


# Under One Sun: Multi-Object Generative Perception of Materials and Illumination

Nobuo Yoshii<sup>1</sup>, Xinran Nicole Han<sup>2</sup>, Ryo Kawahara<sup>1</sup>,  
Todd Zickler<sup>2</sup>, and Ko Nishino<sup>1</sup>

<sup>1</sup> Kyoto University   <sup>2</sup> Harvard University  
<https://vision.ist.i.kyoto-u.ac.jp/research/onesun/>



**Fig. 1:** MultiGP is an ambiguity-aware inverse rendering method that samples the reflectance and texture of each object alongside the global scene illumination from a single image. By leveraging the shared illumination across multiple objects, MultiGP employs a novel end-to-end architecture featuring tailored guidance, axial attention, and ControlNet structures to resolve the otherwise ambiguous radiometric disentanglement.

**Abstract.** We introduce Multi-Object Generative Perception (MultiGP), a generative inverse rendering method for stochastic sampling of all radiometric constituents—reflectance, texture, and illumination—underlying object appearance from a single image. Our key idea to solve this inherently ambiguous radiometric disentanglement is to leverage the fact that while their texture and reflectance may differ, objects in the same scene are all lit by the same illumination. MultiGP exploits this consensus to produce samples of reflectance, texture, and illumination from a single image of known shapes based on four key technical contributions: a cascaded end-to-end architecture that combines image-space and angular-space disentanglement; Coordinated Scheduling for diffusion convergence to a single consistent illumination estimate; Axial Attention applied to facilitate “cross-talk” between objects of different reflectance; and a Texture Extraction ControlNet to preserve high-frequency texture details while ensuring decoupling from estimated lighting. Experimental results demonstrate that MultiGP effectively leverages the complementary spatial and frequency characteristics of multiple object appearances to recover individual texture and reflectance as well as the common illumination.

**Keywords:** Generative Perception · Inverse Rendering · Radiometric Disentanglement

## 1 Introduction

The appearance of an object reveals far more than its 3D shape and semantic identity. A glance at a cup, for example, also reveals its material properties (*e.g.*, ceramic, hard, heavy, perhaps with printed text or graphics) and provides informative reflections of the surrounding illumination. Deciphering this rich, entangled information can be essential for interaction. It can allow an embodied agent to plan a grasp based on perceived materials; predict how the object would look from different vantage points; and predict how other objects would look if placed in the same scene.

However, this radiometric disentanglement is overwhelmingly ill-posed. Surrounding light is reflected with an angular response that is unique to the local geometry and reflectance of each surface point, akin to an angular convolution with a shift-dependent kernel [33]. This deeply entangles object materials with ambient light, and it is inherently lossy because surface reflectance generally attenuates the light’s high angular-frequency details. All of this makes recovering materials and lighting from a single image fundamentally ambiguous [16].

“Generative perception” is a framework for addressing these ambiguities in vision, by using stochastic generative models to sample from the distribution of possible physical explanations. Previous work in this vein has focused on producing samples of material and shape from images [15, 45] or from videos [14, 24], but without producing samples of illumination. In contrast, DRM [10] produces samples of reflectance *and* illumination from an image of a known shape, but it cannot handle objects with texture.

In this paper, we show that stochastic sampling of a richer suite of radiometric constituents—reflectance, texture, and illumination—becomes possible when multiple objects are present in the image. Our key insight is that while objects in a scene differ in their texture and reflectance, they are all lit by the same illumination. We introduce Multi-object Generative Perception (MultiGP), an approach that leverages multi-object consensus to produce samples of reflectance, texture and illumination from a single image of known shapes.

MultiGP is built upon four key technical contributions. First, we propose a cascaded end-to-end architecture that simultaneously disentangles texture in the image domain and reflectance and illumination in the angular (reflectance map) domain. Second, we introduce Coordinated Scheduling, which manages the temporal scheduling of the diffusion process to ensure that illumination estimates across different objects converge to one consistent environment map. Third, we employ Axial Attention to facilitate “cross-talk” between objects in the reflectance map domain, allowing objects with different BRDFs to mutually complement each other’s missing angular frequencies. Fourth, we integrate a Texture Extraction ControlNet that preserves high-frequency surface details while ensuring they remain decoupled from the estimated lighting.

Our extensive experiments on synthetic and real-world data demonstrate that MultiGP achieves state-of-the-art accuracy in texture, reflectance, and illumination estimation. They also show that MultiGP successfully exploits the complementary information entangled in the object appearance of each object

in the spatial and frequency domains. Our results suggest that the radiometric information ingrained in object appearance can be reliably exploited by leveraging the physical constraints naturally occurring in multi-object scenes, providing a path toward robust radiometric scene understanding.

## 2 Related Works

There has been considerable recent progress in appearance decomposition using multiple images with changing viewpoint or illumination. Here we focus on single-image methods, where the ambiguity is substantially greater.

***Deterministic prediction: Inverse rendering and regression.*** Inverse rendering techniques perform gradient descent on a forward rendering operator to recover a set of radiometric constituents that best explain an image, typically using an objective that includes statistical priors or regularization on reflectance and illumination (and sometimes shape when it’s unknown) [6, 8, 27, 31, 44]. Instead of using iterative optimization, regression-based methods train deep networks to deterministically map an input image to one or more of texture, reflectance or illumination [5, 12, 13, 29, 30, 35, 41]. There are also hybrid methods that use regression to predict some radiometric constituents and inverse rendering to predict others [40]. Many of these methods only produce smooth, low-frequency lighting due to their regression to the mean and the intrinsic information loss in forward rendering. Also, all of them have the disadvantage of producing a single “best” estimate, even when there is fundamental ambiguity, where multiple, diverse explanations are equally valid. In contrast, our method stochastically samples sets of explanatory radiometric constituents, including diverse samples of realistic high-frequency lighting environments that are physically consistent with the observed appearance.

***Stochastic illumination.*** Recent generative approaches use diffusion models to hallucinate image-consistent, high-frequency illumination as a panoramic environment map [26, 37] or a virtual mirrored sphere within the scene [25, 32]. Although visually plausible, these methods do not explicitly estimate materials, and they fundamentally bypass physically-based modeling of reflection. In contrast, our goal is to radiometrically disentangle object appearance to recover the texture, reflectance, and illumination. DPI [28] uses differentiable rendering with a diffusion-based illumination prior. DRM [10] considers the special case of textureless objects and reformulates diffusion to generate samples of reflectance and illumination that are consistent with a reflectance map. Our work extends the DRM paradigm by leveraging the common illumination across multiple objects to enable the prediction of texture in addition to reflectance and illumination.

***Stochastic texture and reflectance.*** Texture and reflectance estimation has seen significant progress with diffusion-based intrinsic image decomposition [4, 7, 9, 17, 17, 38, 39, 47]. Kocsis et al. [21] finetune Stable Diffusion to generate

samples of texture and reflectance, and they use them to recover illumination by inverse rendering. RGB-X [45] and DiffusionRender [24] finetune Stable Diffusion to generate samples of texture, reflectance and shape, but they do not provide illumination. Recent work by Han et al. [14] trains a diffusion model from scratch to generate samples of texture, reflectance, and shape without illumination. In our work, we adapt some components of their U-ViT3D architecture.

### 3 Multi-Object Generative Perception

#### 3.1 Problem Formulation

An object’s appearance can be described by its surface radiance  $L_r$  which, ignoring global illumination effects, is determined by the surface normal  $\mathbf{n}$ , the bidirectional reflectance distribution function (BRDF)  $f_r$ , and distant incident illumination  $L_i$  [20]:

$$L_r(\mathbf{x}, \omega_o) = \int_{\omega} f_r(\mathbf{x}, \omega_i, \omega_o) L_i(\omega_i) (\omega_i \cdot \mathbf{n}(\mathbf{x})) d\omega_i, \quad (1)$$

where  $\omega_i$  and  $\omega_o$  are the incident and outgoing light directions in the local hemisphere centered at  $\mathbf{n}(\mathbf{x})$ . We adopt the isotropic Disney BRDF model [3]:

$$f_r(\mathbf{x}, \omega_i, \omega_o; \boldsymbol{\rho}_d(\mathbf{x}), \Psi := \{\rho_s, r, \gamma\}) = (1 - \gamma) \frac{\rho_d(\mathbf{x})}{\pi} (f_{\text{diff}}(\omega_i, \omega_o) + f_{\text{retro}}(\omega_i, \omega_o; r)) + f_{\text{spec}}(\omega_i, \omega_o; \boldsymbol{\rho}_d(\mathbf{x}), \rho_s, r, \gamma), \quad (2)$$

where  $\boldsymbol{\rho}_d(\mathbf{x})$  is the spatially-varying texture and  $\Psi$  is spatially-uniform reflectance consisting of roughness  $r$ , metallicness  $\gamma$ , and specular strength  $\rho_s$ . Most daily objects exhibit homogeneous specular reflection (*e.g.*, due to a uniform coating over a spatially-varying diffuse reflection). Handling fully spatially-varying specular reflection is left for future work.

Inverting Eq. (1) for a single object is fundamentally ambiguous; our key insight is to exploit the complementary radiometric information carried by *multiple* objects to constrain this inversion. Specifically, given an image  $I$  containing  $M$  objects with known shapes  $\mathbf{S} = \{\mathbf{n}(\mathbf{x})_m\}_{m=1}^M$ , we seek the joint posterior over per-object textures  $\mathbf{T} = \{\boldsymbol{\rho}_d(\mathbf{x})_m\}_{m=1}^M$ , per-object reflectances  $\mathbf{R} = \{\Psi^{(m)}\}_{m=1}^M$ , and shared illumination  $L$ :


$$p(\mathbf{T}, \mathbf{R}, L \mid I, \mathbf{S}) \propto p(I \mid \mathbf{S}, \mathbf{T}, \mathbf{R}, L) p(\mathbf{T}, \mathbf{R}, L \mid \mathbf{S}). \quad (3)$$

Computing this posterior exactly is intractable. We approximate it with a cascaded factorization that separates texture from reflectance and illumination:

$$q(\mathbf{T}, \mathbf{R}, L \mid I, \mathbf{S}) = q_\phi(\mathbf{T} \mid I, \mathbf{S}) q_\theta(\mathbf{R}, L \mid I, \mathbf{S}, \mathbf{T}). \quad (4)$$

This factorization directly dictates our architecture (Fig. 2): a texture extraction module  $q_\phi$  (Sec. 3.2), followed by a joint reflectance–illumination module  $q_\theta$  that





**Fig. 2:** Overview of MultiGP. Given a scene with multiple objects, textures are first estimated via a diffusion model  $q_\phi$  that accounts for global light transport. The resulting texture-free appearances are transformed into reflectance maps, from which a multi-object diffusion model  $q_\theta$  estimates a shared illumination and respective reflectances. Finally, a ControlNet refines the textures for physical consistency with the estimated lighting and reflectance through a renderer.

operates on texture-free, shape-invariant reflectance maps and can therefore focus entirely on material and lighting estimation (Sec. 3.3). The estimates are then refined end-to-end by maximizing the physical likelihood with a ControlNet which examines consistency of the disentanglement and the observation (Sec. 3.4).

### 3.2 Texture Extraction Prior ( $q_\phi$ )

The first stage of our factorization isolates the diffuse texture  $\mathbf{T}$  from the input image. We employ a latent diffusion model [36] conditioned on the concatenated latents of the observed appearance and object shapes, denoted as  $\mathbf{c}$ . The intermediate denoising step  $t$  follows standard diffusion dynamics:

$$\mathbf{z}_t = \sqrt{\bar{\alpha}_t} \mathbf{z}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon},$$

with  $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . The network  $\epsilon_\phi$  predicts the added noise to estimate both the texture  $\mathbf{T}$  and the intermediate texture-free appearance  $L_r(\mathbf{x}; \Psi)$ . The base loss is defined as

$$\mathcal{L}_\phi = \mathbb{E}_{\mathbf{z}_0, t, \boldsymbol{\epsilon}} [\|\boldsymbol{\epsilon} - \epsilon_\phi(\mathbf{z}_t, t; \mathbf{c})\|_2^2]. \quad (5)$$

### 3.3 Multi-Object Diffusion Reflectance Maps ( $q_\theta$ )

Conditioned on the extracted texture  $\mathbf{T}$ , we operate on the texture-free appearance of each object to solve  $q_\theta(\mathbf{R}, L \mid \mathbf{I}, \mathbf{S}, \mathbf{T})$ . Each texture-free appearance  $L_r^{(m)}(\mathbf{x})$  is mapped to a shape-invariant raw *reflectance map*  $\tilde{L}_r^{(m)}(\mathbf{n})$ , a Gaussian sphere of surface radiance indexed by surface normal  $\mathbf{n}$ . They are masked due to the visible normal area of the original objects. Then, these raw reflectance maps  $\tilde{L}_r^{(m)}(\mathbf{n})$  are converted to complete reflectance maps  $L_r^{(m)}(\mathbf{n})$  via a learned network (a

multi-object extension of ObsNet [10]). This network leverages axial attention analogous to our multi-object axial attention in the diffusion process to mitigate global light transport effects, such as cast shadows and interreflections.

Building on the Diffusion Reflectance Map (DRM) framework [10], which recovers illumination from a single object via learned inverse diffusion on the rendering equation, we extend this approach to a multi-object setting. Since each surface reflects different spatial extents and frequency bands of the same illumination, determined by its unique BRDF and geometry, we coordinate these complementary signals to jointly recover the shared illumination  $L$  (represented as a mirror reflectance map  $L_r^{(0)}(\mathbf{n})$ ) and individual reflectances  $\mathbf{R}$ .

We index each reflectance map by its diffusion step  $k$ :  $L_r^{(m,K)}(\mathbf{n})$  is the initial observation and  $L_r^{(m,0)}(\mathbf{n}) = \hat{L}_i(\mathbf{n})$  is the target mirror reflectance map. Following DRM, the forward process is not a Markov noise chain, but a deterministic radiometric formation governed by the rendering equation with additive Gaussian noise to account for the inherent ambiguity. Assuming no cross-talk between object reflections (only the Multi-Object Diffusion  $q_\theta$ , the Texture Extraction Prior  $q_\phi$  and the texture-refinement renderer in Sec. 3.4 learn from multi-object path-traced data), the joint forward process factorizes over the  $M$  objects (see the supplementary material for the full derivation):

$$q(L_r^{(1:M,1:K)} | L_r^{(0)}, \Psi^{(1:M,K)}) = \prod_{m=1}^M \prod_{k=1}^K \mathcal{N}(L_r^{(m,k)} | L_r(\Psi^{(m,k)}, L_r^{(0)}), \sigma^2 \mathbf{I}), \quad (6)$$

where  $L_r(\Psi^{(m,k)}, L_r^{(0)})$  denotes the forward rendering operator.


The evidence lower bound (ELBO) decomposes into a sum of per-object terms, yielding the illumination loss:

$$\mathcal{L}_i = \mathbb{E}_{L_i, \Psi, m, k} \|\mu'_{\theta_i, \theta_r} - \{L_r(\Psi^{(m,k-1)}, L_r^{(0)}) - L_r^{(m,k)}\}\|_2^2, \quad (7)$$

where  $\theta_i$  and  $\theta_r$  parameterize the illumination and reflectance estimators, respectively, and  $\mu'_{\theta_i, \theta_r}$  is the expectation of the residual between  $\hat{L}_r^{(m,k-1)}$  and  $L_r^{(m,k)}$ . MultiGP minimizes this objective using two key components: **Coordinated Scheduling**, which guides all object appearances toward a common illumination via reverse diffusion, and **Axial Attention**, which fuses complementary spatio-spectral information across objects.

**Multi-Object Coordinated Scheduling.** Our goal is to convert the  $M$  observed reflectance maps  $L_r^{(m,K)}$ , each representing a different material, back to a single mirror reflectance map  $L_r^{(0)}(\mathbf{n})$  that reveals the environment illumination  $\hat{L}_i(\mathbf{n})$  (Fig. 3a). To achieve this, we schedule the reflectance of each object to evolve linearly from its estimated material  $\hat{\Psi}^{(m,K)}$  to the known mirror reflectance  $\Psi_0$  over  $K$  shared steps:

$$\hat{\Psi}^{(m,k)} = \frac{k}{K} \hat{\Psi}^{(m,K)} + (1 - \frac{k}{K}) \Psi_0. \quad (8)$$



**Fig. 3:** (a) Multi-Object Coordinated Scheduling guides the appearance of diverse objects toward a single shared illumination estimate. At each denoising step, the process is governed by the estimated reflectances of multiple objects alongside a known mirror reflectance. Since the mirror reflectance uniquely corresponds to the illumination, these diverse inputs stochastically converge to a consistent, shared environment map. (b) Multi-Object Axial Attention shares reflectance and spatial information across different reflectance maps. This mechanism enables MultiGP to integrate complementary frequency information which is otherwise attenuated by individual reflectances and aggregate visible lighting directions at every denoising step.

This ensures that all  $M$  objects reach mirror reflectance simultaneously at  $k=0$  while preserving their relative scales, in contrast to DRM [10] which requires infinite steps to converge.

The total number of steps  $K$  is set proportionally to the average distance of the input materials from mirror reflectance:

$$K = \frac{K_{\max}}{M} \sum_{m=1}^M \left\| \frac{\hat{\Psi}^{(m,K)} - \Psi_0}{\sqrt{3}} \right\|^2, \quad (9)$$

where  $K_{\max}$  is the maximum allowed steps and  $\sqrt{3}$  normalizes by the largest possible distance in the 3-D reflectance space  $\hat{\Psi} = \{\text{metallic, roughness, specular}\}$ . Equation (9) guarantees that the most distant reflectance  $\Psi = \{0, 1, 0\}$  can reach mirror state  $\Psi_0 = \{1, 0, 1\}$  in  $K_{\max}$  steps.

The model is trained to predict the reflectance at each step with the loss:

$$\mathcal{L}_r = \mathbb{E}_{L_i, f_r, m, k} [\|\hat{\Psi}_{\theta_r}^{(m,k)} - \Psi^{(m,k)}\|_2^2 + \|\hat{\Psi}_{\theta_r}^{(m,K)} - \Psi^{(m,K)}\|_2^2]. \quad (10)$$

The total loss is defined as  $\mathcal{L} = \lambda_i \mathcal{L}_i + \lambda_r \mathcal{L}_r$ , where  $\lambda_i, \lambda_r$  are hyperparameters.

**Multi-Object Axial Attention.** The texture-free reflectance maps  $\tilde{L}_r^{(m)}(\mathbf{n})$  transformed from each object are spatially sparse. Unless the object is a sphere, its geometry does not densely cover the full hemisphere of camera-facing surface normals. More interestingly, different materials act as distinct frequency filters on the environment lights [33, 34]: a Lambertian-like surface reflects low-pass filtered illumination in its appearance, while more specular objects retain relatively higher-frequency details through band-pass filtering of the illumination spectrum. No

single material behaves as a Dirac delta that preserves all frequencies. The multiple object appearances capture complementary spatio-spectral characteristics of the shared illumination, and we wish to fully leverage this complementarity.

To this end, we propose a Multi-Object Axial Attention mechanism. Inspired by temporal attention in video diffusion [14], this cross-object attention operates across the  $M$  reflectance maps at each corresponding spatial location (*i.e.*, the same surface normal direction). As illustrated in Figure (Fig. 3b), the axial attention applied over multiple partially-observed reflectance maps allows an object with missing spectral-frequency details or unobserved normal directions to borrow information from the other objects in the scene, effectively unifying the multi-object observations into a coherent, complete estimate of the global illumination. We also apply Local Spatial Attention [14] to avoid information integration with Axial Attention solely along the perfectly aligned normal direction so that MultiGP is made robust to geometry errors.

### 3.4 Texture Refinement via ControlNet

While the cascaded factorization produces diverse samples, some may still violate the physics of image formation. We therefore introduce a final refinement step that injects renderer consistency into the sampling process. Given sampled factors  $(\mathbf{T}, \mathbf{R}, L)$  and known geometry  $\mathbf{S}$ , we render the objects  $\hat{L}_r$  with Mitsuba 3 [19] and compute its residual against the observed image  $I$ . We condition  $q_\phi$  on this residual via a ControlNet [46]-style mechanism, effectively steering diffusion sampling toward solutions that better explain the observation. The ControlNet accepts the residual  $\mathbf{c}_f = \mathcal{E}(I - \hat{L}_r)$  as an additional condition, enforcing tighter physical consistency between texture estimates and global light transport:

$$\mathcal{L}_{\phi_c} = \mathbb{E}_{\mathbf{z}_0, t, \epsilon} [\|\epsilon - \epsilon_{\phi_c}(\mathbf{z}_t, t; \mathbf{c}, \mathbf{c}_f)\|_2^2] . \quad (11)$$

This refinement acts as likelihood-guided sampling rather than explicit optimization; it modifies the sampling dynamics using a renderer-based signal to improve physical consistency while preserving generative diversity.

## 4 Experimental Results

We evaluate MultiGP against several state-of-the-art methods [10, 14, 24, 28, 32, 44, 45] on synthetic and real datasets, demonstrating its effectiveness in both accuracy and diversity. Crucially, we derive a canonical framework to evaluate the fundamentally ambiguous disentanglement of inverse rendering components, particularly illumination. We believe this represents a significant contribution to inverse rendering and appearance modeling research at large. Through these extensive evaluations, we show that MultiGP successfully integrates complementary spatio-spectral information from diverse object appearances to achieve superior scene recovery.

#### 4.1 Dataset

**Training.** We train MultiGP on synthetic 3D multi-object scenes with path-traced light transport (not 2D composites) and per-object reflectance-map triplets rendered by Mitsuba3 [19]. We create 96 textures for 1100 Adobe 3D Assets [1] and 24 shape-dependent textures for each of the 4000 shapes of Xu [42] with MV-Adapter [18]. In addition we use unique colored textures. Reflectance properties (*i.e.*, metallic, roughness, and specular) are randomly sampled, with metallicness restricted to 0 or 1 to avoid unrealistic reflectance under Disney BRDF [3], and roughness set above 0.4 for 50% of objects to prevent trivially easy illumination estimation. Illumination is randomly sampled from the Laval Indoor Dataset [11] and the Poly Haven HDRIs [2] totaling 2000 unique environment maps.

**Synthetic data.** For accuracy evaluation, we use the test split of Adobe 3D Assets with MV-Adapter textures and held-out illumination, forming 36 object triplets. For diversity evaluation, we construct pseudo raw reflectance maps from nLMVS-Synth [43] shapes.

**Real-world data.** As no existing dataset contains multiple objects in the same image, we construct multi-object evaluation sets from Stanford-ORB [22] (14 objects across 7 scenes) and nLMVS-Real [43]. We select triplets of objects captured at approximately the same position, viewpoint, and illumination. This yields 23 and 27 object triplets from Stanford-ORB and nLMVS-Real, respectively.

**MultiGP dataset.** We additionally capture 9 real scenes (3 outdoor, 6 indoor) featuring 3 objects under shared illumination. Each object’s shape was scanned with Artec3D to obtain surface normals and aligned to camera coordinates using MegaPose [23].

#### 4.2 Ambiguity-Aware Inverse Rendering Metric

**Standard Metrics.** Following [10, 32], we evaluate recovered illumination at  $128 \times 256$  resolution using scale-invariant logRMSE, PSNR, SSIM, and LPIPS. LogRMSE is computed on HDR images, while other metrics use tone-mapped LDR images. Reflectance is evaluated via scale-invariant logRMSE in the non-parametric MERL BRDF representation [8]. We compare against DPI [28] and DRM [10], as methods generating spatially varying BRDFs [14, 24, 45] are incompatible with our uniform reflectance setting. Following [22], texture is evaluated at  $512 \times 512$  using scale-invariant RMSE, PSNR, SSIM, and LPIPS.

**Ambiguity-Aware Metric.** Standard metrics often falter in inverse rendering because multiple illuminations can produce identical appearances, especially under low-frequency reflectance. To assess how well the estimated distribution captures the ground truth, we utilize Spherical Harmonics (SH). We compute 1089 SH coefficients (up to degree 32) for each illumination sample. By applying PCA

**Table 1:** Quantitative evaluation and ablation study of illumination and reflectance estimates on synthetic test data. For illumination estimation, the sharpest reflectance among multiple objects is shown for both DRM and MultiGP (single). We report the mean of the top 3 out of 10 stochastic predictions. Our full model achieves the highest accuracy and the two novel guidance and attention play crucial roles for this.

|                           | Illumination |              |             |             | Reflectance |
|---------------------------|--------------|--------------|-------------|-------------|-------------|
|                           | logRMSE↓     | PSNR↑        | SSIM↑       | LPIPS↓      | logRMSE↓    |
| DPI [28]                  | 1.64         | 11.96        | 0.37        | <u>0.56</u> | 2.22        |
| DRM [10]                  | 1.48         | 12.62        | 0.34        | 0.61        | 1.85        |
| DiffusionLight [32]       | 1.60         | 11.37        | 0.38        | 0.63        | -           |
| MultiGP (single)          | 1.44         | 12.67        | 0.38        | 0.58        | 2.05        |
| <b>MultiGP</b>            | <b>1.28</b>  | <b>13.54</b> | <b>0.42</b> | <u>0.56</u> | <b>1.81</b> |
| w/o Coordinate Scheduling | <u>1.29</u>  | <u>13.50</u> | <b>0.42</b> | <b>0.54</b> | 1.87        |
| w/o Axial Attention       | 1.37         | 12.97        | <u>0.40</u> | <u>0.56</u> | <u>1.82</u> |
| w/o Texture Refinement    | <u>1.29</u>  | 13.49        | <b>0.42</b> | <u>0.56</u> | 1.87        |
| Global Illumination off   | 1.33         | 13.08        | <b>0.42</b> | <u>0.56</u> | <b>1.81</b> |

(retaining 99% variance) to the coefficients, we model the sample distribution and report the multivariate normal log-likelihood and Mahalanobis distance of the ground truth within this distribution.

**Evaluation Protocol.** For all stochastic methods (except ALP [44]), we sample 10 predictions and report the mean of the top 3 [14]. To evaluate single-object baselines on multi-object scenes, we report results for the object with the sharpest estimated reflectance. DiffusionLight [32] utilizes our reflectance estimate for this selection.

### 4.3 Results on Synthetic Dataset

**Accuracy.** Tab. 1 shows quantitative results of the illumination and the reflectance accuracy. For DRM [10], which processes only a single object, we report scores for the object with the sharpest reflectance (closest to a mirror state). MultiGP (single) follows the same protocol. For DiffusionLight [32], we provide the background consistent with their original assumptions. MultiGP achieves state-of-the-art accuracy across this extensive dataset. Tab. 2 evaluates texture estimates, where MultiGP outperforms DPI [28] among methods assuming known geometry.

**Ablation studies.** Tab. 1 and Tab. 2 demonstrate the impact of Coordinate Scheduling, Axial Attention, and Texture Refinement. Our results confirm that Coordinate Scheduling and Axial Attention are critical for estimating high-fidelity illumination and reflectance. Furthermore, Texture Refinement utilizing the estimated individual reflectances and shared illumination significantly improves

**Table 2:** Quantitative evaluation and ablation study of texture estimation on synthetic test data. MultiGP achieves superior accuracy compared to DPI, which similarly assumes known object geometry.

|                         | RMSE↓        | PSNR↑        | SSIM↑       | LPIPS↓       |
|-------------------------|--------------|--------------|-------------|--------------|
| RGB↔X [45] (LDR)        | 0.11         | 27.37        | <u>0.96</u> | 0.053        |
| RGB↔X [45] (HDR)        | 0.11         | 27.59        | <u>0.96</u> | 0.054        |
| GP Motion [14]          | 0.13         | 25.91        | 0.95        | 0.071        |
| DiffusionRenderer [24]  | 0.11         | 27.37        | <u>0.96</u> | 0.051        |
| DPI [28]                | 0.15         | 24.86        | 0.92        | 0.084        |
| <b>MultiGP</b>          | <b>0.081</b> | <b>29.99</b> | <b>0.97</b> | <b>0.031</b> |
| w/o Texture Refinement  | <u>0.088</u> | <u>29.32</u> | <b>0.97</b> | 0.034        |
| Global Illumination off | 0.089        | 29.35        | <b>0.97</b> | <u>0.033</u> |


the accuracy of texture extraction. Multi-Object Coordinate Scheduling and Axial Attention further attenuate non-shared Global Illumination signals, hence our no-cross-talk assumption in Multi-Object Diffusion. Tab. 1 and Tab. 2 directly ablates GI in the inputs: GI-on yields 1.28 logRMSE vs. 1.33 for GI-off—MultiGP models real GI.

**Runtime comparison.** On one A100 per  $1024 \times 1024$  multi-object image, MultiGP takes 25.3 seconds, DPI [28] 1167.7s, DiffusionLight [32] 1280.1s, and DRM [10] 12.1s (single object). MultiGP is  $\sim 50\times$  faster than DPI and Diffusion-Light.

**Ambiguity-aware accuracy analysis.** Fig. 4 illustrates the distribution of MultiGP estimates. We evaluate MultiGP’s frequency and spatial integration capabilities using two scenarios: (a) objects with varying reflectance (heterogeneous reflectances) and (b) objects with identical reflectance but different geometries (heterogeneous masks).

For objects with distinct frequency characteristics (a), we analyze the Spherical Harmonic (SH) coefficients of the input reflectance maps. As shown in Fig. 4a-1, input appearances exhibit complementary spectral traits; while the shiniest objects often cover a broad spectrum, they do not provide exhaustive information. Thus, multiple objects collectively offer superior visual cues for illumination recovery compared to any single object.

Fig. 4a-2 visualizes the illumination estimates from 100 MultiGP samples. We compare single-object estimation against our multi-object framework. To facilitate visualization, we apply PCA to the SH coefficients and plot the first two principal components. The multi-object distribution covers the ground-truth illumination with significantly higher likelihood than isolated estimates. This confirms the advantage of leveraging complementary material frequencies. Furthermore, it validates our proposed metric: because inverse rendering is inherently ambiguous,



**Fig. 4:** Distribution of illumination samples from MultiGP and MultiGP (single). (a) Heterogeneous reflectances: MultiGP effectively integrates objects with different reflectances. (a-1) shows the complementary Spherical Harmonic (SH) frequency spectra of three input reflectance maps. (a-2) provides a 2D PCA visualization of 100 samples; orange dots represent the joint MultiGP distribution, while other colors represent individual MultiGP (single) estimates. The joint distribution (orange) captures the “ground truth” with the highest density. (b) Heterogeneous masks: Using the same reflectance with different masks demonstrates that MultiGP effectively integrates varied object shapes. Here too, the MultiGP distribution (orange) densely encompasses the “ground truth.”

a probabilistic distribution reflects true accuracy more faithfully than a simple distance to ground truth.

Fig. 4b provides a parallel analysis for scenario (b). The results demonstrate that diverse geometries provide complementary spatial coverage of the environment, which MultiGP successfully exploits to produce more accurate and diverse illumination estimates.

We quantitatively evaluate the distribution of estimates for both cases using synthetic test data (Tab. 3). We randomly paired three object reflectances with 50 test illuminations to render input reflectance maps. For case (b), maps were spatially masked according to visible surface normals from randomly sampled object shapes [43]. For both scenarios, we sampled 1000 illumination outputs per input and performed PCA on the 1089 order-32 SH coefficients (maintaining a 99% cumulative eigen-ratio). Inputs were labeled A, B, and C, corresponding to (a) highest to lowest SH frequency or (b) largest to smallest unmasked area. We then computed the multivariate normal log-likelihood and Mahalanobis distance from the ground-truth illumination within the PCA space.

The results in Tab. 3 demonstrate that the highest accuracy is achieved when integrating all objects. In scenario (a), the multi-object accuracy surpasses that of the highest-frequency object (A) alone. Similarly, for spatial integration (b), using all objects outperforms the single largest-coverage shape (A). These results confirm that MultiGP successfully leverages complementary frequency



**Table 3:** Likelihood and Mahalanobis distance from ground-truth illumination in PCA space (99% SH variance), averaged over 50 test scenes with 1000 samples each. (a) Heterogeneous reflectances: Demonstrates MultiGP’s ability to effectively integrate objects with diverse material properties. (b) Heterogeneous masks: Demonstrates MultiGP’s ability to integrate different object shapes and visible areas. Inputs are labeled A, B, and C, corresponding to (a) highest-to-lowest SH frequency or (b) largest-to-smallest unmasked area. In both cases, the MultiGP “All” configuration outperforms the best single-object input (“A”) across all metrics.

| (a) Heterogeneous reflectances |                                                  |                                      | (b) Heterogeneous masks |                                                  |                                      |
|--------------------------------|--------------------------------------------------|--------------------------------------|-------------------------|--------------------------------------------------|--------------------------------------|
| Input                          | Multivariate Normal<br>Log-Likelihood $\uparrow$ | Mahalanobis<br>Distance $\downarrow$ | Input                   | Multivariate Normal<br>Log-Likelihood $\uparrow$ | Mahalanobis<br>Distance $\downarrow$ |
| All                            | <b>171.9</b>                                     | <b>16.98</b>                         | All                     | <b>61.5</b>                                      | <b>24.18</b>                         |
| A                              | <u>169.2</u>                                     | <u>18.42</u>                         | A                       | <u>-56.8</u>                                     | <u>27.60</u>                         |
| B                              | -137.6                                           | 26.39                                | B                       | -696.9                                           | 38.59                                |
| C                              | -243.6                                           | 28.91                                | C                       | -5211.2                                          | 57.75                                |


characteristics and spatial spans, justifying a multi-object approach for ambiguity-aware inverse rendering.

#### 4.4 Results on Real-World Datasets


**Stanford-ORB dataset.** Fig. 5 shows illumination estimates on the Stanford-ORB dataset. Quantitative evaluations are provided in the supplementary material. Ground truth for reflectance is unavailable, as it is nearly impossible to measure for real objects without an integrating (Ulbricht) sphere. While the dataset provides pseudo-ground-truth textures, we found them too noisy for reliable use. Note that ALP requires multi-view images to pre-estimate texture and reflectance, making it non-comparable to our single-image recovery framework. The results for MultiGP (multiple) demonstrate its ability to effectively integrate object appearances for high accuracy. Fig. 6 provides qualitative texture evaluations, showing that MultiGP removes lighting reflections more accurately than other methods, including DPI, which also utilizes surface normals as input.

**nLMVS-real dataset.** Fig. 7 presents qualitative illumination accuracy on the nLMVS-real dataset. Quantitative results are provided in the supplementary material. All objects in this dataset are textureless, and reflectance ground truth is not available. As before, ALP is shown for reference but is not a competing single-image method, as it requires multi-view images for texture and reflectance pre-estimation. The results demonstrate that MultiGP achieves state-of-the-art accuracy, even for textureless objects, by successfully integrating complementary visual cues of the shared global illumination from surface radiance.

**MultiGP dataset.** Fig. 8 shows a qualitative evaluation of illumination accuracy on our newly captured multi-object dataset. Quantitative evaluations are provided



**Fig. 5:** Illumination estimates on the Stanford-ORB dataset. For a fair comparison, we show scaled illumination results closest to the ground truth. For existing methods, we select the result from the object yielding the best logRMSE score. MultiGP faithfully captures the ground truth illumination structure with high-fidelity.




**Fig. 6:** Texture estimates on the Stanford-ORB dataset. For fair comparison, scaled albedo that best match the reference within the masked object region are shown. MultiGP estimates highly accurate texture (notice the lack of shading).

in the supplementary material. We captured 9 scenes across both indoor and outdoor settings to evaluate multiple objects in a shared physical environment. The results demonstrate that MultiGP generalizes effectively to real-world scenes, successfully handling complex global light transport effects.


## 5 Conclusion

In this paper, we introduced Multi-Object Generative Perception (MultiGP), a framework that resolves the fundamental ambiguities of inverse rendering by leveraging multi-object scene constraints. By treating multiple objects as concurrent probes of the same environment, our method aggregates complementary spatio-spectral information to estimate texture, reflectance, and illumination from a single image. To rigorously evaluate this stochastic approach, we introduced a novel metric for ambiguity-aware inverse rendering, shifting the focus from traditional point estimates to how effectively sampled distributions encompass the physical ground truth. Our results on synthetic and real-world datasets demonstrate that MultiGP significantly outperforms existing methods, providing a robust path toward physics-aware scene understanding.

A primary limitation of our current approach is the requirement for known object geometry, as the radiometric cues are anchored to the 3D surface. Furthermore, our model assumes distant environmental illumination (directional lighting).



**Fig. 7:** Illumination estimates on the nLMVS-real dataset. For all baselines, the result with best logRMSE is shown. MultiGP results capture the “ground truth” light structure most accurately.



**Fig. 8:** Illumination estimates on newly captured MultiGP dataset. MultiGP accurately recovers the illumination structure of the actual environment despite the complex global light transport.

In real-world indoor scenarios, near-field lighting effects, where light source proximity causes spatially varying illumination, can introduce significant deviations from our assumptions required to model the illumination on the reflectance map. Another limitation of the no-cross-talk assumption in the Multi-Object Diffusion could be relaxed by including explicit interreflection augmentation. Spatially varying reflectance can also be handled by extending the proposed framework with multi-patch reflectance maps: given a material-region segmentation, each region becomes its own “object” in our M-input pipeline with its own raw reflectance map, uniform reflectance, and shared illumination. Coordinated Scheduling and Axial Attention operate identically over patches, with more “objects” sharing one illumination. Future work will focus on relaxing the geometry constraint through joint shape estimation and extending the generative framework to represent near-field lighting for more complex scenes.

## Acknowledgements

This work was in part supported by JSPS KAKENHI 21H04893, 26H02534, and JST JPMJAP2305.

## References

1. Adobe Stock. <https://stock.adobe.com/> (Last accessed 23 June 2026)
2. HDRIs: poly haven. <https://polyhaven.com/hdris> (Last accessed 23 June 2026)

3. Burley, B., Studios, W.D.A.: Physically-Based Shading at Disney. In: SIGGRAPH. vol. 2012, pp. 1–7 (2012)
4. Careaga, C., Aksoy, Y.: Colorful Diffuse Intrinsic Image Decomposition in the Wild. ACM TOG **43**(6), 1–12 (2024)
5. Chen, W., Ling, H., Gao, J., Smith, E., Lehtinen, J., Jacobson, A., Fidler, S.: Learning to Predict 3D Objects with an Interpolation-based Differentiable Renderer. NeurIPS **32** (2019)
6. Chen, W., Litalien, J., Gao, J., Wang, Z., Fuji Tsang, C., Khamis, S., Litany, O., Fidler, S.: DIB-R++: Learning to Predict Lighting and Material with a Hybrid Differentiable Renderer. NeurIPS **34**, 22834–22848 (2021)
7. Chen, X., Peng, S., Yang, D., Liu, Y., Pan, B., Lv, C., Zhou, X.: IntrinsicAnything: Learning Diffusion Priors for Inverse Rendering Under Unknown Illumination. In: ECCV. pp. 450–467 (2024)
8. Chen, Z., Nobuhara, S., Nishino, K.: Invertible Neural BRDF for Object Inverse Rendering. IEEE TPAMI **44**(12), 9380–9395 (2021)
9. Chen, Z., Xu, T., Ge, W., Wu, L., Yan, D., He, J., Wang, L., Zeng, L., Zhang, S., Chen, Y.C.: Uni-Renderer: Unifying Rendering and Inverse Rendering Via Dual Stream Diffusion. In: CVPR. pp. 26504–26513 (2025)
10. Enyo, Y., Nishino, K.: Diffusion Reflectance Map: Single-Image Stochastic Inverse Rendering of Illumination and Reflectance. In: CVPR (2024)
11. Gardner, M.A., Sunkavalli, K., Yumer, E., Shen, X., Gambaretto, E., Gagné, C., Lalonde, J.F.: Learning to Predict Indoor Illumination from a Single Image. ACM TOG **9**(4) (2017)
12. Georgoulis, S., Rematas, K., Ritschel, T., Fritz, M., Tuytelaars, T., Van Gool, L.: What Is Around The Camera? In: CVPR. pp. 5170–5178 (2017)
13. Georgoulis, S., Rematas, K., Ritschel, T., Fritz, M., Van Gool, L., Tuytelaars, T.: DeLight-Net: Decomposing Reflectance Maps into Specular Materials and Natural Illumination. arXiv preprint arXiv:1603.08240 (2016)
14. Han, X., Nishino, K., Zickler, T.: Generative Perception of Shape and Material from Differential Motion. NeurIPS (2025)
15. Han, X., Zickler, T., Nishino, K.: Multistable Shape from Shading Emerges from Patch Diffusion. NeurIPS (2024)
16. Hartung, B., Kersten, D.: Distinguishing shiny from matte. Journal of Vision **2**(7), 551–551 (2002)
17. He, Z., Wang, T., Huang, X., Pan, X., Liu, Z.: Neural Lightrig: Unlocking Accurate Object Normal and Material Estimation with Multi-Light Diffusion. In: CVPR. pp. 26514–26524 (2025)
18. Huang, Z., Guo, Y.C., Wang, H., Yi, R., Ma, L., Cao, Y.P., Sheng, L.: MV-Adapter: Multi-View Consistent Image Generation Made Easy. In: ICCV. pp. 16377–16387 (2025)
19. Jakob, W., Speierer, S., Roussel, N., Nimier-David, M., Vicini, D., Zeltner, T., Nicolet, B., Crespo, M., Leroy, V., Zhang, Z.: Mitsuba 3 renderer (2022), <https://mitsuba-renderer.org> (Last accessed 23 June 2026)
20. Kajiya, J.T.: THE RENDERING EQUATION. In: SIGGRAPH. pp. 143–150 (1986)
21. Kocsis, P., Sitzmann, V., Nießner, M.: Intrinsic Image Diffusion for Indoor Single-view Material Estimation. In: CVPR. pp. 5198–5208 (2024)
22. Kuang, Z., Zhang, Y., Yu, H.X., Agarwala, S., Wu, E., Wu, J., et al.: Stanford-ORB: A Real-World 3D Object Inverse Rendering Benchmark. NeurIPS **36**, 46938–46957 (2023)

23. Labbé, Y., Manuelli, L., Mousavian, A., Tyree, S., Birchfield, S., Tremblay, J., Carpentier, J., Aubry, M., Fox, D., Sivic, J.: MegaPose: 6D Pose Estimation of Novel Objects via Render & Compare. In: CoRL (2022)
24. Liang, R., Gojcic, Z., Ling, H., Munkberg, J., Hasselgren, J., Lin, Z.H., Gao, J., Keller, A., Vijaykumar, N., Fidler, S., Wang, Z.: DiffusionRenderer: Neural Inverse and Forward Rendering with Video Diffusion Models. In: CVPR (2025)
25. Liang, R., Gojcic, Z., Nimier-David, M., Acuna, D., Vijaykumar, N., Fidler, S., Wang, Z.: Photorealistic Object Insertion with Diffusion-Guided Inverse Rendering. In: ECCV. pp. 446–465 (2024)
26. Liang, R., He, K., Gojcic, Z., Gilitschenski, I., Fidler, S., Vijaykumar, N., Wang, Z.: LuxDiT: Lighting Estimation with Video Diffusion Transformer (2025)
27. Lombardi, S., Nishino, K.: Reflectance and Natural Illumination from a Single Image. In: ECCV. pp. 582–595 (2012)
28. Lyu, L., Tewari, A., Habermann, M., Saito, S., Zollhöfer, M., Leimkühler, T., Theobalt, C.: Diffusion Posterior Illumination for Ambiguity-aware Inverse Rendering. ACM TOG **42**(6) (2023)
29. Maximov, M., Leal-Taixé, L., Fritz, M., Ritschel, T.: Deep Appearance Maps. In: ICCV. pp. 8729–8738 (2019)
30. Meka, A., Maximov, M., Zollhoefer, M., Chatterjee, A., Seidel, H.P., Richardt, C., Theobalt, C.: LIME: Live Intrinsic Material Estimation. In: CVPR. pp. 6315–6324 (2018)
31. Nishino, K.: Directional Statistics BRDF Model. In: ICCV. pp. 476–483 (2009)
32. Phongthawee, P., Chinchuthakun, W., Sinsunthithet, N., Jampani, V., Raj, A., Khungurn, P., Suwajanakorn, S.: DiffusionLight: Light Probes for Free by Painting a Chrome Ball. In: CVPR. pp. 98–108 (2024)
33. Ramamoorthi, R., Hanrahan, P.: A Signal-Processing Framework for Inverse Rendering. In: SIGGRAPH (2001)
34. Ramamoorthi, R., Hanrahan, P.: An Efficient Representation for Irradiance Environment Maps. In: SIGGRAPH. pp. 497–500 (2001)
35. Rematas, K., Ritschel, T., Fritz, M., Gavves, E., Tuytelaars, T.: Deep Reflectance Maps. In: CVPR. pp. 4508–4516 (2016)
36. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-Resolution Image Synthesis with Latent Diffusion Models. In: CVPR. pp. 10684–10695 (2022)
37. Shen, S., Bao, Z., Xu, W., Xiao, C.: IllumiDiff: Indoor Illumination Estimation from a Single Image with Diffusion Model. IEEE TVCG (2025)
38. Sun, S., Wang, Y., Zhang, H., Xiong, Y., Ren, Q., Fang, R., Xie, X., You, C.: Ouroboros: Single-step Diffusion Models for Cycle-consistent Forward and Inverse Rendering. In: ICCV. pp. 10386–10397 (2025)
39. Wang, H., Wang, Z., Long, X., Lin, C., Hancke, G., Lau, R.W.: MAGE: Single Image to Material-Aware 3D via the Multi-View G-Buffer Estimation Model. In: CVPR (2025)
40. Wang, L., Tran, D.M., Cui, R., TG, T., Dahl, A.B., Bigdeli, S.A., Frisvad, J.R., Chandraker, M.: Materialist: Physically Based Editing Using Single-Image Inverse Rendering. arXiv preprint arXiv:2501.03717 (2025)
41. Wei, X., Chen, G., Dong, Y., Lin, S., Tong, X.: Object-based Illumination Estimation with Rendering-aware Neural Networks. In: ECCV. pp. 380–396 (2020)
42. Xu, Z., Sunkavalli, K., Hadap, S., Ramamoorthi, R.: Deep Image-Based Relighting from Optimal Sparse Samples. ACM TOG **37**(4), 1–13 (2018)
43. Yamashita, K., Enyo, Y., Nobuhara, S., Nishino, K.: nLMVS-Net: Deep Non-Lambertian Multi-View Stereo. In: WACV. pp. 3037–3046 (2023)

- 44. Yu, H.X., Agarwala, S., Herrmann, C., Szeliski, R., Snavely, N., Wu, J., Sun, D.: Accidental Light Probes. In: CVPR. pp. 12521–12530 (2023)
- 45. Zeng, Z., Deschaintre, V., Georgiev, I., Hold-Geoffroy, Y., Hu, Y., Luan, F., Yan, L.Q., Hašan, M.: RGB $\leftrightarrow$ X: Image decomposition and synthesis using material- and lighting-aware diffusion models. In: SIGGRAPH Conference Papers (2024)
- 46. Zhang, L., Rao, A., Agrawala, M.: Adding Conditional Control to Text-to-Image Diffusion Models. In: ICCV. pp. 3836–3847 (2023)
- 47. Zheng, R., Zhang, Q., Long, C., Zheng, W.S.: DNF-Intrinsic: Deterministic Noise-Free Diffusion for Indoor Inverse Rendering. In: ICCV. pp. 10342–10352 (2025)