

VideoNeuMat: Neural Material Extraction from Generative Video Models

BOWEN XUE*, The University of Manchester, United Kingdom

SAEED HADADAN*, NVIDIA, USA

ZHENG ZENG, University of California Santa Barbara, USA and NVIDIA, USA

FABRICE ROUSSELLE, NVIDIA, Switzerland

ZAHRA MONTAZERI, The University of Manchester, United Kingdom

MILOS HASAN, NVIDIA, USA




Fig. 1. VideoNeuMat generates realistic and diverse neural materials, optionally tileable, from text or image prompts. The resulting material assets can be used on any surfaces under arbitrary lighting. We use the NeuMIP [Kuznetsov et al. 2021] neural representation, which approximates parallax through neural offset (visualized as red and blue for negative and positive U-offset; V-offset looks similar). See supplementary materials for animated results.

Creating photorealistic materials for 3D rendering requires exceptional artistic skill. Generative models for materials could help, but are currently limited by the lack of high-quality training data. While recent video generative models effortlessly produce realistic material appearances, this knowledge remains entangled with geometry and lighting. We present VideoNeuMat, a two-stage pipeline that extracts reusable neural material assets from video diffusion models. First, we finetune a large video model (Wan 2.1 14B) to generate material sample videos under controlled camera and lighting trajectories, effectively creating a "virtual gonioreflectometer" that preserves the model's material realism while learning a structured measurement pattern. Second, we reconstruct compact neural materials from these videos through a Large Reconstruction Model (LRM) finetuned from a smaller Wan 1.3B

video backbone. From 17 generated video frames, our LRM performs single-pass inference to predict neural material parameters that generalize to novel viewing and lighting conditions. The resulting materials exhibit realism and diversity far exceeding the limited synthetic training data, demonstrating that material knowledge can be successfully transferred from internet-scale video models into standalone, reusable neural 3D assets. Code and data for this paper are at: <https://bowenxueai.github.io/VideoNeuMat/>.

CCS Concepts: • **Computing methodologies** → **Reflectance modeling**.

ACM Reference Format:

Bowen Xue, Saeed Hadadan, Zheng Zeng, Fabrice Rousselle, Zahra Montazeri, and Milos Hasan. 2026. VideoNeuMat: Neural Material Extraction from Generative Video Models. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '26)*, July 19–23, 2026, Los Angeles, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3799902.3811190>

1 Introduction

Creating realistic materials is critical for 3D rendering of believable virtual worlds, but requires time-consuming effort by highly skilled artists. There are few people in the world who can author truly convincing photo-real materials using classical tools such as procedural graphs, microfacet reflectance models, displacement mapping and layering; furthermore, the resulting complex layerings and graphs are often too costly for real-time rendering. Neural materials [Kuznetsov et al. 2021; Zeltner et al. 2024] are a promising

*Equal contribution.

Authors' Contact Information: Bowen Xue, The University of Manchester, Manchester, United Kingdom, bowen.xue@manchester.ac.uk; Saeed Hadadan, NVIDIA, Santa Clara, USA, shadadan@nvidia.com; Zheng Zeng, University of California Santa Barbara, Santa Barbara, USA and NVIDIA, Santa Clara, USA, zhengzeng@ucsb.edu; Fabrice Rousselle, NVIDIA, Zurich, Switzerland, frousselle@nvidia.com; Zahra Montazeri, The University of Manchester, Manchester, United Kingdom, zahra.montazeri@manchester.ac.uk; Milos Hasan, NVIDIA, Santa Clara, USA, mhasan@nvidia.com.



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGGRAPH Conference Papers '26, Los Angeles, CA, USA*
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2554-8/26/07
<https://doi.org/10.1145/3799902.3811190>

material deployment technology, relaxing some of the limitations of classical models, but they also lack straightforward solutions for photo-real content authoring.

On the other hand, recent image and video generative AI models effortlessly deliver virtually unlimited photo-real imagery and clearly understand material realism. To deliver a leap forward in the materials used in 3D rendering, we hypothesize that AI-assisted material generation can generate the quality of materials that currently only the best film material artists in the world can create. Ideally, this generative approach should directly produce neural representations, as the alternative of producing an open-ended number of layers using a large palette of classical BRDF lobes appears harder, more constrained and less suitable for real-time applications.

However, the root problem is data. The straightforward approach of supervising the desired generative model directly by high quality materials hits a barrier: as noted by recent work on generating materials, both classical [Zhou et al. 2026] and neural [Raghavan et al. 2025], there are effectively zero datasets of materials at the realism level required. How do we overcome this data scarcity, and get to generative photo-real materials?

One approach would be to hire a team of the best material artists in the world, or build the largest and highest-quality material measurement lab. However, we believe there may a less brute-force solution. Video generative models already produce realistic material appearance, which suggests the possibility of "extracting" this material knowledge into standalone neural material assets, reusable on new shapes in new scenes. Of course, there is a technical obstacle: the materials in generated videos are entangled with shapes and lighting, and presented as final RGB frames.

In this paper, we show that neural materials can indeed be bootstrapped out of video models. Our solution is to teach a video model to generate videos of materials samples under a moving light and camera, and then train another model (also built on a video backbone) to turn such videos into neural materials.

First, we finetune a recent video model, Wan 2.1 14B [Wan et al. 2025], to generate text- or image-conditioned videos of material samples under a specific light and camera trajectory. In a sense, we turn the video model into a "virtual gonioreflectometer". The finetuning is based on videos rendered using synthetic data from MatSynth [Vecchio and Deschaintre 2024], with full displacement and multi-bounce path tracing. This teaches the video model the desired light/camera trajectory and the geometry of the sample, while preserving much of its material knowledge and realism. Therefore, we are able to generate videos of materials with appearances that go well beyond MatSynth in realism and diversity.

Second, we reconstruct a material in a neural representation previously used by Kuznetsov et al. [2021] and Raghavan et al. [2025], consisting of a parallax module and main module, each defined by a feature texture, combined with a small neural network (MLP). The key challenge is that the videos generated in the first step are very small slices of the 6D reflectance space that neural materials model. While direct optimization of feature textures and MLPs to match the training views does give usable results in some cases, we also introduce a more robust and much faster reconstruction approach based on the idea of *large reconstruction models* (LRMs) [Hong et al. 2023]. We finetune the transformer layers of a smaller

video model, Wan 2.1 1.3B, to map 17-frame material videos into the feature textures of the corresponding material (with universal MLPs) in a single forward inference. The training uses a rendering loss under new view and lighting conditions, not known by the LRM, thus forcing it to produce neural materials that look plausible under a range of cameras and lights. In summary, our contributions are as follows:

- Teaching a large generative video model to become a "virtual gonioreflectometer", producing videos of material samples under a fixed light/camera trajectory, while preserving much of its prior knowledge about material realism.
- Reconstructing neural materials from the above videos, usable on new shapes in new scenes, thus solving the challenge of disentangling the generative model's material knowledge from shape and lighting.
- Introducing the first large reconstruction model (LRM) for producing neural materials (and materials in general) from sparse views, and also the first LRM finetuned from a video backbone, achieving significantly higher robustness and performance than direct optimization.

Our two-stage pipeline generates high-quality neural materials from text prompts, with realism and diversity far exceeding the small synthetic data used to train the method, thus breaking through the available data barrier (see Figure 1).

2 Related Work

Material representations. Realistic material appearance has been modeled through increasingly expressive representations, from analytic BRDF models [Cook and Torrance 1982] to measured BTFs [Dana et al. 1999] and spatially-varying BRDFs [Lensch et al. 2003]. Neural representations have emerged as a compact alternative. Rainer et al. [2019] were the first to demonstrate that neural networks can faithfully fit and compress measured BTFs. NeuMIP [Kuznetsov et al. 2021] and follow-up work [Xue et al. 2024b; Zeltner et al. 2024] encodes view- and light-dependent appearance in learned latent textures decoded by small MLPs, achieving real-time performance independent of underlying material model complexity. We adopt the NeuMIP representation as our reconstruction target, as it balances expressiveness with efficient rendering, though any differentiable material representation would work with our LRM approach.

Material capture. High-fidelity material capture traditionally requires specialized hardware such as gonioreflectometers [Matusik et al. 2023] or light stages [Debevec et al. 2023]. The concept of virtual gonioreflectometry dates back to at least Cabral et al. [1987] and Westin et al. [1992], who simulated light interactions with a synthetic microstructure. Single-image-based inverse rendering methods offer a more accessible alternative, learning neural priors to predict explicit SVBRDF maps [Deschaintre et al. 2018; Sartor and Peers 2023; Shi et al. 2020; Ying et al. 2025]. A recent line of work relights the input image into new configurations to aid SVBRDF reconstruction [Bieron et al. 2023], more recently leveraging diffusion priors [Xing et al. 2025]. We take a different approach: instead of capturing real materials, we teach a video model to simulate the acquisition process, generating measurement-like sequences that

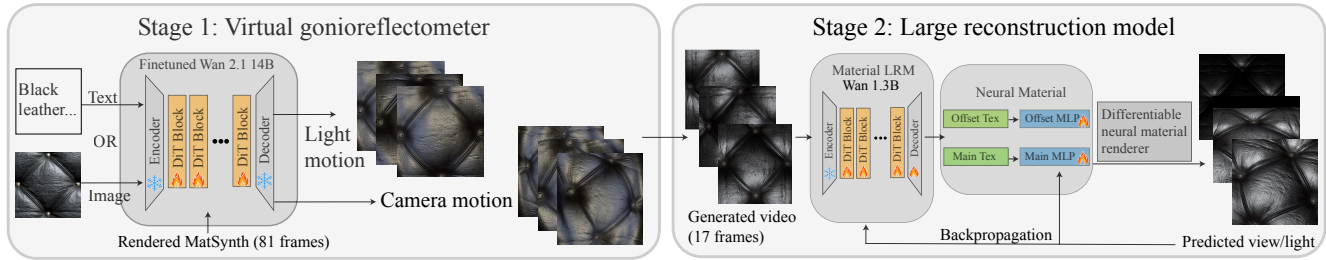


Fig. 2. Our method has two stages. First, we finetune a video diffusion model into a virtual gonio-reflectometer that generates structured material videos from text or image prompts. Second, a feed-forward LRM infers a NeuMIP-style material from 17 frames using a rendering loss under novel views and lights. The resulting material supports relighting and novel shapes.

can be converted to reusable neural assets, with no assumptions on analytic SVBRDF models.

Material generation. Classical approaches synthesize materials through procedural graphs or example-based texture synthesis [Wei and Levoy 2000]. More recently, GAN-based methods were introduced to generate directional and spatial material variation [Guo et al. 2020; Kuznetsov et al. 2019]. Recent diffusion-based methods generate SVBRDF maps from text or images [Vecchio et al. 2024a,b; Xue et al. 2024a], but remain constrained by training data: existing material datasets such as MatSynth [Vecchio and Deschaintre 2024] or Adobe Substance [Adobe Inc. 2025] lack the realism and diversity of natural imagery. PhotoMat [Zhou et al. 2023] trains a material generator from a dataset of real images collected by the authors, but the size of such datasets is necessarily limited. The closest related work to ours [Raghavan et al. 2025] generates neural materials in the NeuMIP representation using a diffusion model trained *from scratch* and supervised directly by neural versions of materials from Adobe Substance; this data is not publicly available and is still limited. Our work sidesteps this data limitation by extracting material knowledge from video generative models pretrained on internet-scale data.

Video diffusion models. Large-scale video diffusion models have demonstrated remarkable ability to synthesize temporally coherent, photorealistic sequences [Blattmann et al. 2023; Kong et al. 2024; Wan et al. 2025]. Recent work repurposes these models for 3D tasks: SV3D [Voleti et al. 2024] generates multi-view images for 3D reconstruction, while others extract motion priors for animation [Hu et al. 2023]. VideoMat [Munkberg et al. 2025] generates PBR materials for a given 3D mesh, leveraging a finetuned video diffusion model to produce a camera orbit, similar to our first stage but under uncontrolled natural lighting.

Large reconstruction models. Feed-forward reconstruction using LRMs [Hong et al. 2023] has gained traction for 3D assets: many follow-up works predict neural radiance fields [Li et al. 2023], meshes [Xu et al. 2024] or Gaussian splats [Zhang et al. 2024] from sparse views in a single forward pass. To our knowledge, we introduce the first LRM for material reconstruction, mapping short material videos directly to neural material textures without iterative optimization; ours is also the first LRM that is trained from a strong video prior instead of from scratch.

3 Structured Material Video Generation

Extracting a neural material representation from a pretrained video generative model is a challenging task that we divide into two steps: structured video generation (this section) and neural material reconstruction (section 4).

Many material capture and generation pipelines start from a single photograph, under unknown or flash illumination. This is motivated by making material capture accessible to users without complex measurement setups. As a result, view- and lighting-dependent effects (specular highlight motion, shadowing changes, and relief-induced parallax) are ambiguous: they are frequently entangled into a single 2D appearance map or suppressed by priors that favor view- and illumination-invariant outputs. Since we are generating instead of capturing, we have the opportunity to produce multiple views and lighting conditions, simulating a more advanced capture setup. We choose to represent a material sample as a short video, in which viewpoint and lighting changes become temporal variation. This allows a pretrained video diffusion model to contribute its built-in temporal coherence and preserve its strong priors.

The rest of this section answers two questions: which video sequences are effective for making the subsequent reconstruction feasible, and how can we teach the model to generate such sequences without compromising its existing knowledge of realistic materials?

3.1 Structured Video Trajectory

To make the generated path interpretable, we render training videos under controlled trajectories that separate the two sources of appearance variation: lighting and view. Each sequence is composed of two consecutive, symmetric phases, as shown in Figure 4a: (1) Light-motion phase: the camera is fixed directly above the sample while the light moves along a circular path at an elevation angle of 56.31° ($= \arctan(3/2)$, corresponding to a height of 1.5 units and radius of one unit above a unit-width sample); (2) Camera-motion phase: the light is fixed directly above while the camera orbits along a circular path at the same elevation angle.

By separating the two phases, we provide clean supervision that is easy for the video model to learn: the light-motion phase teaches illumination response, while the camera-motion phase teaches view-dependent effects and parallax.

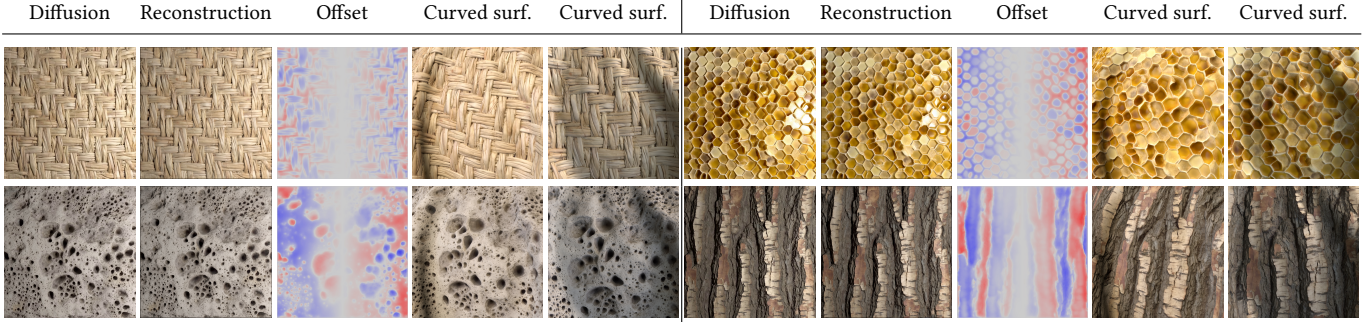


Fig. 3. For a few example materials, we show the first generated frame, the LRM-reconstructed rendering of the same frame, the U-offset map, and renderings of the material on a curved surface under different environment illuminations. The offset amount (for the neural parallax mapping effect of NeuMIP) is shown along the U-axis in texture space. A red pixel means texture look-up a few pixels to the left, while blue means offset to the right). Not only our approach can reconstruct the generated frames faithfully, but also it generates meaningful offset maps, indicating a correct understanding of the material geometry.




Fig. 4. Camera/light trajectories used for training. (a) Two-phase gonioreflectometer trajectory with 81 frames. (b) Additional 81 LRM training samples formed by cross-combining 9 camera and 9 light positions.

Perspective projection and near-field lighting. Traditional BTF acquisition uses near-orthographic cameras and typically rectifies data such that each pixel consistently maps to the same surface point across all view / light conditions. This layout simplifies data processing, but such imagery does not exist in natural videos and falls far outside the pretraining distribution of video diffusion models. We therefore adopt standard perspective projection and a small near-field area light, which produces natural patterns that better match real video statistics. Although this sacrifices per-pixel alignment in the view variation phase, we show in section 6 that the preserved prior leads to higher-quality generated materials and better downstream reconstruction.

3.2 From Pretrained Video Model to Material Model

Our intent is not to rebuild video realism from scratch, but to steer a strong pretrained model towards producing material sample videos under our controlled measurement trajectories. Our training data (MatSynth renderings) is limited in realism and diversity compared to what the base model has implicitly learned; successful finetuning should carefully retain the pretrained material realism while only teaching the base model the new trajectory and sample geometry.

Synthetic data generation. We render training videos from MatSynth using Mitsuba 3, utilizing over 5,600 materials at 1024×1024 resolution. To capture realistic parallax and self-shadowing, we apply the height maps as true geometric displacement on a 1024×1024

mesh. We augment the dataset with 90° rotations of each material. We apply a distance-based light intensity correction to the synthetic training data so that overall brightness remains constant.

Caption generation. Since MatSynth does not provide text captions, we generate them from rendered training videos using Qwen2-VL [Wang et al. 2024]. Since the captioner naturally describes visible lighting variations, we explicitly prevent lighting-related text content through two strategies:

- (1) Filtered captions: we prompt the model to describe the video freely, then post-process to remove sentences containing standalone lighting-related words (e.g., "light", "lighting", "illumination");
- (2) Material-only captions: we directly instruct the captioner to focus only on intrinsic material attributes and ignore all lighting information.

During training, we mix both caption types equally.


Full finetuning and LoRA. We choose a recent open-source model, Wan2.1-14B, due to its strong generative prior and clean software implementation in DiffSynth-Studio [ModelScope 2024]. We finetune the model using the standard diffusion objective of the base model, without introducing task-specific losses. We explored two finetuning strategies with different prior-preservation characteristics: LoRA and full finetuning, both of which lead to realistic results. LoRA (rank 32) updates only low-rank adapters, which in theory better preserve the pretrained prior. However, this comes at a trade-off: LoRA achieves weaker trajectory compliance (see Table 3b), suggesting the measurement trajectory requires broader weight modifications. For full finetuning, we apply classifier-free guidance dropout (probability 0.1) to preserve unconditional generation capability, ensuring negative prompts remain effective; LoRA inherently retains this capability without explicit dropout. We use both model types in our results.

Image conditioning. We also finetune an image-conditioned variant that generates from a reference photograph. For training, we render MatSynth materials under natural illumination sampled from

Table 1. Ablation study results. **Best**, **2nd best**, **2nd worst**, **worst**. "Rand. init." is short for random initialization.

Experiment	Features				Metrics			
	Pretained WAN	Pretained MLP	Trainable MLP	VAE upsampler	LPIPS ↓	MAPE ↓	MSE ↓	PSNR ↑
Full method	✓	✓	✓	✓	0.3017	0.0407	0.005667	24.29
Linear upsampler	✓	✓	✓	✓	0.2919	0.0454	0.005767	24.25
Frozen MLP	✓	✓	✓	✓	0.3155	0.0461	0.006105	23.96
Rand. init. MLP	✓	✓	✓	✓	0.3341	0.0565	0.009961	22.13
Rand. init. LRM		✓	✓	✓	0.3543	0.0563	0.008769	22.45
<i>Network size (base: 64×4)</i>								
128×6	✓		✓	✓	0.3211	0.0483	0.007285	23.22
32×2	✓		✓	✓	0.3325	0.0501	0.007287	23.17
<i>LRM input views (base: 17 frames)</i>								
5 inputs	✓	✓	✓	✓	0.3352	0.0517	0.006763	23.45
41 inputs	✓	✓	✓	✓	0.3009	0.0435	0.006151	24.04

Note: Checkmarks indicate features present. Full method uses 64×4 MLP and LRM 17 input frames.



(a) Linear layer

(b) VAE decoder

Fig. 5. **LRM upsampler**. A linear layer as the upsampler of the DiT tokens processes each token separately from the others, causing patch artifacts, while the convolutional layers in the VAE decoder resolve this issue.

500 environment maps (normalized to equivalent irradiance), encouraging disentanglement of intrinsic material appearance from incident lighting (Figure 9).

Training details. We finetune Wan2.1-14B updating only the diffusion transformer (DiT) while keeping the variational autoencoder and text encoder (UMT5-XXL) frozen. We train models at 512×512 with 81-frame sequences (the native video length of the Wan model), using AdamW with a learning rate of 1×10^{-5} . Training is conducted on 4 NVIDIA A100 GPUs. For the image-conditioned variant, we finetune Wan2.1-I2V-14B-480P with identical hyperparameters, conditioning on a reference material photograph.

Tileable material variant. For applications requiring tileable materials, we further finetune our model on the tileable subset. This restricts the training set to roughly 3,000 materials with periodic boundary conditions. The resulting model generates outputs that are largely tileable, though with occasional imperfect samples. We also apply a latent-space post-processing step, which modifies a narrow band along the boundaries to reduce seam visibility. See the supplementary document for details.

3.3 Inference and Evaluation

Inference with novel prompts. Although training captions originate from synthetic MatSynth renderings, at inference the finetuned model can be driven by prompts that are outside the training caption distribution, producing material appearances that go beyond MatSynth in both diversity and realism.

Prompt expansion. We use an LLM (Claude Opus 4.5 [Anthropic 2026]) to expand short material prompts into detailed captions matching the training format, adding surface micro-geometry, weathering, color variation, and other material-specific cues. Examples are provided in the supplementary material.

Flexible inference resolution. Thanks to the DiT architecture of Wan2.1, our finetuned model exhibits resolution flexibility beyond the training resolution. We train model at 512×512 , but inference at resolutions up to 1920×1920 without architectural modification, enabling high-resolution material videos with fine detail suitable for production use cases.

BTF sequence quality metric. Since no existing metric quantifies BTF sequence quality, we propose *RPC (Residual Parallax Coherence)*. The key insight is that for a perfectly rectified BTF sequence, residual optical flow should correlate monotonically with camera baseline: larger pose changes induce proportionally larger parallax for non-planar surface details.

We first obtain rectified frames by unwarping rendered images to a canonical UV space using reference-plane position AOV to remove perspective distortion. On Phase 2 (the camera-motion phase), we form frame pairs $(t, t + \Delta)$ with multiple temporal lags Δ to cover a wide range of pose baselines. For each pair, we compute a composite baseline combining rotation and normalized translation:

$$b_{t,\Delta} = \sqrt{\Delta\theta_{t,\Delta}^2 + \left(\frac{\|\Delta\mathbf{t}_{t,\Delta}\|}{s}\right)^2}, \quad (1)$$

where $\Delta\theta_{t,\Delta}$ is the relative rotation angle, $\Delta\mathbf{t}_{t,\Delta}$ is the translation difference, and s is the world-space extent of the material sample. We then measure the median optical-flow magnitude $m_{t,\Delta}$ over a robust intensity mask that excludes the 10th/90th intensity percentiles after bilateral filtering. RPC loss penalizes weak monotonic dependence:

$$L_{\text{RPC}} = 1 - \max(0, \rho_s(\{b_{t,\Delta}\}, \{m_{t,\Delta}\})), \quad (2)$$

where ρ_s is Spearman correlation, chosen over Pearson for robustness to outliers and to capture monotonic relationships. Optical flow is computed using Dual TV-L1 [Zach et al. 2007].

4 Neural Material Reconstruction

In this section, we discuss ways to obtain a neural representation from the videos of our gonireflectometer. Specifically, the goal is to obtain the parameters of a given neural material model (in this case NeuMIP) from a set of frames under known geometry, camera pose, and illumination. Namely, we will discuss *direct optimization*, and a data-driven approach inspired by *Large Reconstruction Models (LRMs)* [Hong et al. 2023]. Despite our focus on NeuMIP, the approach is applicable to any neural (or even classical) material model with a differentiable rendering operator available.

4.1 NeuMIP representation

Instead of the original NeuMIP formulation [Kuznetsov et al. 2021], we use the recent modification by Raghavan et al. [2025], which does not use a multi-resolution pyramid, and uses universal instead of per-material MLPs. Our representation can be easily converted into the original one by distillation if desired.

This architecture has four components: a per-material offset feature texture T_{off} , a per-material reflectance feature texture T_{rgb} , a universal offset MLP M_{off} , and a universal reflectance MLP M_{rgb} . These can be thought of as an offset module (responsible for the parallax effect) and a reflectance module (returning the final RGB reflectance value).

Given a material-space location \mathbf{u} , incoming (light) direction ω_i and outgoing (camera) direction ω_o , the final reflectance value $f_{\phi, \theta}(\mathbf{u}, \omega_i, \omega_o)$, where ϕ are the MLP parameters and θ are the feature texture parameters, is evaluated as follows:

$$\Delta \mathbf{u} = M_{\text{off}}(T_{\text{off}}[\mathbf{u}], \omega_o) \in [-1, 1]^2, \quad (3)$$

$$f_{\phi, \theta}(\mathbf{u}, \omega_i, \omega_o) = M_{\text{rgb}}(T_{\text{rgb}}[\mathbf{u} + \Delta \mathbf{u}], \omega_i, \omega_o) \in \text{RGB}. \quad (4)$$

This evaluates the offset module at \mathbf{u} , computes an offset $\Delta \mathbf{u}$, and evaluates the reflectance module at $\mathbf{u} + \Delta \mathbf{u}$. Feature texture queries are by bilinear interpolation. See Raghavan et al. [2025] (Section 3) for more details. We require a rendering operator (differentiable in ϕ, θ) that evaluates the model on a flat plane with a perspective camera and a point light, which is straightforward to implement.

4.2 Direct optimization

A straightforward approach to obtain a neural representation from a sequence of frames under known geometry and lighting is to directly optimize the parameters of the neural material in an inverse rendering framework. This involves differentially rendering candidate images of the neural material and comparing them to the input frames using an image-based loss. The parameters of the neural model are then iteratively updated using gradient descent. Our experiments show direct optimization can sometimes produce good results, but with three significant limitations.

Limitations. Direct optimization requires perfect multi-view consistency; otherwise, conflicting information prevents convergence to a sharp solution. Therefore, it works best on traditionally rendered videos. Since the videos generated by the diffusion model have a degree of view/light inconsistency, the reconstruction has significant artifacts as shown in Figure 7. Worse, direct optimization motivates the neural model to only explain the input video sequence. Since the video covers a limited slice of the 4D space (ω_i, ω_o) , and the optimization does not make use of any prior knowledge, the resulting neural material can fail to robustly generalize to unseen camera/light locations (Figure 7). Lastly, direct optimization is slow due to iterative rendering and optimization per-material (1.5 hours to fit 81 frame videos at a resolution of 1024).

4.3 Large Reconstruction Model

To overcome the limitations of direct optimization, we propose a feed-forward Large Reconstruction Model (LRM) that replaces per-material, iterative optimization with a single forward pass: given a material video in the gonioreflectometer’s camera/light trajectory

as input, the LRM outputs the latent textures of our neural representation. Our LRM uses only 17 input frames, a subset of the 81 generated by the video model. Despite the sparse inputs, the LRM is much more robust compared to direct optimization, both for synthetic and especially for diffusion-generated materials (Figure 7). It is also much faster than direct optimization (roughly 4 seconds per material at resolution 1024). Figure 3 shows the LRM approach can successfully reconstruct the material appearance and geometry (through the offset module).

Training. We do not directly supervise the LRM output using a set of ground truth latent textures, as we do not have access to such a dataset (Raghavan et al. [2025] did not release their data). Instead, we use a rendering loss: we render images using the latent textures for new view / lighting conditions unknown to the LRM, and compute an L2 loss against the ground truth MatSynth rendering. This rendering loss is calculated not only on the gonioreflectometer’s trajectory (81 frames), but also on new combinations of light and camera locations (81 frames), Figure 4. The latter set consists of camera and light locations placed on the hemisphere on center texel, ensuring 3 evenly spaced samples in each dimension of the 4D space of w_i, w_o , thus $3^4 = 81$ frames.

LRM architecture. We adopt the Wan2.1-1.3B video diffusion model and repurpose it into a non-generative regressor that converts a material video into NeuMIP latent textures. The DiT operates in Wan’s VAE latent space, which we reuse. We encode the input frames using the frozen VAE encoder before feeding them to the DiT. The DiT output tokens corresponding to the first frame of the video are then fed to the VAE decoder (with modified final layers) to obtain the NeuMIP latent textures T_{off} and T_{rgb} (12 channels each). All other output tokens are ignored (i.e., do not affect loss computation). The DiT and VAE decoder are trained end-to-end.

We initialize both the DiT and VAE decoder weights using the pretrained Wan weights as opposed to random initialization. As shown in Table 1, this produces better results. We attribute this to the pretrained model’s ability to parse temporal appearance variation from input videos, a skill that transfers directly to material reconstruction. The only layers that are randomly initialized are the DiT and VAE heads (final layers), due to the domain gap.

NeuMIP MLP decoders. The textures T_{off} and T_{rgb} produced by the LRM are fed into the offset decoder M_{off} and reflectance decoder M_{rgb} to render images under novel view / lighting. The MLPs have 4 hidden layers with 64 neurons and are trained end-to-end with the LRM. The MLPs can be randomly initialized, or pretrained on a set of materials. The latter produces better results, as shown in Table 1.

Input frame count. Video diffusion VAEs typically apply a temporal downsampling factor of 4 after an independent keyframe, requiring frame counts of the form $1 + 4n$, hence our choices of 5, 17, 41, and 81 frames. Our two-phase trajectory naturally suggests a minimum sampling density: enough frames to capture both illumination response (light-motion phase) and view-dependent effects (camera-motion phase). We find 17 uniformly-spaced frames optimal; reducing to 5 frames undersamples the light-motion phase and

degrades perceptual quality (LPIPS), while 41 frames offers no improvement and slightly hurts generalization, likely due to overfitting to redundant observations.

Training details. We pretrain universal NeuMIP MLPs on 512 random MatSynth materials, then jointly train with the LRM on MatSynth renderings. We train the LRM on MatSynth renderings at resolution 256^2 for 75 hours at fp32 precision, followed by 60 hours at bf16 at resolution 1024^2 , using 64 A100 GPUs. In addition to L2 pixel-wise loss, we used LPIPS with weight 0.1.

5 Results

Figure 10 shows the final output of our pipeline, i.e., generated neural materials from various prompts, rendered on flat and curved geometries. See supplementary for more results.

Qualitative comparison. We compare against three baselines: GNM [Raghavan et al. 2025], RealMat [Zhou et al. 2026], and Reflectance-Fusion [Xue et al. 2024a] in Figure 8. We use examples from their original publications and generate results with the same text prompts. Note that viewing and lighting conditions differ across methods. Our approach consistently achieves higher output resolution and more pronounced displacement effects in all comparisons.

Single-image-to-material comparison. We further evaluate image-conditioned material generation on 7 materials from Adobe Substance [Adobe Inc. 2025], not used in training. Given a single input image, we compare against CHORD [Ying et al. 2025] and MatFusion [Sartor and Peers 2023]. As shown in Figure 6, our method better preserves the material structure and relief cues across diverse materials. We report quantitative results in Table 2, evaluating both the displayed frame and additional test views from the same materials. High-relief materials are particularly challenging for competing SVBRDF methods, especially since the predicted normal maps were not integrated into a height field to capture parallax effects. Note that our approach is a generative model conditioned on the input image, so some differences in color and tone are expected; still, our method achieves the lowest error across all metrics.

Quantitative evaluation. A key hypothesis of our work is that finetuning the video model on synthetic, measurement-like trajectories can steer the generator toward our protocol *without* collapsing to a synthetic-rendered regime, i.e., it retains the video model’s realistic prior. We therefore evaluate two complementary signals (Table 3a). We compute FVD to OpenVid-1M [Nan et al. 2024] to measure distributional closeness to real videos in a standard video embedding space. The FVD serves as an indication of whether finetuning preserves real-video statistics rather than drifting toward synthetic-rendered artifacts. Our generated clips (783 samples) achieve a lower FVD (118.04) than MatSynth renderings (138.86). We additionally report RealMatScore [Zhou et al. 2026], a material-focused perceptual metric, where our results improve over MatSynth (0.6289 vs. 0.5463). These results suggest that finetuning retains the pretrained model’s real-video appearance prior, rather than collapsing toward a purely synthetic-rendered regime, as further supported by the nearest-neighbor analysis in the supplementary material.

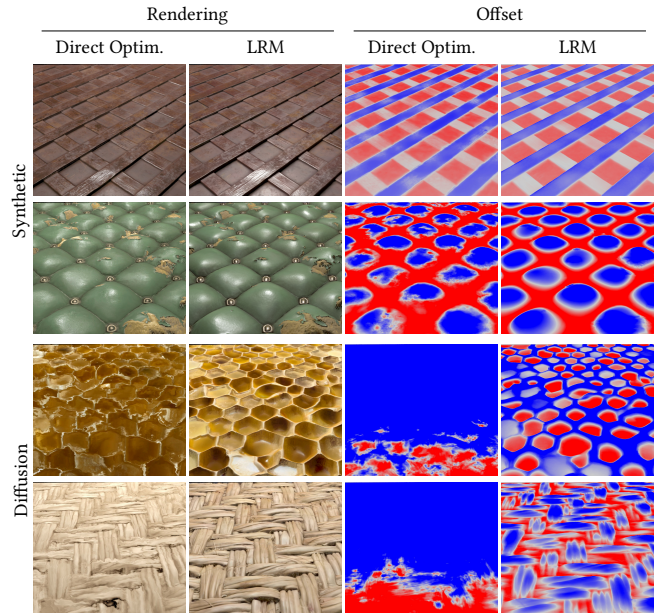


Fig. 7. We compare direct optimization with LRM by showing renderings of unseen views during training. Direct optimization fails to reach the correct offset in diffusion materials, likely due to the misalignments in the generated videos. Even for synthetic materials with no misalignments, the LRM has better generalization to unseen views, since it was trained on more camera/light settings. Full videos are included in the supplementary video.

6 Ablation Study

Perspective vs. rectified projection (Table 3b). Traditional BTF acquisition employs orthographic cameras with rectified sensing, ensuring pixel-to-surface correspondence across frames. We compare this against our perspective projection approach by training two video models on the same material set, differing only in projection mode. The perspective model preserves material realism at the cost of correspondence ambiguity, yet still achieves competitive reconstruction quality on generated videos.

Video model finetuning strategies (Table 3b). We provide two finetuning variants with different resource-quality trade-offs. Our LoRA variant (rank 32) requires less GPU memory and training time while still producing realistic materials that follow the measurement trajectory. Our full finetuning variant updates all DiT parameters and achieves the best trajectory coherence (RPC 0.2106 vs. 0.3102) and perceptual realism (RealMatScore 0.6289 vs. 0.6190, though this difference is very small). User can choose based on their computational budget and quality requirements.

LRM ablations (Table 1). The table shows that not using pretrained weights for either NeuMIP MLP or Wan DiT/VAE decoder has the lowest scores, as explained in subsection 4.3. The pretrained MLP was trained on 512 randomly selected MatSynth materials. It also reveals freezing the pre-trained MLP during LRM training leads to degraded results. Additionally, utilizing a smaller/larger NeuMIP MLP, or using fewer or more input views do not change the metrics




Fig. 6. Comparison on material generation from a single input photograph. The input image is shown in the first column. The outputs are rendered from an oblique viewpoint, significantly different from the input, to better evaluate reconstruction quality. “Our full” denotes the complete pipeline: image-conditioned video generation followed by LRM reconstruction. “Our LRM” applies the LRM directly to synthetic rendered videos, isolating the reconstruction quality from the video generation step. Additional examples are shown in supplementary.

Method	MSE↓	LPIPS↓	PSNR↑
<i>Displayed frame</i>			
CHORD	0.0469	0.465	13.29
MatFusion	0.0271	0.594	15.67
Ours	0.0142	0.292	18.47
<i>Test views</i>			
CHORD	0.0332	0.400	14.80
MatFusion	0.0254	0.572	15.96
Ours	0.0105	0.260	19.81

Table 2. Quantitative comparison on the image-to-texture test materials. “Displayed frame” evaluates the frame shown in the figure; “Test views” evaluates additional held-out views.



Fig. 8. **Qualitative comparison with baseline methods:** GNM [Raghavan et al. 2025], RealMat [Zhou et al. 2026] and ReflectanceFusion [Xue et al. 2024a]. All three previous methods are based on strong diffusion backbones, and produce meaningful results. In our opinion, our results are the most successful at extracting realistic material appearance from the backbone generative model. Credits: GNM panels © Generative Neural Materials project authors; RealMat panels © 2026 Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd.; ReflectanceFusion panels © 2024 The Authors.



Fig. 9. **Image-to-video generation results.** Given a single material photograph as input conditioning (first row), our fine-tuned model generates structured material videos adhering to the learned camera and lighting trajectory. The middle row shows first frame of the video-generated appearance under novel viewing conditions, and the last row shows the LRM-reconstructed neural materials.

Image credits: Input condition images 1-8 from left to right are Adobe Stock [Adobe Stock 2026] images, © BubbleSnap, Earthen Stock, Anna Klimchuk, mr_studioo, Elena, mila103, Cleopatra and shushan1974 / stock.adobe.com; the last input condition image designed by Magnific, www.magnific.com.

Table 3. (a) Appearance Evaluation. FVD to OpenVid-1M measures distributional closeness to real videos and is used as a proxy for preserving the pretrained video prior (not as a direct material-correctness metric). RealMatScore measures material-focused perceptual appearance. (b) Ablation study of training the video model compares finetuning strategies and projection modes.

(a) Appearance			(b) Ablation			
	MatSynth	Ours		Ours	Rectified	
			LoRA	Finetune	Projection	
FVD↓	138.86	118.04	RPC↓	0.3102	0.2106	0.2146
RealMat↑	0.5463	0.6289	RealMat↑	0.6190	0.6289	0.5992

dramatically. Lastly, using a linear layer to upsample the DiT tokens instead of the VAE decoder leads to lower scores, except for LPIPS. Comparing renderings of either upsamplers in Figure 5 reveals the patchy artifacts of linear upsampler as each upsampled token is processed independently of others.

7 Limitations and Future Work

The NeuMIP representation has limits in its ability to approximate true 3D effects; a more sophisticated neural representation can be designed, as our LRM design is representation-agnostic. Our pipeline is finetuned solely on MatSynth, which lacks complexity such as translucent, layered or fiber materials. While the video diffusion model can still generate these complex materials due to its generative prior, augmenting MatSynth, or training the LRM to predict generated video frames from other generated frames could further improve the complexity of the generated materials. Lastly, further improvements on our tileability solution may be possible.

8 Conclusion

We presented VideoNeuMat, a pipeline that extracts reusable neural materials from video diffusion models. By finetuning a pretrained

video model as a "virtual gonioreflectometer" and introducing a large reconstruction model for single-pass neural material inference from such generated videos, we successfully transfer material knowledge from internet-scale video data into standalone material assets, usable on new surfaces under new views and lighting. Our results demonstrate higher realism than synthetic training data and previous diffusion-based solution, breaking through the data scarcity barrier inherent to the material generation problem.

Acknowledgments

We would also like to thank Aaron Lefohn for his support, and NVIDIA for supporting the work through an NVIDIA academic partnership.

References

- Adobe Inc. 2025. *Adobe Substance 3D*. <https://www.adobe.com/products/substance3d.html>
- Adobe Stock. 2026. Adobe Stock. <https://stock.adobe.com/>
- Anthropic. 2026. Claude (Opus 4.5). <https://www.anthropic.com/news/claude-opus-4-5>. Large language model.
- James Bieron, Xin Tong, and Pieter Peers. 2023. Single Image Neural Material Relighting. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) (SIGGRAPH '23). Association for Computing Machinery, New York, NY, USA, Article 80, 11 pages. doi:10.1145/3588432.3591515
- Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. 2023. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127* (2023).
- Brian Cabral, Nelson Max, and Rebecca Springmeyer. 1987. Bidirectional reflection functions from surface bump maps. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '87)*. Association for Computing Machinery, New York, NY, USA, 273–281. doi:10.1145/37401.37434
- R. L. Cook and K. E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.* 1, 1 (Jan. 1982), 7–24. doi:10.1145/357290.357293
- Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. 1999. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.* 18, 1 (Jan. 1999), 1–34. doi:10.1145/300776.300778
- Paul Debevec, Tim Hawkins, Chris Tchou, Westley Sarokin, and Mark Sagar. 2023. *Acquiring the Reflectance Field of a Human Face* (1 ed.). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3596711.3596762>



Fig. 10. **Reconstruction results.** A selection of materials generated by our pipeline from text prompts, shown on a flat and curved surface under different illuminations. Note the realism of the results and the ability to handle non-trivial geometry (leaves, fur, fabric) that cannot be represented by heightfields. Please see the extensive supplementary materials for animated results, showing parallax effects.

- Valentin Deschaintre, Miika Aittala, Fr'edo Durand, George Drettakis, and Adrien Bousseau. 2018. Single-Image SVBRDF Capture with a Rendering-Aware Deep Network. *ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)* 37, 128 (aug 2018), 15. <http://www.sop.inria.fr/revues/Basilic/2018/DADDB18>
- Yu Guo, Cameron Smith, Miloš Hašan, Kalyan Sunkavalli, and Shuang Zhao. 2020. MaterialGAN: reflectance capture using a generative SVBRDF model. *ACM Trans. Graph.* 39, 6, Article 254 (Nov. 2020), 13 pages.
- Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. 2023. LRM: Large Reconstruction Model for Single Image to 3D. *ArXiv abs/2311.04400* (2023). <https://api.semanticscholar.org/CorpusID:265050698>
- Li Hu, Xin Gao, Peng Zhang, Ke Sun, Bang Zhang, and Liefeng Bo. 2023. Animate Anyone: Consistent and Controllable Image-to-Video Synthesis for Character Animation. *arXiv preprint arXiv:2311.17117* (2023).
- Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. 2024. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603* (2024).
- Alexandr Kuznetsov, Miloš Hašan, Zexiang Xu, Ling-Qi Yan, Bruce Walter, Nima Khademi Kalantari, Steve Marschner, and Ravi Ramamoorthi. 2019. Learning generative models for rendering specular microgeometry. *ACM Trans. Graph.* 38, 6, Article 225 (Nov. 2019), 14 pages.
- Alexandr Kuznetsov, Krishna Mullia, Zexiang Xu, Miloš Hašan, and Ravi Ramamoorthi. 2021. NeuMIP: multi-resolution neural materials. *ACM Trans. Graph.* 40, 4, Article 175 (July 2021), 13 pages. doi:10.1145/3450626.3459795
- Hendrik P. A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. 2003. Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph.* 22, 2 (April 2003), 234–257. doi:10.1145/636886.636891
- Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. 2023. Instant3D: Fast Text-to-3D with Sparse-View Generation and Large Reconstruction Model. *arXiv:2311.06214 [cs.CV]* <https://arxiv.org/abs/2311.06214>
- Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. 2023. *A Data-Driven Reflectance Mode* (1 ed.). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3596711.3596750>
- ModelScope. 2024. DiffSynth-Studio. <https://github.com/modelscope/DiffSynth-Studio>.
- J. Munkberg, Z. Wang, R. Liang, T. Shen, and J. Hasselgren. 2025. VideoMat: Extracting PBR Materials from Video Diffusion Models. *Computer Graphics Forum* 44, 4 (2025), e70180. <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.70180> doi:10.1111/cgf.70180
- Kejan Nan, Rui Xie, Penghao Zhou, Tiehan Fan, Zhenheng Yang, Zhijie Chen, Xiang Li, Jian Yang, and Ying Tai. 2024. OpenVid-1M: A Large-Scale High-Quality Dataset for Text-to-video Generation. *arXiv preprint arXiv:2407.02371* (2024).
- Nithin Raghavan, Krishna Mullia, Alexander Trevithick, Fujun Luan, Miloš Hašan, and Ravi Ramamoorthi. 2025. Generative Neural Materials. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25)*. Association for Computing Machinery, New York, NY, USA, Article 162, 11 pages. doi:10.1145/3721238.3730746
- Gilles Rainer, Wenzel Jakob, Abhijeet Ghosh, and Tim Weyrich. 2019. Neural BTF Compression and Interpolation. *Computer Graphics Forum* 38, 2 (2019), 235–244. <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13633> doi:10.1111/cgf.13633
- Sam Sartor and Pieter Peers. 2023. MatFusion: A Generative Diffusion Model for SVBRDF Capture. In *SIGGRAPH Asia 2023 Conference Papers (Sydney, NSW, Australia) (SA '23)*. Association for Computing Machinery, New York, NY, USA, Article 86, 10 pages. doi:10.1145/3610548.3618194
- Liang Shi, Beichen Li, Miloš Hašan, Kalyan Sunkavalli, Tamy Boubekeur, Radomir Mech, and Wojciech Matusik. 2020. MATch: Differentiable Material Graphs for Procedural Material Capture. *ACM Trans. Graph.* 39, 6 (Dec. 2020), 1–15.
- Giuseppe Vecchio and Valentin Deschaintre. 2024. MatSynth: A Modern PBR Materials Dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Giuseppe Vecchio, Rosalie Martin, Arthur Roullier, Adrien Kaiser, Romain Rouffet, Valentin Deschaintre, and Tamy Boubekeur. 2024a. ControlMat: A Controlled Generative Approach to Material Capture. *ACM Trans. Graph.* 43, 5, Article 164 (Sept. 2024), 17 pages. doi:10.1145/3688830
- Giuseppe Vecchio, Renato Sortino, Simone Palazzo, and Concetto Spampinato. 2024b. MatFuse: Controllable Material Generation with Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4429–4438.
- Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitrii Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. 2024. SV3D: Novel Multi-view Synthesis and 3D Generation from a Single Image using Latent Video Diffusion. In *European Conference on Computer Vision (ECCV)*.
- Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruiqi Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wentu Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. 2025. Wan: Open and Advanced Large-Scale Video Generative Models. *arXiv preprint arXiv:2503.20314* (2025).
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. Qwen2-VL: Enhancing Vision-Language Model's Perception of the World at Any Resolution. *arXiv preprint arXiv:2409.12191* (2024).
- Li-Yi Wei and Marc Levoy. 2000. Fast texture synthesis using tree-structured vector quantization. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '00)*. ACM Press/Addison-Wesley Publishing Co., USA, 479–488. doi:10.1145/344779.345009
- Stephen H. Westin, James R. Arvo, and Kenneth E. Torrance. 1992. Predicting reflectance functions from complex surfaces. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 255–264. doi:10.1145/142920.134075
- Youxin Xing, Zheng Zeng, Youyang Du, Lu Wang, and Beibei Wang. 2025. Diffusion-Guided Relighting for Single-Image SVBRDF Estimation. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25)*. Association for Computing Machinery, New York, NY, USA, Article 162, 10 pages. doi:10.1145/3757377.3763809
- 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).
- Bowen Xue, Giuseppe Claudio Guarnera, Shuang Zhao, and Zahra Montazeri. 2024a. ReflectanceFusion: Diffusion-based text to SVBRDF Generation. In *Eurographics Symposium on Rendering*, Eric Haines and Elena Garces (Eds.). The Eurographics Association. doi:10.2312/sr.20241152
- Bowen Xue, Shuang Zhao, Henrik Wann Jensen, and Zahra Montazeri. 2024b. A Hierarchical Architecture for Neural Materials. *Computer Graphics Forum* 43, 6 (2024), e15116. <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.15116> doi:10.1111/cgf.15116
- Zhi Ying, Boxiang Rong, Jingyu Wang, and Maoyuan Xu. 2025. Chord: Chain of Rendering Decomposition for PBR Material Estimation from Generated Texture Images. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25)*. Association for Computing Machinery, New York, NY, USA, Article 164, 11 pages. doi:10.1145/3757377.3763848
- C. Zach, T. Pock, and H. Bischof. 2007. A Duality Based Approach for Realtime TV-L1 Optical Flow. In *Pattern Recognition*, Fred A. Hamprecht, Christoph Schnörr, and Bernd Jähne (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 214–223.
- Tizian Zeltner, Fabrice Rousselle, Andrea Weidlich, Petrik Clarberg, Jan Novák, Benedikt Bitterli, Alex Evans, Tomáš Davidovič, Simon Kallweit, and Aaron Lefohn. 2024. Real-time Neural Appearance Models. *ACM Trans. Graph.* 43, 3, Article 33 (June 2024), 17 pages. doi:10.1145/3659577
- Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. 2024. GS-LRM: Large Reconstruction Model for 3D Gaussian Splatting. *European Conference on Computer Vision* (2024).
- X. Zhou, P. Figueiredo, M. Hašan, V. Deschaintre, P. Guerrero, Y. Hu, and N. Khademi Kalantari. 2026. RealMat: Realistic Materials With Diffusion and Reinforcement Learning. *Computer Graphics Forum* (2026), e70307. <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.70307> doi:10.1111/cgf.70307
- Xilong Zhou, Milos Hasan, Valentin Deschaintre, Paul Guerrero, Yannick Hold-Geoffroy, Kalyan Sunkavalli, and Nima Khademi Kalantari. 2023. PhotoMat: A Material Generator Learned from Single Flash Photos. In *ACM SIGGRAPH 2023 Conference Proceedings (Los Angeles, CA, USA) (SIGGRAPH '23)*. Association for Computing Machinery, New York, NY, USA, Article 49, 11 pages. doi:10.1145/3588432.3591535