# TactStyle: Generating Tactile Textures with Generative AI for Digital Fabrication


Faraz Faruqi
MIT CSAIL
Cambridge, MA, USA
ffaruqi@mit.edu

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

Jaskaran Singh Walia
Vellore Institute of Technology
Chennai, India
karanwalia2k3@gmail.com

Yunyi Zhu
MIT CSAIL
Cambridge, MA, USA
yunyizhu@mit.edu

Shuyue Feng
Zhejiang University
Hangzhou, China
shuyuefeng@zju.edu.cn

Donald Degraen*
HIT Lab NZ, University of Canterbury
Christchurch, New Zealand
donald.degraen@canterbury.ac.nz

Stefanie Mueller
MIT CSAIL
Cambridge, MA, USA
stefanie.mueller@mit.edu


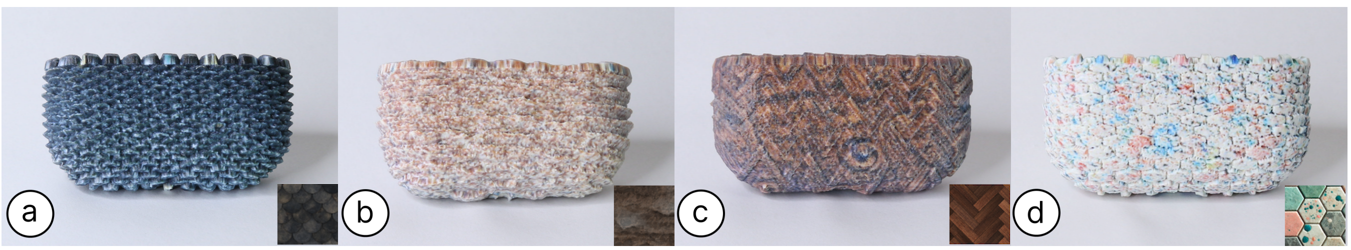


**Figure 1: TactStyle allows creators to stylize 3D models with image input while incorporating the tactile properties of the texture in addition to its color. Here, we show four different textures applied to the same 3D model, an Airpods cover, with the image stylization prompt shown on the bottom right. The different textures used are: a) round stone roof, b) layered brown rock, c) herringbone wood, and d) colorful hexagonal tiles.**


## Abstract

Recent work in Generative AI enables the stylization of 3D models based on image prompts. However, these methods do not incorporate tactile information, leading to designs that lack the expected tactile properties. We present TactStyle, a system that allows creators to stylize 3D models with images while incorporating the expected tactile properties. TactStyle accomplishes this using a modified image-generation model fine-tuned to generate heightfields for given surface textures. By optimizing 3D model surfaces to embody a generated texture, TactStyle creates models that match the desired style and replicate the tactile experience. We utilize a large-scale dataset of textures to train our texture generation model. In a psychophysical experiment, we evaluate the tactile qualities of a set of 3D-printed original textures and TactStyle's generated textures. Our results show that TactStyle successfully generates a wide range of tactile features from a single image input, enabling a novel approach to haptic design.


## CCS Concepts

• **Human-centered computing → Human computer interaction (HCI)**.


## Keywords

Personal Fabrication; Digital Fabrication; 3D Printing; Generative AI.

**ACM Reference Format:**
Faraz Faruqi, Maxine Perroni-Scharf, Jaskaran Singh Walia, Yunyi Zhu, Shuyue Feng, Donald Degraen, and Stefanie Mueller. 2025. TactStyle: Generating Tactile Textures with Generative AI for Digital Fabrication. In *CHI Conference on Human Factors in Computing Systems (CHI '25), April 26–May 01, 2025, Yokohama, Japan.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3706598.3713740


## 1 Introduction

With the growing popularity of 3D printing research within HCI [35], there is also increasing interest in developing tools that enable users to customize 3D models. Open-source repositories, such as Thingiverse [4], are a useful resource for ready-to-print 3D models. However, their customization is limited to changing predefined parameters [1]. Recent advances in Generative AI allow users to more freely customize their 3D models using text prompts or images as user-provided style descriptions [11, 34]. However, these existing frameworks for stylizing 3D models primarily focus

*Also with HCI Group, University of Duisburg-Essen.

on modifying the model to match a desired visual appearance as described via the provided text or image prompt [43].

One underexplored area of model customization is texture—specifically, the 'tactile feedback' of printed structures, such as whether a surface feels smooth, rough, or potentially reminiscent of materials like wood grain or stone. Our ability to sense these textures through touch plays a crucial role in our interactions with the physical world, shaping not only how we perceive and manipulate objects but also influencing our emotional and cognitive responses to them [16]. Therefore, augmenting printed structures with appropriate tactile properties can enrich interaction with physical objects, especially when mimicking materials that differ from the 3D printing material.

Recent advances in computer vision have proposed methods for capturing high-fidelity visual properties from images, enabling digital replication of textures from real-world surfaces [21, 56]. However, these techniques are limited to digital replication, as highlighted in TextureDreamer [56], where they do not optimize normal maps to avoid details that are inconsistent with the target mesh. In the field of digital fabrication, researchers have proposed techniques to capture surface microgeometry [7] as a heightfield and use this data to replicate the texture using fabrication methods such as 3D printing. However, they currently require sophisticated equipment such as photometric sensing techniques [25] to capture the surface microgeometry of each texture, limiting their usability. Thus, image-based replication is currently limited to replicating visual elements of textures, and creators are currently limited to replicating textures through associated digital surface microgeometry data. We hypothesize that by learning a correlation between a texture's visual image, and its heightfield (surface microgeometry), we can replicate the tactile properties of textures, directly from an input image.

We present TactStyle, a system that allows creators to stylize 3D models with texture images while incorporating the expected tactile properties. TactStyle accomplishes this by separating the visual and the geometry stylization, and augmenting the process with a novel geometry stylization module that replicates the tactile properties of textures based on user input. The novel geometry stylization module uses a fine-tuned variational autoencoder (VAE) [27], that translates the user provided visual image of a texture into a surface microgeometry or heightfield. The model then uses this heightfield to manipulate the geometry to create the tactile properties on the 3D model surface. Separately, the visual appearance of the 3D model is optimized with a method [11] that has been shown to accurately replicate visual qualities. Thus, by optimizing 3D model surfaces to embody both the color and tactile properties of a given input image, TactStyle allows creators to generate stylized models that not only visually match the desired style but also replicate the tactile experience.

# 2 Related Work

To situate our question and findings, we draw upon previous research in personalizing open-source designs, 3D printed haptics, and tactile surface reconstruction to develop our proposed system.

## 2.1 Personalizing Open-Source Designs

As 3D printing has become the preferred digital fabrication tool for both expert and amateur makers alike [10, 29], makers have increasingly shared their 3D models and designs through online open repositories, such as Thingiverse. Alcock et al. [1] propose that such repositories serve as ideal training platforms for novice makers. However, several challenges remain in allowing creators to personalized open-source designs. Numerous studies have investigated novice makers [22, 37] in the process of modifying shared 3D models. These studies consistently find that users face significant challenges when attempting to modify existing designs. Editing 3D printable meshes requires advanced expertise of computer-aided design workflows [46] which are often incompatible with the skill levels of novice users. Moreover, as highlighted by Oehlberg et al. [38], even when designs are customizable, they often fall short of the scope of modifications makers desire. For instance, customizing 3D models to support varying tactile properties is rarely supported.

## 2.2 3D Printed Haptics & Tactile Surfaces

3D printing haptics has been an increasing area of interest in HCI. Design tools have been developed that provide users with the ability to customize their designs for a variety of haptic properties, ranging from desired heat dissipation (Thermal Comfort [58]), to customizable stiffness (X-Bridges [48]) and softness (OmniSoft [26]), and personalized force feedback (Shape-Haptics [59]).

To 3D print surfaces with desired tactile properties, researchers developed a range of new 3D printing techniques. For instance, 3D printed Hair [31] creates hair-like structures and bristles through FDM 3D-printing by exploiting the stringing phenomena inherent to the 3D-printing process. Cillia [40] creates high-resolution tactile surfaces by modifying the input to the SLA printing process. Such hair-like structures have been shown to influence texture perception, especially when combined with visual augmentation [8]. On-the-Fly Fine Texture 3D Printing [55], Thickness Control [49] and ExtruderTurtle [42] modify the G-code underlying FDM 3D printers to create varying types of tactile surface textures. Another approach to create 3D printed tactile surfaces is to print mechanical structures that can transform into various surface structures. For instance, Metamaterial Textures [24] demonstrated how varying tactile surface can be created within a single 3D print, and such structural approaches have been shown to directly influence perception during tactile fingertip exploration [13]. Most closely aligned with our approach is HapticPrint [50], which modifies the "feel" of 3D-printed objects by providing a tool to automatically generate heightfields by grayscaling raster images to create tactile textures on arbitrary 3D geometries. This approach gives an approximate texture from the image, which can then be applied to the 3D model. However, a grayscale image and a heightfield differ since grayscale only encodes luminance, not the local height of the surface. Since not all color changes in an image are related to changes in depth or height [56], taking a grayscale version would fail to distinguish between color patterns caused by surface detail and those caused by lighting. In this work, we modify an image-generation model to generate heightfield data, by fine-tuning a diffusion-based model on pairs of texture-heightfield data.

## 2.3 Reconstructing Tactile Surfaces

Researchers have also explored augmenting a 3D model's geometry to enable accurate tactile perceptions. Haptography [28] introduces an approach that uses sensors to capture the haptic properties of real objects and recreate them in virtual environments, while Metareality [13] designed adaptable metamaterial structures that can alter their hardness and roughness upon compression. Degraen et al [7] showed how a real-world texture's tactile properties can be replicated by using its microgeometry processed with a photometric sensing technique [25]. Existing Generative AI-based stylization methods such as Style2Fab [11] and Text2Mesh [34] allow users to stylize their 3D models based on text and image prompts. These methods perform iterative refinement of the mesh, making small changes on vertex and color channels of the 3D model and estimating its similarity to the goal text or image prompt provided by the user. However, since these methods are based on image-based losses, replicating the surface microgeometry becomes a challenge [12].

TactStyle extends this line of work by proposing a system that allows creators to stylize 3D models using images as input, optimizing not only the texture's appearance but also its expected tactile properties.

# 3 Formative Study

We hypothesize that current stylization frameworks that leverage latent representations [11] such as CLIP [44] are efficient at replicating the visual appearance of a texture but ineffective at replicating its tactile properties. We first test this hypothesis by performing a formative study. Although stylization frameworks allow both text and image-based stylization, we consider only image-based stylization for our experiments. This is because text-based stylization methods require a captioning technique to generate textual descriptions of textures, which may not express all its details. This limitation was also highlighted by TextureDreamer [56]. Thus, in this formative study, we focus on testing the stylization of a 3D model based on image prompts.

## 3.1 Dataset and Stylization Baseline

To investigate the accuracy of texture replication, we use a large-scale dataset of PBR (Physically Based Rendering) textures from CGAxis [3]. This dataset contains both visual and heightfield information about textures. We collect a total of 500 textures which contain textures for '*Parquets*', '*Wood*', '*Rocks*', '*Walls*', and '*Roofs*'. For each of these 500 textures, we take the visual texture and its associated heightfield as ground-truth pairs.

For the stylization framework, we use Style2Fab [11], which allows users to personalize 3D models based on text prompts. We modified Style2Fab's system to take image prompts instead of text by changing the hyperparameters in the stylization module.

## 3.2 Procedure

For consistency, we perform stylization of a single tile of size 5×5×1 $cm^3$. To create the ground truth set of textures, we apply the heightfield from our dataset on the tile surface following the technique from Degraen et al. [7]. We take 50 random textures from our dataset (10% of the dataset size) and stylize the tile with the texture image as the prompt. We subdivide the tile surface to 25k resolution for accurate texture generation and run the stylization process for 1500 iterations, as specified in Style2Fab [11]. We apply stylization to only one face of the tile, the same as that of the ground truth textures, and freeze the geometry on the remaining faces, retaining a flat surface. This allows for a consistent comparison. Stylization iteratively modifies the geometry and color channel of the 3D model, and using the CLIP loss to assess the stylization quality. At the end of the study, we have 50 modified 3D tiles created using 50 random textures from our dataset.

## 3.3 Results:

To quantitatively assess the fidelity of the stylized textures in replicating the ground-truth textures, we compare the Root Mean Square (RMS) values of the textures' heightfields as it has been shown to correlate to surface roughness [7]. We take the 50 heightfields associated with the texture images used to stylize the 3D tiles. To extract the heightfield from the stylized tile, we take the boolean difference of original unstylized tile, and then map the displacement of the modified vertices onto the grayscale range(0 - 255). The RMS values capture the overall surface variation, allowing us to evaluate the differences between the original textures and the stylized outputs.

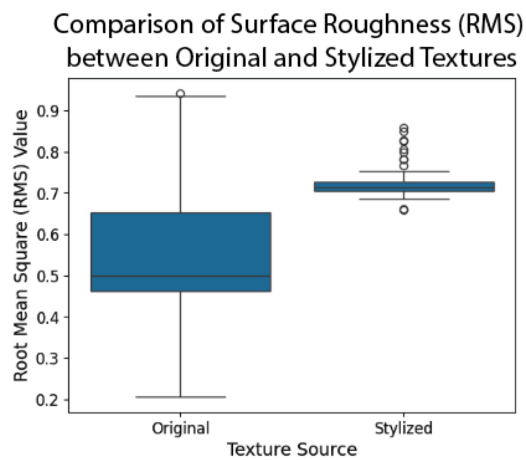


**Figure 2: The boxplot shows the distribution of Root Mean Square (RMS) values for original and stylized textures, representing surface roughness. The original textures exhibit a wider range of RMS values, indicating higher variability in surface roughness. In contrast, the stylized textures have consistently higher RMS values with less variability, indicating rougher and more uniform surfaces as a result of the stylization process.**

Figure 2 presents a boxplot comparing the RMS distributions for the original textures and their stylized counterparts. We observe that the stylized textures generally exhibit higher RMS values compared to the original textures. RMS values can be interpreted as a metric for surface roughness [7] suggesting that the stylization process results in rougher surfaces. Moreover, the RMS values for the original surface textures have a wider range of values, showing higher variability, whereas the stylized surfaces have lesser variability indicating more uniformity.

To determine the statistical significance of the observed differences, we perform a Welch's t-test between the RMS values of the original and stylized textures. The test reveals a statistically significant difference between the two groups ($t = 11.89$, $p < 0.0001$), indicating that the stylized textures have significantly different RMS values compared to the original heightfields.

This result suggests that Style2Fab and similar stylization strategies do not accurately modify the surface geometry to replicate a specific texture's heightfield. Further refinement in the stylization process could enhance the replication of texture variation for more accurate texture replication in 3D models. In the next sections, we present TactStyle, a system that allows creators to accurately replicate the tactile properties via a new geometry stylization approach.

## 4 System Overview

Prior work [11, 32, 34] has shown that Generative-AI based stylization methods closely approximate the user's style visually. However, our formative study found that such geometry modifications do not accurately replicate the desired texture represented by its surface microgeometry. We designed TactStyle to enable the replication of a surface's microgeometry, and by extension, its tactile properties.

TactStyle augments existing stylization methods with a new approach to modifying the geometry of 3D models that replicates the tactile properties of the texture described by the user. It accomplishes this by co-optimizing the geometry and the color channels separately. We call these two modules: (1) the color stylization module and (2) the geometry stylization module. The focus of this paper is on the geometry stylization module.

Our main challenge was to design a geometry stylization module that modifies a 3D model's geometry to replicate the tactile properties of a texture. We leverage the fact that heightfields can be represented as images, and thus TactStyle accomplishes this goal by fine-tuning an image generation model to generate heightfields based on visual images of a texture. This heightfield is then used to modify the 3D model's geometry using an approach based on UV mapping. Thus, the color and geometry stylization modules work in tandem, stylizing the color and geometry of the 3D model to replicate both the visual appearance and tactile properties of a given texture (Figure 3a).

## 5 Heightfield Generation Technique

In this section, we describe our novel heightfield generation model. This model takes a texture image as a prompt and generates the associated heightfield. For this purpose, we fine-tune a trained diffusion-based Image-to-Image model and integrate it into the TactStyle system. In the following subsections, we describe the modified architecture of the diffusion model and the dataset used to train and test the system.

### 5.1 Diffusion Model

We approach this problem as an image generation task and use a modified version of the Stable Diffusion model [45], a popular open-source image generation model. Specifically, we use an image-to-image generation model proposed in SDEdit [33]. This deep-learning model uses a diffusion model to synthesize new realistic images. Given an input image along with a user prompt in the form of text or image, SDEdit first adds noise to the input, then subsequently denoises the resulting image to generate a modified image based on the user prompt. At the core of this diffusion-based generative model is a variational autoencoder [27] (VAE), which encodes images into a latent representation, and decodes that latent representation into an image.

The VAE is trained to encode an image into a latent representation, a compact high-dimensional representation that can then be 'decoded' using another network called a 'Decoder' to generate another image. More details on the architecture and training approach are available in Meng et al. [33] and Kingma et al. [27]. Our goal with this model was to generate a heightfield given an image of a texture. Since heightfields are traditionally represented as grayscale images, we (1) modify the VAE architecture to generate representative grayscale images, i.e., heightfields, (2) and fine-tune the trained model on our texture image-heightfield pairs.

### 5.2 Modified Model Architecture

We use a trained open-source Image-to-Image Generation model available through the Diffusers library [52]. As described above, this model's essential component is the VAE, which encodes an image into a latent representation and then decodes it into another image. This VAE is structured to generate images in 3 (RGB) channels. We modify the architecture's decoder module by adding 4 additional layers to learn heightfield features and modify the final layer to output single-channel grayscale images. This approach was motivated by the fact that the pre-trained model was trained to generate colored images, and there are additional features that the model would need to learn to generate heightfield-specific features.

In fine-tuning our modified image generation model, our goal was both to maximize the similarity in intensity between the target and generated heightfield and minimize their perceptual difference. For comparing overall intensities, we use the Mean Squared Error Loss (MSE), a standard in regression and image generation tasks. For the perceptual similarity metric, we use the Structural Similarity Index Measure (SSIM) [53]. These two loss functions serve two different purposes. MSE calculates an average of per-pixel similarity that provides a guide towards a similar intensity in generated images. However, independent training with MSE does not generate high-quality heightfields because it assumes pixel-wise independence. For instance, blurred images can have a large perceptual difference but a small MSE loss. SSIM on the other hand, takes into account the luminance, contrast, and structure of the two images being compared, highlighting local structural differences. Thus, a combination of these two loss functions allows us to generate heightfields that are similar in both overall intensity (MSE) and local structural features (SSIM). In training our model, we use these both loss measures.

### 5.3 Training Methodology

The Variational Auto Encoder (VAE) is fine-tuned for generating accurate heightfields using the PBR Dataset consisting of texture

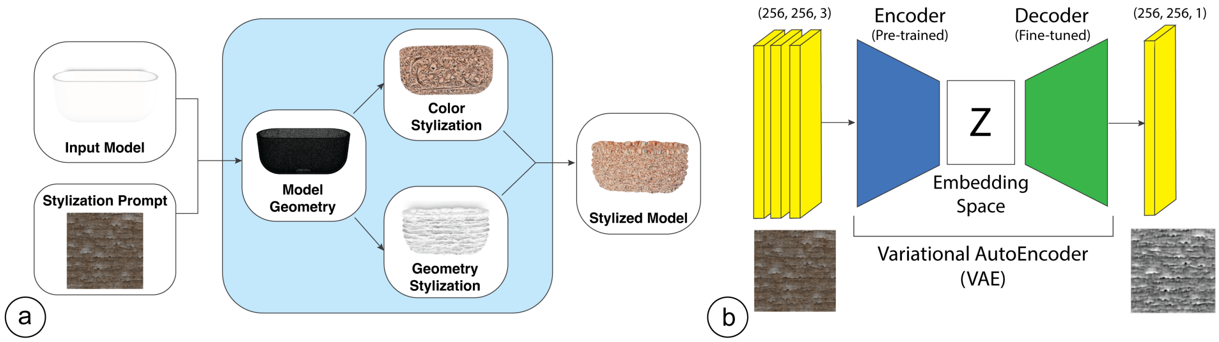


**Figure 3: TactStyle augments traditional 3D model stylization techniques by introducing a novel geometry stylization module that replicates the tactile properties of textures based on user input. (a) The system takes an input model and a stylization prompt (e.g., an image of a texture) and applies two separate stylization processes: (1) Color Stylization and (2) Geometry Stylization. The color stylization modifies the model's visual appearance, while the geometry stylization alters its surface to reflect tactile properties. The two modules operate in tandem, creating a stylized 3D model that replicates both the visual and tactile aspects of the texture. (b) The geometry stylization module uses a variational autoencoder (VAE) to generate heightfields from texture images, which are then applied to modify the model's surface geometry, enabling co-optimization of geometry and color for a unified tactile and visual experience.**

image-heightfield pairs. The associated heightfields serve as ground truth representations of the tactile features, and our model learns the correlation between visual appearance and tactile properties.

We fine-tune the model updating the decoder parameters over 60 epochs, using a batch size of 10 images with an RMSprop optimizer. We use a lower learning rate ($1e^{-5}$) for fine-tuning existing layers in the VAE model, and a higher learning rate ($1e^{-3}$) for the newer layers. This was done because the original layers are already trained and need small adjustments, whereas the newer layers are randomly initialized and require larger changes. Since our goal was to modify the 'Decoder' module of the VAE, we froze the weights for the encoder module, training the weights for only the decoder module.

### 5.4 Dataset

To train our model on realistic textures, we utilize the CGAxis repository [3], which contains a wide range of textures designed to provide accurate real-world simulations of materials in 3D environments. We collected 500 pairs of texture images and corresponding heightfields in 4k resolution. The dataset contains 5 different material types: '*Parquets*', '*Wood*', '*Rocks*', '*Walls*', and '*Roofs*', containing 100 textures each. This allows for a diverse set of textures to train our model. For each of these 500 textures, we collect the visual texture and its associated heightfield as ground-truth pairs. These heightfields represent the tactile features of the textures and are critical for learning the correlation between visual appearance and haptic properties.

We split our dataset into a train and test set, using a 90% - 10% split, resulting in 450 textures to train our model and 50 textures to test it. We also augment the train set by rotating each image-heightfield pair by 90 degrees three times, effectively generating four variations for each texture and resulting in a total of 1,800 textures in our train set. This augmentation allows us to increase the diversity of the data, providing a more comprehensive set of examples for training our model without introducing synthetic artifacts. This enables the model to learn more robust and invariant representations of visual and tactile features, improving its ability to generalize across different orientations of textures.

### 5.5 Texture Application

To apply our textures to 3D models, we apply the heightmap by displacing vertices along their normals based on the corresponding height values from a UV map normalized to fit the texture map, producing a texturized object ready for 3D printing. This process creates a final texturized object that is ready to be 3D printed, as shown in geometry stylization step of Figure 3a.

## 6 User Interface and Workflow

TactStyle has been implemented as a plugin for the open-source 3D design software tool Blender [5] to allow easy integration with makers' existing workflows. Figure 4 shows a view of the interface. To stylize a model with TactStyle, the user (1) loads their model, (2) uploads the image prompt of the desired texture, and (3) clicks the stylize button. TactStyle then processes the model and stylizes it using the integrated color and geometry simulation modules. The stylized model is rendered next to the original model, which the user can export for fabrication.

*6.0.1 Preprocessing.* Once the user has loaded an OBJ file of their 3D mesh into the plugin, the model is automatically pre-processed for stylization. The model is first standardized to a unit-sized cube for stylization. Next, we use Pymeshlab [36] to increase the model's resolution by subdivision to 25k faces following the standardization protocol from Style2Fab [11]. This enables accurate stylization of the model by increasing the number of vertices on the model, which are then modified both in color and geometry to approximate the style desired by the user.

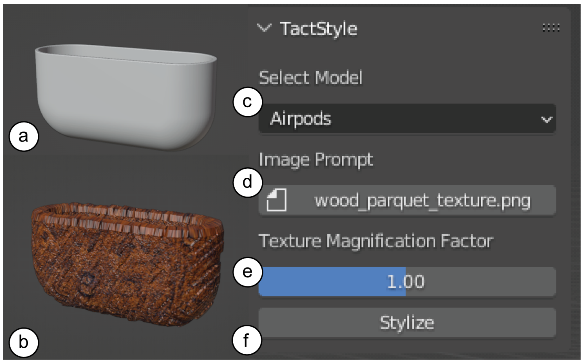


**Figure 4: TactStyle's user interface, implemented as a Blender plugin, allows users to load (a) original 3D model and (b) stylize with image prompts. In order to use TactStyle, the user (c) loads the model, (d) uploads an image of their desired texture, (e) optionally adjust the Texture Magnification Factor to control the level of height displacement applied on the 3D model. (f) Finally, the user clicks the "Stylize" button, which starts the stylization process using TactStyle's integrated color and geometry stylization modules.**

*6.0.2 Stylization.* TactStyle used two modules — the color stylization module and the geometry stylization module to optimize both visual and tactile properties. As shown in Figure 3, TactStyle uses Style2Fab [11] for iterative color optimization. Here the model's geometry is frozen, and the generative AI model modifies the color channels of the vertices to approximate the style in the image. Next, the geometry stylization module uses the modified image generation model to generate a heightfield using the texture image prompt provided by the user. This heighfield is applied on the model using the technique described in section 5.5. The completed model is rendered alongside the original model for review. Furthermore, the segmentation tool from Style2Fab [11] has been integrated into TactStyle. This allows the user to have multiple textures on the same model. For this, the user can segment the model through Style2Fab's segmentation, and then apply TactStyle on individual segments.

*6.0.3 Fine-Tuning and Export.* Users can iterate on this process and apply new styles using new image prompts as needed. Users can also optionally increase the amount of height displacement to magnify their texture by changing the 'Texture Magnification Factor' slider shown in Figure 4e which can exaggerate or diminish the texture applied. This factor is by default at 1.0, which corresponds to the value used in our study following the height displacement values from Degraen et al. [7].

Finally, the user can export the stylized model and fabricate it.

## 7 Technical Evaluation

To validate the effectiveness of our system, we conducted a quantitative evaluation. We evaluate TactStyle's performance on its ability to replicate the surface micro-geometry, represented by the ground truth heightfield. We evaluate TactStyle's results using two metrics: (1) the RMS Error, which represents the difference in surface roughness, (2) the Mean Squared Error (MSE) which calculates the average error in per-pixel intensity between the textures. In the following subsections, we discuss the results of the quantitative evaluation.

*7.0.1 Analyzing Root Mean Square Error:* We evaluate the Root Mean Square (RMS) values of the generated heightfields, which allow us to compare the overall surface roughness of the generated textures [7]. Figure 5a shows the comparison between the Original, Stylized, and TactStyle textures. The original textures exhibit a wide range of RMS values, reflecting the inherent variability in surface roughness across different textures.

To evaluate the differences in RMS values between the Original, Stylized, and TactStyle textures, we performed a Welch's ANOVA test which indicated a statistically significant difference ($F = 47.58$, $p < 0.0001$). Next we conducted a Games-Howell post-hoc analysis. We found a significant difference between Stylized and Original textures ($T = 6.79$, $p < 0.0001$) and TactStyle and Stylized ($T = 14.34$, $p < 0.0001$) textures. However, we found no significant difference between the TactStyle and Original textures ($p > 0.05$). These results suggest that stylization process results in textures with significantly higher surface roughness compared to the original and TactStyle's generated textures.

*7.0.2 Analyzing Mean Squared Error:* MSE measures per-pixel intensity differences, where 0 indicates identical per-pixel intensities, and 1 indicates completely different intensities. As shown in Figure 5b, TactStyle exhibits lower MSE (M = 0.03, std-dev = 0.03) compared to the stylized method (M = 0.10,std-dev = 0.05), indicating more accurate texture replication. To evaluate statistical significance, we conducted a Welch's t-test, and found that the MSE Loss for TactStyle's results was significantly lower than that of Stylized results ($F = 6.79$, $p < 0.0001$).

## 8 Perception Study

In order to evaluate TactStyle's accuracy at replicating the tactile feedback of textures, we performed a psycho-physical experiment used to evaluate texture replication techniques [7, 41]. The goal of our study was to understand if the reconstructed heightfield from TactStyle creates similar tactile perceptions to the original heightfield. In addition, our second goal was to evaluate if the tactile perception of TactStyle's heightfields are similar to the expected tactile perception from just looking at a visual image of the texture. To compare tactile properties, we take a representative set of descriptors from Degraen et al. [7].

### 8.1 Conditions

We created four distinct conditions to evaluate the tactile and visual characteristics of textures. These conditions were designed to isolate specific aspects of perception, allowing us to better understand how each modality contributes to the overall experience of texture replication. The conditions were:

(1) **Visual (No Heightfield):** The texture image was printed with no heightfield on glossy paper and pasted on flat tiles.
(2) **Original Heightfield:** The texture was printed with the heightfield originally provided with the texture (groundtruth).

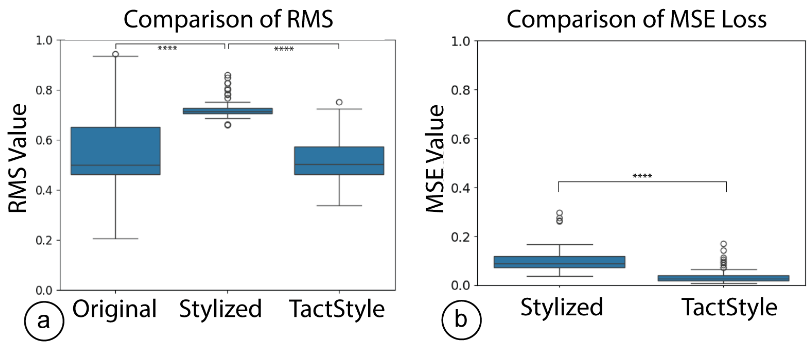


**Figure 5: Quantitative Comparison of TactStyle and Stylized Textures, and original heightfields. (a) Comparison of RMS values for Original, Stylized, and TactStyle textures, demonstrating that TactStyle replicates surface roughness more closely to the Original textures. (b) Box plot of MSE loss between the Original textures and the Stylized and TactStyle textures. TactStyle exhibits significantly lower MSE compared to the stylized method, indicating more accurate texture replication. (****p < 0.0001)**

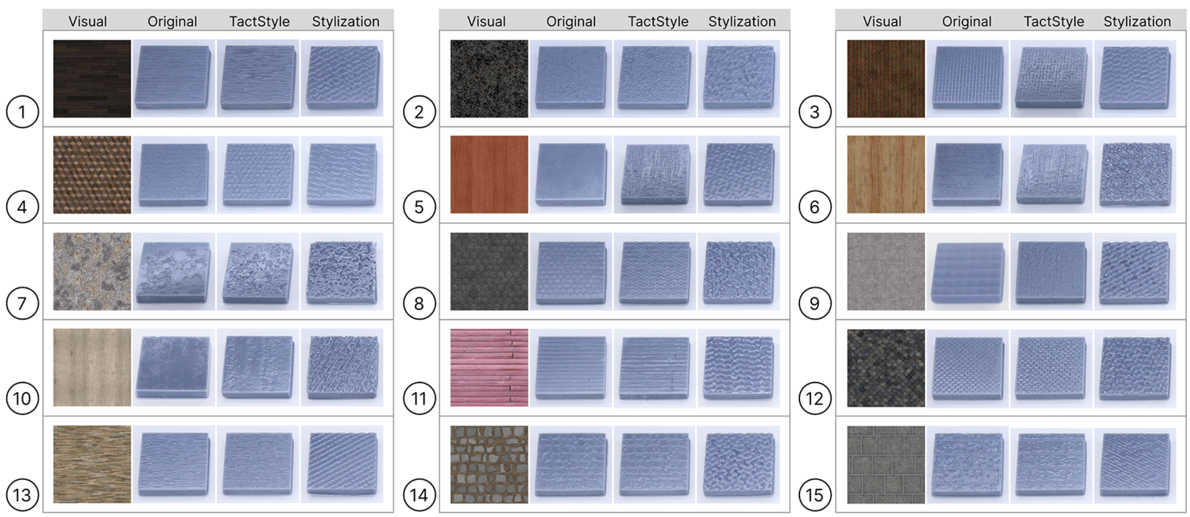


**Figure 6: 3D-Printed samples of 15 textures from our test set used in perception study: We created four different sets for our perception study: the 'visual set', 'original set', 'TactStyle set', and the 'stylization set'. The original set was created with the heightfield associated with the texture and served as the groundtruth. The reconstructed set was created using TactStyle, with the texture image as input. The visual set was created using printed texture images. Finally, the stylization (baseline) set was created using Style2Fab, using the texture image as input.**

(3) **TactStyle:** The texture was printed with the heightfield created from TactStyle with the texture image as input.
(4) **Stylization:** The texture was printed with the heightfield generated from Style2Fab [11] using the texture image as input.

This separation of conditions was motivated by the literature on visuo-haptic stimuli integration by humans. When humans explore objects with their hands, vision and touch both provide information for estimating the properties of the object [9]. Vision frequently dominates the integrated visual–haptic percept [14, 15]. To address this, we kept the visual and tactile conditions separate to mitigate cross-modality influence and isolate modality-specific effects to assess tactile perception. Here, condition 4 (stylization) is our baseline to be compared with TactStyle's results.

In this experiment, we investigate the following research questions:

(1) **RQ1:** How accurately does TactStyle replicate the tactile properties of a texture, as represented by its original heightfield?
(2) **RQ2:** To what extent do TactStyle-generated textures align with user expectations based on their visual appearance?

(3) **RQ3:** How do the tactile expectations derived from a texture's visual appearance differ from its actual tactile properties, as represented by the heightfield?

## 8.2 Dataset

We collected 15 random samples from our test set, matching the size of the sample set used by Degraen et al. [7] in their study on the perceptual similarity between real textures and their digital replicas. As these textures were not used to train the model, they can be used to evaluate TactStyle's ability to replicate unseen textures. Figure 6 shows the models used in our study. This gave us a total set of 60 models (4 conditions for each of the 15 textures). All the conditions were presented to the user on tiles of the same size - 5cm x 5cm x 1cm. Our models were printed using an SLA printer, namely, Elegoo Saturn 3 Ultra. In order to keep our printed objects comparable to the previous studies, we used the Elegoo Resin Standard 2.0 – Grey, which has a Shore Hardness of 80-86 (Scale D). For comparison, Degraen et al. [7] used a material with Shore Hardness of 83-86 (Scale D). We printed all the samples at a layer resolution of 30 $\mu$m.

## 8.3 Study Design

We used a within-subjects experimental design. To control for carry-over effects, we counter-balanced conditions using round-robin ordering between participants. Our study was structured as a self-assessment test in which participants compared and recorded perceptual attributes of the 3D-printed texture samples from the different conditions.

During the study, each participant recorded their ratings of the sample in terms of hardness, roughness, bumpiness, stickiness, scratchiness, uniformity, and how isotropic the surface is, each on a 1-to-9 Likert scale, 1 indicating a low assessment and 9 indicating a high assessment of the respective variable. To rate these dimensions, participants were asked the following questions:

Q1: How **hard** does this surface feel? (1 meaning extremely soft, 9 meaning extremely hard)
Q2: How **rough** does this surface feel? (1 meaning extremely smooth, 9 meaning extremely rough)
Q3: How **bumpy** does this surface feel? (1 meaning extremely flat, 9 meaning extremely bumpy)
Q4: How **sticky** does this surface feel? (1 meaning extremely slippery, 9 meaning extremely sticky)
Q5: How **scratchy** does this surface feel? (1 meaning extremely dull, 9 meaning extremely scratchy)
Q6: How **uniform** does this surface feel? (1 meaning extremely irregular, 9 meaning extremely uniform)
Q7: How **isotropic** does this surface feel? (1 meaning extremely anisotropic, 9 meaning extremely isotropic)

Hardness, Roughness and Stickiness are motivated by related work indicating these are the base dimensions of tactile discrimination [19, 20, 39, 51, 57]. Bumpiness and Scratchiness are informed by the notion that roughness can be divided into respectively macro and micro dimensions [39]. The inclusion of Uniformity and Isotropy stems from the fact that our textures embed some directionality and localized variations, which affect perception during tactile exploration [7]. While other works have considered the additional dimension of Hairiness [7, 8], we excluded this descriptor since none of our textures in our dataset were representative of it.

## 8.4 Apparatus

Our apparatus was built to limit visual cues and ensure accurate recording of purely tactile perceptual attributes of the textures. Participants were positioned in front of a screen that separated them from the experimenter, as shown in Figure 7. A small opening in the screen, covered by a piece of cloth, allowed participants to reach through and access the samples, placed by the experimenter. On the other side, the experimenter arranged the samples for the participants to explore. The samples were held in place with a laser-cut wooden frame.

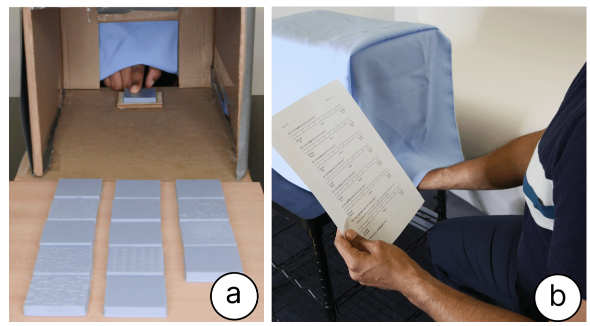


**Figure 7: Experimental Setup for perception study: a) Experimenter side, b) Participant side**

## 8.5 Participants

A total of 15 participants (6 female, 9 male, 22 - 38 years, $M$ = 27.3 years, $SD$ = 5.4) were recruited for our study. When asked about their hand dominance, 14 participants indicated to be right-handed with 1 participant indicating ambidexterity. All participants chose to use their right-hand index finger for the study. They were informed that they could only use this finger throughout the study for consistency in perception. All participants indicated that they do not suffer from any impairment to haptic perception to best of their knowledge. Participants were compensated with $20 an hour for the 90-minutes long study.

## 8.6 Study Procedure

The total study duration was set to be 90 minutes. At the start, participants were asked to fill in a consent form, and a short survey about their demographic data. Next, to ensure that participants clearly understood the perceptual descriptors used in the study, we conducted a short training session where participants were allowed to explore exemplar textures separate from the set being investigated. Once they were confident in their understanding of the descriptors, we proceeded to the evaluation stage.

During this stage, the experimenter placed one sample at a time at a fixed location behind a screen. The participant could insert their hand into the screen and feel the texture but were not allowed to see the texture itself. The participant was then asked to explore the texture and rate its tactile properties based on the 7 descriptors.

During the visual perception stage, the samples were placed on a board next to the screen where the participants could see the texture but were not allowed to touch them. This allowed us to isolate visual and tactile perception for textures evaluated in the study, and mitigate cross-modal influence [14, 15]. Prior work has shown that human fingers are particularly sensitive in perceiving and distinguishing textures [47]. Since all the participants chose their right hand for the study, all participants were requested to use only their right index finger throughout their study for consistency. The interaction window was limited to 5 seconds per sample so that the participants' first impressions could be communicated. All participants answered the descriptor questions for all 60 textures.

Ethical approval for this study was obtained from the Ethical Review Board of the author's institute.

## 9 Results

In the following section, we describe the analysis and the obtained results from our texture perception study.

### 9.1 Comparing Visual and Tactile Ratings

In this section, we present the results from the perception study. To analyze the individual tactile assessments, we conducted Friedman tests with post-hoc analysis using Wilcoxon signed-ranks tests and Bonferroni-Holm correction. Figure 8 shows box plots for each assessment. The assessment data for each descriptor is provided in Appendix A.1.

*9.1.1 Hardness.* Users perceived significant differences in Hardness between Original vs. Visual, Original vs. Stylization, Visual vs. TactStyle, Visual vs. Stylization, and TactStyle vs. Stylization. However, users did not perceive significant differences in Original vs. TactStyle.

*9.1.2 Roughness.* Users perceived significant differences in Roughness in all comparisons except between the Visual and the TactStyle conditions.

*9.1.3 Bumpiness.* Users perceived no significant differences in bumpiness between Original vs. Visual, or Visual vs. TactStyle. However, they perceived significant differences between Original vs. TactStyle, Original vs. Stylization, Visual vs. Stylization, and TactStyle vs. Stylization.

*9.1.4 Scratchiness.* Users perceived significant differences in Scratchiness between Original vs. Visual, Original vs. Stylization, Visual vs. TactStyle, Visual vs. Stylization, and TactStyle vs. Stylization. However, they did not perceive any significant difference between Original vs. TactStyle.

*9.1.5 Stickiness.* Users perceived no significant difference in Stickiness between Original vs. Visual, Visual vs. TactStyle, or TactStyle vs. Stylization conditions. However, significant differences were found between Original vs. TactStyle, Original vs. Stylization, and Visual vs. Stylization.

*9.1.6 Uniformity.* Users perceived no significant differences in uniformity between Original vs. Visual, Original vs. Stylization, Original vs. TactStyle, or TactStyle vs. Stylization. However, significant differences were found between Visual vs. TactStyle and Visual vs. Stylization.

*9.1.7 Isotropy.* Users perceived no significant difference in isotropy between Original vs. Visual, Original vs. TactStyle, or Visual vs. TactStyle samples. However, they perceived significant differences between Original vs. Stylization, Visual vs. Stylization, and TactStyle vs. Stylization.

### 9.2 Perceptual Correlations

To uncover relationships between different tactile perceptions in our samples, we performed Spearman's rank-order correlation analysis. For each descriptor, we evaluate the relationship between different texture descriptors across the Original, TactStyle, Visual, and Stylization samples. This analysis helped determine whether the tactile ratings of the textures were correlated across different conditions. Figure 8 shows correlation plots for each assessment. The assessment data for each descriptor is provided in Appendix A.2.

*9.2.1 Hardness.* There were significant correlations in perception of hardness between all pairs of textures. This can be explained by the fact that the samples were printed with the same material. However, as described in Section 9.1.1, users perceived significant differences between all pairs except Original vs TactStyle samples. This suggests that TactStyle effectively replicates the tactile properties of hardness from the Original textures, while stylization does not.

*9.2.2 Roughness.* There were significant correlations in perception of roughness between Original vs Visual, Original v TactStyle, and Visual vs TactStyle. We found in Section 9.1.2 that users perceived significant differences in roughness perception in all comparisons except Visual vs TactStyle condition. Thus, TactStyle effectively replicates tactile perception of roughness from visual expectations of a texture.

*9.2.3 Bumpiness.* There were significant correlations in perception of bumpiness between Original vs Visual, Original vs TactStyle and Visual vs TactStyle. However, we found significant differences in comparing Original vs TactStyle (Section 9.1.3), but not in Visual vs TactStyle, and Original vs Visual. Thus, TactStyle effectively replicates bumpiness of textures from their visual expectations.

*9.2.4 Scratchiness.* There were significant correlations in perception of scratchiness between Original and Visual, and Original and TactStyle. We found that users did not perceive significant differences between Original and TactStyle samples (Section 9.1.4). Thus, TactStyle effectively replicates Scratchiness of Original textures.

*9.2.5 Stickiness.* There were significant correlations in perceived stickiness of textures between Original and Visual, Original and TactStyle, Visual and TactStyle, Visual and Stylization. Since users did not perceive any significant differences in perceived stickiness (Section 9.1.5) for Visual and TactStyle, this shows that TactStyle effectively replicates perceived stickiness of textures from their visual expectations.

*9.2.6 Uniformity.* There were significant correlations in perceived uniformity, between Original and Visual, Original and TactStyle,

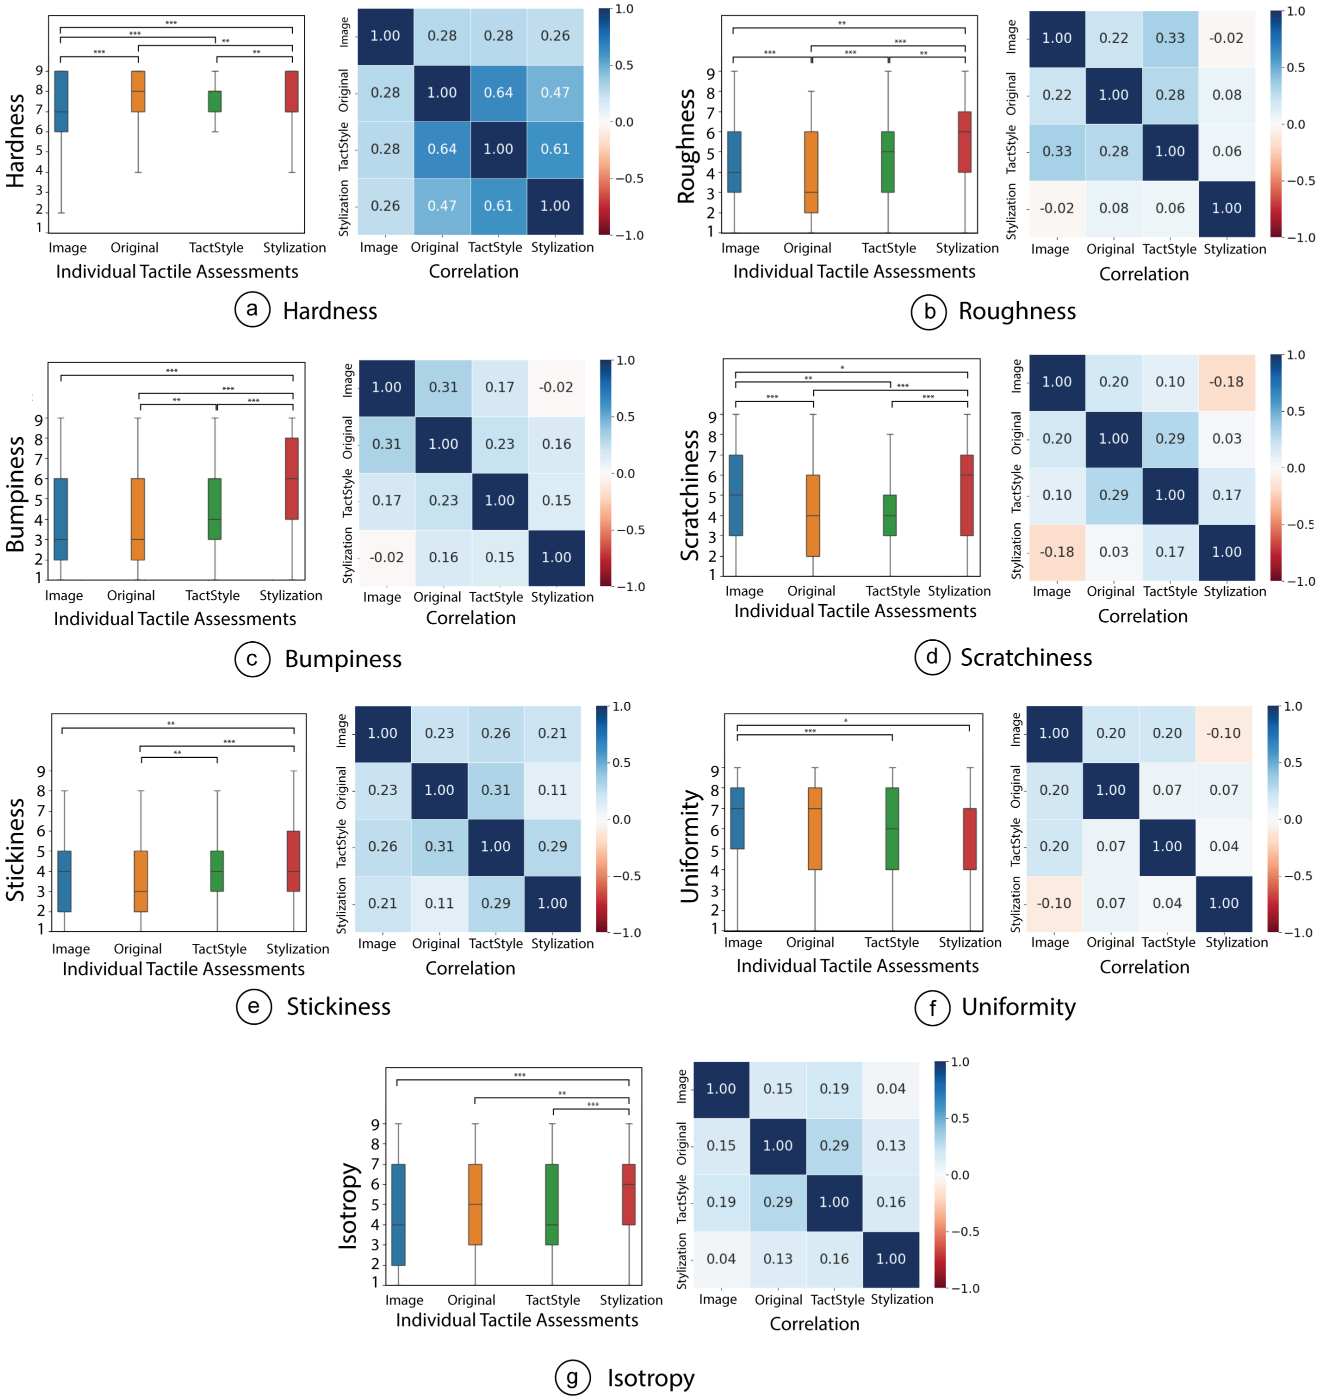


**Figure 8: Box Plots showing the individual assessments on Hardness, Roughness, Bumpiness, Stickiness, Scratchiness, Uniformity, and Isotropy. Tactile Correlations are shown as heatmaps showing the correlations between the 4 conditions for each descriptor. (*p < 0.05, **p < 0.01, ***p < 0.001)**

and Visual and TactStyle conditions. Since we found that users did not perceive any significant differences in perceived uniformity between Original and TactStyle (Section 9.1.6), TactStyle effectively replicates perceived uniformity of textures from their Original textures.

*9.2.7 Isotropy.* There were significant correlations in perceived isotropy between Original and TactStyle, Original and Visual, and Visual and TactStyle samples. Since users did not perceive any significant differences between Original vs TactStyle, Original vs Visual, and Visual vs TactStyle (Section 9.1.7), this shows us that TactStyle is able to effectively replicate perceived isotropy from both Visual expectations and Original textures.

## 9.3 Discussion

We conducted a comparative analysis of the 4 texture sets - Visual, Original, TactStyle, and Stylization, identifying which textures are significantly different on various tactile descriptors, and which of them are correlated. We found that TactStyle effectively replicates visual and Original textures for several key descriptors, and outperforms the baseline stylization method. Based on our analysis, we were able to answer all our research questions stipulated in Section 8.1.

*9.3.1 RQ1: TactStyle accurately replicates several Original Texture Descriptors:* TactStyle effectively replicated the tactile experiences expected from visual cues for several descriptors. Hardness, Scratchiness, Uniformity, and Isotropy are correlated between Original and TactStyle samples without exhibiting significant perceptual differences. This suggests that TactStyle is capable of closely replicating the tactile sensations associated with these descriptors from the original textures.

The analysis also showed that Stylization does not perform well in replicating the tactile features of Original textures. In between Stylization and Original samples, significant differences across most descriptors such as Hardness, Roughness, Bumpiness, Stickiness, and Scratchiness, indicate poor alignment between the tactile experiences replicated by Stylization compared to the Original textures.

*9.3.2 RQ2: TactStyle accurately replicates several Visual Texture Descriptors:* TactStyle effectively replicates several tactile features based on the visual textures, as demonstrated by significant correlations and the absence of perceptual differences in descriptors such as Roughness, Bumpiness, Stickiness, and Isotropy. This indicates that TactStyle can reproduce the tactile experiences expected from visual cues based on these descriptors.

In contrast, stylization does not reliably replicate tactile expectations derived from visual samples. Stylization does not effectively replicate the tactile expectations from Visual samples, with most descriptors showing significant differences.

*9.3.3 RQ3: Differences Between Tactile Expectations from Visual Textures and Actual Tactile Perceptions:* We also compared the expected tactile perceptions of visual textures, and the tactile perceptions of the original textures. We found that most descriptors such as Bumpiness, Stickiness, Uniformity, and Isotropy align closely between visual perceptions and actual original surface tactile experiences. In contrast, descriptors like Hardness, Roughness, and Scratchiness show significant perceptual differences between Visual and Original samples, although moderate correlations indicate some level of consistency. These results suggest that while certain tactile experiences can be anticipated based on visual cues alone, others may require direct tactile interaction for accurate perception.

# 10 Applications

In this section, we showcase how TactStyle's stylization technique allows users to stylize 3D models with accurate tactile properties for fabrication. We demonstrate five application scenarios across four categories: home decor, personal accessories, tactile learning tools, and personalized health applications. All 3D models shown in Figure 9 were stylized with TactStyle using textures not present in the training dataset and printed on a Stratasys J55 printer.

## 10.1 Home Decor

TactStyle can be used to apply textures to objects downloaded from platforms like Thingiverse, enabling users to enhance the tactile experience of 3D-printed items at home. This allows individuals to create customized, textured versions of everyday objects, adding both aesthetics and functionality. We illustrate two applications of TactStyle in applying textures to functional home objects. Figure 9a we shows a wood-parquet textured phone stand, demonstrating how organic textures can be applied to enhance the visual appeal and usability of frequently handled items. Figure 9b, shows a granite-textured vase. By combining TactStyle and digital fabrication techniques, users can now personalize their objects, or prototype specific tactile properties in addition to the aesthetics of their home decor objects.

## 10.2 Personalizing Accessories

Personal accessories are a popular domain for personalized fabrication. TactStyle enables creators to replicate both the 'look' and the 'feel' of textures based on image input, allowing creators to create customized versions of their accessories with specific textures and fabricate them with digital fabrication. In Figure 9c we showcase an AirPods case, stylized with two different textures: a round stone roof texture taken from an image of a 'round stone roof' (top), and another stylized with an image of a 'layered brown rock'. These textures not only provide visual distinction but also have different surface microgeometry, associated with the texture.

## 10.3 Tactile Learning Tools

TactStyle has the potential to create educational tools that enhance learning in subjects such as geometry, topography (e.g., the texture of different terrains), and biology (e.g., the texture of animal skins). To exemplify this concept, we present two examples in Figure 9d: the top surface features a 'volcanic rock texture', and the bottom surface replicates the texture of stone from 'the Grand Canyon'. Both textures are not present in the dataset, however are samples of a class that TactStyle is trained on. Thus, TactStyle is able to effectively generalize over different texture classes provided they were represented in the training data. Tangible learning materials are well-known to improve educational outcomes, particularly by engaging multiple senses [18]. TactStyle offers a new way to create such materials, allowing educators to bring textures and surfaces to life in the classroom. By giving students the ability to physically interact with these textures, TactStyle could potentially help them better understand the tactile properties of objects, making abstract concepts more concrete and accessible.


**Figure 9: Application Examples for TactStyle: a) a phone stand stylized with a wood-parquet texture, b) a granite-textured vase, c) an airpods case stylized with a texture of 'round stone roof' and 'layered brown rock', d) two tiles, one styled with a volcanic rock texture and the other styled with stone from the Grand Canyon, e) walking stick handle stylized with a rough rock texture.**

### 10.4 Customizable Assistive Devices

In the field of "Medical Making" [30] and "DIY Assistive Technology" [2] personalized fabrication by nontechnical experts is becoming an emerging and critical domain. TactStyle can be employed to customize assistive devices with specific textures, enhancing grip, comfort, or usability tailored to the unique needs of the user. Figure 9e illustrates this by applying a 'rough rock' texture to the handle of a walking stick. This texture generates a rough heightfield, which post-fabrication, significantly increases surface friction, thereby improving grip and stability for the user. Such tactile enhancements are particularly valuable for assistive devices, where safety and ease of use are critical, offering a practical solution that can be tailored to the specific requirements of individuals with mobility challenges.

## 11 Discussion and Future Work

TactStyle demonstrates an ability to replicate both visual and tactile features from an image input, allowing creators to stylize their 3D models for both accurate color replication, and expected tactile properties. In this section, we discuss TactStyle's current limitations, and its possible extensions in the future.

### 11.1 Opportunities for Richer Datasets

TactStyle's performance and robustness to a diversity of textures is dependent the quality and diversity of its training dataset. Currently, the model utilizes the CGAxis repository [3], which provides 500 texture-heightfield pairs across five material categories: *Parquets*, *Wood*, *Rocks*, *Walls*, and *Roofs*. The high quality of the available images, and their corresponding heightfields allowed us to train the image-generation model with a high accuracy. While this dataset offers a diverse selection of real-world textures, this does not cover all types of textures encountered by humans. Additional material categories such as fabrics, metals, and organic surfaces could enhance the model's generalizability, and allow personalization of tactile surfaces in fashion, automotive design, etc. Moreover, expanding the dataset to include dynamic material properties, such as elasticity, thermal responsiveness, or friction, could enable TactStyle to model textures with more complex interactions, further improving the reproducibility of their tactile properties.

### 11.2 Cross-Modal Texture Design

TactStyle is able to replicate specific tactile descriptors from both expected tactile features extracted from the visual texture, and the perceived tactile properties correlated to the original heightfield. This combined replication of expected and perceived tactile properties allows for a cross-model design of textures. Recent work in VR and Haptics [6] have explored novel ways to design and map user-defined tactile properties in virtual reality, such as voice. TactStyle approaches a similar problem, but in the fabrication domain, allowing users to apply textures that have 'expected' tactile properties. In the future, this approach can be extended to text prompts, allowing users to describe their expected tactile response, and fabricate a 3D model with such tactile properties.

### 11.3 Incorporating Material Properties of Textures

The material properties of textures play a critical role in defining their tactile experience. Hardness and Scratchiness, for instance, are closely related to the rigidity and resistance of a material, directly affecting how a surface feels when touched. Currently, TactStyle operates by taking images as input to generate tactile features. While this method effectively aligns visual and tactile experiences, there is potential to enhance its accuracy by incorporating material descriptors as input. In this study, we standardized the material used for all texture samples for consistency, to evaluate the accuracy of generated heightfields in replicating tactile perception. However, tactile perception is also influenced by material-specific properties such as compliance, thermal conductivity, and surface friction [7]. Future work could explore how these properties can be replicated by predicting material types that approximate their tactile properties. Additionally, integrating novel approaches like metamaterials [23] could provide new avenues for tailoring and enhancing texture replication across diverse applications.

## 11.4 Analyzing Visuo-Haptic Properties Together

TactStyle currently evaluates visual and tactile perceptions separately, identifying key differences between expected tactile properties based on visual cues and the actual tactile perceptions of textures. These findings highlight an interplay between visual and haptic modalities in shaping texture perception [47]. Future work could explore this property of tactile perception, and leverage visuo-haptic mismatches to create novel experiences, such as "*impossible materials*" that visually appear soft but feel rigid, defying conventional expectations. Additionally, photochromic materials have been used in prior work to create re-programmable multi-color surfaces. Such materials offer opportunities to dynamically link visual and tactile feedback to create novel dynamic textures.

## 11.5 3D Model Generation with accurate texture information

Recent Generative AI methods have enabled users to generate novel 3D models from scratch based on image and text prompts [17, 54]. However, while current systems excel in generating visual representations of textures, they often lack the capacity to accurately generate the tactile properties on these materials. Since TactStyle works with image modality as well, an extension of TactStyle could allow creators to provide an image of description of a novel object and its expected tactile properties, allowing creators to not only create novel digital artifacts but also fabricatable designs with accurate texture information. By extending generative tools to encode material properties, these models could also propose materials to fabricate the object such that the tactile experience is closely approximated.

## 12 Conclusion

In this paper, we present TactStyle, a system that allows users to stylize 3D models using image prompts, replicating both visual appearance and tactile properties. By extending generative AI techniques, TactStyle generates tactile features as a heightfield and applies them to 3D models. A quantitative study demonstrates significant improvements over traditional stylization methods. In a psychophysical experiment with 15 participants, we evaluate TactStyle's ability to create textures perceived as similar to both visually expected tactile properties and the original texture's tactile features.

Our findings show that TactStyle successfully aligns visual and tactile properties, enabling more realistic 3D model personalization. This work opens up new possibilities in cross-modal design, and future work can expand TactStyle by incorporating material descriptors to further enhance its tactile accuracy.

## Acknowledgments

We would like to extend our sincere gratitude to the MIT-Google Program for Computing Innovation for their generous support, which made this research possible. Furthermore, we thank Varun Jampani from Stability AI, Vrushank Phadnis, Federico Tombari, and Fabian Manhardt from Google Research, and Martin Nisser from University of Washington for their valuable insights and feedback on this research.

## References

[1] 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* (Sanibel Island, Florida, USA) *(GROUP '16)*. Association for Computing Machinery, New York, NY, USA, 195–199. doi:10.1145/2957276.2957301

[2] Erin Buehler, Stacy Branham, Abdullah Ali, Jeremy J. Chang, Megan Kelly Hofmann, Amy Hurst, and Shaun K. Kane. 2015. Sharing is Caring: Assistive Technology Designs on Thingiverse. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 525–534. doi:10.1145/2702123.2702525

[3] CGAxis. [n. d.]. PBR 20 Parquets. https://cgaxis.com/cgaxis_preview/PBR_20_Parquets/. Accessed: 2024-09-10.

[4] Thingiverse. com. 2024. Thingiverse - Digital Designs for Physical Objects. https://www.thingiverse.com [Online; accessed 4. Apr. 2024].

[5] Blender Online Community. 2018. *Blender - a 3D modelling and rendering package.* Blender Foundation, Stichting Blender Foundation, Amsterdam. http://www.blender.org

[6] Donald Degraen, Bruno Fruchard, Frederik Smolders, Emmanouil Potetsianakis, Seref Güngör, Antonio Krüger, and Jürgen Steimle. 2021. Weirding haptics: In-situ prototyping of vibrotactile feedback in virtual reality through vocalization. In *The 34th Annual ACM symposium on user interface software and technology.* 936–953.

[7] Donald Degraen, Michal Piovarči, Bernd Bickel, and Antonio Krüger. 2021. Capturing Tactile Properties of Real Surfaces for Haptic Reproduction. In *The 34th Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) *(UIST '21)*. Association for Computing Machinery, New York, NY, USA, 954—-971. doi:10.1145/3472749.3474798

[8] Donald Degraen, André Zenner, and Antonio Krüger. 2019. Enhancing Texture Perception in Virtual Reality Using 3D-Printed Hair Structures. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300479

[9] Marc O Ernst and Martin S Banks. 2002. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415, 6870 (2002), 429–433.

[10] 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 preprint arXiv:2109.14722* (2021).

[11] 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.

[12] Faraz Faruqi, Yingtao Tian, Vrushank Phadnis, Varun Jampani, and Stefanie Mueller. 2024. Shaping Realities: Enhancing 3D Generative AI with Fabrication Constraints. *arXiv preprint arXiv:2404.10142* (2024).

[13] Martin Feick, Donald Degraen, Fabian Hupperich, and Antonio Krüger. 2023. Metareality: enhancing tactile experiences using actuated 3D-printed metamaterials in virtual reality. *Frontiers in Virtual Reality* 4 (2023), 1172381.

[14] Roland W Fleming. 2014. Visual perception of materials and their properties. *Vision research* 94 (2014), 62–75.

[15] Roland W Fleming. 2017. Material perception. *Annual review of vision science* 3, 1 (2017), 365–388.

[16] A Gallace. 2014. *In touch with the future: The sense of touch from cognitive neuroscience to virtual reality.* Oxford University Press.

[17] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. 2022. GET3D: A Generative Model of High Quality 3D Textured Shapes Learned from Images. In *Advances In Neural Information Processing Systems.*

[18] Carina S González-González, María D Guzmán-Franco, and Alfonso Infante-Moro. 2019. Tangible technologies for childhood education: A systematic review. *Sustainability* 11, 10 (2019), 2910.

[19] Mark Hollins, Sliman Bensmaïa, Kristie Karlof, and Forrest Young. 2000. Individual differences in perceptual space for tactile textures: Evidence from multidimensional scaling. *Perception & psychophysics* 62 (2000), 1534–1544.

[20] Mark Hollins, Richard Faldowski, Suman Rao, and Forrest Young. 1993. Perceptual dimensions of tactile surface texture: A multidimensional scaling analysis. *Perception & Psychophysics* 54, 6 (01 Nov 1993), 697–705. doi:10.3758/BF03211795

[21] Yi-Hua Huang, Yan-Pei Cao, Yu-Kun Lai, Ying Shan, and Lin Gao. 2023. NeRF-texture: Texture synthesis with neural radiance fields. In *ACM SIGGRAPH 2023 Conference Proceedings.* 1–10.

[22] Nathaniel Hudson, Celena Alcock, and Parmit K. Chilana. 2016. Understanding Newcomers to 3D Printing: Motivations, Workflows, and Barriers of Casual Makers. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing

Machinery, New York, NY, USA, 384–396. doi:10.1145/2858036.2858266
[23] Alexandra Ion, Johannes Frohnhofen, Ludwig Wall, Robert Kovacs, Mirela Alistar, Jack Lindsay, Pedro Lopes, Hsiang-Ting Chen, and Patrick Baudisch. 2016. Metamaterial mechanisms. In *Proceedings of the 29th annual symposium on user interface software and technology*. 529–539.
[24] Alexandra Ion, Robert Kovacs, Oliver S. Schneider, Pedro Lopes, and Patrick Baudisch. 2018. Metamaterial Textures. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3173574.3173910
[25] Micah K Johnson and Edward H Adelson. 2009. Retrographic sensing for the measurement of surface texture and shape. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 1070–1077.
[26] Jeeeun Kim, Qingnan Zhou, Amanda Ghassaei, and Xiang 'Anthony' Chen. 2021. OmniSoft: A Design Tool for Soft Objects by Example. In *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction* (Salzburg, Austria) *(TEI '21)*. Association for Computing Machinery, New York, NY, USA, Article 15, 13 pages. doi:10.1145/3430524.3440634
[27] Diederik P Kingma. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
[28] Katherine J Kuchenbecker, Joseph Romano, and William McMahan. 2011. Haptography: Capturing and recreating the rich feel of real surfaces. In *Robotics Research: The 14th International Symposium ISRR*. Springer, 245–260.
[29] Stacey Kuznetsov and Eric Paulos. 2010. Rise of the Expert Amateur: DIY Projects, Communities, and Cultures. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries (NordiCHI '10)*. ACM, 295–304. doi:10.1145/1868914.1868950
[30] Udaya Lakshmi, Megan Hofmann, Stephanie Valencia, Lauren Wilcox, Jennifer Mankoff, and Rosa I. Arriaga. 2019. "Point-of-Care Manufacturing": Maker Perspectives on Digital Fabrication in Medical Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 91 (nov 2019), 23 pages. doi:10.1145/3359193
[31] Gierad Laput, Xiang'Anthony' Chen, and Chris Harrison. 2015. 3D printed hair: Fused deposition modeling of soft strands, fibers, and bristles. In *Proceedings of the 28th annual ACM symposium on user interface software & technology*. 593–597.
[32] Yiwei Ma, Xiaoqing Zhang, Xiaoshuai Sun, Jiayi Ji, Haowei Wang, Guannan Jiang, Weilin Zhuang, and Rongrong Ji. 2023. X-Mesh: Towards Fast and Accurate Text-driven 3D Stylization via Dynamic Textual Guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 2749–2760.
[33] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. 2021. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073* (2021).
[34] Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. 2022. Text2mesh: Text-driven neural stylization for meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13492–13502.
[35] Stefanie Mueller. 2017. 3D printing for human-computer interaction. *interactions* 24, 5 (2017), 76–79.
[36] Alessandro Muntoni and Paolo Cignoni. 2021. *PyMeshLab*. doi:10.5281/zenodo.4438750
[37] Behnaz Norouzi, Marianne Kinnula, and Netta Iivari. 2021. Making Sense of 3D Modelling and 3D Printing Activities of Young People: A Nexus Analytic Inquiry. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 481, 16 pages. doi:10.1145/3411764.3445139
[38] Lora Oehlberg, Wesley Willett, and Wendy E. Mackay. 2015. Patterns of Physical Design Remixing in Online Maker Communities. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 639–648. doi:10.1145/2702123.2702175
[39] Shogo Okamoto, Hikaru Nagano, and Yoji Yamada. 2012. Psychophysical dimensions of tactile perception of textures. *IEEE Transactions on Haptics* 6, 1 (2012), 81–93.
[40] Jifei Ou, Gershon Dublon, Chin-Yi Cheng, Felix Heibeck, Karl Willis, and Hiroshi Ishii. 2016. Cilllia: 3D Printed Micro-Pillar Structures for Surface Texture, Actuation and Sensing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 5753–5764. doi:10.1145/2858036.2858257
[41] Alvaro G Perez, Daniel Lobo, Francesco Chinello, Gabriel Cirio, Monica Malvezzi, Jose San Martin, Domenico Prattichizzo, and Miguel A Otaduy. 2016. Optimization-based wearable tactile rendering. *IEEE transactions on haptics* 10, 2 (2016), 254–264.
[42] Franklin Pezutti-Dyer and Leah Buechley. 2022. Extruder-Turtle: A Library for 3D Printing Delicate, Textured, and Flexible Objects. In *Proceedings of the Sixteenth International Conference on Tangible, Embedded, and Embodied Interaction* (Daejeon, Republic of Korea) *(TEI '22)*. Association for Computing Machinery, New York, NY, USA, Article 6, 9 pages. doi:10.1145/3490149.3501312
[43] Michal Piovarci, Michael Foshey, Vahid Babaei, Szymon Rusinkiewicz, Wojciech Matusik, and Piotr Didyk. 2020. Towards spatially varying gloss reproduction for 3D printing. *ACM transactions on graphics* 39, 6 (2020).
[44] 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 (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 8748–8763. https://proceedings.mlr.press/v139/radford21a.html
[45] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. arXiv:2112.10752 [cs.CV]
[46] Ryan Schmidt and Karan Singh. 2010. Meshmixer: An Interface for Rapid Mesh Composition. In *ACM SIGGRAPH 2010 Talks*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/1837026.1837034
[47] Lisa Skedung, Martin Arvidsson, Jun Young Chung, Christopher M Stafford, Birgitta Berglund, and Mark W Rutland. 2013. Feeling small: exploring the tactile perception limits. *Scientific reports* 3, 1 (2013), 2617.
[48] Lingyun Sun, Jiaji Li, Junzhe Ji, Deying Pan, Mingming Li, Kuangqi Zhu, Yitao Fan, Yue Yang, Ye Tao, and Guanyun Wang. 2022. X-Bridges: Designing Tunable Bridges to Enrich 3D Printed Objects' Deformation and Stiffness. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR, USA) *(UIST '22)*. Association for Computing Machinery, New York, NY, USA, Article 20, 12 pages. doi:10.1145/3526113.3545710
[49] Haruki Takahashi and Homei Miyashita. 2016. Thickness Control Technique for Printing Tactile Sheets with Fused Deposition Modeling. In *Adjunct Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology* (Tokyo, Japan) *(UIST '16 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 51–53. doi:10.1145/2984751.2985701
[50] Cesar Torres, Tim Campbell, Neil Kumar, and Eric Paulos. 2015. HapticPrint: Designing feel aesthetics for digital fabrication. In *Proceedings of the 28th annual ACM symposium on user interface software & technology*. 583–591.
[51] Yasemin Vardar, Christian Wallraven, and Katherine J. Kuchenbecker. 2019. Fingertip Interaction Metrics Correlate with Visual and Haptic Perception of Real Surfaces. In *2019 IEEE World Haptics Conference (WHC)*. IEEE. doi:10.1109/whc.2019.8816095
[52] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, Dhruv Nair, Sayak Paul, William Berman, Yiyi Xu, Steven Liu, and Thomas Wolf. 2022. Diffusers: State-of-the-art diffusion models. https://github.com/huggingface/diffusers.
[53] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
[54] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. 2024. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191* (2024).
[55] Xin Yan, Lin Lu, Andrei Sharf, Xing Yu, and Yulu Sun. 2021. Man-made by Computer: On-the-Fly Fine Texture 3D Printing. In *Proceedings of the 6th Annual ACM Symposium on Computational Fabrication* (Virtual Event, USA) *(SCF '21)*. Association for Computing Machinery, New York, NY, USA, Article 6, 10 pages. doi:10.1145/3485114.3485119
[56] Yu-Ying Yeh, Jia-Bin Huang, Changil Kim, Lei Xiao, Thu Nguyen-Phuoc, Numair Khan, Cheng Zhang, Manmohan Chandraker, Carl S Marshall, Zhao Dong, et al. 2024. Texturedreamer: Image-guided texture synthesis through geometry-aware diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4304–4314.
[57] T. Yoshioka, S. J. Bensmaïa, J. C. Craig, and S. S. Hsiao. 2007. Texture perception through direct and indirect touch: An analysis of perceptual space for tactile textures in two modes of exploration. *Somatosensory & Motor Research* 24, 1-2 (2007), 53–70. doi:10.1080/08990220701318163
[58] Xiaoting Zhang, Guoxin Fang, Chengkai Dai, Jouke Verlinden, Jun Wu, Emily Whiting, and Charlie C.L. Wang. 2017. Thermal-Comfort Design of Personalized Casts. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (Québec City, QC, Canada) *(UIST '17)*. Association for Computing Machinery, New York, NY, USA, 243–254. doi:10.1145/3126594.3126600
[59] Clement Zheng, Zhen Zhou Yong, Hongnan Lin, HyunJoo Oh, and Ching Chiuan Yen. 2022. Shape-Haptics: Planar & Passive Force Feedback Mechanisms for Physical Interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 171, 15 pages. doi:10.1145/3491102.3501829

# A Perception Study Assessment Readings

## A.1 Comparing Visual and Tactile Readings

*A.1.1 Hardness.* The ratings of hardness significantly differed depending on the presented texture ($\chi^2(3) = 29.398$, $p < 0.001$).

For the Original samples, significant differences were found when compared to the Visual ($W = 2476.5$, $p < 0.001$) and Stylization samples ($W = 1161.5$, $p < 0.01$), but no significant difference was observed with the TactStyle samples ($W = 927.5$, $p > 0.05$). For the Visual samples, significant differences were found with TactStyle ($W = 2978.5$, $p < 0.001$) and Stylization samples ($W = 4319.5$, $p < 0.001$). Additionally, the TactStyle samples differed significantly from Stylization samples ($W = 1080.5$, $p < 0.01$).

*A.1.2 Roughness.* The ratings of roughness differed significantly depending on the presented texture ($\chi^2(3) = 46.48$, $p < 0.001$). For the Original samples, significant differences were found with Visual ($W = 5768.0$, $p < 0.001$), TactStyle ($W = 4511.0$, $p < 0.001$) and Stylization samples ($W = 4347.5$, $p < 0.001$). For the Visual samples, significant differences were also found with Stylization samples ($W = 6526.0$, $p < 0.01$), but no significant difference with the TactStyle samples ($W = 8503.0$, $p > 0.05$). The TactStyle samples differed significantly from the Stylization samples ($W = 7110.0$, $p < 0.01$).

*A.1.3 Bumpiness.* The ratings of bumpiness significantly differed depending on the presented texture ($\chi^2(3) = 95.72$, $p < 0.001$). For the Original samples, no significant difference was observed with the Visual samples ($W = 7070.0$, $p > 0.05$), but significant differences were found with both the TactStyle ($W = 6257.5$, $p < 0.01$) and Stylization ($W = 2296.0$, $p < 0.001$). For the Visual samples, no significant difference was observed with the TactStyle ($W = 8242.5$, $p > 0.05$), but a significant difference was found with the Stylization ($W = 3901.5$, $p < 0.001$). Additionally, the TactStyle samples significantly differed from the Stylization samples ($W = 3326.0$, $p < 0.001$).

*A.1.4 Scratchiness.* The ratings of scratchiness significantly differed depending on the presented texture ($\chi^2(3) = 45.20$, $p < 0.001$). For the Original samples, significant differences were found when compared to both the Visual ($W = 6383.5$, $p < 0.001$) and Stylization samples ($W = 4940.5$, $p < 0.001$), but no significant difference was observed with the TactStyle samples ($W = 6907.0$, $p > 0.05$). For the Visual samples, significant differences were found with both the TactStyle ($W = 6379.0$, $p < 0.01$) and Stylization samples ($W = 8198.0$, $p < 0.05$). Additionally, the TactStyle samples significantly differed from the Stylization samples ($W = 4379.0$, $p < 0.001$).

*A.1.5 Stickiness.* The ratings of stickiness significantly differed depending on the presented texture ($\chi^2(3) = 24.53$, $p < 0.001$). For the Original samples, significant differences were observed with both the TactStyle ($W = 3726.0$, $p < 0.01$) and Stylization samples ($W = 3685.5$, $p < 0.001$), but no significant difference was found compared to the Visual samples ($W = 5656.5$, $p > 0.05$). For the Visual samples, no significant difference was found with the TactStyle samples ($W = 6985.0$, $p > 0.05$), but a significant difference was found with the Stylization samples ($W = 5679.0$, $p < 0.01$). Additionally, no significant difference was found between TactStyle and Stylization samples ($W = 5316.0$, $p > 0.05$).

*A.1.6 Uniformity.* The ratings of uniformity significantly differed depending on the presented texture ($\chi^2(3) = 25.02$, $p < 0.001$). For the Original samples, no significant difference was observed compared to the Visual ($W = 7504.0$, $p > 0.05$), TactStyle ($W = 6783.0$, $p > 0.05$), or Stylization samples ($W = 7393.5$, $p > 0.05$). For the Visual samples, significant differences were found with both the TactStyle ($W = 5534.0$, $p < 0.001$) and Stylization samples ($W = 7579.0$, $p < 0.05$). Additionally, no significant difference was observed between the TactStyle and Stylization samples ($W = 8565.5$, $p > 0.05$).

*A.1.7 Isotropy.* The ratings of isotropy significantly differed depending on the presented texture ($\chi^2(3) = 27.13$, $p < 0.001$). For the Original samples, no significant difference was observed compared to the Visual ($W = 7527.5$, $p > 0.05$) or TactStyle samples ($W = 7359.5$, $p > 0.05$), but a significant difference was found with the Stylization samples ($W = 6651.5$, $p < 0.01$). For the Visual samples, no significant difference was found with the TactStyle samples ($W = 7825.0$, $p > 0.05$), but a significant difference was observed with the Stylization samples ($W = 5726.0$, $p < 0.001$). For TactStyle samples, a significant difference was found with the Stylization samples ($W = 5085.5$, $p < 0.001$).

## A.2 Perceptual Correlations

*A.2.1 Hardness.* The ratings of hardness were found to significantly correlate depending on the presented texture. For the Original samples, significant correlations were found with Visual samples ($r = 0.28$, $p < 0.001$), the TactStyle samples ($r = 0.64$, $p < 0.001$), and the Stylization samples ($r = 0.47$, $p < 0.001$). For the Visual samples, significant correlations were found with the TactStyle samples ($r = 0.28$, $p < 0.001$) and the Stylization samples ($r = 0.25$, $p < 0.001$). Finally, for the TactStyle samples, significant correlations were found with the Stylization samples ($r = 0.61$, $p < 0.001$).

*A.2.2 Roughness.* The ratings of roughness were found to significantly correlate depending on the presented texture. For the Original samples, significant correlations were found with the Visual samples ($r = 0.22$, $p < 0.01$) and the TactStyle samples ($r = 0.28$, $p < 0.001$), but no significant correlation was observed with the Stylization samples ($r = 0.08$, $p > 0.05$). For the Visual samples, a significant correlation was found with the TactStyle samples ($r = 0.33$, $p < 0.001$), but no significant correlation was observed with the Stylization samples ($r = -0.023757$, $p > 0.05$). Finally, for the TactStyle samples, no significant correlation was observed with the Stylization samples ($r = 0.055$, $p > 0.05$).

*A.2.3 Bumpiness.* The ratings of bumpiness were found to significantly correlate depending on the presented texture. For the Original samples, significant correlations were found with the Visual samples ($r = 0.31$, $p < 0.001$) and the TactStyle samples ($r = 0.23$, $p < 0.01$), but no significant correlation was observed with the Stylization samples ($r = 0.16$, $p > 0.05$). For the Visual samples, a significant correlation was found with TactStyle samples ($r = 0.17$, $p < 0.05$), but not with Stylization samples ($r = -0.021$, $p > 0.05$). Finally, for the TactStyle samples, no significant correlation was observed with the Stylization samples ($r = 0.14$, $p > 0.05$).

*A.2.4 Scratchiness.* The ratings of scratchiness were found to significantly correlate depending on the presented texture. For the Original samples, a significant correlation was found with the Visual samples ($r = 0.20$, $p < 0.05$) and the TactStyle samples ($r = 0.29$,

$p < 0.001$), but no significant correlation was observed with the Stylization samples ($r = 0.03$, $p > 0.05$). For the Visual samples, no significant correlation was observed with the TactStyle samples ($r = 0.10$, $p > 0.05$) or the Stylization samples ($r = -0.18$, $p > 0.05$). Finally, for the TactStyle samples, no significant correlation was found with the Stylization samples ($r = 0.17$, $p > 0.05$).

*A.2.5 Stickiness.* The ratings of stickiness were found to significantly correlate depending on the presented texture. For the Original samples, significant correlations were found with the Visual samples ($r = 0.23$, $p < 0.01$) and the TactStyle samples ($r = 0.31$, $p < 0.001$), but no significant correlation was observed with the Stylization samples ($r = 0.11$, $p > 0.05$). For the Visual samples, significant correlations were found with both the TactStyle samples ($r = 0.26$, $p < 0.001$) and the Stylization samples ($r = 0.21$, $p < 0.01$). Finally, for the TactStyle samples, a significant correlation was found with the Stylization samples ($r = 0.29$, $p < 0.001$).

*A.2.6 Uniformity.* For the Original samples, a significant correlation was observed with the Visual samples ($r = 0.20$, $p < 0.05$), and with TactStyle samples ($r = 0.074$, $p < 0.05$), but not the Stylization samples ($r = 0.072$, $p > 0.05$). For the Visual samples, a significant correlation was found with the TactStyle samples ($r = 0.20$, $p < 0.05$), but no significant correlation was observed with the Stylization samples ($r = -0.10$, $p > 0.05$). Finally, for the TactStyle samples, no significant correlation was found with the Stylization samples ($r = 0.04$, $p > 0.05$).

*A.2.7 Isotropy.* The ratings of isotropy were found to significantly correlate depending on the presented texture. For the Original samples, a significant correlation was observed with the TactStyle samples ($r = 0.29$, $p < 0.001$), and with the Visual samples ($r = 0.15$, $p < 0.05$), but not with the Stylization samples ($r = 0.13$, $p > 0.05$). For the Visual samples, a significant correlation was found with the TactStyle samples ($r = 0.19$, $p < 0.05$), but not with the Stylization samples ($r = 0.04$, $p > 0.05$). Finally, for the TactStyle samples, no significant correlation was found with the Stylization samples ($r = 0.16$, $p > 0.05$).