply lightweight LoRA fine-tuning [34], a computationally efficient fine-tuning approach that injects small, trainable weights into the frozen model, using 100 Creative Commons–licensed images of 3D prints. We will release this dataset of images upon publication of this manuscript.

In Figure 11, we show the effect of this fine-tuning on the SDXL model alone, by presenting generated images for the prompt “A 3D printed cup” before and after our fine-tuning approach. The cups appear more realistic and faithful to true 3D prints after fine-tuning.


#### 4.8 Implementation

The system is implemented in PyTorch, using `diffusers` for SDXL inference and fine-tuning and `rembg` for background removal. Our core model is `StableDiffusionXL-ControlNet-Inpaint-Pipeline` with two ControlNet branches (depth and edges). Depth maps are computed with DPT [64]. Appearance is guided by IP-Adapter XL.

Optional LoRA fine-tuning runs 10,000 steps at learning rate 0.01. Default ControlNet scales are 0.9 (depth and edges); inference uses 70 steps with seed 42. The *Slicing Influence* slider linearly adjusts edge vs. depth weights ( $w_{edge}$ ,  $w_{depth}$ ), initialized at 0.9 each and clipped to  $[0, 1]$ . Preprocessing uses masks eroded with a  $5 \times 5$  kernel. Color correction applies histogram matching with contrast weight 0.5 and blend factor 0.8. Experiments ran on an NVIDIA A100 GPU.

#### 4.9 Scope and Assumptions

*VisiPrint* is designed primarily for FDM prints, but its approach generalizes to other technologies such as SLA. In the Supplemental, we show additional examples using resin and other materials. The system focuses on visual appearance: it previews what a print is expected to look like *if fabrication is successful*. It does not estimate printability, mechanical feasibility, or the likelihood of failure. Our goal is to give users a clearer sense of surface finish, material look, and color outcome—not to predict whether a print will succeed technically.



**Figure 12: User study results. Top row: mean scores for the NASA-TLX and extra questions (left), task completion rate (middle), and completion times for successful trials (right). Bottom rows: distribution of non-neutral responses for individual TLX dimensions and realism/confidence ratings. *VisiPrint* consistently reduced workload and increased success compared to baselines, and its improvement was statistically significant across all metrics other than physical demand.**

## 5 User Study

We evaluate *VisiPrint* as an appearance-first preview aid by asking whether, before printing, it helps participants reach a preview that they judge as *faster and more reliable*, requiring *lower workload*, and offering *higher realism* than baselines. We report task completion, time-to-preview, perceived realism, and workload. To compare against existing 3D print preview approaches (baseline), we set up three different tools during the experiment:

- **Ultimaker Cura Slicer Preview Tab:** A standard slicer visualization that displays the layer-by-layer slicing pattern.
- **Blender-Based Rendering:** A manually configured rendering pipeline using Blender to simulate realistic materials and lighting.
- ***VisiPrint*:** Our approach, which uses generative models to transfer material appearance onto sliced 3D models.

The study examined trade-offs among ease of use, photorealism, visual fidelity, and task efficiency. While Cura and Blender do not cover the full landscape of preview tools, they are popular and accessible in printing and modeling communities. We conducted a one-hour user study in a lab environment with fifteen participants. Users were given five minutes to complete various tasks, and the completion rate for *VisiPrint* was 100% (30/30 tasks across all participants) but only 63% (19/30 tasks) for Cura and 13% (4/30) for Blender. Participants consistently rated *VisiPrint* as significantly easier to use, less demanding, and more photorealistic than the other tools. In a perception study, previews generated by *VisiPrint* were rated as more realistic than those from the baseline and nearly on par with actual photographs. Beyond measuring task completion time, this workflow study was designed to probe users' perceived

workload, confidence, and trust when interacting with different preview tools, reflecting the cognitive demands of real 3D-printing decision workflows rather than simply comparing interface speed.

This study received institutional approval. Participants gave informed consent and were able to withdraw at any time.

### 5.1 Participant Demographics

We recruited fifteen participants (self reported genders of ten men and five women) aged 18 to 37 (with a mean age of 26.27). The participants were university students with different majors and backgrounds. All but one participant had prior experience with printing and modeling tools.

### 5.2 Pre-survey Questionnaire

**Process.** Before starting the tasks, participants completed a background questionnaire covering:

- Prior experience with printers and modeling tools.
- Importance of visualization before printing (1 = Not at all, 7 = Essential).
- Current methods or strategies used to preview prints.
- Prior experience with different print materials (e.g., colors, plastics).
- Willingness to spend time on visualization before printing.
- Prior experience with Blender and Cura.

**Results.** We assessed participants' familiarity with each tool using a 7-point Likert scale (1 = novice, 7 = expert). The mean familiarity

scores were 2.63 ( $SD = 1.40$ ) for Blender and 2.60 ( $SD = 2.26$ ) for Cura, indicating that participants were generally knowledgeable amateurs. Participants rated print visualization as highly important, with an average score of 6.40 ( $SD = 0.63$ ). The time they were willing to spend on visualization averaged 14.93 minutes, with a mode of 5 minutes. Five participants reported frequent experience with aesthetic materials, six occasional, two infrequent, and two none.

### 5.3 Evaluating Software Interaction Experience

For each tool, participants completed two tasks:

**Task 1:** Visualize the boat in silver PLA.

**Task 2:** Visualize the whistle in clear PLA.

Participants' completion times were recorded, with a maximum task duration of 5 minutes. If 5 minutes had elapsed, they were prompted to stop and asked if they felt they had finished. If yes, time was recorded as 5 minutes; if not, the task was marked incomplete. Tool order was rotated with a Latin-Square to reduce order effects.

To isolate software interaction from asset preparation (studied separately in the Supplemental), some resources were pre-provided:

- Blender: curated texture maps for silver and clear PLA or built-in materials.
- *VisiPrint*: sliced screenshots and images of the silver and clear PLA references.

### Results.

*Completion Rates and Times.* Completion rates per software tool were 30/30 for *VisiPrint*, 4/30 for Blender, and 19/30 for Cura. For completed tasks, average completion times were 66.4 s ( $SD = 10.0$ ) for *VisiPrint*, 158.5 s ( $SD = 64.9$ ) for Cura, and 250.2 s ( $SD = 34.2$ ) for Blender (Figure 12). These results show a sizable improvement for *VisiPrint* over both baselines in both completion rate and task time.

*NASA-TLX.* After each task, participants completed NASA-TLX [30] and two extra 7-point Likert questions: (i) "How photorealistic was the result?" (1 = Not at all, 7 = Extremely) and (ii) "How confident are you that the preview matches the real object?" (1-7). For *VisiPrint* specifically, they also answered: "Did the automatic visualization help with previewing the print?" (1-7).

We used non-parametric tests for within-subjects ordinal data. For each subjective metric, a Friedman test assessed overall differences across the three tools. If the Friedman test was significant, we ran Wilcoxon signed-rank tests for pairwise comparisons with Holm-Bonferroni correction for multiple comparisons. We report statistical significance at  $\alpha = 0.05$ , with all raw p-values listed in Table 1. Significant results at  $p < .05$  (after Holm-Bonferroni correction) are shown in boldface. *VisiPrint* was rated as significantly lower-effort, less mentally demanding, less hurried, and more successful than both Blender and Cura. Participants also reported feeling less insecure, perceiving higher photorealism, and expressing greater confidence in the resulting prints.

For *VisiPrint*, the extra question "Did the automatic visualization help with previewing the print?" had a mean score of 6.20 out of 7.

**Table 1: Friedman and Wilcoxon test results. We report raw p-values; boldface indicates significance after Holm-Bonferroni correction ( $p < .05$ ).**

Metric	Friedman	VisiPrint vs Blender	VisiPrint vs Cura
Mental Demand	<b>.000016</b>	<b>.000650</b>	<b>.003483</b>
Hurried	<b>.000036</b>	<b>.001576</b>	<b>.017679</b>
Success	<b>.000006</b>	<b>.000655</b>	<b>.001024</b>
Effort	<b>.000009</b>	<b>.000647</b>	<b>.001498</b>
Insecure	<b>.000049</b>	<b>.000682</b>	<b>.003483</b>
Photorealism	<b>.000005</b>	<b>.000974</b>	<b>.000686</b>
Confidence	<b>.000123</b>	<b>.001149</b>	<b>.008519</b>

### 5.4 User Perception Study


**Process.** Participants held the actual 3D-printed object (a silver boat and a clear whistle) while viewing side-by-side previews from each tool (Fig. 13) and photographs of that same object. In this setup, the physical print serves as the ground-truth baseline: all similarity ratings are made with the real 3D-printed artifact in hand, rather than relative to another preview. Each set included:

(a) Blender uniform, (b) Blender PBR, (c) Cura yellow, (d) *VisiPrint*, (e) photo of the printed object, and (f) Cura silver (boat) or Cura clear (whistle). Backgrounds were removed. Participants rated each image on 7-point Likert scales (1 = no resemblance, 7 = exact match) for three metrics: overall similarity, texture similarity, and slicing similarity. They also provided full rankings for each metric.

**Results.** We collected Likert-scale ratings and rankings for each model and metric. Figure 14 shows a visualization of both Likert-scale ratings and average rankings for the boat model. Notably, the profile of *VisiPrint* closely tracks that of a photo across most conditions. For slicing similarity, Cura previews sometimes matched or outperformed *VisiPrint* and the photo, which is expected given Cura's primary function as a slicing visualization tool. For overall and texture similarity, however, Cura consistently lagged behind *VisiPrint*. In a handful of cases, participants ranked *VisiPrint* ahead of the real photograph, suggesting that lighting conditions in photos may bias perception. The perception study was structured not merely to rank previews, but to examine when users treat *VisiPrint* previews as interchangeable with photographs of the actual object, and how this judgment shifts across filaments, materials, and baseline methods—providing insight into expectation-setting, rather than simply demonstrating that *VisiPrint* performs well.

*Overall Appearance.* Both whistle and boat conditions showed significant differences across methods (Whistle:  $\chi^2 = 36.53$ ,  $p < .001$ ,  $W = .459$ ; Boat:  $\chi^2 = 47.38$ ,  $p < .001$ ,  $W = .602$ ). For the whistle, *VisiPrint* was significantly higher than Cura-Yellow and comparable to both the photo and Cura-Clear, while all three outperformed the Blender baselines. For the boat, *VisiPrint* and the photo formed the top tier, rated higher than most baselines and not statistically differently from one another. Cura-Silver fell in between the top tier and the weaker baselines.

*Texture Similarity.* Texture ratings again showed significant differences (Whistle:  $\chi^2 = 46.42$ ,  $p < .001$ ,  $W = .602$ ; Boat:  $\chi^2 = 53.58$ ,  $p < .001$ ,  $W = .690$ ). For both objects, *VisiPrint* and the photo were rated highly and were not significantly different from each other. Both consistently outperformed Blender-Uniform, Blender-PBR, and Cura-Yellow. For the whistle, Cura-Clear also achieved strong



**Figure 13: Left: Perception study setup.** The participant holds the real 3D-printed object and compares it to six different previews on a screen, recording their results. **Right: Example row of options shown in the perception study.** From left to right we show the options participants were given: Blender uniform texture, Blender PBR texture, Cura default color, *VisiPrint*, a real-world photograph of the actual object and Cura “clear PLA”.

ratings, significantly better than the Blender baselines, while for the boat, Cura-Silver received moderate scores.

**Slicing Similarity.** Slicing comparisons also yielded significant effects (Whistle:  $\chi^2 = 44.46$ ,  $p < .001$ ,  $W = .534$ ; Boat:  $\chi^2 = 51.15$ ,  $p < .001$ ,  $W = .611$ ). For both objects, *VisiPrint* and the photo were not significantly different from each other and outperformed the Blender baselines. As expected, slicing-focused previews from Cura performed competitively in this dimension: Cura-Clear for the whistle and Cura-Silver for the boat matched or surpassed some baselines, reflecting their emphasis on visualizing slicing patterns.

**Summary.** Across both objects, *VisiPrint* achieved perceptual similarity on par with photographs and significantly outperformed slicer previews. Blender PBR improved shape and overall appearance but generally lagged behind *VisiPrint* and photos.

## 5.5 Post Survey

Participants completed a post-study survey covering: (1) preferred tool for photorealistic print preview (Blender, Cura, or *VisiPrint*) with justification; (2) overall *VisiPrint* experience (7-point Likert); (3) willingness to use *VisiPrint* again and reasons; (4) extent to which *VisiPrint* aided print preview (1–7) with explanation; (5) anticipated frequency of using *VisiPrint* for prototyping (1–7) with example cases; (6) perceived reliability of visualizations (1–7) with elaboration; (7) suggested improvements; (8) applications where *VisiPrint* would be most helpful; (9) reflections on what they learned; and (10) any additional comments.

Twelve out of the fifteen participants said they would prefer to use *VisiPrint* instead of just Cura alone or Blender alone for print visualization, and three participants said they would prefer to use just Cura. The mean experience score was 6.33 out of 7, the mean answer to how much *VisiPrint* helped was 5.33 out of 7, and the mean answer to how reliable the previews were was 6.27.

Fourteen out of the fifteen participants expressed that they would use *VisiPrint* again in the future, with a mean score of 4.60 when answering how often they would use it when doing 3D-print preview. Participants proposed a range of applications for *VisiPrint*, including product design where appearance matters (e.g., P4 said “*Product design focused applications, where visual appearance matters*”), exploring new materials (e.g., P2: “*It would be nice to visualize new materials on print jobs*”), and fine-tuning models late in the design process (e.g., P3: “*I would use it mostly at the end of a modeling process to fine-tune details and confirm material selection*”). Some

noted limited use for everyday prints (e.g., P6: “*Probably won’t need to preview appearances, I imagine product designers need it*”).


Participants found *VisiPrint* to be highly accurate and visually realistic (e.g., P7: “*I thought it was photorealistic*,” P8: “*They look like the material I’m trying to use*,” P9: “*The boat and print were very accurate to each other*”), and appreciated its ease of use (e.g., P1: “*It was very easy to use*,” P8: “*Cool stuff*”). However, several participants noted limitations such as the lack of a 3D preview or orbiting view (P3: “*Would be cool to rotate the object and view it from different angles*”), and a desire for more interactivity and integration (P3: “*I would like to provide a 3D model instead of a single image*”). Trust and understanding of the system also came up (P2: “*An understanding of the algorithm would better ground the trust I place in the visualizations*”).

We mitigate learning and order effects via a within-subjects design with counterbalanced condition order and standardized task scripts. Hawthorne effects remain possible in any tool-comparison study; we reduce them by avoiding feedback during tasks and by using identical prompts across conditions. Our findings center on appearance-related decisions rather than print success, consistent with *VisiPrint*’s intended use.

## 6 Adhering to Responsible Use of AI for Fabrication Previews

In our study (Section 5), participants’ confidence in *VisiPrint* previews was lower before seeing the real object in question than it was after inspecting the real physical prints, suggesting that they treated images as provisional rather than guarantees. Nonetheless, as preview realism improves, as with any pipeline that incorporates computer-vision material transfer or generative models, we acknowledge the risk of over-trusting the outputs of the model may grow. Prior work shows that explicitly communicating the limitations and fallibility of AI systems helps calibrate users’ expectations and reduces both over- and under-reliance on automated advice [41]. Following guidance on calibrated trust [9], we implemented the following safeguards into *VisiPrint*:

**System safeguards:** Unlike some generative models such as GPT-5, *VisiPrint* previews are tailored with a custom computer-vision algorithm, conditioned on slicer screenshots with geometry and layer cues preserved to keep outputs tied to toolpaths and prevent geometry hallucination; the UI warns if strong text prompting may reduce faithfulness; a single *Slicing Influence* slider exposes



**Figure 14: Perception study (boat). Top row: distribution of non-neutral Likert responses for overall, texture, and slicing similarity. Bottom row: average participant rankings. *VisiPrint* and photos were statistically indistinguishable and ranked highest, with Cura silver mid-tier and Blender/Cura yellow lowest.**

the trade-off between edges (lamination detail, vents/holes) and depth (global shape/shading); and all outputs are labeled “Simulated Preview” to set expectations.

**User guidelines:** We also provide the following guidelines for users, which are included in a pop-up to the User Interface: “Use *VisiPrint* only for visible appearance (sheen, hue tendencies, layer-line direction, seam/banding) not for printability, strength, tolerance, or safety, and always pair it with slicer checks and for high-stakes parts, a small test print. Know that *VisiPrint* is designed for previewing sliced 3D prints and tested primarily on FDM printing. And do not use *VisiPrint* previews as an indicator of likely print success for safety-critical parts (e.g., load-bearing or clinical), keeping records of slicer view, model/weights (incl. LoRA), conditioning scales, seed, and exemplar reference for auditability.” [81].

These safeguards and guidelines aim to keep trust calibrated: previews reveal slicing patterns and material appearance before printing, without being mistaken for evidence that a print will succeed, and our findings thus complement broader HCI work on expectation-setting and AI-assisted creativity, which emphasizes transparency and user control when integrating generative tools into designers’ workflows [41, 80].

## 7 Technical Evaluation

We evaluate *VisiPrint* using both a user study and technical benchmarks. Users provide the most meaningful signal of utility; quantitative metrics (PSNR, LPIPS, CLIP) are included as additional checks that compare each preview against a photograph of the corresponding actual 3D-printed object.

**Input preparation.** All methods used the same models and material exemplars. Aspect ratios were matched to avoid distortion; outputs were foreground-masked for quantitative metrics. Diffusion-based methods (*VisiPrint*, *ZeST*, IP-Adapter) used identical screen-shots, exemplars, defaults, and seeds. Cura previews hid non-print moves with a neutral theme. As using *GPT-5* is non-deterministic


even for a fixed prompt, we prompted the model twice with the same prompt to show variations across two distinct attempts.

### 7.1 Dataset

We built a dataset of twelve objects: 8 popular Thingiverse models plus 4 primitives (sphere, torus, cylinder, cube). All were printed in four filaments—solid pink, solid blue, clear white, and metallic silver, chosen to span translucency, specular, and matte vs. glossy finishes. Ground-truth photos were captured under diffuse lighting with fixed camera distance and neutral backgrounds, with the same setup as described in Section 4.1, and object orientation matched slicer views. This qualitative alignment was performed by manually rotating the physical object to visually resemble the slicer view, not via any automatic registration, but was held constant across baselines. Thus, for each filament, the exemplar sphere and the corresponding evaluation photos of printed objects share the same illumination environment. Printer/slicer settings were held constant within each color (full details in the Supplemental).

### 7.2 Qualitative Evaluation

We compared previews from: (i) ground-truth photos, (ii) *VisiPrint*, and baselines. All methods used the same models and exemplars. Figure 15 shows representative results. Across all models, only *VisiPrint* and Cura consistently preserved shape and slicing; *ZeST* introduced blocking artifacts, and *GPT-5* hallucinated geometry. On specular materials (like the whistle and octopus), *VisiPrint*, *ZeST*, and *GPT-5* captured shininess, but only *VisiPrint* and Cura preserved slicing patterns. Unlike *GPT-5*, *VisiPrint* and *ZeST* support arbitrary aspect ratios without stretching, and are geometry-aware and deterministic. Overall, *VisiPrint* qualitatively best balances geometry, slicing, and material (Table 2).



**Figure 15:** Left: ground-truth photographs of 3D printed objects from Thingiverse [22, 38, 45, 55, 57], with sliced geometries and material samples. Right: previews from *VisiPrint* and baselines. *VisiPrint* best combines material fidelity, slicing, and geometry.

### 7.3 Quantitative Evaluation

We compared *VisiPrint* against *ZeST*, *IP-Adapter*, *MatSwap* [50], and Single-Image SVBRDF estimation [17], using their official code releases, with results shown in Table 3. Evaluation metrics include PSNR [33], LPIPS (VGG) [87], and CLIP similarity [62], computed against photographs of real 3D-printed objects.

For the SVBRDF baseline, we used “sliced” STLs with the ground-truth slicing patterns extracted from GCode using *GCode2STL* [56] to ensure a fair geometric comparison with our method. We estimated material maps (diffuse albedo, specular albedo, roughness, normals) from each exemplar using [17]. While single-image SVBRDF capture assumes planar material samples, we use the same sphere exemplar images as in other baselines to keep the inputs consistent across methods. We crop to the central 50% of the image

**Table 2: Feature support across different tools.**

	<i>VisiPrint</i>	<i>ZeST/IP-Adapter</i>	<i>Cura</i>	<i>GPT-5</i>
<b>Geometry</b>	Yes	Yes	Yes	No
<b>Slicing</b>	Yes	No	Yes	No
<b>Material</b>	Yes	Yes	No	Yes

**Table 3: Quantitative comparison of *VisiPrint* with mean scores and standard deviations of different baselines (lower LPIPS is better; higher CLIP is better).**

Method	PSNR $\uparrow$	LPIPS <sub>VGG</sub> $\downarrow$	CLIP $\uparrow$
<i>VisiPrint</i> (Ours)	15.587 $\pm$ 3.470	0.278 $\pm$ 0.078	0.911 $\pm$ 0.040
<i>ZeST</i>	14.567 $\pm$ 3.165	0.285 $\pm$ 0.074	0.887 $\pm$ 0.043
<i>IP-Adapter</i>	12.446 $\pm$ 2.532	0.306 $\pm$ 0.070	0.823 $\pm$ 0.056
<i>MatSwap</i>	11.506 $\pm$ 2.220	0.298 $\pm$ 0.080	0.865 $\pm$ 0.041
<i>SVBRDF</i>	11.518 $\pm$ 3.226	0.324 $\pm$ 0.092	0.887 $\pm$ 0.043


to prevent background whitespace from appearing in the material maps, which reduces the amount of curvature in the image. However, as the surfaces are not flat, the resulting maps may be suboptimal. See the Supplemental for preprocessing and reproducibility details as well as a discussion of how the material images were adapted for use with the SVBRDF code, including material maps and details on alignment results.

Camera elevation and azimuth were manually specified per object to closely match the ground-truth photograph viewpoint. Ambient and directional light intensities were calibrated once by minimizing RGB error on a held-out reference photograph, then fixed for all renders. While *MatCap* is conceptually similar—mapping a photographed sphere to arbitrary geometry—we do not include *MatCap* as a baseline in Table 3 because there is no standard, open-source implementation that can be applied directly to slicer screenshots and our dataset. Backgrounds were removed and metrics averaged across 48 object–material pairs.

*VisiPrint* achieved the highest PSNR (15.587), LPIPS (.278), and CLIP (.911), indicating that its performance quantitatively exceeds the closest baselines in pixel-level fidelity, perceptual similarity, and semantic alignment with the ground-truth appearance. This is a demanding benchmark, since we compare against real-world photographs whose lighting and other capture conditions can influence scores, and the CLIP differences are modest; we do not claim statistical significance. While the quantitative results highlight *VisiPrint*’s capabilities, we treat these metrics as complementary to—but less decisive than—the qualitative user study comparisons.

### 7.4 Ablation Study

We evaluate the contribution of each component in the *VisiPrint* pipeline across 12 target objects and 4 material exemplars (48 test pairs), with results shown in Fig. 16. Starting from a depth-only



**Figure 16: Ablation study results. Top: Component ablation showing gains from adding Canny conditioning, color harmonization, and BLIP prompts to a depth-only baseline; color harmonization yields the largest improvement. Bottom: Sweeps over Canny weight (left) and LoRA scale (right). Higher Canny weights improve semantic consistency (CLIP ↑), while LoRA fine-tuning provides minor perceptual benefits.  $y$ -axes are zoomed to data range; arrows indicate preferred directions. Error bars show 95% confidence intervals.**

*ControlNet* baseline ( $w_d=0.8$ ,  $w_c=0$  for Depth and Canny weights), adding Canny conditioning ( $w_c=0.8$ ), color harmonization, and BLIP prompts produces progressively better results. Color harmonization delivers the largest improvement, increasing PSNR by 1.93 and CLIP by 0.018 relative to the Canny-only configuration, and the full pipeline achieves the strongest overall performance (PSNR = 14.99, CLIP = 0.912). To assess sensitivity to Canny conditioning, we sweep  $w_c \in \{0.0, 0.4, 0.6, 0.8, 1.0\}$  while holding  $w_d=0.8$ . Increasing the Canny weight improves semantic consistency (CLIP) with minimal impact on PSNR or LPIPS, with  $w_c=0.8$  offering a good balance of edge fidelity and generative flexibility. Finally, varying the LoRA finetuned-weights scale  $\{0, 0.2, 0.5, 0.8, 1.0\}$  produces modest perceptual gains (LPIPS from 0.715 to 0.712) while leaving PSNR and CLIP largely unchanged, indicating that LoRA contributes incremental refinement atop an already capable base model.

## 8 Limitations

While *VisiPrint* improves the visual fidelity of 3D print previews, several limitations remain. First, the method relies on monocular depth estimation and Canny edges extracted from slicer screenshots; these can introduce artifacts in hollow or detailed regions. Second, preview quality depends on the exemplar photo: poor lighting, low contrast, or misalignment can reduce realism, particularly for user-provided inputs. Our photo-slicer pose alignment is manual rather than optimization-based, reusing a visually matched pose across all methods instead of performing automatic camera registration.

As with *MatCap* previews and exemplar-driven methods such as *ZeST* and *MatSwap*, *VisiPrint* inherits the exemplar’s lighting and does not support arbitrary relighting. Previews may therefore diverge under different illumination, though users can provide

exemplars captured in alternate environments. The system thus offers expectation-setting under a specific lighting setup rather than the relighting flexibility of physically based rendering.

We intentionally favor an exemplar-driven pipeline over SVBRDF or inverse-rendering methods, which require controlled capture, calibrated cameras, and dedicated renderers. *VisiPrint* instead operates from a single uncalibrated photo and a slicer screenshot, trading physical interpretability for speed and accessibility while maintaining competitive fidelity relative to *MatSwap* and single-image SVBRDF baselines. *VisiPrint* is currently optimized for FDM printing with standard PLA filaments; materials with complex optical properties (metallic, translucent, multicolor) may not be fully captured, and we have not evaluated performance on SLA/resin prints. Importantly, the method does not model physical feasibility or support-structure artifacts and is not a substitute for slicer validation. Instead, it complements existing tools by offering a visual approximation of expected appearance.

## 9 Conclusion

We introduced *VisiPrint*, a photorealistic preview tool for 3D printing that incorporates both material appearance and slicing structure, enabling users to visualize how prints will look before fabrication. Our method operates as a post-slicing step and is agnostic to printer or slicer type. By bridging the gap between input 3D models and the actual physical output, *VisiPrint* lays the groundwork for future intelligent fabrication tools that support faster, more informed, and aesthetics-driven prototyping.

## Acknowledgments

We thank Leonardo Hernández Cano for their support, and gratefully acknowledge the support of the MIT MAD (Morningside Academy for Design) Fellowship, and the MIT MathWorks Fellowship.

## References

- [1] 2013. <https://community.ultimaker.com/topic/3486-color-becomes-darker-in-some-layers/>. [Accessed 11-09-2025].
- [2] 2022. Tips and tricks for easy printing with transparent PLA. <https://www.bcn3d.com/tips-and-tricks-for-easy-printing-with-transparent-pla/>. [Accessed 11-09-2025].
- [3] 2025. *Seam painting*. [https://help.prusa3d.com/article/seam-painting\\_168620](https://help.prusa3d.com/article/seam-painting_168620)
- [4] 2025. *Ultimaker Cura 5.8 beta release notes*. <https://ultimaker.com/learn/ultimaker-cura-5-8-beta-release-notes/>
- [5] Miika Aittala, Tim Weyrich, and Jaakko Lehtinen. 2015. Two-shot SVBRDF capture for stationary materials. *ACM Trans. Graph.* 34, 4 (2015), 110–1.
- [6] Yuval Alaluf, Daniel Garibi, Or Patashnik, Hadar Averbuch-Elor, and Daniel Cohen-Or. 2024. Cross-image attention for zero-shot appearance transfer. In *ACM SIGGRAPH 2024 Conference Papers*. 1–12.
- [7] Celena Alcock, Nathaniel Hudson, and Parmit K Chilana. 2016. Barriers to using, customizing, and printing 3D designs on thingiverse. In *Proceedings of the 2016 ACM International Conference on Supporting Group Work*. 195–199.
- [8] Amal Alfaraj and Wei-Shao Lin. 2025. Color reproduction trueness of 3D-printed full-color dental casts with scans derived from an intraoral scanner. *Journal of Prosthodontics* 34, 2 (2025), 196–203.
- [9] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fournay, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [10] Oğuz Arslan, Artun Akdoğan, and Mustafa Doga Dogan. 2025. TinkerXR: In-Situ, Reality-Aware CAD and 3D Printing Interface for Novices. In *Proceedings of the ACM Symposium on Computational Fabrication*. 1–19.
- [11] Alan Brunton, Can Ates Arikian, Tejas Madan Tanksale, and Philipp Urban. 2018. 3D printing spatially varying color and translucency. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–13.
- [12] John Canny. 1986. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-8, 6 (1986), 679–698.
- [13] Ta-Ying Cheng, Prafull Sharma, Andrew Markham, Niki Trigoni, and Varun Jampani. 2024. ZeST: Zero-shot material transfer from a single image. In *European Conference on Computer Vision*. Springer, 370–386.
- [14] Clevertize Blog. 2025. How Pixar’s animation pipeline has evolved with AI. <https://clevertize.com/blog/how-pixars-animation-pipeline-has-evolved-with-ai/>.
- [15] Jorge Condor, Michal Piovarci, Bernd Bickel, and Piotr Didyk. 2023. Gloss-aware color correction for 3D printing. In *ACM SIGGRAPH 2023 Conference Proceedings*. 1–11.
- [16] CreativeTools. 2015. 3DBenchy. <https://www.thingiverse.com/thing:763622>. Accessed: 2026-01-23.
- [17] Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. 2018. Single-image SVBRDF capture with a rendering-aware deep network. *ACM Transactions on Graphics (ToG)* 37, 4 (2018), 1–15.
- [18] Valentin Deschaintre, George Drettakis, and Adrien Bousseau. 2020. Guided fine-tuning for large-scale material transfer. *Computer Graphics Forum* 39, 4 (2020), 91–105.
- [19] Mustafa Doga Dogan, Steven Vidal Acevedo Colon, Varnika Sinha, Kaan Akşit, and Stefanie Mueller. 2021. SensiCut: Material-aware laser cutting using speckle sensing and deep learning. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. 24–38.
- [20] Mustafa Doga Dogan, Faraz Faruqi, Andrew Day Churchill, Kenneth Friedman, Leon Cheng, Sriram Subramanian, and Stefanie Mueller. 2020. G-ID: Identifying 3D prints using slicing parameters. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. ACM, 1–13.
- [21] Kyriakos Efsthathiou and Anthony McGarry. 2023. 3D printed cosmetic covers for lower limb prosthetics. *Canadian Prosthetics & Orthotics Journal* 6, 2 (2023), 42176.
- [22] eggnote. 2017. Curved honeycomb vase. <https://www.thingiverse.com/thing:2376777>. Accessed: 2026-01-23.
- [23] Willemijn Elkhuisen, Tessa Essers, Yu Song, Jo Geraedts, Clemens Weijkamp, Joris Dik, and Sylvia Pont. 2019. Gloss, color, and topography scanning for reproducing a painting’s appearance using 3D printing. *Journal on Computing and Cultural Heritage (JOCCCH)* 12, 4 (2019), 1–22.
- [24] Faraz Faruqi, Kenneth Friedman, Leon Cheng, Michael Wessely, Sriram Subramanian, and Stefanie Mueller. 2021. SliceHub: Augmenting shared 3D model repositories with slicing results for 3D printing. *arXiv preprints arXiv: 2109.14722* (2021).
- [25] Jiri Filip and Martin Vitek. 2024. Psychophysical insights into anisotropic highlights of 3D printed objects. *Workshop on Material Appearance Modeling, MANER Conference* (2024).
- [26] Manuela Galati, Paolo Minetola, Giovanni Marchiandi, Eleonora Atzeni, Flaviana Calignano, Alessandro Salmi, and Luca Iuliano. 2019. A methodology for evaluating the aesthetic quality of 3D printed parts. *Procedia CIRP* 79 (2019), 95–100.
- [27] Duan Gao, Xiao Li, Yue Dong, Pieter Peers, Kun Xu, and Xin Tong. 2019. Deep inverse rendering for high-resolution SVBRDF estimation from an arbitrary number of images. *ACM Trans. Graph.* 38, 4 (2019), 134–1.
- [28] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2016. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2414–2423.
- [29] Elizabeth George, Peter Liacouras, Frank J Rybicki, and Dimitrios Mitsouras. 2017. Measuring and establishing the accuracy and reproducibility of 3D printed medical models. *Radiographics* 37, 5 (2017), 1424–1450.
- [30] Sandra G Hart. 2006. NASA-task load index (NASA-TLX); 20 years later. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 50. Sage publications, Sage CA: Los Angeles, CA, 904–908.
- [31] Jess Hartzher-O’Brien, Jeremy Evers, and Erik Tempelman. 2019. Surface roughness of 3D printed materials: Comparing physical measurements and human perception. *Materials Today Communications* 19 (2019), 300–305.
- [32] Raed Hlayhel, Simon Hviid Del Pin, Takahiko Horiuchi, Sony George, Jon Yingve Hardeberg, and Aditya Suneel Sole. 2025. Naturalness perception of 3D prints with highlighted features. *Journal of Perceptual Imaging* 8 (2025), 1–14.
- [33] Alain Hore and Djemel Ziou. 2010. Image quality metrics: PSNR vs. SSIM. In *2010 20th International Conference on Pattern Recognition*. IEEE, 2366–2369.
- [34] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- [35] Syed Fouzan Iftekar, Abdul Aabid, Adibah Amir, and Muneer Baig. 2023. Advancements and limitations in 3D printing materials and technologies: a critical review. *Polymers* 15, 11 (2023), 2519.
- [36] Lan Jiang, Guihua Cui, Manuel Melgosa, Kaida Xiao, and Suchitra Suceprasana. 2021. Color-difference evaluation for 3D objects. *Optics Express* 29, 15 (2021), 24237–24254.
- [37] Yongcheng Jing, Yezhou Yang, Zunlei Feng, Jingwen Ye, Yizhou Yu, and Mingli Song. 2019. Neural style transfer: A review. *IEEE Transactions on Visualization and Computer Graphics* 26, 11 (2019), 3365–3385.
- [38] jzisa. 2015. V29. <https://www.thingiverse.com/thing:1179160>. Accessed: 2026-01-23.
- [39] Hyunyoung Kim, Aluna Everitt, Carlos Tejada, Mengyu Zhong, and Daniel Ashbrook. 2021. MorphesPlug: A toolkit for prototyping shape-changing interfaces. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [40] Jeeun Kim, Anhong Guo, Tom Yeh, Scott E Hudson, and Jennifer Mankoff. 2017. Understanding uncertainty in measurement and accommodating its impact in 3D modeling and printing. In *Proceedings of the 2017 Conference on Designing Interactive Systems*. 1067–1078.
- [41] Rafal Kocielnik, Saleema Amershi, and Paul N Bennett. 2019. Will you accept an imperfect AI? Exploring designs for adjusting end-user expectations of AI systems. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [42] Peter Kocsis, Vincent Sitzmann, and Matthias Nießner. 2024. Intrinsic image diffusion for indoor single-view material estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5198–5208.
- [43] Nahyun Kwon, Tong Steven Sun, Yuyang Gao, Liang Zhao, Xu Wang, Jeeun Kim, and Sungsoo Ray Hong. 2024. 3DPFlix: Improving remote novices’ 3D printing troubleshooting through human-AI collaboration design. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–33.
- [44] Seulhee Kwon and Dongwook Hwang. 2025. Understanding and Resolving 3D Printing Challenges: A Systematic Literature Review. *Processes* 13, 6 (2025), 1772.
- [45] Lalo\_Solo. 2017. USB SD and MicroSD holder for wide USB sticks. <https://www.thingiverse.com/thing:2637487>. Accessed: 2026-01-23.
- [46] Beichen Li, Yiwei Hu, Paul Guerrero, Milos Hasan, Liang Shi, Valentin Deschaintre, and Wojciech Matusik. 2024. Procedural material generation with reinforcement learning. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–14.
- [47] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*. PMLR, 12888–12900.
- [48] Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. 2017. Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. *ACM Transactions on Graphics (ToG)* 36, 4 (2017), 1–11.
- [49] Marco Livesu, Stefano Ellero, Jonas Martinez, Sylvain Lefebvre, and Marco Attene. 2017. From 3D models to 3D prints: An overview of the processing pipeline. *Computer Graphics Forum* 36, 2 (2017), 537–564.

- [50] I. Lopes, V. Deschaintre, Y. Hold-Geoffroy, and Raoul de Charette. 2025. MatSwap: Light-aware material transfers in images. *Computer Graphics Forum* 44, 4 (2025), e70168. <https://doi.org/10.1111/cgf.70168>
- [51] Ivan Lopes, Fabio Pizzati, and Raoul de Charette. 2024. Material Palette: Extraction of materials from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4379–4388.
- [52] Luo M. Ronnier, Guihua Cui, and Bryan Rigg. 2001. The development of the CIE 2000 colour-difference formula: CIEDE2000. *Color Research & Application: Endorsed by Inter-Society Color Council, The Colour Group (Great Britain), Canadian Society for Color, Color Science Association of Japan, Dutch Society for the Study of Color, The Swedish Colour Centre Foundation, Colour Society of Australia, Centre Français de la Couleur* 26, 5 (2001), 340–350.
- [53] Zehua Ma, Hang Zhou, and Weiming Zhang. 2023. AnisoTag: 3D printed tag on 2D surface via reflection anisotropy. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [54] Dario Magri, Paolo Cignoni, Marco Tarini, et al. 2016. Anisotropic MatCap: Easy capture and reproduction of anisotropic materials. In *Smart Tools and Apps for Graphics-Eurographics Italian Chapter Conference*. The Eurographics Association.
- [55] MakerBot. 2017. Headphone Stand. <https://www.thingiverse.com/thing:2050885>. Accessed: 2026-01-23.
- [56] Matter and Form. 2018. GCode2STL. <https://github.com/Matter-and-Form/GCode2STL>. GCode parser and STL generator for 3D printing.
- [57] McGybeer. 2019. Cute Mini Octopus. <https://www.thingiverse.com/thing:3495390>. Accessed: 2026-01-23.
- [58] Ilan E Moyer, Samuelle Bourgault, Devon Frost, and Jennifer Jacobs. 2024. Don't mesh around: Streamlining manual-digital fabrication workflows with domain-specific 3D scanning. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–16.
- [59] Stefanie Mueller, Anna Seufert, Huaishu Peng, Robert Kovacs, Kevin Reuss, François Guimbretière, and Patrick Baudisch. 2019. FormFab: Continuous interactive fabrication. In *Proceedings of the Thirteenth International Conference on Tangible, Embedded, and Embodied Interaction*. 315–323.
- [60] Huaishu Peng, Jimmy Briggs, Cheng-Yao Wang, Kevin Guo, Joseph Kider, Stefanie Mueller, Patrick Baudisch, and François Guimbretière. 2018. RoMA: Interactive fabrication with augmented reality and a robotic 3D printer. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [61] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. 2024. SDXL: Improving latent diffusion models for high-resolution image synthesis. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*.
- [62] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, Vol. 139. PMLR, 8748–8763.
- [63] Nerrolyn Ramstrand, Maria Riveiro, Lars Eriksson, and Michael Ceder. 2025. Effects of conventional versus 3D-printed cosmetic covers on user satisfaction and psychosocial well-being in lower limb prostheses users: A randomised crossover trial. *Journal of Rehabilitation and Assistive Technologies Engineering* 12 (2025), 2055683251330996.
- [64] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. 2021. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 12179–12188.
- [65] Erik Reinhard, Michael Ashikhmin, Bruce Gooch, and Peter Shirley. 2001. Color transfer between images. *IEEE Computer Graphics and Applications* 21, 5 (2001), 34–41.
- [66] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10684–10695.
- [67] Thijs Jan Roumen, Willi Müller, and Patrick Baudisch. 2018. Grafter: Remixing 3D-printed machines. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [68] Rosemary Seelaus and Gerald T Grant. 2024. 3D printing and digital design for maxillofacial prosthetics. In *3D Printing at Hospitals and Medical Centers: A Practical Guide for Medical Professionals*. Springer, 165–184.
- [69] Prafull Sharma, Varun Jampani, Yuanzhen Li, Xuhui Jia, Dmitry Lagun, Fredo Durand, Bill Freeman, and Mark Matthews. 2024. Alchemist: Parametric control of material properties with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 24130–24141.
- [70] Pitchaya Sithi-Amorn, Javier E Ramos, Yuwang Wang, Joyce Kwan, Justin Lan, Wenshou Wang, and Wojciech Matusik. 2015. MultiFab: A machine vision assisted platform for multi-material 3D printing. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–11.
- [71] M Šljivic, A Pavlovic, Milija Krašnik, and Jovica Ilić. 2019. Comparing the accuracy of 3D slicer software in printed enduse parts. In *IOP conference series: materials science and engineering*, Vol. 659. IOP Publishing, 012082.
- [72] Evgeny Stemasov, Simon Demharter, Max Rädler, Jan Gugenheimer, and Enrico Rukzio. 2024. pARam: Leveraging parametric design in extended reality to support the personalization of artifacts for personal fabrication. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–23.
- [73] Evgeny Stemasov, Tobias Wagner, Jan Gugenheimer, and Enrico Rukzio. 2020. Mix&Match: Towards omitting modelling through in-situ remixing of model repository artifacts in mixed reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [74] Stratasys. 2025. Advanced Multi-Color and Full-Color 3D Printing. <https://www.stratasys.com/en/resources/blog/full-color-3d-printing/>.
- [75] Blair Subbaraman, Nathaneal Bursch, and Nadya Peek. 2025. It's not the shape, it's the settings: Tools for exploring, documenting, and sharing physical fabrication parameters in 3D printing. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [76] Daniel Subias and Manuel Lagunas. 2023. In-the-wild material appearance editing using perceptual attributes. *Computer Graphics Forum* 42, 2 (2023), 333–345.
- [77] TechRadar. 2025. Adobe Max London 2025 Live. <https://www.techradar.com/news/live/adobe-max-london-2025-live>.
- [78] Alexander Teibrich, Stefanie Mueller, François Guimbretière, Robert Kovacs, Stefan Neubert, and Patrick Baudisch. 2015. Patching physical objects. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. 83–91.
- [79] Lezhong Wang, Duc Minh Tran, Ruiqi Cui, Thomson TG, Anders Bjorholm Dahl, Siavash Arjomand Bigdeli, Jeppe Revall Frisvad, and Manmohan Chandraker. 2025. Materialist: Physically based editing using single-image inverse rendering. *arXiv preprint arXiv:2501.03717* (2025).
- [80] Wen-Fan Wang, Chien-Ting Lu, Nil Ponsa i Campanya, Bing-Yu Chen, and Mike Y Chen. 2025. Aldeation: Designing a human-AI collaborative ideation system for concept designers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–28.
- [81] Yue Wang and Sai Ho Chung. 2022. Artificial intelligence in safety-critical systems: A systematic review. *Industrial Management & Data Systems* 122, 2 (2022), 442–470.
- [82] Karl DD Willis, Cheng Xu, Kuan-Ju Wu, Golan Levin, and Mark D Gross. 2010. Interactive fabrication: New interfaces for digital fabrication. In *Proceedings of the Fifth International Conference on Tangible, Embedded, and Embodied Interaction*. 69–72.
- [83] Jiangping Yuan, Guangxue Chen, Hua Li, Hartmut Prautzsch, and Kaida Xiao. 2021. Accurate and computational: A review of color reproduction in full-color 3D printing. *Materials & Design* 209 (2021), 109943.
- [84] Steve Zelinka, Hui Fang, Michael Garland, and John C Hart. 2005. Interactive material replacement in photographs. In *Proceedings of Graphics interface 2005*. 227–232.
- [85] Lianghao Zhang, Fangzhou Gao, Li Wang, Minjing Yu, Jiamin Cheng, and Jiawan Zhang. 2023. Deep SVBRDF estimation from single image under learned planar lighting. In *ACM SIGGRAPH 2023 conference proceedings*. 1–11.
- [86] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 3813–3824.
- [87] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 586–595.
- [88] Xiaoting Zhang, Xinyi Le, Athina Panotopoulou, Emily Whiting, and Charlie CL Wang. 2015. Perceptual models of preference in 3D printing direction. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 1–12.
- [89] Xilong Zhou, Miloš Hašan, Valentin Deschaintre, Paul Guerrero, Kalyan Sunkavalli, and Nima Khademi Kalantari. 2023. A semi-procedural convolutional material prior. *Computer Graphics Forum* 42, 6 (2023), e14781.
- [90] Liyuan Zhu, Shengqu Cai, Shengyu Huang, Gordon Wetzstein, Naji Khosravan, and Iro Armeni. 2025. ReStyle3D: Scene-level appearance transfer with semantic correspondences. *CoRR* (2025).