

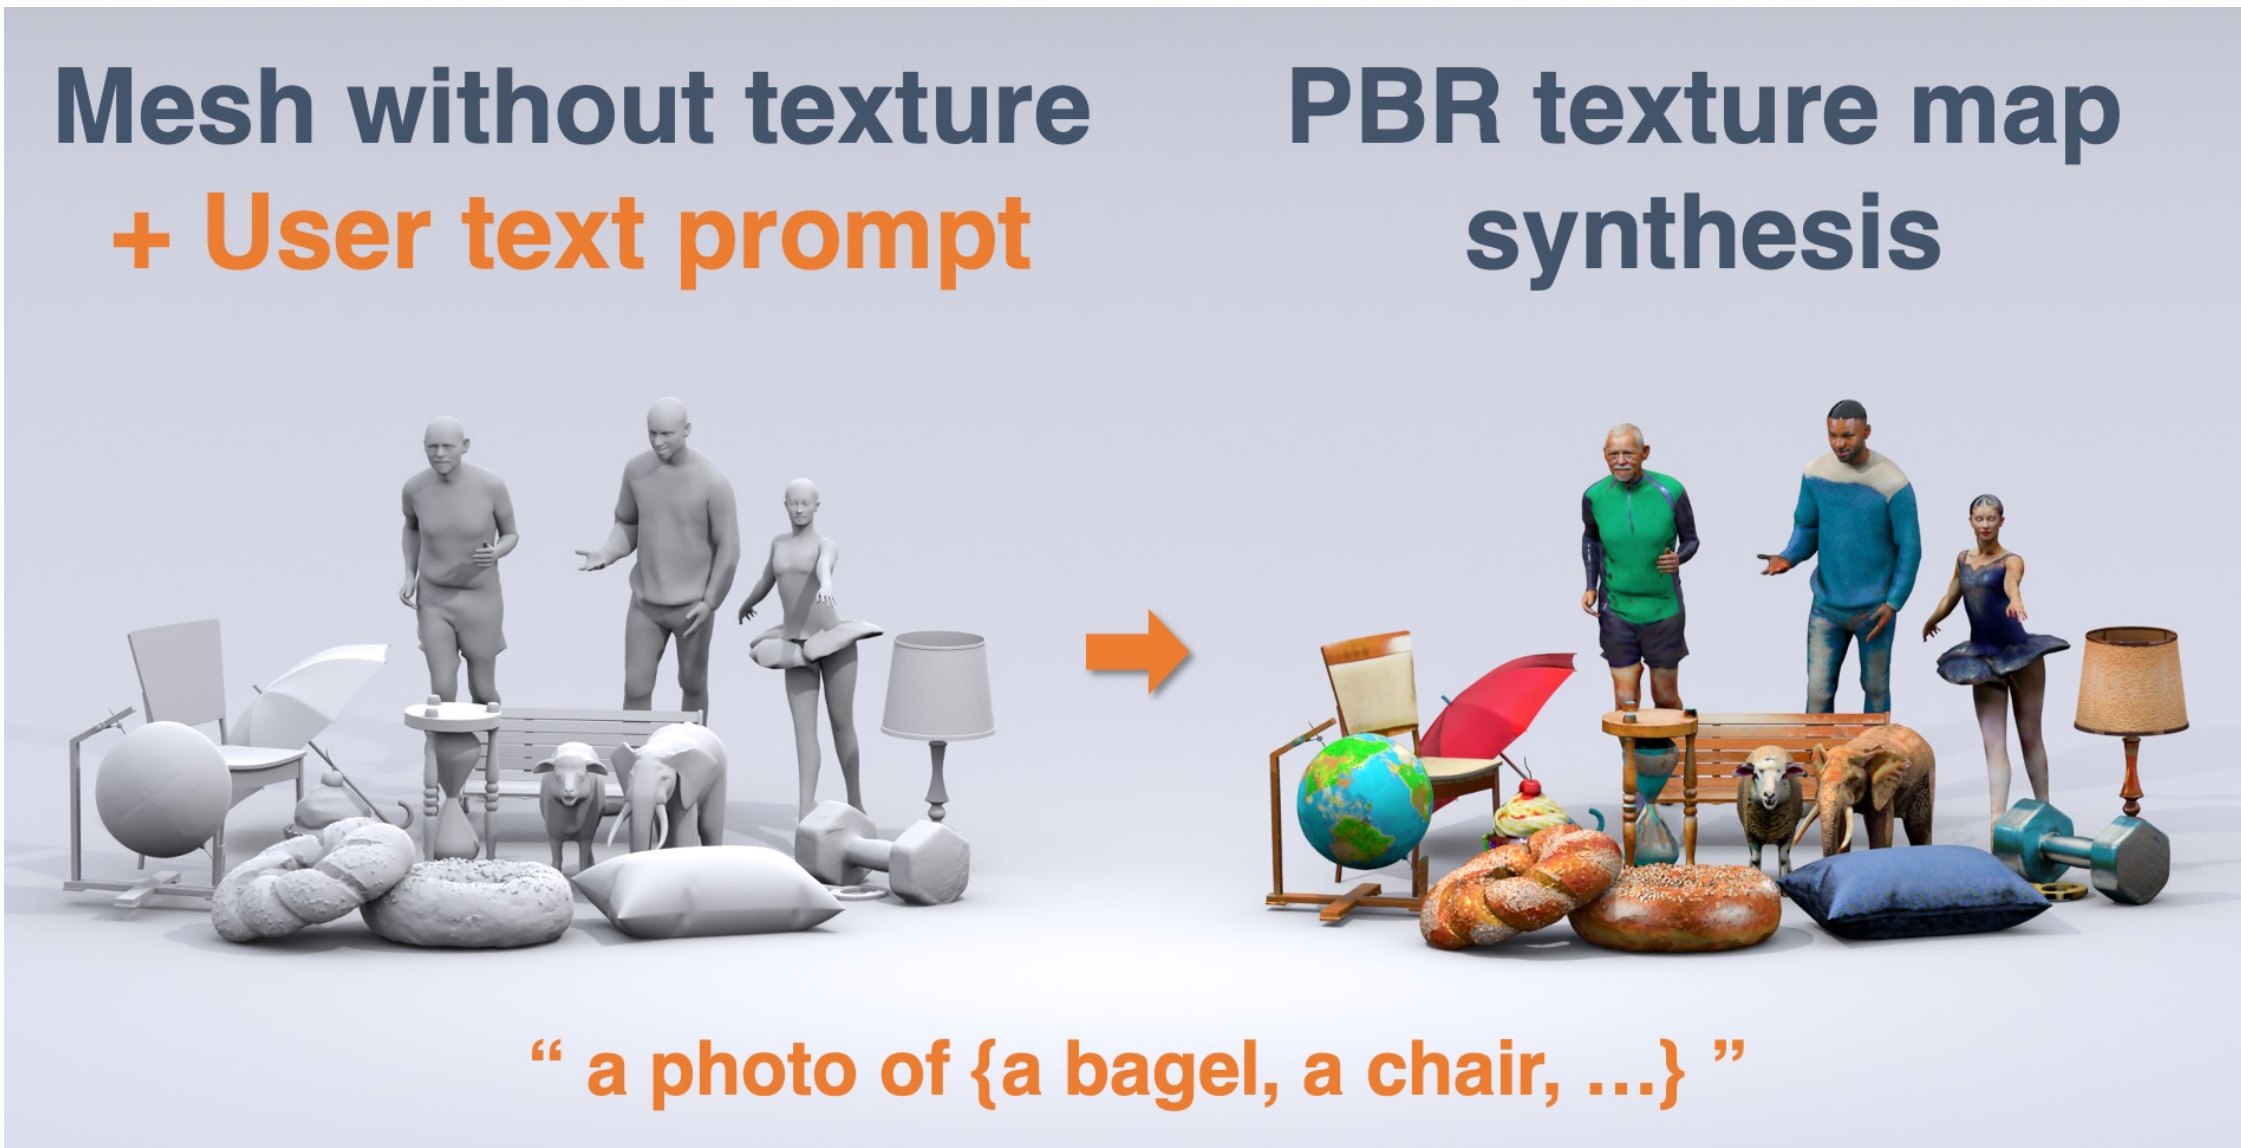
# Paint-it: Text-to-Texture Synthesis via Deep Convolutional Texture Map Optimization and Physically-Based Rendering

Kim Youwang<sup>1</sup> Tae-Hyun Oh<sup>1</sup> Gerard Pons-Moll<sup>2</sup>  
<sup>1</sup>POSTECH <sup>2</sup>University of Tübingen

<https://kim-youwang.github.io/paint-it>

## Overview

**TL;DR** Convolutional parameterization of PBR texture maps helps high-quality text-to-texture synthesis!

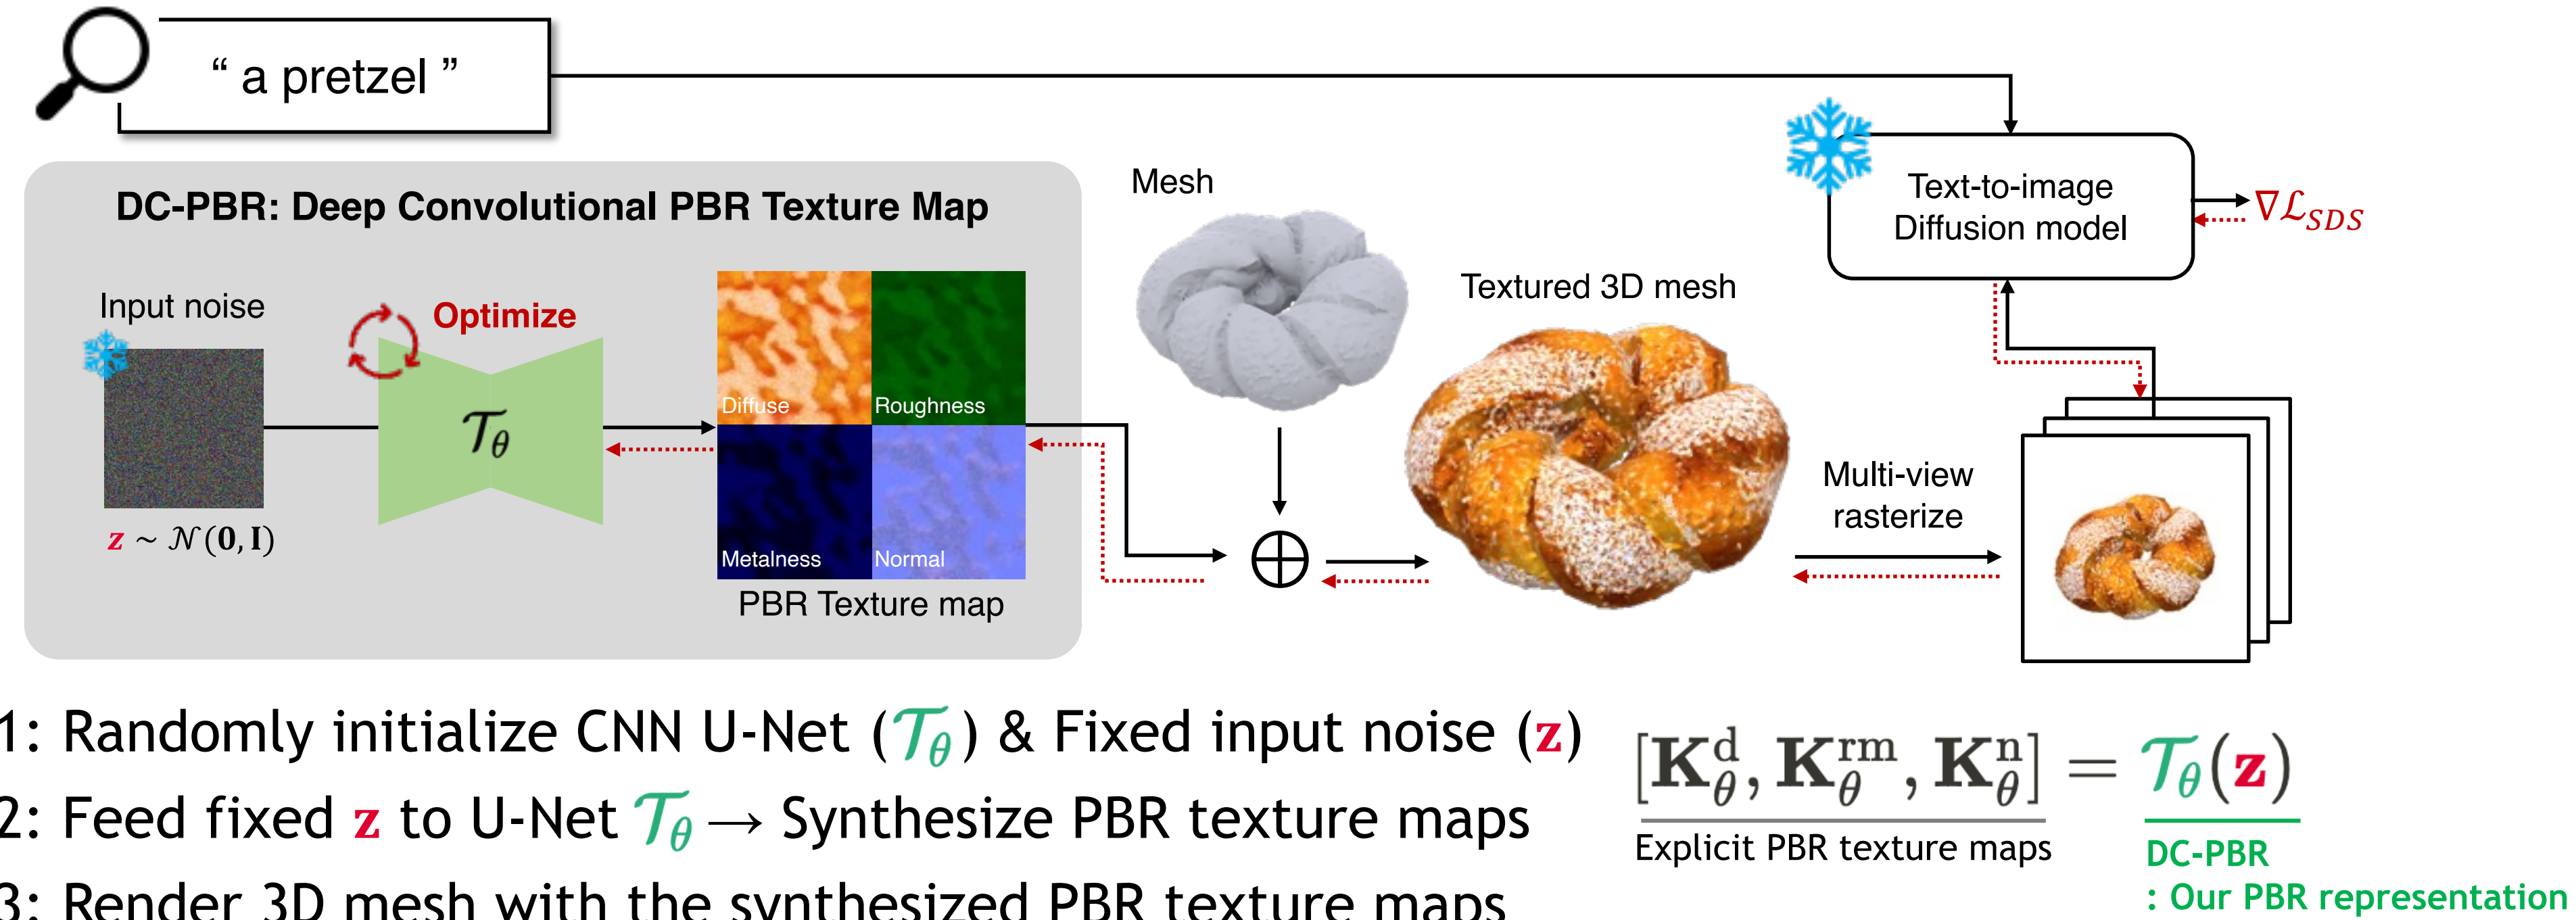


## Contributions

- Paint-it:** Text  $\rightarrow$  3D mesh texture synthesis method, generating PBR maps compatible with graphics engines
- DC-PBR:** CNN re-parameterization of PBR texture maps guides frequency-scheduled texture synthesis



## Paint-it: Text-Driven PBR Texture Map Synthesis



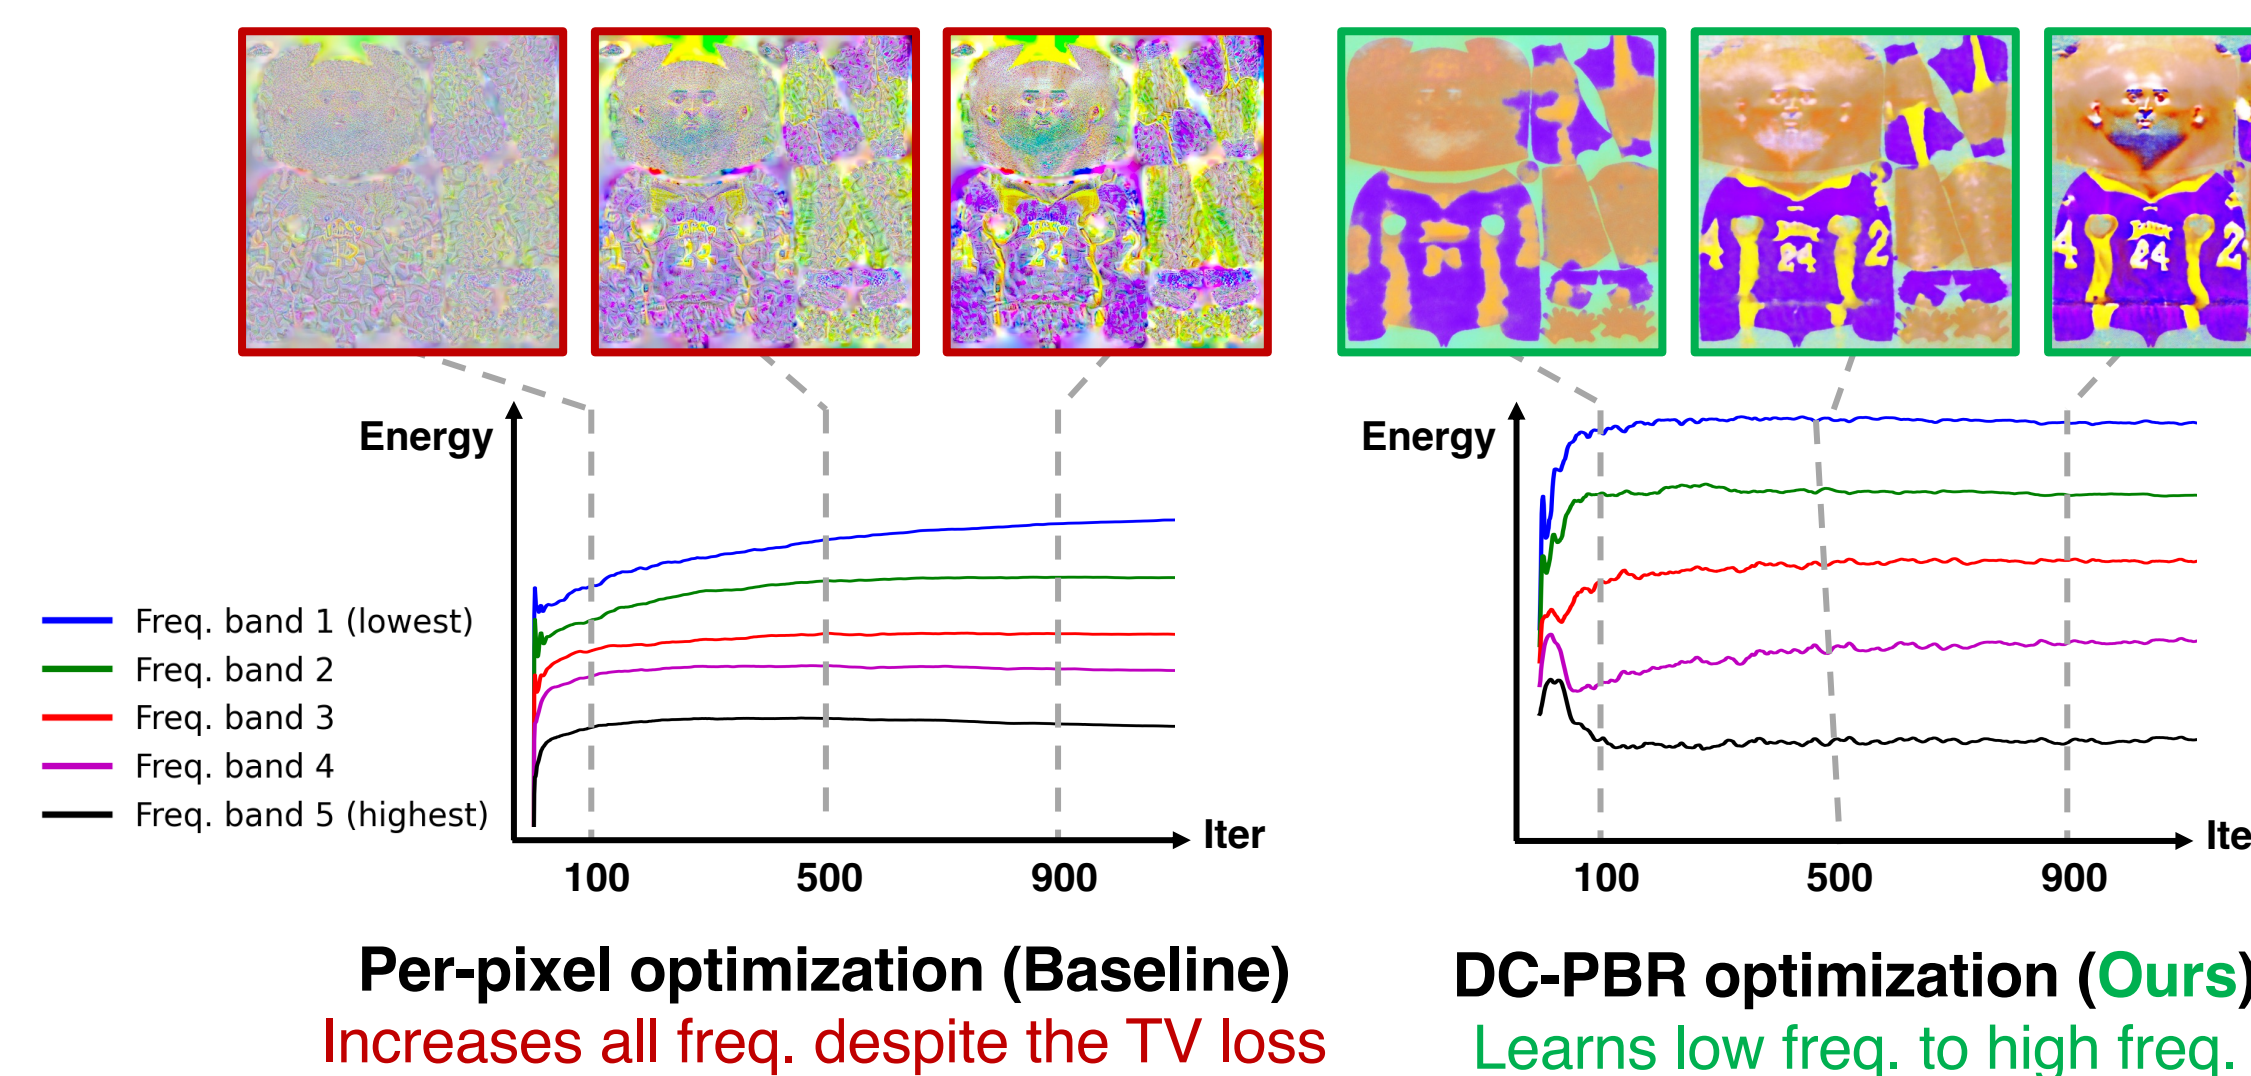
- Step 1: Randomly initialize CNN U-Net ( $\mathcal{T}_\theta$ ) & Fixed input noise ( $\mathbf{z}$ )
- Step 2: Feed fixed  $\mathbf{z}$  to U-Net  $\mathcal{T}_\theta \rightarrow$  Synthesize PBR texture maps
- Step 3: Render 3D mesh with the synthesized PBR texture maps
- Step 3: Update CNN kernels  $\theta$ , using Score-Distillation Sampling (SDS) gradient

$$[\mathbf{K}_\theta^d, \mathbf{K}_\theta^{rm}, \mathbf{K}_\theta^n] = \mathcal{T}_\theta(\mathbf{z})$$

Explicit PBR texture maps      DC-PBR : Our PBR representation

## Deep Convolutional PBR Texture Map Re-param. (DC-PBR)

**Our observation:** Re-parameterizing PBR texture maps with CNN kernels naturally induces “coarse-to-fine frequency scheduling” for texture synthesis & noise robustness



### (Baseline) Per-pixel PBR optimization

$$[\mathbf{K}^{d*}, \mathbf{K}^{rm*}, \mathbf{K}^{n*}] = \arg \min_{\mathbf{K}^d, \mathbf{K}^{rm}, \mathbf{K}^n} \mathbb{E}_{t, \epsilon} [\|\hat{\epsilon}_\phi(\mathcal{R}_t^M(\mathbf{K}^d, \mathbf{K}^{rm}, \mathbf{K}^n); \mathbf{y}, t) - \epsilon\|_2^2] + \mathcal{L}_{TV}$$

### (Ours) DC-PBR re-parameterized optimization

$$\theta^* = \arg \min_\theta \mathbb{E}_{t, \epsilon} [\|\hat{\epsilon}_\phi(\mathcal{R}_t^M(\mathbf{K}_\theta^d, \mathbf{K}_\theta^{rm}, \mathbf{K}_\theta^n); \mathbf{y}, t) - \epsilon\|_2^2],$$

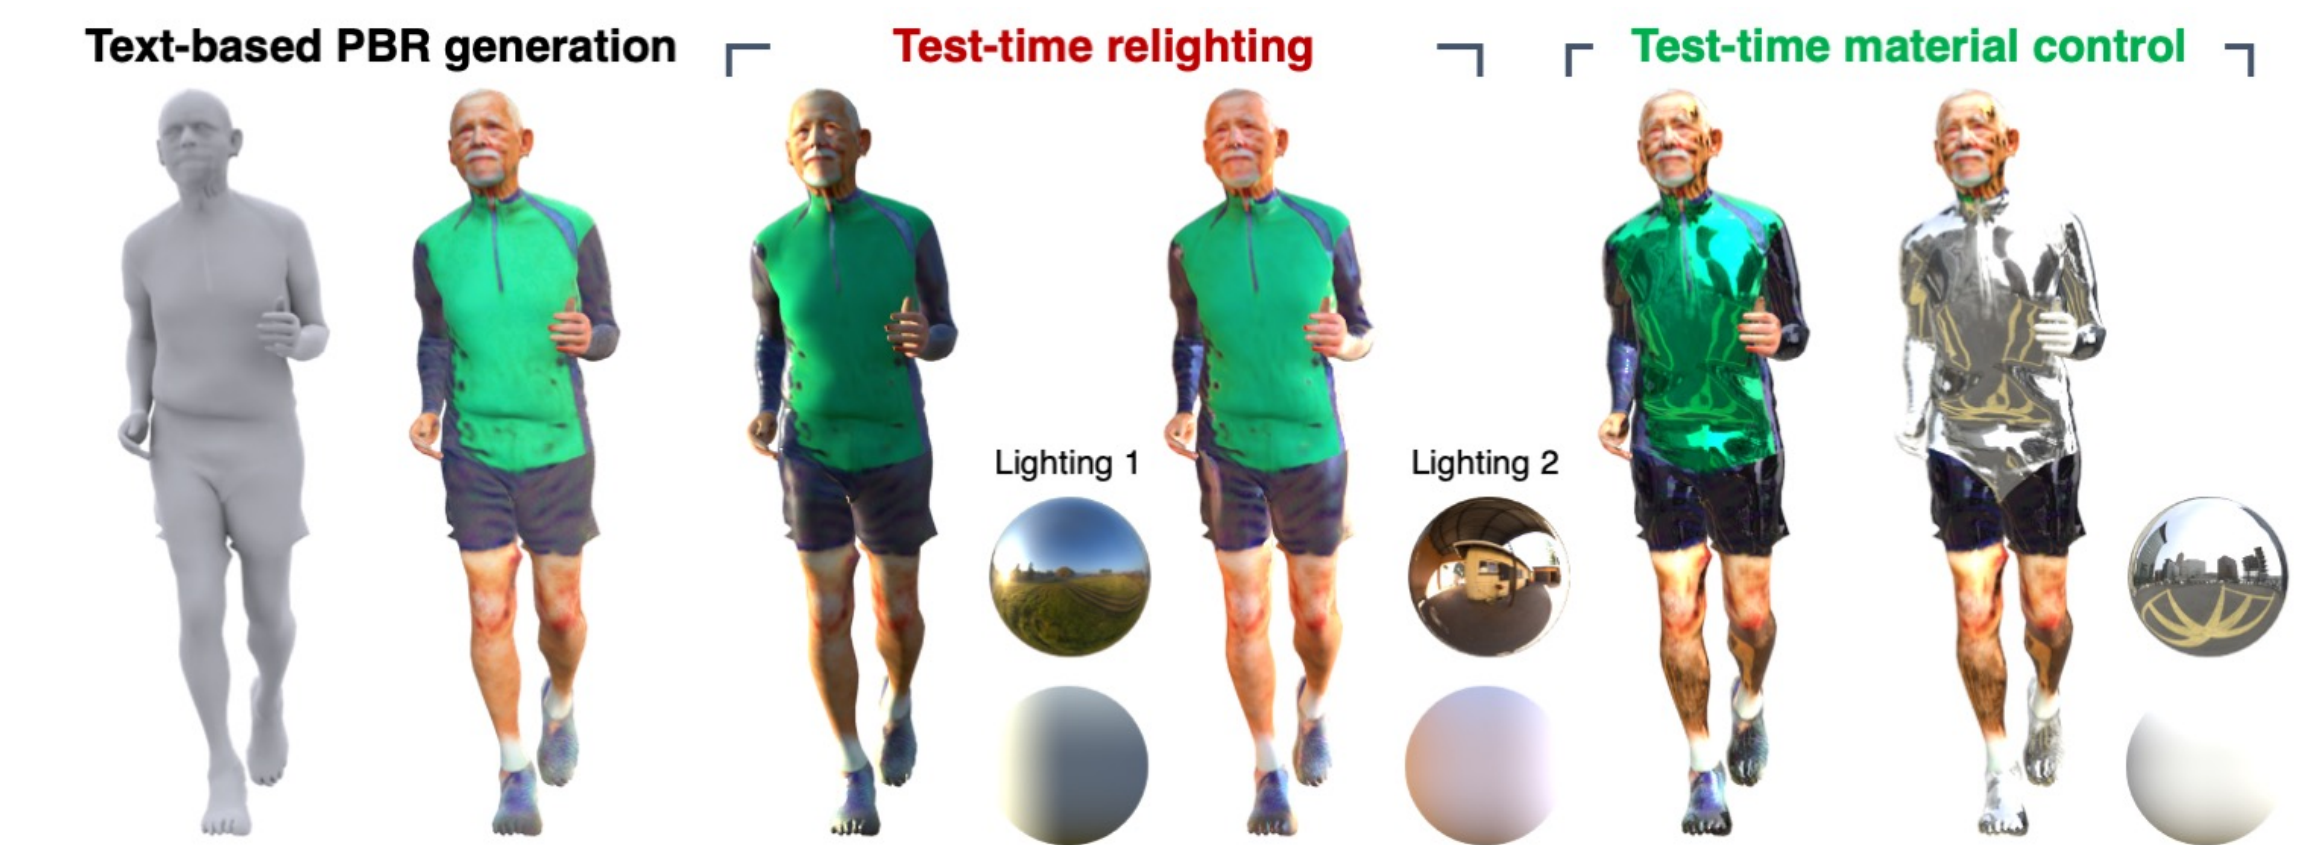
$$[\mathbf{K}_\theta^d, \mathbf{K}_\theta^{rm}, \mathbf{K}_\theta^n] = \mathcal{T}_\theta(\mathbf{z})$$

## PBR Texture Synthesis Results

### Qualitative results - Objaverse, RenderPeople



### Test-time applications with CG engines



## Take-home message

The SDS loss is known to be noisy. We show a better PBR texture representation (parameterization) sufficiently stabilizes the SDS loss and its empirical mechanism.