

# DealMaTe: Multi-Dimensional Material Transfer via Diffusion Transformer

NISHA HUANG, Tsinghua University, China and Pengcheng Laboratory, China

YIZHOU LIN, Tsinghua University, China

JIE GUO, Pengcheng Laboratory, China

XIU LI<sup>†</sup>, Tsinghua University, China

TONG-YEE LEE<sup>†</sup>, National Cheng-Kung University, Taiwan

ZITONG YU<sup>†</sup>, Great Bay University, China and Dongguan Key Laboratory for Intelligence and Information Technology, China



Fig. 1. **DealMaTe** is a material transfer method that can transform materials from a single real-world image without any prior knowledge. This method is not only capable of successfully extracting texture information from antiques with thousands of years of history but also handles novel virtual materials generated by computer graphics images. It excels across diverse scenarios such as product design, antique restoration, and furniture design, thereby robustly facilitating design tasks.

Recently, diffusion-based material transfer methods rely on image fine-tuning or complex architectures with auxiliary networks but face challenges such as text dependency, additional computational costs, and feature misalignment. To address these limitations, we propose **DealMaTe**, using **depth**, **normal**, and **lighting** images for **material transfer**. DealMaTe is a simplified diffusion framework that eliminates text guidance and reference networks. We design a lightweight 3D information injection method, Multi-Dim 3D Shader LoRA, which, without modifying the base model weights, enables compatible control conditions and achieves harmonious and stable results. Additionally, we optimize the attention mechanism with Shader Causal Mutual Attention and key-value (KV) caching to reduce inference latency caused by multiple conditions, improve computational efficiency, and achieve high-quality material transfer results with low architectural complexity. Extensive experiments covering a wide variety of objects and

lighting conditions consistently demonstrate that DealMaTe achieves remarkable high-fidelity material transfer under arbitrary input materials.

Additional Key Words and Phrases: Material transfer, diffusion transformer, shader multi-modal attention

## ACM Reference Format:

Nisha Huang, Yizhou Lin, Jie Guo, Xiu Li, Tong-Yee Lee, and Zitong Yu. 2026. DealMaTe: Multi-Dimensional Material Transfer via Diffusion Transformer. 1, 1 (May 2026), 15 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 Introduction

Material transfer is a technique that precisely maps the properties of a specific material sample onto the surface of a target object, as shown in Fig. 1. Due to its broad application prospects in digital content creation and industrial design, this technology has garnered attention in recent years. Traditional material transfer frameworks primarily rely on parametric modeling paradigms, such as optical reflection models based on bidirectional scattering distribution function (BRDF) [9, 20, 25] or procedural texture generation algorithms [24, 59]. These frameworks limit the diversity of the generated results due to the limited size of the basic material library,

<sup>†</sup>Corresponding author

Authors' Contact Information: Nisha Huang, Tsinghua University, China and Pengcheng Laboratory, China, [hns24@mails.tsinghua.edu.cn](mailto:hns24@mails.tsinghua.edu.cn); Yizhou Lin, Tsinghua University, China, [yz-lin24@mails.tsinghua.edu.cn](mailto:yz-lin24@mails.tsinghua.edu.cn); Jie Guo, Pengcheng Laboratory, China, [guoj01@pcl.ac.cn](mailto:guoj01@pcl.ac.cn); Xiu Li, Tsinghua University, China, [li.xiu@sz.tsinghua.edu.cn](mailto:li.xiu@sz.tsinghua.edu.cn); Tong-Yee Lee, National Cheng-Kung University, Taiwan, [tonylee@ncku.edu.tw](mailto:tonylee@ncku.edu.tw); Zitong Yu, Great Bay University, China and Dongguan Key Laboratory for Intelligence and Information Technology, China, [zitong.yu@ieee.org](mailto:zitong.yu@ieee.org).

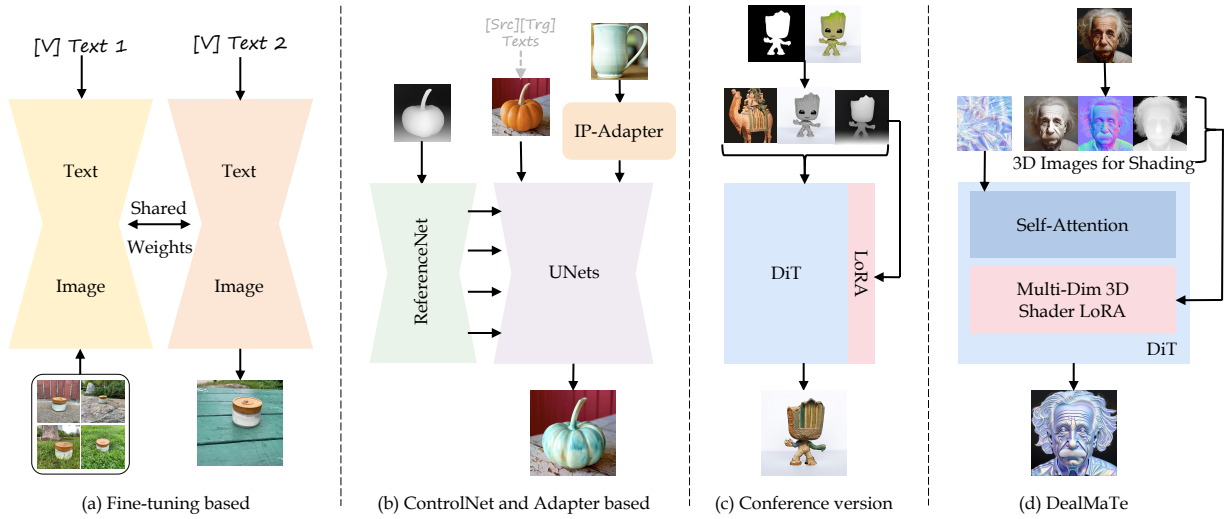


Fig. 2. Simplified structure comparison of different kinds of material transfer methods. Our approach neither relies on fine-tuning image sets/individual images nor requires additional image encoding by IP-Adapters. It only requires basic image information as input, rather than complex text guidance, to obtain high-quality material transfer results.

failing to meet the customization needed for non-uniform composite materials in digital artistic creation.

In recent years, inspired by the breakthroughs of diffusion models in conditional generation tasks, diffusion-based material transfer methods [37, 50, 57] have achieved remarkable improvements in high-quality material transfer effects. Current state-of-the-art (SOTA) approaches often employ fine-tuning of diffusion models with sample sets bound to text identifiers [49, 57, 64] (in Fig. 2 (a)), implicitly encoding material features through conceptual semantic binding. Although these methods address some shortcomings of traditional pipelines, their heavy reliance on text prompts restricts fine-grained control of material properties. Moreover, the full-parameter fine-tuning paradigm significantly increases training costs and the risk of overfitting. The latest studies [6, 7, 18] have also introduced pre-trained general image encoders IP-Adapters [58] to extract material features, as well as ControlNet [62] to inject depth information, as shown in Fig. 2 (b). However, these methods frequently induce hierarchical decoupling between material and structural information (instead of seamless fusion) during generation, while also suffering from prolonged inference times. Fig. 2 (c) eliminates the use of additional image encoders in material transfer tasks. MaTe prioritizes the semantic alignment of cross-modal features and their interaction in a shared latent space, mapping materials, depth, and grayscale images to a unified latent representation. However, our previous work [27] and other SOTAs [6, 7, 18] only employ Depth LoRA for structural control and use grayscale images as initial sampling images for lighting control, resulting in insufficient preservation of 3D information.

Therefore, to solve the problem, DealMaTe (in Fig. 2 (d)) designs the Multi-Dim 3D Shader LoRA: depth, normal, and lighting. Specifically, the depth LoRA is capable of precisely capturing the depth information of the target object, providing an accurate

three-dimensional spatial reference for material transfer. The normal LoRA focuses on controlling surface curvature, ensuring that material textures closely conform to the geometric shape of the target object, thereby effectively avoiding the general problem [6, 7, 18] of material-geometry misalignment. To the best of our knowledge, we pioneer the fine-tuning of the Lighting LoRA in the image generation field. It finely adjusts the direction, intensity, and color of the lighting, enabling specular highlights to accurately follow the direction of the light source and significantly improving the consistency of lighting and reflection in the material transfer results.

We suggest a lightweight conditional injection LoRA module at the model architecture level to address the problem of multi-condition collaborative control. Our essential innovation is the solitary insertion of conditional signals. By means of a parallel-branch design, the module is grafted onto the frozen backbone so that low-rank projection is restricted to the conditional-branch tokens, whereas the prompt and noise pathways stay locked and unchanged. This advanced architecture facilitates that the Multi-Dim 3D Shader LoRA works seamlessly with the customized model while also allowing for harmonic and strong zero-shot generalization.

Regarding the inference latency caused by multiple conditions, at the attention mechanism level, we replace the traditional full attention with our innovative Shader Causal Mutual Attention and combine it with KV caching technology to achieve the transformation. At the initial diffusion step, the system computes the key-value tensors for every conditioning feature once and stores them. These cached tensors are then reused in all subsequent steps, cutting out repeated computation and delivering significant savings.

In summary, we highlight our contributions as follows:

- We propose DealMaTe, an efficient and concise material transfer architecture that can infer all necessary information using only image conditions, without the need for fine-tuning or manual text settings, and supports zero-shot generalization.




Fig. 3. Comparison with our conference version work MaTe. DealMaTe is capable of faithfully transferring material characteristics while ensuring enhanced structural consistency with the input image.

- We introduce Multi-Dim 3D Shader LoRA to address material-geometry misalignment and lighting-reflection inconsistency. To the best of our knowledge, we are the first to fine-tune the Lighting LoRA. We also design a lightweight conditional injection LoRA module that isolates the injection of conditional signals, integrates seamlessly with pre-trained models, and enhances the ability for multi-condition collaborative control.
- We optimize the attention mechanism through Shader Causal Mutual Attention and KV caching to reduce inference latency caused by various conditions, improve computational efficiency, and reduce computational resources, achieving high-quality material transfer results with lower architectural complexity.

This work extends our recent conference paper MaTe published in ICCV 2025 [27]. The material transfer framework has been enhanced by adjusting the injection method of Low-Rank Adaptation (LoRA) [23] and proposing Multi-Dim 3D Shader LoRA, which introduces brand-new control conditions for surface curvature and lighting. Furthermore, to improve the harmonious stability of multi-condition collaboration, we have adopted Shader Causal Mutual Attention and KV caching to optimize the attention mechanism. The extended version has been thoroughly evaluated through comprehensive experiments, which show that our DealMaTe can greatly improve the quality of findings for existing arbitrary image material transfer models. For visual comparisons between MaTe and DealMaTe, please refer to Fig. 3.

## 2 Related Work

*Material and texture extraction.* Material acquisition has always been an ongoing challenge [19]. Traditional material capture often relies on costly multi-view [1] or polarized light [11] equipment.

Many researchers have used synthetic datasets to train networks that estimate SVBRDF from a single viewpoint [9], often in conjunction with additional single-view data [17, 40] or specific training strategies [10, 33, 54]. However, it is important to note that all the aforementioned studies require a frontal view of the material, meaning the camera must be parallel to the material surface, which is often difficult to achieve in practical applications. For material analysis, UMat [46] employs a single image taken using a flatbed scanner. Rather than depicting the substance itself, TexSynth [12] offers a guided texture modification method. Material Palette [37] suggests a way to extract materials from a single image without previous knowledge. MaterialPicker [38] uses a DiT-based text-to-video generation methodology to generate material maps. Moreover, a common way to extract textures from real-world images includes clustering the textures and scaling them up to full resolution [34, 48] or applying dataset distillation techniques [3].

*Image-guided generation.* Image-guided generation is a challenging yet promising field, leveraging visual guidance for content creation. The emergence of diffusion models [28, 43, 47] has significantly propelled advancements in this field, enabling SOTA performance across a wide range of generative visual tasks, such as image-to-image translation [26, 58], subject-driven image generation [21, 49], etc. DreamBooth [49] and Textual Inversion [16] use transfer learning in text-to-image (T2I) diffusion models to generate customized concepts through parameter fine-tuning or word vector optimization. To enhance the controllability of image-guided generation, adapter-based architectures have emerged as bridges between external control signals (e.g., sketches) and diffusion models. ControlNet [62] validates an adapter that can be trained to capture task-specific input conditions, whereas T2I-adapter [41] uses a lightweight adapter to achieve fine-grained control in the color and structure of the produced images. IP-Adapter [58] allows for more flexible and intuitive control of the generation process, expanding the capabilities of image-guided generation. These methods demonstrate image-conditioned generation strategies with varying specificity, suggesting potential pathways for addressing material transfer challenges.

*Material acquisition and transfer.* Material acquisition and transfer represent a realm of research considering illumination conditions, object geometry, and physical properties of materials. Traditional 3D material transfer methods, such as Text2tex [4], TEXTure [45], and TextureDreamer [60], rely on 3D geometric shapes and lighting estimation, followed by careful adjustment of material properties. As a result, the quality and diversity are restricted, presenting limited and unsatisfactory results. In contrast, 2D-to-2D material transfer, bypassing ground-truth 3D data, is challenging yet highly practical. Prevalent approaches based on fine-tuning, including DreamBooth [49], Material Palette [37], Prospect [64], and U-VAP [57] fine-tune diffusion models on small sample sets associated with text identifiers. While existing methods mitigate conventional pipeline limitations, text prompts reliance restricts control and fine-tuning risks computational costs, and overfitting. Another category of methods, such as MaterialFusion [18] and ZeST [7], uses extra encoder modules like IP-Adapter [58] or ControlNet [62] to extract material features. When multiple image conditions are injected into the

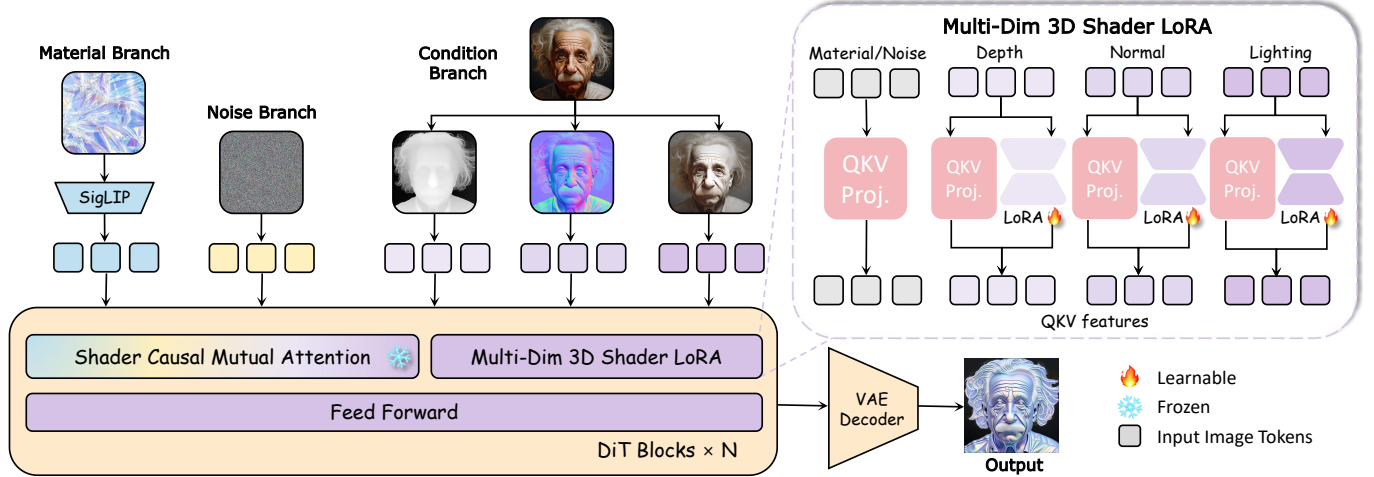


Fig. 4. Our method achieves high-quality material transfer by feeding depth, normal, and lighting inputs into respective shader LoRAs and then passing the material and 3D condition tokens into the shader causal mutual attention for interaction, ensuring that they remain in the same feature space throughout the diffusion process. It remarkably eliminates the need for unnecessary image fine-tuning and text guidance, resulting in a streamlined inference process.

network in parallel, the material and structure information in the results are separated into two independent layers, lacking semantic integration (see Fig. 6 (d)-(g)). The latest SOTA method Material-Fusion [18] applies DDIM inversion [51] to both the material and the input images, which increases the sampling steps and time by more than double. Different from methods that rely on text-assisted fine-tuning or complex architectures with additional reference networks, our method DealMaTe projects material, content, and depth images into the same latent space and performs semantically aligned generation.

### 3 Method

Building on the Diffusion Transformer architecture introduced in Section 3.1, we maintain spatial consistency through rotary position encoding and integrate multiple conditional signals via the multimodal attention mechanism. For conditional feature enhancement, the Multi-Dim 3D Shader LoRA proposed in Section 3.2 selectively boosts conditional representations while preserving the base model’s material representation capabilities. During inference, the Shader Causal Mutual Attention mechanism designed in Section 3.3 efficiently fuses multiple conditions and effectively prevents cross-condition interference. The training process employs the flow-matching loss function detailed in Section 3.4, and inference efficiency is optimized through Key-Value (KV) caching techniques that eliminate redundant computations across denoising steps. Additionally, the data collection process outlined in Section 3.5 leverages Marigold-based pipelines to obtain high-quality depth, normal, and lighting estimates from natural images. The simultaneous estimation of these three intrinsic properties requires 0.3 seconds for per input image.

#### 3.1 Preliminary

**3.1.1 Rectified-Flow Models.** Generative models seek to establish a correspondence between samples  $x_1$  originating from a noise

distribution  $p_1$  and instances  $x_0$  derived from a target distribution  $p_0$ , where  $p_0$  corresponds to real images. According to [35, 36], rectified flows establish a forward procedure that constructs linear pathways across distributions  $p_0$  and  $p_1$ , with  $p_1$  defined as  $\mathcal{N}(0, 1)$  in Eq. 1. This method is time-reliant given the presence of the timestep  $t$ .

$$x_t = (1 - t)x_0 + t\epsilon, \quad \epsilon \sim \mathcal{N}(0, 1). \quad (1)$$

The acquisition of this transformation is achieved by optimizing a neural model with parameters  $\theta$  to estimate the rectified flow velocity  $v$ , denoted  $v_\theta$ . Through the parameterization approach from [14], this network that estimates flow velocity can alternatively operate as a model for predicting noise,  $\epsilon_\theta$ , which is trained with the Conditional Flow Matching (CFM) objective detailed in Eq. 2.

$$\mathcal{L}_{CFM} = -\frac{1}{2} \mathbb{E}_{t \sim \mathcal{U}(t), \epsilon \sim \mathcal{N}(0, 1)} [w_t \lambda'_t \|\epsilon_\theta(x_t, t) - \epsilon\|^2]. \quad (2)$$

$\lambda'_t$  denotes the redefined signal-to-noise ratio and  $w_t$  represents a weighting factor that varies with time.

**3.1.2 Multimodal Attention.** The DiT model [43] is used in designs such as FLUX.1 [32], Stable Diffusion 3 [14], and PixArt [5] iteratively refine noise image tokens using a transformer denoising network. The DiT architecture operates on two token categories, namely noise image tokens  $X \in \mathbb{R}^{N \times d}$  and text conditioning tokens  $C_T \in \mathbb{R}^{M \times d}$ . Here  $d$  indicates embedding dimensionality, with  $N$  and  $M$  representing the respective counts of visual and textual tokens. Throughout the network, these tokens maintain consistent shapes as they pass through multiple transformer blocks. FLUX.1’s DiT blocks include layer normalization, multimodal attention (MMA) [42], and rotary position embedding (RoPE) [52] to encode spatial information.

Next, the multimodal attention module takes the tokens with positional encodings as its input and generates distinct query  $Q$ , key  $K$ , and value  $V$  representations. It facilitates the calculation of

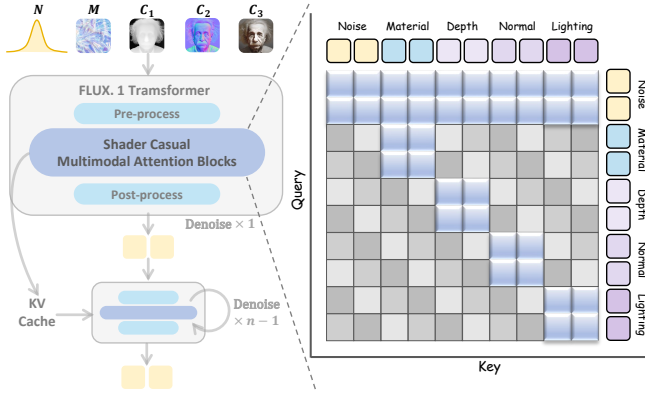


Fig. 5. To optimize the inference performance of the material transfer task, the framework is novelly designed with a shader causal multimodal attention mechanism. This mechanism aligns with the task’s requirements for feature correlation processing. Besides, it can directly enable key-value caching, ultimately reducing inference latency and improving efficiency effectively.

attention among all tokens:

$$\text{MMA}([X; C_T]) = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V, \quad (3)$$

where  $[X; C_T]$  represents the concatenation of image and text tokens. This concept supports bidirectional attention. Building on the DiT architecture with FLUX.1 as the implementation basis, we aim to develop DealMaTe, a framework that balances superior performance with minimalist design, guiding material transfer exclusively via visual conditions.

### 3.2 Multi-Dim 3D Shader LoRA

To efficiently integrate 3D structural conditional signals while preserving the generalization ability of the pre-trained model and the material representation capability of the diffusion backbone, we extend and optimize the FLUX architecture [32]. Instead of adopting the traditional approach of adding independent control modules, which often directly embeds control modules into the main network layer and can interfere with the model original feature representation, as seen in the method proposed by Tan et al. [53], we introduce an additional conditional branch. This allows for the seamless and natural integration of 3D structural conditional information into the existing architecture. Moreover, we deploy a dedicated Multi-Dim 3D Shader LoRA only on this new conditional branch. This strategy not only achieves the intended functionality but also avoids the generation of redundant parameters and the increase in additional computational costs. This separation is crucial, as it ensures precise structural control while preventing degradation of the model’s core generative capabilities, and this key advantage ultimately enables the high-fidelity synthesis of both structure and material.

In the Transformer architecture, the input features are first mapped to three types of features: queries ( $Q$ ), keys ( $K$ ), and values ( $V$ ), which are then processed by the self-attention mechanism. When the input feature representations correspond to material, noise, and 3D

shader conditions, respectively, their corresponding input representations are  $Z_m$ ,  $Z_n$ , and  $Z_c$ . The standard query, key, and value transformations can be defined as:

$$Q_i, K_i, V_i = W_Q Z_i, W_K Z_i, W_V Z_i, \quad i \in \{m, n, c\}. \quad (4)$$

The projection matrices  $W_Q, W_K, W_V$  are shared between branches to share parameters, however this approach does not maximize the representation of conditional signals. To solve this constraint, we creatively insert the LoRA into the conditional path to adaptively boost the expressiveness of conditional features while guaranteeing that the parameters and structures of the other branches are unaffected:

$$\Delta Q_c = B_Q A_Q Z_c, \quad \Delta K_c = B_K A_K Z_c, \quad \Delta V_c = B_V A_V Z_c. \quad (5)$$

In the low-rank adaptation matrices,  $A_Q, A_K, A_V \in \mathbb{R}^{r \times d}$  and  $B_Q, B_K, B_V \in \mathbb{R}^{d \times r}$ , where  $r \ll d$ . Under this setting, the query, key, and value in the conditional branch are updated to:

$$Q'_c = Q_c + \Delta Q_c, \quad K'_c = K_c + \Delta K_c, \quad V'_c = V_c + \Delta V_c. \quad (6)$$

Meanwhile, the material and noise branches remain unchanged:

$$Q'_i = Q_i, \quad K'_i = K_i, \quad V'_i = V_i, \quad i \in \{m, n\}. \quad (7)$$

By applying LoRA-based adaptation solely to the 3D shader conditional branch, we ensure that the conditional signal is efficiently integrated into the model without interfering with the material representation. This focused adjustment allows the model to flexibly incorporate conditional information while preserving the integrity of its original feature space, thereby enabling more controllable and high-fidelity material transfer, as demonstrated in Sec. 4.5.1.

### 3.3 Shader Casual Multimodal Attention

**3.3.1 Shader Casual Attention.** Shader Casual Attention is a unidirectional attention mechanism whose core objective is to ensure strict adherence to temporal causality by restricting the flow of information within sequence models. Specifically, it stipulates that each position in the sequence can only attend to itself and the preceding positions. This rule is achieved by applying a mask composed of 0 and  $-\infty$  to the attention logits before the softmax operation, with the following mathematical expression:

$$Q = [Q'_m; Q'_n; Q'_c], \quad K = [K'_m; K'_n; K'_c], \quad V = [V'_m; V'_n; V'_c], \quad (8)$$

$$\text{MMA}([N; M; C_c]) = \text{Softmax}(QK^T / \sqrt{d_k} + M)V. \quad (9)$$

In this mechanism,  $M$  is responsible for ensuring the adherence to causality, while QKV is composed of features from the material, noise, and 3D condition branches. To improve efficiency and integrate multiple condition signals during inference, we have developed a new causal attention mechanism called Shader Casual Mutual Attention (SCMA). By employing different masking strategies, this mechanism can precisely control the flow of information, ensuring that while multiple condition signals are integrated, the independence between each condition signal is maintained.

**3.3.2 Shader Causal Mutual Attention (SCMA).** However, during multi-condition inference, the interaction between condition tokens and denoising tokens stays normal, but interference across conditions occurs due to the model’s lack of training on cross-condition token interactions. The following mechanism, with its unique architecture, intelligently combines several condition signals while avoiding interference during the inference phase. We specifically describe the input sequence for multi-condition inference as follows:

$$Z = [Z_n; Z_m; Z_{c_1}; Z_{c_2}; Z_{c_3}]. \quad (10)$$

In this mechanism,  $Z_n$  and  $Z_m$  are used to identify the tokens related to noise and material, respectively. While  $Z_{c_1}$ ,  $Z_{c_2}$ , and  $Z_{c_3}$  correspond to the depth, normal, and lighting tokens, respectively.

To precisely control the flow of attention, we introduce a special attention mask  $M \in \{0, -\infty\}^{n \times n}$ , whose specific construction is as follows:

$$M_{ij} = \begin{cases} -\infty, & \text{if } i \notin n_{t\&n} \text{ and } j \in n_{t\&n}, \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

Under this setup, the total sequence length is given by  $n = n_{m\&n} + \sum_{i=1}^m n_{c_i}$ . This carefully designed masking strategy ensures that while the image tokens aggregate information from all conditions, the distinct conditions remain isolated from one another, maintaining their independence. We present experimental validation in Sec. 4.5.2.

## 3.4 Training and Inference

**3.4.1 Loss Function.** The flow-matching loss is used as our loss function, defined as

$$L_{RF} = E_{t, \epsilon \sim N(0, I)} \|v_\theta(z, t, c_i) - (\epsilon - x_0)\|_2^2. \quad (12)$$

In this formula,  $\epsilon$  represents the predicted noise,  $v_\theta$  denotes the velocity field,  $z$  is the image feature at time  $t$ ,  $c_i$  indicates the input condition, and  $x_0$  is the original image feature.

**3.4.2 Efficient Inference via KV Cache.** To achieve efficient inference, we employ a novel key-value (KV) caching strategy specifically designed for multi-conditional diffusion models. Unlike previous approaches in diffusion inference optimization [39, 56] that primarily focus on reducing sampling steps or model distillation, our method focuses on decoupling conditional feature computation. The core of this approach is to design the conditioning branch as a computation module independent of the denoising timesteps, leveraging the properties of the causal attention mechanism to make the conditioning branch agnostic to the timesteps.

During inference, we calculate and store the key-value pairs of all conditional features in the cache dictionary  $\mathcal{D}$  just once at the initial timestep, as the calculation of the conditioning branch is independent of individual timesteps. These cached KV pairs are subsequently utilized in all following denoising processes, preventing duplicate computation of similar conditional features. Specifically, for 3 conditional features  $\{\text{cond}_i\}_{i=1}^3$ , we compute at the initial timestep:

$$K_{C_i}, V_{C_i} = f_\theta(\text{cond}_i), \quad \forall i \in \{1, 2, 3\}, \quad (13)$$

where  $\text{cond}_1$ ,  $\text{cond}_2$ , and  $\text{cond}_3$  correspond to depth, normals, and lighting conditions, respectively. We store these in the cache  $\mathcal{D}[i] \leftarrow (K_{C_i}, V_{C_i})$ . In each subsequent denoising step  $t$ , we directly retrieve the cached key-value pairs  $\{(K_{C_i}, V_{C_i})\}_{i=1}^3$ , and compute the query, key, and value vectors corresponding to the current noise and text features:

$$Q_{\text{denoising}}, K_{\text{denoising}}, V_{\text{denoising}} = f_\theta(x_t, t). \quad (14)$$

Subsequently, we concatenate the cached conditional key-value pairs with the current computed key-value pairs:

$$\begin{aligned} Q &= Q_{\text{denoising}}, \\ K &= \text{Concat}(K_{\text{denoising}}, K_{C_1}, K_{C_2}, K_{C_3}), \\ V &= \text{Concat}(V_{\text{denoising}}, V_{C_1}, V_{C_2}, V_{C_3}). \end{aligned} \quad (15)$$

The attention output is then computed using the standard attention mechanism. This method significantly reduces inference latency by eliminating  $N$ -fold redundant computations (corresponding to  $N$  denoising steps), while maintaining generation quality and model flexibility. See Sec. 4.5.5 for relevant experimental verification.

## 3.5 Dataset Collection

For the spatial control LoRA training of depth, lighting, and normals, we primarily utilized 2,400 natural images from Unsplash [29] that encompass arbitrary subject information. The corresponding depth, lighting, and normals for these natural images were obtained through Marigold [30, 31]. All of the images are preprocessed, including cropping and alignment, to guarantee that our training inputs are consistent and accurate. This resulted in four subsets of data, each containing corresponding real, depth, lighting, and normal images.

**Depth.** Previous material transfer methods [6, 7, 18] predominantly use depth as the structural condition, and we adopt and further strengthen this paradigm. Given an image  $x$ , we use the Stable Diffusion VAE to encode it into a latent code  $z^{(x)}$ . We then concatenate it with the depth latent code  $z_t^{(d)}$  and input it to the fine-tuned UNet for iterative denoising. After  $T$  steps of the schedule, the depth latent code  $z_0^{(d)}$  is decoded into an image, and the average of its 3 channels is taken to obtain the initial depth estimate  $\hat{d}$ .

To improve the quality of the depth-guiding results, we propose a novel test-time ensemble scheme that leverages the stochasticity of the inference process. For each input, we run  $N$  inferences to obtain a set of predictions  $\{\hat{d}_1, \dots, \hat{d}_N\}$ . We align these predictions by optimizing for the scales  $\hat{s}_i$  and offsets  $\hat{t}_i$ , with the objective function given by:

$$\min_{\substack{s_1, \dots, s_N \\ t_1, \dots, t_N}} \left( \sqrt{\frac{1}{b} \sum_{i=1}^{N-1} \sum_{j=i+1}^N \|\hat{d}'_i - \hat{d}'_j\|_2^2} + \lambda \mathcal{R} \right), \quad (16)$$

where  $\hat{d}' = \hat{d} \times \hat{s} + \hat{t}$ ,  $\mathcal{R} = |\min(\mathbf{m})| + |1 - \max(\mathbf{m})|$ , and  $b = \binom{N}{2}$ . During optimization, the merged depth  $\mathbf{m}$  is computed via pixel-wise median, and it serves as the final ensemble prediction. This scheme does not require real data and balances efficiency and quality by adjusting  $N$ .

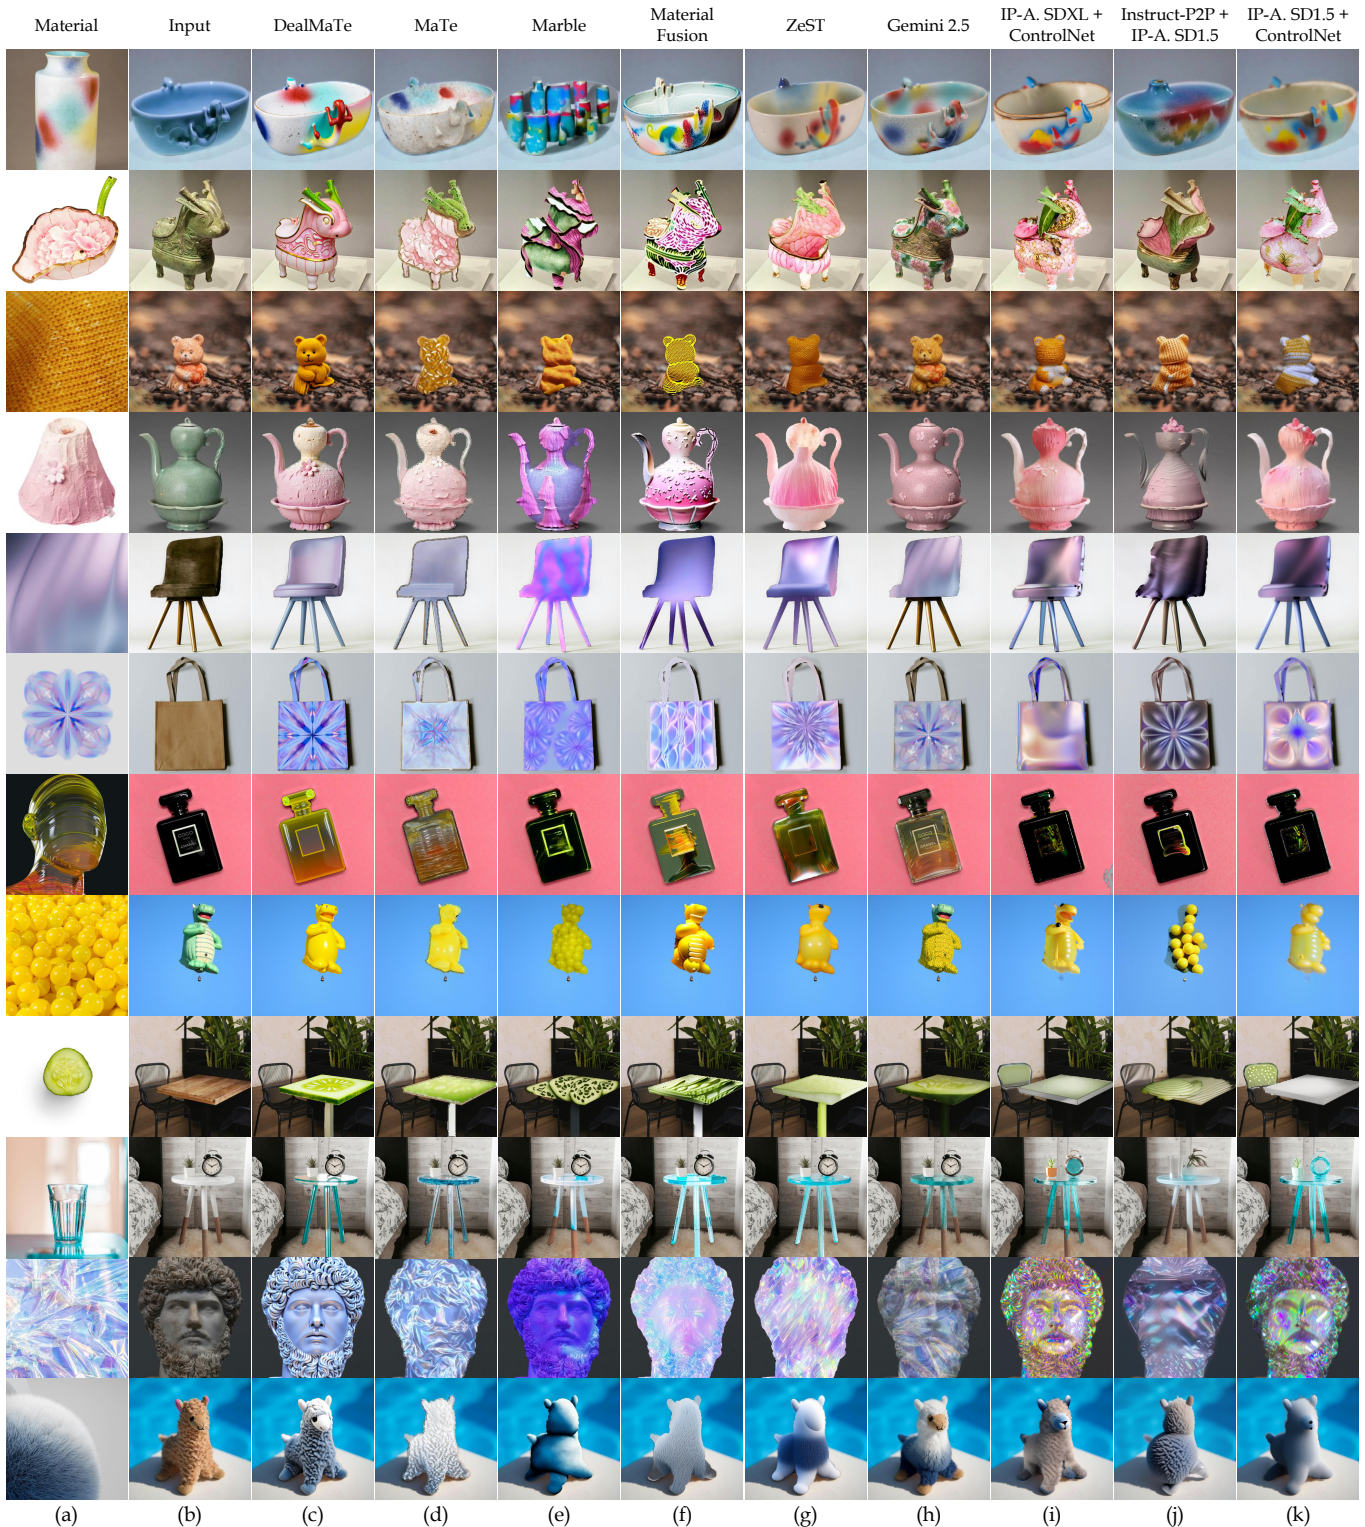


Fig. 6. Qualitative comparisons with other methods. Column (d) shows the result of our conference version work [XXX]. Columns (e)-(g) show a qualitative comparison of DealMaTe with several state-of-the-art material transfer methods. Column (h) is one of the most advanced image generation agents that can understand complex editing commands. Column (i)-(k) show the pipeline we build using the IP-Adapter and reference networks.

Table 1. Quantitative comparisons with other methods are presented. The results that represent the current SOTA baselines for material transfer are highlighted in gray, with the best results shown in bold and the second-best results underlined.

	DealMaTe	MaTe [ICCV 2025]	Marble [CVPR 2025]	Material Fusion	ZeST [ECCV 2024]	Gemini 2.5	IP-A. SDXL + ControlNet	Inst-P2P + IP-Adapter	IP-A. SD1.5 + ControlNet
<b>SSIM</b> ↑	<b>0.8906</b>	0.8176	0.7038	0.8263	0.7231	<u>0.8825</u>	0.8152	0.7093	0.7489
<b>LPIPS</b> ↓	<b>0.1285</b>	0.1548	0.2970	0.1565	0.2430	<u>0.1317</u>	0.1876	0.2743	0.2537
<b>CLIP</b> ↑	<b>0.8927</b>	<u>0.8834</u>	0.7816	0.8521	0.8627	0.8692	0.8358	0.7904	0.8271
<b>PSNR</b> ↑	<b>16.755</b>	14.987	15.337	<u>16.308</u>	15.540	15.0323	9.4670	11.482	8.9032
<b>DreamSim</b> ↓	<b>0.3527</b>	0.3841	0.4352	0.4038	0.6020	<u>0.3832</u>	0.6089	0.6897	0.6676




Fig. 7. When 3D-SCB is removed, the surface textures become contaminated with artifacts and irregular bumps, and naive condition interference becomes evident.

**Normals.** Surface normals are closely related to depth estimation, as both aim to recover 3D geometric information. Depth estimation outputs scalar values, while normals are represented as three-dimensional unit vectors. Real normal annotations are difficult to collect and are often generated from depth, which might result in noise in flat areas and excessive smoothing at edges. Marigold-Normals bridges the gap between simulation and reality through the Stable Diffusion prior. We first run inference  $N$  times with different noise initializations. The resulting normal predictions  $\{\hat{\mathbf{n}}_1, \dots, \hat{\mathbf{n}}_N\}$  are averaged into  $\bar{\mathbf{n}}$  and normalized. Finally, for each pixel  $(u, v)$ , we select the prediction with the highest cosine similarity to the mean normal  $\bar{\mathbf{n}}^{(u,v)}$ :

$$\arg \max_i \bar{\mathbf{n}}^{(u,v)} \cdot \hat{\mathbf{n}}_i^{(u,v)}. \quad (17)$$

**Lighting.** We obtain the lighting image using Marigold-IID-Lighting by first feeding the input image  $I$  into the model, which outputs three decomposition components: albedo  $A$ , diffuse shading  $l$ , and non-diffuse residual  $R$ . The diffuse shading component  $l$  corresponds directly to the lighting image, representing illumination color. The decomposition follows the formula:

$$I = A \cdot l + R. \quad (18)$$




Fig. 8. Visual comparison of the SCMA ablation. With SCMA, the generated structures remain clean; without it, structures are partly lost, and the different structural conditions mutually corrupt, yielding blurred details.

By extracting the predicted  $l$  component from the model, we obtain the lighting image. This formulation gives an organized approach to distinguishing between reflectance and shading while taking into consideration complicated lighting processes.

## 4 Experiments

### 4.1 Implementation Details

We base the method on FLUX.1 [32], a latent rectified flow transformer for image generation. In our experiments, we set the inference steps to 25 and the image generation size to  $1024 \times 1024$  pixels. All the inference experiments are conducted on a single NVIDIA A100 GPU. Each 3D Shader LoRA takes roughly three days to train for 500k steps on 4 NVIDIA A100 GPUs. We employ a batch size of one and a learning rate of  $1 \cdot 10^{-4}$ .

For comparison experiments settings, we compare DealMaTe with Marble [6], Material Fusion [18], ZeST [7], Gemini 2.5 [8], IP-Adapter SDXL [58], Instruct-P2P [2], and IP-Adapter SD1.5 [58] in the task of material transfer. To enhance the fairness of the experiments, we added the corresponding version of the ControlNet depth control models [13, 61] to the pipelines for IP-Adapter SD1.5 and IP-Adapter SDXL. Additionally, we incorporated mask segmentation operations for Gemini 2.5, IP-Adapter SDXL / SD 1.5 + ControlNet, and Instruct-P2P + IP-Adapter SD 1.5.

### 4.2 Qualitative Comparison

As shown in Fig. 6, our method excels in preserving the structure of the input images, faithfully rendering the structural characteristics of the original images. Material transfer based solely on image




Fig. 9. We evaluate the necessity of each control signal by removing exactly one condition (depth, lighting, or normal) from our 3D Shader Conditional Branch.

Table 2. The user study results show the percentage of participants who preferred the outcomes of the corresponding methods in different aspects. The results that represent the current SOTA baselines for material transfer are highlighted in gray, with the best results shown in bold and the second-best results underlined.

	DealMaTe	Marble [CVPR 2025]	Material Fusion	ZeST [ECCV 2024]	Gemini 2.5	IP-A. SDXL + ControlNet	Instruct-P2P + IP-Adapter	IP-A. SD1.5 + ControlNet
<b>Material</b> ↑	<b>68.77%</b>	0.88%	2.63%	6.49%	<u>9.82%</u>	8.42%	1.40%	1.58%
<b>Structure</b> ↑	<b>79.30%</b>	0.70%	2.28%	1.75%	<u>7.54%</u>	4.39%	1.23%	2.81%
<b>Overall</b> ↑	<b>85.79%</b>	0.53%	1.75%	1.93%	<u>4.91%</u>	2.98%	0.88%	1.23%

conditions has been unattainable with previous methods [6, 7, 18]. In contrast, the DealMaTe method proposed in this paper (see Fig. 6 (c)) achieves balanced alignment between structural preservation and material representation by optimizing the extraction mechanism of appearance information from material images and accurately capturing the mapping relationship between the subject structure and the target material. These results demonstrate that our method makes a meaningful step forward in tackling current challenges, offering an effective new paradigm for material transfer studies.


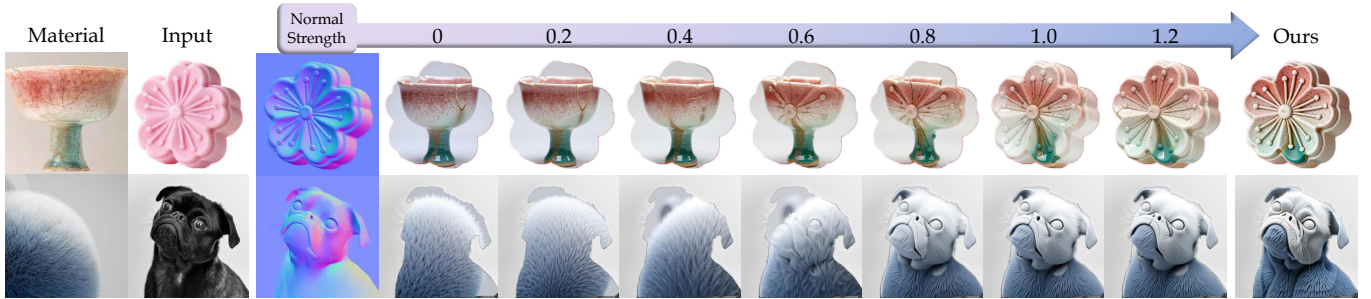
When compared to our conference version MaTe [27] (see Fig. 6 (d)), DealMaTe demonstrates significant qualitative improvements in maintaining rigorous geometric consistency. While the baseline MaTe achieves general texture alterations, it frequently struggles with structural fidelity and fine-grained object details, issues that our enhanced architecture successfully resolves. Moreover, methods specifically designed for material transfer tasks, such as Marble, Material Fusion, and ZeST, have significant shortcomings in preserving structural information. As shown in Fig. 6 (e)-(g), these methods often struggle to effectively retain such features, whether it is the morphological details of antiques, the contour features of toys, or the structural information of human limbs and faces. This further indicates that existing advanced material transfer methods still face technical bottlenecks in preserving the 3D structural information of complex objects.

From the comparative experimental results, it is also evident that the mainstream Gemini 2.5 method, despite its representative performance, fails to completely strip away the original features of the input images, resulting in incomplete material transfer (see Fig. 6 (h)). Taking the “toy bear” as an example, the result processed by this method shows almost no noticeable change in material properties.

Furthermore, to verify the effectiveness of the method, this paper also constructed three types of composite processing pipelines for supplementary experiments (see Fig. 6 (i)-(k)). Among them, the combination of ControlNet and IP-Adapter (based on SDXL or SD1.5) has improved the overall structural expression, but the preservation effect of key facial details such as eyes, nose, and mouth is still not satisfactory. On the other hand, the combination of Instruct-P2P and IP-Adapter shows a certain uniqueness in material reproduction but has significant deviations in the accuracy of structural information.

### 4.3 Quantitative Comparison

We quantitatively validate the approach on the MTB benchmark [27] with regard to state-of-the-art material transfer methods. The Structural Similarity (SSIM) [55] index and Peak Signal-to-Noise Ratio (PSNR) [22] are used to assess structural and geometric consistency between input images and outputs. The Learned Perceptual Image Patch Similarity (LPIPS) [63] and the recent DreamSim [15] metric are used to evaluate the consistency of content detail and mid-level perceptual features between input and output images. We apply the Contrastive Language-Image Pre-training (CLIP) [44] to evaluate the semantic similarity between the materials and the outcomes. Table 1 provides a complete presentation of the experimental results. Our methodology consistently outperforms existing state-of-the-art approaches across evaluated metrics. Notably, compared to the conference version [27], DealMaTe achieves significant performance gains due to its improved capability in disentangling complex material properties and enforcing geometric constraints. This shows that our model is highly effective and well-suited for tackling the specific task of material transfer.

Fig. 10. Ablation experiment on the varying depth control parameter  $\lambda$ .Fig. 11. Ablation analysis of the varying normal control parameter  $\gamma$ .

#### 4.4 User Study

Subjective evaluation consists of three main dimensions: material consistency, structure preservation, and overall quality. Participants were asked to select their most preferred method from each dimension. We assembled a group of 19 experts in the fields of vision and graphics to conduct the assessment. Through 30 comparative experiments (each involving 8 different methods), we collected a total of 1,710 voting results from these three perspectives. The percentage data displayed in Table 2 reveals the preferences of the participants, demonstrating that our proposed DealMaTe method is more favored across all three dimensions. By comparing the quantitative experimental results with the user study outcomes, we recognize the differences between objective and subjective evaluations. While objective evaluations typically consider each dimension independently, participants may synthesize information across dimensions when making their choices, even though they are presented with separate options. The preferences of the participants indicate that our generated results achieve a better balance in material texture, structure preservation, and overall visual appeal. For more information, please see the supplementary materials.

#### 4.5 Ablation Study

**4.5.1 3D shader conditional branch (3D-SCB).** In our ablation study we analyze the impact of removing the 3D-SCB. Specifically, as illustrated by the ‘Single Condition’ column in Fig. 7, this ablation setting retains only a single condition signal and removes the independent LoRA structure of the conditional branch. Instead, the

condition is directly injected into the main network using a naive LoRA structure. Without the 3D-SCB, the material details deviate from the source because such a generic condition LoRA injection corrupts the material latent space, as shown in the fifth column of Fig. 7. For metallic and plastic samples spurious patterns emerge on the surface, while for the blue-and-white porcelain unexpected bumps appear. The 3D-SCB preserves micro-textures while strictly adhering to depth, normal, and lighting constraints, delivering a faithful appearance. This benefit is quantitatively confirmed in Table 3 where the CLIP score of 0.8828 is rendered in the deepest purple in its column, underscoring the positive role of 3D-SCB for material fidelity.

**4.5.2 Shader Casual Mutual Attention (SCMA).** We conduct an ablation study by removing the SCMA module and replacing it with standard multimodal attention (MMA, as formulated in Eq. 3) to evaluate its contribution. As illustrated in Fig. 8, the replacement of SCMA with standard MMA disables the independent guidance from different structural conditions and results in negligible structural refinement. Quantitative results in Table 3 confirm that eliminating SCMA leads to clear degradation in both SSIM and LPIPS. These findings demonstrate that SCMA is essential for maintaining the independence and mutual non-interference of each structural condition, which jointly enables superior structural control.

**4.5.3 Single-Factor Control Ablation.** To strictly verify the necessity of each spatial control signal in our conditional branch, we conducted a leave-one-out ablation study, as shown in Fig. 9. In this experiment, we remove exactly one input condition to observe the


Fig. 12. Ablation experiment on the varying lighting control parameter  $\mu$ .

Table 3. Ablation studies on 3D-SCB and SCMA on quantitative metrics. ‘Avg.’ denotes the arithmetic mean across each ablation configuration. Cells outperforming the average are shaded teal, while underperforming ones are shaded maroon; darker hues indicate larger deviations.

Method	SSIM $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
Avg.	0.8704	0.1314	0.8872
w/o 3D-SCB	0.8634	0.1309	0.8828
w/o SCMA	0.8572	0.1348	0.8862
Ours	0.8906	0.1285	0.8927

specific degradation in the generated results. First, without depth guidance, the depth of field and separation between the foreground and background are imprecise, leading to a lack of 3D volume. Second, without lighting guidance, the generated objects appear extremely flat, and critical semantic details, such as the clear structure of the dog’s eyes and nose, are severely degraded or washed out. Finally, without normal guidance, the model struggles to map complex materials accurately onto uneven surfaces, causing severe texture discontinuities and a loss of fine-grained structural fidelity. Overall, these ablation results show that depth, lighting, and normal control different and complementary spatial properties. Removing any single factor harms the structural consistency of material transfer.

**4.5.4 Multi-Dim 3D Shader LoRA Effect Strength.** The effects of each LoRA’s intensity on transfer quality are shown in Figs. 10, 11, and 12. In practice, setting these parameters involves a general trade-off between rigidly adhering to the target object’s spatial and lighting constraints and preserving the source material’s intrinsic appearance. For the depth LoRA ( $\lambda$ ), lower strength ( $\lambda < 1.0$ ) helps retain the material’s inherent characteristics, such as surface undulations and micro-texture details, but may result in insufficient adherence to the global geometric structure. When  $\lambda \geq 1.0$ , the depth constraint is fully activated, ensuring the material accurately conforms to the target object’s contours. However, pushing beyond this threshold yields no further structural improvement and risks overly constraining the material’s natural displacement, leading to rigid appearance. Similarly, the normal LoRA ( $\gamma$ ) balances surface curvature alignment against the material’s inherent geometric features. Lower strength ( $\gamma < 1.2$ ) better preserves the material’s

Table 4. Ablation studies on KV cache. Calculation of the inference time is performed for images at 1024×1024 resolution.

Method	Sample steps	Inference time (s) $\downarrow$	Speedup $\uparrow$	CLIP Score $\uparrow$
w/o KV Cache	25	73	1.00×	0.8941
w/ KV Cache	25	28	2.61×	0.8927

original physical texture and gloss expression but fails to wrap the material accurately around the object’s curved surfaces. At  $\gamma = 1.2$ , harmonious integration is achieved by fitting the target’s surface normals without completely flattening the material’s intrinsic appearance. Yet higher strength may weaken the material-specific gloss and translucency characteristics. For the lighting LoRA ( $\mu$ ), lower strength helps preserve the material’s base color and inherent reflective properties (such as glossiness and color purity) but fails to correct the lighting direction, resulting in inconsistent shading. At  $\mu = 0.8$ , a balance is struck between lighting consistency and material fidelity. Excessive strength ( $\mu > 0.8$ ), while improving lighting consistency, leads to overexposure, negatively impacting accurate color presentation and potentially destroying fine optical effects such as translucency. Therefore, to achieve an optimal balance among structure preservation, material fidelity, and lighting consistency,  $\lambda$ ,  $\gamma$ , and  $\mu$  are empirically set to 1.0, 1.2, and 0.8, respectively.

**4.5.5 KV Cache.** As shown in Table 4, the proposed KV cache mechanism significantly improves inference efficiency while maintaining output quality. Experimental results demonstrate that with 25 sampling steps, enabling KV cache reduces inference time from 73 seconds to 28 seconds, achieving a 2.61× speedup ratio. This confirms that our method effectively eliminates redundant computation of conditional features during the denoising process, thereby accelerating the generation procedure. Notably, the performance improvement does not come at the cost of output quality degradation. The essentially unchanged CLIP score indicates that the visual quality and semantic alignment of the generated 1024×1024 resolution images are preserved.




Fig. 13. Various downstream applications of our material transfer method. From top to bottom: (a) product packaging, (b) garment fabric, and furniture design. Our approach demonstrates robust generalization and high-fidelity texture synthesis across diverse object categories and complex real-world scenarios.

## 5 Applications and Discussions

### 5.1 Applications

To verify the practicality and versatility of DealMaTe in material transfer, we applied it to fields such as product packaging, garment




Fig. 14. Limitations. When encountering the extreme cases shown in the figure, DealMaTe produce failure and ambiguous material transfer results.

fabrics, and furniture design, demonstrating its adaptability and flexibility under various design requirements.

**Product packaging.** In product packaging design, the visual effect of materials is crucial for attracting consumers. As shown in Fig. 13(a), our method can transfer the textures of metals or gemstones to packaging, achieving a high-end visual effect at a lower cost. By adjusting model parameters, we can simulate the reflective and refractive properties of different materials, allowing for a preview of the final packaging effect during the design phase. This approach saves the cost of physical sample production and accelerates the design iteration process.

**Garment fabric.** In the garment industry, the texture and color of fabrics are decisive for the final appearance of clothing. DealMaTe (see Fig. 13(b)) can transfer the smooth texture of silk to cotton fabric, giving it a high-end silk appearance visually. This not only enhances the innovation in garment design but also allows for the prediction of the effects of different material combinations before production, reducing material waste and production costs.

**Furniture design.** In furniture design, the choice of materials directly affects the durability and aesthetics of the furniture. Our model can transfer the natural texture of plastic to synthetic materials, giving them a natural plastic appearance visually. This not only enhances the aesthetic appeal of furniture design but also allows for the prediction of the effects of different material combinations before production, optimizing material selection and the manufacturing process, as shown in Fig. 13(c).

### 5.2 Limitations and Discussions

Although our proposed material transfer method generally achieves satisfactory results, it still has some limitations under specific conditions. First, we formally define a limitation called geometric-material mismatch. This conflict happens when a material requires a specific physical texture that contradicts the fixed shape of the input image. For instance, rendering matcha powder requires a granular and uneven physical structure. Our framework tightly restricts the generation process using the original depth and normal conditions. This restriction strongly suppresses the fine physical bumps needed for the powder effect. As Fig. 14 (a) illustrates, the model successfully transfers the green hue. Yet, it leaves the surface smooth and fails to capture the realistic physical nature of the powder. Second, due to the nature of zero-shot learning, some detail features in the texture may be overemphasized, as shown in Fig. 14 (b). Furthermore, the structural consistency of the results generated by DealMaTe is




Fig. 15. Failure cases due to inaccurate geometric conditions estimation.




Fig. 16. Discussions. The results when the material diagram contains multiple types of materials or includes information beyond the materials.

fundamentally bottlenecked by the accuracy of the underlying 3D geometric condition estimator [31]. As shown in Fig. 15, the estimator may fail to capture fine geometric changes, such as vertical (first row) or horizontal (second row) ridges, as well as complex ripple structures (third row). Therefore, the transferred results lack these complex structures and appear overly smooth or flattened in these areas. Addressing the robustness of pre-processing condition extraction for highly complex geometries remains an important direction for future research.

Additionally, if the material map contains multi-material information (as shown in Fig. 16(a)), multiple types of material elements are all transferred to the result. Finally, if the material includes particularly large special shapes (as illustrated in Fig. 16(b)), these shapes may also be transferred. Furthermore, as shown in Fig. 16(c), when the reference material is illuminated by colored light, the model tends to interpret this illumination color as an intrinsic property of the material, resulting in the colored lighting being transferred to the generated object. To address these issues, future work will focus on optimizing the model’s precision in identifying materials in images to better handle these special cases and improve the overall quality and reliability of material transfer. Future work will focus on enhancing the model’s accuracy in enabling region-aware or category-aware materials selection within one image to better handle these special cases and improve the overall quality and reliability of material transfer.

## 6 Conclusion

Existing methods rely on image fine-tuning or complex architectures with reference networks, facing challenges such as text dependency, high computational costs, and insufficient structure preservation. To solve these issues, we propose the promotion framework DealMaTe. It establishes semantic alignment across material, depth, normal, and lighting images, allowing for material transfer from any 2D image to the target item. The lightweight Multi-Dim 3D Shader LoRA module provides stable 3D information support without modifying the base model weights, ensuring harmonious and stable transfer results. Additionally, shader causal attention and KV caching optimize the attention mechanism, reducing inference latency and maintaining low architectural complexity. Extensive tests demonstrate that DealMaTe provides cutting-edge performance with a compact and efficient architecture. In the future, we will further extend the 2D material transfer architecture to 3D objects and scenes.

## Acknowledgments

This work is supported by the Science and Technology Innovation 2030-Major Projects (Grant No. 2021ZD0201404), the Shenzhen Key Laboratory of New Generation Interactive Media Technology Innovation (Grant No. ZDSYS20210623092001004), the National Science and Technology Council, Taiwan (Grant No. 114-2221-E-006-114-MY3), the National Natural Science Foundation of China (Grant Nos. 62306061 and 62576076), and is sponsored by the CCF-Tencent Rhino-Bird Open Research Fund. Nisha Huang is supported by the Doctoral Student Program of the Young Science and Technology Talents Cultivation Project, China Association for Science and Technology.

## References

- [1] Louis-Philippe Asselin, Denis Laurendeau, and Jean-Francois Lalonde. 2020. Deep SVBRDF estimation on real materials. In *International Conference on 3D Vision (3DV)*. IEEE, 1157–1166.
- [2] Tim Brooks, Aleksander Holynski, and Alexei A Efros. 2023. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 18392–18402.
- [3] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A Efros, and Jun-Yan Zhu. 2022. Wearable ImageNet: Synthesizing tileable textures via dataset distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2278–2282.
- [4] Dave Zhenyu Chen, Yawar Siddiqui, Hsin-Ying Lee, Sergey Tulyakov, and Matthias Nießner. 2023. Text2tex: Text-driven texture synthesis via diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 18558–18568.
- [5] Junsong Chen, YU Jincheng, GE Chongjian, Lewei Yao, Enze Xie, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. 2024. Pixart- $\alpha$ : Fast training of diffusion transformer for photorealistic text-to-image synthesis. In *International Conference on Learning Representations*.
- [6] Ta Ying Cheng, Prafull Sharma, Mark Boss, and Varun Jampani. 2025. MARBLE: Material Recomposition and Blending in CLIP-Space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13061–13071.
- [7] 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.
- [8] Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261* (2025).

- [9] 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.
- [10] Valentin Deschaintre, George Drettakis, and Adrien Bousseau. 2020. Guided fine-tuning for large-scale material transfer. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 91–105.
- [11] Valentin Deschaintre, Yiming Lin, and Abhijeet Ghosh. 2021. Deep polarization imaging for 3d shape and svbrdf acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15567–15576.
- [12] Olga Diamanti, Connelly Barnes, Sylvain Paris, Eli Shechtman, and Olga Sorkine-Hornung. 2015. Synthesis of complex image appearance from limited exemplars. *ACM Transactions on Graphics (TOG)* 34, 2 (2015), 1–14.
- [13] Diffusers. 2023. controlnet-depth-sdxl-1.0. <https://huggingface.co/diffusers/controlnet-depth-sdxl-1.0>.
- [14] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. 2024. Scaling Rectified Flow Transformers for High-Resolution Image Synthesis. In *International Conference on Machine Learning*. PMLR, 12606–12633.
- [15] Stephanie Fu, Netanel Tamir, Shobhita Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and Phillip Isola. 2023. DreamSim: Learning New Dimensions of Human Visual Similarity using Synthetic Data. *Advances in Neural Information Processing Systems* 36 (2023), 50742–50768.
- [16] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit Haim Bermano, Gal Chechik, and Daniel Cohen-or. 2023. An Image is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion. In *International Conference on Learning Representations*.
- [17] 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 Transactions on Graphics (TOG)* 38, 4 (2019), 134–1.
- [18] Kamil Garifullin, Maxim Nikolaev, Andrey Kuznetsov, and Aibek Alanov. 2025. MaterialFusion: High-Quality, Zero-Shot, and Controllable Material Transfer with Diffusion Models. *arXiv preprint arXiv:2502.06606* (2025).
- [19] D. Guarnera, G.C. Guarnera, A. Ghosh, C. Denk, and M. Glencross. 2016. BRDF Representation and Acquisition. *Computer Graphics Forum* 35, 2 (2016), 625–650. doi:10.1111/cgf.12867
- [20] Philipp Henzler, Valentin Deschaintre, Niloy J Mitra, and Tobias Ritschel. 2021. Generative modelling of BRDF textures from flash images. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–13.
- [21] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-or. 2023. Prompt-to-Prompt Image Editing with Cross-Attention Control. In *International Conference on Learning Representations*.
- [22] Alain Hore and Djemel Ziou. 2010. Image quality metrics: PSNR vs. SSIM. In *2010 20th international conference on pattern recognition*. IEEE, 2366–2369.
- [23] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- [24] Yiwei Hu, Julie Dorsey, and Holly Rushmeier. 2019. A novel framework for inverse procedural texture modeling. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–14.
- [25] Yiwei Hu, Chengan He, Valentin Deschaintre, Julie Dorsey, and Holly Rushmeier. 2022. An inverse procedural modeling pipeline for svbrdf maps. *ACM Transactions on Graphics (TOG)* 41, 2 (2022), 1–17.
- [26] Nisha Huang, Weiming Dong, Yuxin Zhang, Fan Tang, Ronghui Li, Chongyang Ma, Xiu Li, Tong-Yee Lee, and Changsheng Xu. 2025. CreativeSynth: Cross-Art-Attention for Artistic Image Synthesis With Multimodal Diffusion. *IEEE Transactions on Visualization and Computer Graphics* (2025).
- [27] Nisha Huang, Henglin Liu, Yizhou Lin, Kaer Huang, Chubin Chen, Jie Guo, Tongyee Lee, and Xiu Li. 2025. MaTe: Images Are All You Need for Material Transfer via Diffusion Transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 15117–15126.
- [28] Nisha Huang, Yuxin Zhang, Fan Tang, Chongyang Ma, Haibin Huang, Weiming Dong, and Changsheng Xu. 2024. Diffstyler: Controllable dual diffusion for text-driven image stylization. *IEEE Transactions on Neural Networks and Learning Systems* (2024).
- [29] Unsplash Inc. 2025. Unsplash. <https://unsplash.com/>.
- [30] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. 2024. Repurposing Diffusion-Based Image Generators for Monocular Depth Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [31] Bingxin Ke, Kevin Qu, Tianfu Wang, Nando Metzger, Shengyu Huang, Bo Li, Anton Obukhov, and Konrad Schindler. 2025. Marigold: Affordable Adaptation of Diffusion-Based Image Generators for Image Analysis. *arXiv:2505.09358 [cs.CV]*
- [32] Black Forest Labs. 2024. FLUX. <https://github.com/black-forest-labs/flux>.
- [33] 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.
- [34] Xueting Li, Xiaolong Wang, Ming-Hsuan Yang, Alexei A Efros, and Sifei Liu. 2022. Scraping textures from natural images for synthesis and editing. In *European Conference on Computer Vision*. Springer, 391–408.
- [35] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. 2023. Flow Matching for Generative Modeling. In *International Conference on Learning Representations*.
- [36] Xingchao Liu, Chengyue Gong, et al. 2023. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow. In *International Conference on Learning Representations*.
- [37] 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.
- [38] Xiaohe Ma, Valentin Deschaintre, Miloš Hašan, Fujun Luan, Kun Zhou, Hongzhi Wu, and Yiwei Hu. 2025. MaterialPicker: Multi-Modal DiT-Based Material Generation. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–12.
- [39] Xinyin Ma, Gongfan Fang, Michael Bi Mi, and Xinchao Wang. 2024. Learning-to-cache: Accelerating diffusion transformer via layer caching. *Advances in Neural Information Processing Systems* 37 (2024), 133282–133304.
- [40] Rosalie Martin, Arthur Roullier, Romain Rouffet, Adrien Kaiser, and Tamy Boubekeur. 2022. MaterIA: Single Image High-Resolution Material Capture in the Wild. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 163–177.
- [41] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. 2024. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 4296–4304.
- [42] Zexu Pan, Zhaojie Luo, Jichen Yang, and Haizhou Li. 2020. Multi-Modal Attention for Speech Emotion Recognition. In *Proceedings of the Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 364–368.
- [43] William Peebles and Saining Xie. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4195–4205.
- [44] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR, 8748–8763.
- [45] Elad Richardson, Gal Metzger, Yuval Alaluf, Raja Giryes, and Daniel Cohen-Or. 2023. Texture: Text-guided texturing of 3d shapes. In *ACM SIGGRAPH Conference Proceedings*. 1–11.
- [46] Carlos Rodriguez-Pardo, Henar Dominguez-Elvira, David Pascual-Hernandez, and Elena Garces. 2023. Umat: Uncertainty-aware single image high resolution material capture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5764–5774.
- [47] 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.
- [48] Amir Rosenberger, Daniel Cohen-Or, and Dani Lischinski. 2009. Layered shape synthesis: automatic generation of control maps for non-stationary textures. *ACM Transactions on Graphics (TOG)* 28, 5 (2009), 1–9.
- [49] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 22500–22510.
- [50] 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.
- [51] Jiaming Song, Chenlin Meng, and Stefano Ermon. 2021. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*.
- [52] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing* 568 (2024), 127063.
- [53] Zhenxiang Tan, Songhua Liu, Xingyi Yang, Qiaochu Xue, and Xinchao Wang. 2025. Ominicontrol: Minimal and universal control for diffusion transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14940–14950.
- [54] Giuseppe Vecchio, Simone Palazzo, and Concetto Spampinato. 2021. SurfaceNet: Adversarial svbrdf estimation from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12840–12848.
- [55] 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.
- [56] Felix Wimbauer, Bichen Wu, Edgar Schoenfeld, Xiaoliang Dai, Ji Hou, Zijian He, Artsiom Sanakoyeu, Peizhao Zhang, Sam Tsai, Jonas Kohler, et al. 2024. Cache me if you can: Accelerating diffusion models through block caching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6211–6220.

- [57] You Wu, Kean Liu, Xiaoyue Mi, Fan Tang, Juan Cao, and Jintao Li. 2024. U-VAP: User-specified Visual Appearance Personalization via Decoupled Self Augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9482–9491.
- [58] Hu Ye, Jun Zhang, Sibao Liu, Xiao Han, and Wei Yang. 2023. IP-Adapter: Text Compatible Image Prompt Adapter for Text-to-Image Diffusion Models. *arXiv preprint arxiv:2308.06721* (2023).
- [59] Chih-Kuo Yeh, Zhanping Liu, I-Hsuan Lin, Eugene Zhang, and Tong-Yee Lee. 2020. WYSIWYG Design of Hypnotic Line Art. *IEEE Transactions on Visualization and Computer Graphics* 28, 6 (2020), 2517–2529.
- [60] 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.
- [61] Lvmin Zhang. 2023. Controlnet-v1.1-depth. [https://huggingface.co/lllyasviel/control\\_v11f1p\\_sd15\\_depth](https://huggingface.co/lllyasviel/control_v11f1p_sd15_depth).
- [62] 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*. 3836–3847.
- [63] 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/CVF Conference on Computer Vision and Pattern Recognition*. 586–595.
- [64] Yuxin Zhang, Weiming Dong, Fan Tang, Nisha Huang, Haibin Huang, Chongyang Ma, Tong-Yee Lee, Oliver Deussen, and Changsheng Xu. 2023. Prospect: Prompt spectrum for attribute-aware personalization of diffusion models. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–14.