

LumiX: Structured and Coherent Text-to-Intrinsic Generation

Xu Han^{1,2◊}

Biao Zhang²

Xiangjun Tang²
¹HUST ²KAUST

Xianzhi Li^{1†}

Peter Wonka^{2†}



Figure 1. **LumiX for Text-to-Intrinsic Generation.** Given a text prompt, LumiX jointly generates a coherent set of intrinsic maps, including RGB color, albedo, irradiance, depth, and normal. Built on a powerful diffusion prior, it produces diverse and physically grounded intrinsic images, and can also perform image-conditioned intrinsic decomposition even though it is trained with text-only conditioning.

Abstract

We present *LumiX*, a structured diffusion framework for coherent text-to-intrinsic generation. Conditioned on text prompts, *LumiX* jointly generates a comprehensive set of intrinsic maps (e.g., albedo, irradiance, normal, depth, and final color), providing a structured and physically consistent description of an underlying scene. This is enabled by two key contributions: 1) *Query-Broadcast Attention*, a mechanism that ensures structural consistency by sharing queries across all maps in each self-attention block. 2) *Tensor LoRA*, a tensor-based adaptation that parameter-efficiently models cross-map relations for efficient joint training. Together, these designs enable stable joint diffusion training and unified generation of multiple intrinsic properties. Experiments show that *LumiX* produces coherent and physically meaningful results, achieving 23% higher alignment and a better preference score (0.19 vs. -0.41) compared to the state of the art, and it can also perform image-conditioned intrinsic decomposition within the same framework.

1. Introduction

Recent advances in text-to-image diffusion models [39] have made it possible to generate realistic and detailed images

from natural language. However, these models still produce a single RGB image and do not reveal the underlying structure of the scene, such as geometry, lighting, or material properties. For many vision and graphics tasks, a single image is not enough. We often need a structured representation that separates these intrinsic factors, including albedo, irradiance, normal, depth, and the final color. Together, these maps describe both the physical and semantic aspects of a scene. This raises an important question: can we directly generate a coherent set of intrinsic maps from text? In this work, we study this problem under the term *text-to-intrinsic generation* and aim to build a unified model that can jointly generate scene content and physical attributes.

Most existing studies focus on intrinsic image decomposition [6, 10, 15, 30], where a model receives a rendered or captured image and predicts its intrinsic components such as albedo, shading, or normal. This formulation is effective for analysis but inherently limited, as it depends on a given image and cannot generate new scenes from language. *Text-to-intrinsic generation* removes this restriction. It starts from text and produces both the scene content and its physical attributes within a single generative process. The same diffusion backbone can also switch to image-conditioned decomposition during inference, enabling both generation and understanding within one unified framework. Therefore, text-to-intrinsic generation is not merely a generalization of intrinsic decomposition, but a broader step toward unified multi-attribute generation.

[†] Corresponding authors.

[◊] This work was done during Xu Han’s internship at KAUST.




Figure 2. **Overview of LumiX.** Our goal is to generate a coherent set of intrinsic maps from text. **Left: Training.** Multiple intrinsic images are encoded into the latent space and concatenated along the batch dimension. We introduce *Query-Broadcast Attention* (Sec. 4.1) to ensure pixel alignment across properties, and *Tensor LoRA* (Sec. 4.2) to efficiently finetune the KV projections for each property. Different timesteps are assigned to different properties for flexible conditioning. **Right: Inference.** Given a text or image input, LumiX jointly outputs all intrinsic maps in a single forward pass, supporting both text-to-intrinsic generation and intrinsic decomposition.

The main challenge in text-to-intrinsic generation is maintaining cross-map structural consistency under text-only conditioning. In image-conditioned settings such as intrinsic decomposition or RGB-to-X translation [7, 51], all outputs share the same input image, so spatial alignment is naturally preserved. In contrast, each intrinsic map here is generated from its own noise sample and is only weakly coupled through the shared text embedding. Without an explicit spatial anchor, structures can drift across maps: an object may appear in the color map but disappear in the normal map, or the geometry in the depth map may not match the shading in the irradiance map. Ensuring consistent scene structure across all intrinsic maps without image supervision is therefore the central problem of this task.

To address the consistency challenge across multiple images, related works have explored two strategies. One strategy [4] trains each intrinsic map independently, using separate models or LoRA [18] branches. With strong diffusion priors, these models can produce high-quality albedo, normal, or depth maps, but the results often lack semantic alignment across maps. Another strategy [23] builds a fully coupled framework that concatenates features from all maps and applies cross-intrinsic attention to enforce global interaction. This improves consistency but makes training unstable, and computation scales quadratically with the number of maps. Overall, these directions reveal a clear trade-off between consistency and efficiency, motivating a more structured sharing mechanism that maintains semantic coherence while preserving property-specific independence.

We present a systematic analysis of the pipeline, examining it from two critical perspectives: the *forward pass* and the *training/finetuning* process. In the forward pass, we ensure structural consistency by broadcasting identical queries to all intrinsic maps (a mechanism we term **Query-Broadcast Attention**). This design is grounded in prior research on content and style separation [13, 17], which

establishes that queries primarily capture content, whereas keys and values contribute to style in self-attentions. For the training/finetuning process, we present several possible low-rank-based (LoRA) finetuning strategies. We extend the conventional matrix-based LoRA to a tensor-based formulation. We demonstrate that our proposed **Tensor LoRA** not only yields the best quality results but also effectively preserves structural consistency and operates with high efficiency. Our main contributions:

- We present LumiX, a structured diffusion framework for coherent text-to-intrinsic generation that jointly produces multiple physically consistent intrinsic maps.
- We design Query-Broadcast Attention to align scene content across maps by sharing queries while keeping attribute projections independent.
- We propose Tensor LoRA, a tensor-based low-rank adaptation that models global cross-map relations in a parameter-efficient and scalable way.

2. Related Work

2.1. Intrinsic and Multi-map Generation

Intrinsic image decomposition has long focused on separating an image into physically meaningful factors such as albedo, shading, and materials [6, 7, 10, 19, 27, 30, 40, 43]. Most methods are image-conditioned and rely on predefined physical priors, such as albedo–shading consistency or inverse–forward rendering cycles [11, 12, 31, 41, 57]. While effective for analysis, they cannot synthesize new scenes or handle open-vocabulary text input. To move beyond decomposition, later works predict multiple complementary maps such as depth [20], normal [49], or segmentation [55]. Feed-forward [59, 60] and GAN-based models [2] support multi-map prediction but require dense supervision and often lose structural coherence.

Diffusion-based methods [22, 35, 51] extend text-to-

image backbones or add task tokens for multi-map generation, but sharing one backbone across maps easily leads to interference as the number of maps grows. Recent designs add per-map LoRA adapters [33] or communication layers [4, 42], though they do not clearly separate shared structure from property-specific variation. More recently, text-to-X diffusion models [14, 23, 47] explore intrinsic generation without image supervision. Among them, IntrinsicX [23] improves consistency with a single fused attention space, but this limits scalability and interpretability. In contrast, our method adopts a structured formulation that preserves pixel alignment across intrinsic maps without heavy fusion, enabling scalable and coherent text-to-intrinsic generation.

2.2. Cross-map Attention and Adaptation

Attention layers are central to diffusion models, as they integrate conditional information. Many works manipulate the query, key, or value projections to control generation [5, 8, 44, 52]. Queries typically capture spatial and semantic content, while keys and values encode appearance or modality cues [17, 50]. Attend-and-Excite [8] modifies Q to highlight object presence, and MasaCtrl [5] mixes Q, K, and V across images to improve spatial consistency. Other methods introduce separate K/V branches [52, 53] or adapter modules [25, 45, 50] to inject conditions such as segmentation, depth, or reference images [1, 17]. Recent works extend these ideas to multi-map diffusion by aligning K/V projections across maps or by learning property-specific branches [4, 58]. These methods focus on external conditioning or reference, rather than ensuring structural consistency between maps generated from text. Our Query-Broadcast Attention reframes attention for intrinsic map alignment by sharing queries across maps while keeping keys and values independent, allowing shared scene structure and property-specific variation to be preserved.

2.3. Parameter-efficient and Tensorized Adaptation

Low-rank adaptation methods such as LoRA [18, 24, 32, 54] fine-tune large models efficiently by adding small low-rank updates to frozen weights [16]. In multi-map generation, assigning an independent LoRA to each map [4, 33, 42] ignores cross-map relations and causes quadratic parameter growth as the number of maps increases. To improve efficiency, tensor decomposition methods such as tensor-train [37] and tensor-ring [56] provide a compact way to represent high-dimensional parameters. Following this idea, we design Tensor LoRA, which encodes all LoRA updates in one structured tensor and factorizes it into shared cores and per-map components. This formulation captures cross-map relations while keeping the parameter cost close to linear.

3. Preliminaries

Finetuning with LoRA. In a denoising diffusion network, each self-attention layer accepts a hidden feature $H \in \mathbb{R}^{d \times L}$ as input,

$$\text{SA}(H) = \text{softmax}\left(QK^\top/\sqrt{d}\right)V, \quad (1)$$

where $Q = W_Q H$, $K = W_K H$, and $V = W_V H$. The learnable parameters are $\{W_Q \in \mathbb{R}^{d \times d}, W_K \in \mathbb{R}^{d \times d}, W_V \in \mathbb{R}^{d \times d}\}$. Low-Rank Adaptation (LoRA) is applied for efficient finetuning,

$$W \leftarrow W + \Delta, \quad \Delta = AB^\top, \quad (2)$$

where $A \in \mathbb{R}^{d \times R}$, $B \in \mathbb{R}^{d \times R}$ are learnable low-rank adapters. Usually, $R \ll d$. For a single token $h \in \mathbb{R}^d$, the calculation is often done by

$$Wh + A(B^\top h), \quad (3)$$

which is a more efficient version of $(W + AB^\top)h$.

Flow matching training. To train flow matching on a dataset $\mathcal{D}$, we minimize the loss function,

$$\min_{\theta} \mathbb{E}_{t \in [0,1], \epsilon \sim \mathcal{N}(0,1), \mathbf{z} \sim \mathcal{D}} \|v_{\theta}(\mathbf{z}_t, t, \mathcal{C}) - (\epsilon - \mathbf{z})\|, \quad (4)$$

where $\mathcal{C}$ denotes the conditioning input, and v_{θ} is the denoising network parameterized by θ . Most models, such as Stable Diffusion [39] and FLUX [26], are trained in the latent space $\mathbf{z}$ instead of the original pixel space $\mathbf{x}$.

4. Method

Given a text prompt condition $\mathcal{C}$, our goal is to jointly generate pixel-aligned intrinsic maps of the same scene: color $\mathbf{x}^{(c)}$, albedo $\mathbf{x}^{(a)}$, irradiance $\mathbf{x}^{(i)}$, depth $\mathbf{x}^{(d)}$, and normal $\mathbf{x}^{(n)}$. We denote their corresponding latent representations, obtained via a pretrained autoencoder: $\mathbf{z}^{(c)}$, $\mathbf{z}^{(a)}$, $\mathbf{z}^{(i)}$, $\mathbf{z}^{(d)}$, and $\mathbf{z}^{(n)}$, respectively.

These maps share a common spatial layout but capture different intrinsic properties, including appearance, reflectance, illumination, and geometry. Our task has two primary objectives: (1) **consistency**: defined as structural consistency across properties to ensure shared content and pixel alignment, and (2) **quality**: defined as realism within each property to preserve its distinct characteristics.

We address this task by finetuning existing text-to-image diffusion models. As a starting point for our discussion, we consider the naive solution to train $M = 5$ independent models, one for each map (color, albedo, irradiance, depth, and normal). This approach promises good per-map *quality*, but fails to maintain structural *consistency* across the maps, as they are generated in isolation.

To ensure *consistency*, we introduce two core components that link the different maps together: Query-Broadcast Attention in Sec. 4.1 and a Tensor LoRA in Sec. 4.2. The former modifies the forward pass to enable information exchange across all maps, while the latter improves training efficiency by proposing a lightweight finetuning strategy.

4.1. Forward with Query-Broadcast Attention

A baseline approach is to compute the self-attention for each model $m \in \mathcal{M} = \{c, a, i, d, n\}$ (color, albedo, irradiance, depth, and normal) independently,

$$H^{(m)} \leftarrow \text{softmax} \left(Q^{(m)} K^{(m)\top} / \sqrt{d} \right) V^{(m)}. \quad (5)$$

The design, hereafter referred to as **Vanilla Attention**, evidently ignores potential interactions between the different models in $\mathcal{M}$.

IntrinsiX [23] proposed concatenating K and V from all models,

$$H^{(m)} = \text{softmax} \left(Q^{(m)} K^{(\mathcal{M})\top} / \sqrt{d} \right) V^{(\mathcal{M})}, \quad (6)$$

where $K^{(\mathcal{M})} = \text{Concat}([K^{(c)}, K^{(a)}, \dots])$, $V^{(\mathcal{M})} = \text{Concat}([V^{(c)}, V^{(a)}, \dots])$. This approach, termed **Cross-Intrinsic Attention** in [23], improves consistency by exchanging information between models. However, its computational complexity is M times larger than the vanilla approach (Eq. (5)).

To address this issue, we propose a more efficient design. Our approach is inspired by [13, 17, 58], which argue that the K and V matrices capture modality-specific *style* (e.g., reflectance, illumination, or surface-normal variations), and that the Q matrix encodes the scene’s *content*. Based on this insight, we broadcast the color model’s query matrix (i.e., $Q^{(c)}$) to all models. Specifically, for each model m ,

$$H^{(m)} \leftarrow \text{softmax} \left(Q^{(c)} K^{(m)\top} / \sqrt{d} \right) V^{(m)} \quad (7)$$

We refer to the design as **Query-Broadcast Attention**. In this formulation, all models utilize the shared query $Q^{(c)}$ from the color map model. We empirically demonstrate that our strategy in Eq. (7) outperforms IntrinsiX in Eq. (6) in later sections, while being significantly more efficient.

4.2. Finetuning with Tensor LoRA

As discussed above, $Q^{(c)}$ is broadcast to all models. Consequently, we do not fine-tune the query projection W_Q . The learnable parameters for a single attention layer are

$$\left\{ (W_K^{(c)}, W_V^{(c)}), (W_K^{(a)}, W_V^{(a)}), (W_K^{(i)}, W_V^{(i)}), \dots \right\}, \quad (8)$$

which results in $2 \times M \times d \times d$ parameters in total. When finetuning with LoRA, we do not directly update the original

weight matrices W . Instead, we only consider the low-rank updates,

$$\underbrace{Wh}_{\text{base model}} + \underbrace{AB^\top h}_{\text{LoRA update}} \quad (9)$$

In the following, we discuss different designs of this LoRA update.

S Separate LoRA. For a single weight matrix $W^{(m)}$ and its update $\Delta^{(m)}$, we apply LoRA in Eq. (2) to obtain the parameters $A^{(m)} \in \mathbb{R}^{d \times R}$ and $B^{(m)} \in \mathbb{R}^{d \times R}$. The resulting update for each map m ’s activation is calculated as $A^{(m)}(B^{(m)\top} h^{(m)})$. While this approach is efficient, it ignores any interaction between the M models, and thus fails to maintain structural consistency. This independent update process can be visualized in a compact block-diagonal form for all M models,

$$\begin{bmatrix} \Delta^{(c)} & 0 & 0 & 0 & 0 \\ 0 & \Delta^{(a)} & 0 & 0 & 0 \\ 0 & 0 & \Delta^{(i)} & 0 & 0 \\ 0 & 0 & 0 & \Delta^{(d)} & 0 \\ 0 & 0 & 0 & 0 & \Delta^{(n)} \end{bmatrix} \begin{bmatrix} h^{(c)} \\ h^{(a)} \\ h^{(i)} \\ h^{(d)} \\ h^{(n)} \end{bmatrix} \quad (10)$$

F Fused LoRA. A straightforward solution is to fuse all the activations $h^{(m)} \in \mathbb{R}^d$ from all models into a single, concatenated vector $h^{(\mathcal{M})} \in \mathbb{R}^{Md}$. A single LoRA update is then applied on a large update matrix $\Delta^{(\mathcal{M})} \in \mathbb{R}^{Md \times Md}$, which is further decomposed into low-rank matrices $A^{(\mathcal{M})} \in \mathbb{R}^{Md \times R}$ and $B^{(\mathcal{M})} \in \mathbb{R}^{Md \times R}$. The fused output $A^{(\mathcal{M})}(B^{(\mathcal{M})\top} h^{(\mathcal{M})})$, is subsequently split back into M separate models. Unlike **Separate LoRA**, this approach creates a dense update matrix where the off-diagonal blocks are non-zero, explicitly modeling the interactions between all models:

$$\underbrace{\begin{bmatrix} \Delta^{(cc)} & \Delta^{(ca)} & \Delta^{(ci)} & \Delta^{(cd)} & \Delta^{(cn)} \\ \Delta^{(ac)} & \Delta^{(aa)} & \Delta^{(ai)} & \Delta^{(ad)} & \Delta^{(an)} \\ \Delta^{(ic)} & \Delta^{(ia)} & \Delta^{(ii)} & \Delta^{(id)} & \Delta^{(in)} \\ \Delta^{(dc)} & \Delta^{(da)} & \Delta^{(di)} & \Delta^{(dd)} & \Delta^{(dn)} \\ \Delta^{(nc)} & \Delta^{(na)} & \Delta^{(ni)} & \Delta^{(nd)} & \Delta^{(nn)} \end{bmatrix}}_{\Delta^{(\mathcal{M})} \in \mathbb{R}^{Md \times Md}} \underbrace{\begin{bmatrix} h^{(c)} \\ h^{(a)} \\ h^{(i)} \\ h^{(d)} \\ h^{(n)} \end{bmatrix}}_{h^{(\mathcal{M})} \in \mathbb{R}^{Md}} \quad (11)$$

H Hybrid LoRA. Inspired by the block-matrix form of **Separate LoRA** in Eq. (10) and **Fused LoRA** in Eq. (11), we propose a hybrid approach. We enhance **Separate LoRA** with non-zero off-diagonal entries. Specifically, we use higher rank (R_1) decomposition in **diagonal** entries and lower rank ($R_2 < R_1$) in **off-diagonal** entries,

Table 1. **Comparison of attention and LoRA designs for text-to-intrinsic generation on our Hypersim test set [38].** #P: trainable parameters (M) per attention block; FLOPs: FLOPs (G) of LoRA and attention per block; Align.: cross-map consistency; Other columns: human preference scores [21, 46] across intrinsic maps. [†]Without the first training stage in IntrinsiX [23]. [‡]Official IntrinsiX repository.

Method	LoRA	#P↓	FLOPs↓	Align.↑	Color↑		Albedo↑		Irradiance↑		Avg.↑
			LoRA/Attn		ImageReward	PickScore	ImageReward	PickScore	ImageReward	PickScore	IR/PS
Vanilla Attention (FLUX)											
Separate	S	2.95	9.1 / 145.1	2.40	0.25	20.89	-0.33	19.47	0.26	20.75	0.06 / 20.37
Fused	F	2.95	9.1 / 145.1	5.91	-0.28	19.94	-0.69	19.30	-0.61	19.81	-0.53 / 19.68
Hybrid	H	5.90	18.1 / 145.1	6.86	-0.10	20.36	-0.62	19.24	-0.36	19.98	-0.36 / 19.86
Tensor	T	2.46	14.1 / 145.1	7.16	-0.01	20.16	-0.44	19.40	-0.38	19.87	-0.28 / 19.81
Cross-Intrinsic Attention [23]											
IntrinsiX [†]	S	2.95	9.1 / 724.7	3.17	-0.91	18.87	-1.54	18.24	-1.21	18.50	-1.22 / 18.54
IntrinsiX [‡]	S	-	-	6.73	-0.37	19.83	-0.44	19.73	-	-	-0.41 / 19.78
IntrinsiX [†]	T	2.46	14.1 / 724.7	7.65	-0.43	19.62	-0.62	19.31	-0.51	19.64	-0.52 / 19.52
IntrinsiX	T	2.46	14.1 / 724.7	7.98	-0.05	20.01	-0.48	19.28	-0.30	19.85	-0.28 / 19.71
Query-Broadcast Attention (Ours)											
Separate	S	2.16	6.7 / 145.1	4.40	-0.06	20.46	-0.14	19.76	-0.16	20.40	-0.12 / 20.21
Fused	F	2.16	6.7 / 145.1	6.82	0.05	20.28	-0.47	19.45	-0.29	20.06	-0.24 / 19.93
LumiX	H	4.03	12.4 / 145.1	8.21	0.40	20.32	-0.04	19.74	0.17	20.31	0.18 / 20.12
LumiX	T	2.34	12.1 / 145.1	8.30	0.45	21.02	0.04	19.88	0.09	20.66	0.19 / 20.52

$$\begin{bmatrix} \Delta(cc) & \Delta(ca) & \Delta(ci) & \Delta(cd) & \Delta(cn) \\ \Delta(ac) & \Delta(aa) & \Delta(ai) & \Delta(ad) & \Delta(an) \\ \Delta(ic) & \Delta(ia) & \Delta(ii) & \Delta(id) & \Delta(in) \\ \Delta(dc) & \Delta(da) & \Delta(di) & \Delta(dd) & \Delta(dn) \\ \Delta(nc) & \Delta(na) & \Delta(ni) & \Delta(nd) & \Delta(nn) \end{bmatrix} \begin{bmatrix} h^{(c)} \\ h^{(a)} \\ h^{(i)} \\ h^{(d)} \\ h^{(n)} \end{bmatrix} \quad (12)$$

T Tensor LoRA. In our experiments, we found that both **Fused LoRA** and **Hybrid LoRA** exhibit better consistency but suffer from inefficiency due to the large dimensionality of their update matrices. Inspired by the element-wise form of LoRA ($\Delta = AB^T$),

$$\Delta[i, j] = \sum_{\alpha=1}^r A[i, \alpha] B[j, \alpha], \quad (13)$$

we propose to represent $\Delta^{(\mathcal{M})} \in \mathbb{R}^{M \times M \times d}$ using a tensor decomposition. Specifically, we reshape the matrix into a 4th-order tensor $\Delta^{(\mathcal{M})} \in \mathbb{R}^{N \times d_{\text{out}} \times M \times d_{\text{in}}}$, where M is the number of input models and N is the number of output models. As a general case, we allow $M \neq N$ and $d_{\text{in}} \neq d_{\text{out}}$. The element-wise form of our proposed decomposition is

$$\Delta^{(\mathcal{M})}[i, j, k, l] = \sum_{\alpha_1=1}^{R_1} \sum_{\alpha_2=1}^{R_2} A[i, j, \alpha_1] B[i, k, \alpha_2] C[i, l, \alpha_1, \alpha_2], \quad (14)$$

where $A \in \mathbb{R}^{N \times d_{\text{out}} \times R_1}$, $B \in \mathbb{R}^{N \times M \times R_2}$, $C \in \mathbb{R}^{N \times d_{\text{in}} \times R_1 \times R_2}$. This is related to tensor-train (TT) decom-

positions [37]. The output is calculated using tensor contractions (einsum is a built-in function in many libraries such as NumPy, PyTorch and Jax):

```
# N d R1 R2, M d -> N M R1 R2
Ch = einsum('ndrs, md -> nmrs', C, h)
# N M R2, N M R1 R2 -> N R1
BCh = einsum('nms, nmrs -> nr', B, Ch)
# N d R1, N R1 -> N d
ABCh = einsum('ndr, nr -> nd', A, BCh)
```

In practice, we set $M = N$, $d_{\text{in}} = d_{\text{out}}$, and $R_1 = R_2$.

5. Experiment

5.1. Implementation Details

Training Details. We fine-tune the FLUX.1-dev model [26] on a subset of the Hypersim dataset [38] containing about 3K images, which is sufficient to produce stable and high-quality results. We use BLIP-2 [29] to generate captions from color images as text inputs. For Tensor LoRA, we set the rank to 8, resulting in about 133.1M trainable parameters in total. We use the Prodigy optimizer [36] with a learning rate of 1.0 and a batch size of 16, and train the model for 10K steps on four NVIDIA A100 (80GB) GPUs, taking about 40 hours. All images are resized to 512×512 resolution using ratio-preserving scaling and random cropping for data augmentation, following RGB↔X [51].

Disentangled Timestep Sampling. Inspired by prior works [4, 9], we adopt a *timestep-disentangled* scheme that assigns each intrinsic property an independent diffusion timestep. This introduces different noise levels across



Figure 3. **Visual comparison of Attention and LoRA designs.** Using vanilla attention with separate LoRA leads to the weakest alignment. Replacing it with our Tensor LoRA improves consistency and quality. IntrinsicX without its first training stage becomes unstable, while substituting Tensor LoRA alleviates collapse but still struggles to distinguish different modality characteristics. Our Query-Broadcast Attention combined with Hybrid or Tensor LoRA achieves the best results, producing consistent and high-quality intrinsic maps.

properties, acting as a soft mask that encourages flexible conditioning. Although the model is trained with text-only conditioning, it naturally supports image-conditioned generation at inference by keeping one property clean while denoising others. Combined with Query-Broadcast Attention, this ensures structural consistency between generated maps and the conditioning image, enabling intrinsic decomposition as well as diverse conditional and editing tasks.

Baselines. We compare LumiX with IntrinsicX [23], the first model targeting text-to-PBR generation. IntrinsicX produces albedo, material, and normal maps, but does not generate irradiance, depth, or full-color images. For fairness, we reproduce its stage-2 training using separate LoRA and vanilla attention, and evaluate albedo, irradiance, and color outputs. Since IntrinsicX cannot perform intrinsic image decomposition, we additionally include RGB $\leftrightarrow$ X and Colorful Shading [7], two state-of-the-art decomposition baselines trained

on large-scale datasets, as references for albedo quality and cross-map consistency. Finally, we conduct ablations on different attention and LoRA designs to analyze their effect on image quality and structural alignment.

Evaluation Metrics. We use two types of quantitative evaluation. For text-to-intrinsic generation, evaluating the quality of intrinsic maps is challenging because no ground-truth supervision is available. We therefore use human preference score models ImageReward [46] and PickScore [21] to assess perceptual quality. To measure cross-map consistency, we use Qwen3-VL [48], a strong vision-language evaluator. We sample 200 images from the Hypersim test set, generate captions with BLIP-2 [29], and use them as both model inputs and evaluation prompts. For intrinsic decomposition, we follow prior work [7] and perform zero-shot evaluation on ARAP [3], reporting RMSE and SSIM for albedo reconstruction. To further examine generalization, we also



Figure 4. **Text-to-Intrinsic Generation Comparison.** Both models are built upon FLUX. While IntrinsiX tends to overfit to specific indoor scenes, our method preserves FLUX’s strong prior and produces consistent, high-quality intrinsic maps even for out-of-domain prompts.



Figure 5. **Intrinsic Decomposition Comparison.** Our method performs intrinsic decomposition on in-the-wild data, producing albedo maps with less embedded lighting and generating consistent, high-quality intrinsic maps across all properties.

Table 2. **Zero-shot albedo evaluation on the synthetic ARAP Dataset [3].** We evaluate intrinsic decomposition quality using RMSE and SSIM. T2I indicates whether each method supports text-to-image generation. * denotes non-zero-shot methods. Our method is the only one that supports text-to-intrinsic generation.

Method	T2I	Base	RMSE↓	SSIM↑
<i>Without Diffusion Prior</i>				
Chromaticity	✗	-	0.193	0.710
Constant Shading	✗	-	0.264	0.693
Lettry et al. [28]	✗	-	0.193	0.732
NIID-Net [34]*	✗	-	0.129	0.788
Zhu et al. [59]	✗	-	0.184	0.729
Ordinal Shading [6]	✗	-	0.162	0.751
Colorful Shading [7]	✗	-	0.149	0.796
<i>With Diffusion Prior</i>				
IntrinsicAnything [11]	✗	SD-CLIP _I	0.171	0.692
IID [22]	✗	SD2-Depth	0.160	0.738
RGB↔X [51]	✗	SD2.1	0.238	0.564
LumiX (Ours)	✓	FLUX	0.165	0.753

evaluate on 50 in-the-wild photographs collected from the Internet, covering diverse indoor and outdoor scenes.

5.2. Text-to-Intrinsic Generation

Comparisons on the Hypersim dataset. Tab. 1 and Fig. 3 compare different attention and LoRA designs for text-to-intrinsic generation. The vanilla FLUX attention with Sepa-

rate LoRA produces high-quality individual maps but lacks cross-map coherence, resulting in the lowest alignment score. Fused LoRA and Hybrid LoRA partially address this issue through shared updates, yet they either compromise fidelity or incur a higher parameter cost. In contrast, our Tensor LoRA achieves a better trade-off between quality, consistency, and efficiency. We also evaluate the Cross-Intrinsic Attention from IntrinsiX. Without stage-1 initialization using vanilla attention, it tends to collapse since all properties rely on globally shared K and V . Stage-1 pretraining improves stability, but replacing Separate LoRA with our Tensor LoRA enables robust joint training even without such initialization. Finally, integrating our Query-Broadcast Attention improves all LoRA variants. Fused LoRA achieves moderate alignment gains, while Hybrid LoRA enhances structural coherence at the cost of more parameters. Combining Query-Broadcast Attention with Tensor LoRA yields the best overall performance. As shown in Fig. 4, our model produces visually coherent intrinsic maps with high perceptual quality and accurate geometric consistency.

5.3. Intrinsic Image Decomposition

Albedo quality on ARAP dataset. Although our model is trained only with text conditioning, it can also take an image as input during inference to perform intrinsic decomposition. We follow the Colorful Shading [7] setup and evaluate albedo quality on the ARAP dataset, as shown in Tab. 2. Models without diffusion priors are usually trained from scratch on large datasets and can achieve strong re-




Figure 6. **Image-Conditioned Intrinsic Decomposition.** Given a color image, our model performs intrinsic decomposition on Hypersim and in-the-wild data, and can further take albedo or irradiance as additional conditions. Conditioned images are shown with blue boxes.

Table 3. **Zero-shot albedo evaluation on our in-the-wild dataset.** We evaluate intrinsic decomposition quality using ImageReward and PickScore across different variants of our method.

Method	ImageReward $\uparrow$	PickScore $\uparrow$
RGB $\leftrightarrow$ X [51]	-0.20	20.01
Colorful Shading [7]	0.06	20.03
LumiX (Ours)	0.14	20.16

sults. Diffusion-based methods typically build on Stable Diffusion U-Net architectures and are trained specifically for image decomposition. In contrast, our model is based on the FLUX DiT architecture and is trained without any image conditioning or ControlNet [53] design. Even so, it achieves comparable or better performance than state-of-the-art diffusion-based baselines. Moreover, it is the only model that can also generate intrinsic maps directly from text.

Generalization to in-the-wild data. We evaluate our model on an in-the-wild dataset containing diverse real-world photographs from both indoor and outdoor scenes. As shown in Tab. 3, our method outperforms state-of-the-art intrinsic decomposition models, achieving the highest preference scores (ImageReward and PickScore) despite being trained on a much smaller dataset (3K vs. 900K). Qualitative comparisons in Fig. 5 further confirm that our approach generalizes well to real-world imagery. Our generated albedo maps contain less baked-in lighting, and the predicted normals remain geometrically consistent under complex illumination.

5.4. Ablation Study

Tab. 4 reports ablations on the Hypersim dataset. Fine-tuning the Query projection W_Q in our Query-Broadcast Attention increases parameters and FLOPs but degrades intrinsic qual-

Table 4. **Ablation of different variants of our method.** We report the averaged ImageReward and PickScore. #P: trainable parameters (M) per attention block; #F: FLOPs (G) of LoRA per block.

Method	#P $\downarrow$	#F $\downarrow$	Align. $\uparrow$	Average	
				ImageReward $\uparrow$	PickScore $\uparrow$
LumiX	2.34	12.1	8.30	0.19	20.52
+ Tune W_Q	2.46	14.1	7.14	-0.09	20.04
$R = 4$	0.68	4.7	7.86	-0.18	19.79
$R = 8$	2.34	12.1	8.30	0.19	20.52
$R = 12$	4.98	22.5	8.10	0.14	20.29

ity, as it disrupts the diffusion prior and confuses learning objectives. Varying the rank $R_1 = R_2 = R$, we find $R = 8$ gives the best balance of quality and efficiency, while $R = 4$ remains competitive with minimal parameters.

6. Conclusion

We presented LumiX, a structured diffusion framework for text-to-intrinsic generation. By combining Query-Broadcast Attention, which enforces pixel-level alignment across intrinsic maps, with Tensor LoRA, which provides lightweight yet expressive parameter adaptation, LumiX produces coherent intrinsic images with strong structural consistency and high visual fidelity. Our experiments show that stable multi-map generation arises not only from scaling up supervision, but also from structured parameter sharing. We plan to scale LumiX to a broader set of intrinsic properties and larger datasets in the future, moving toward models that achieve unified understanding and generation of scene geometry, materials, and illumination.

References

- [1] Yuval Alaluf, Amir Hertz, Yael Vinker, and Daniel Cohen-Or. Cross-image attention for zero-shot appearance transfer. In *ACM SIGGRAPH 2024 Conference Papers*, 2024. 3
- [2] Anand Bhattad, Daniel McKee, Derek Hoiem, and David Alexander Forsyth. Stylegan knows normal, depth, albedo, and more. *arXiv preprint arXiv:2306.00987*, 2023. 2
- [3] Nicolas Bonneel, Balazs Kovacs, Sylvain Paris, and Kavita Bala. Intrinsic decompositions for image editing. In *Computer graphics forum*, pages 593–609. Wiley Online Library, 2017. 6, 7
- [4] Kwon Byung-Ki, Qi Dai, Lee Hyoseok, Chong Luo, and Tae-Hyun Oh. Jointdit: Enhancing rgb-depth joint modeling with diffusion transformers. *arXiv preprint arXiv:2505.00482*, 2025. 2, 3, 5
- [5] Mingdeng Cao, Youwei Wang, Kan Zhang, Yongming Chen, Jie Chen, and Biao Wang. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12446–12456, 2023. 3
- [6] Chris Careaga and Yağız Aksoy. Intrinsic image decomposition via ordinal shading. *ACM Transactions on Graphics (TOG)*, 2023. 1, 2, 7
- [7] Chris Careaga and Yağız Aksoy. Colorful diffuse intrinsic image decomposition in the wild. *ArXiv abs/2409.13690*, 2024. 2, 6, 7, 8
- [8] Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–10, 2023. 3
- [9] Boyuan Chen, Diego Martí Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. Diffusion forcing: Next-token prediction meets full-sequence diffusion. *Advances in Neural Information Processing Systems*, 37:24081–24125, 2024. 5
- [10] Qifeng Chen and Vladlen Koltun. A simple model for intrinsic image decomposition with depth cues. In *Proceedings of the IEEE international conference on computer vision*, pages 241–248, 2013. 1, 2
- [11] Xi Chen, Sida Peng, Dongchen Yang, Yuan Liu, Bowen Pan, Chengfei Lv, and Xiaowei Zhou. Intrinsicanything: Learning diffusion priors for inverse rendering under unknown illumination. In *European Conference on Computer Vision*, pages 450–467. Springer, 2024. 2, 7
- [12] Zhifei Chen, Tianshuo Xu, Wenhong Ge, Leyi Wu, Dongyu Yan, Jing He, Luozhou Wang, Lu Zeng, Shunsi Zhang, and Ying-Cong Chen. Uni-renderer: Unifying rendering and inverse rendering via dual stream diffusion. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26504–26513, 2025. 2
- [13] Jiwoo Chung, Sangeek Hyun, and Jae-Pil Heo. Style injection in diffusion: A training-free approach for adapting large-scale diffusion models for style transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2, 4
- [14] Alara Dirik, Tuanfeng Wang, Duygu Ceylan, Stefanos Zafeiriou, and Anna Frühstück. Prism: A unified framework for photorealistic reconstruction and intrinsic scene modeling. *arXiv preprint arXiv:2504.14219*, 2025. 3
- [15] Roger Grosse, Micah K Johnson, Edward H Adelson, and William T Freeman. Ground truth dataset and baseline evaluations for intrinsic image algorithms. In *2009 IEEE 12th International Conference on Computer Vision*, pages 2335–2342. IEEE, 2009. 1
- [16] Zeyu Han, Chao Gao, Jinyang Liu, Jeff Zhang, and Sai Qian Zhang. Parameter-efficient fine-tuning for large models: A comprehensive survey. *Transactions on Machine Learning Research*, 2024. 3
- [17] Amir Hertz, Yuval Alaluf, Yael Vinker, and Daniel Cohen-Or. Style aligned image generation via shared attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4775–4785, 2024. 2, 3, 4
- [18] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. 2, 3
- [19] Michael Janner, Jiajun Wu, Tejas D Kulkarni, Ilker Yildirim, and Josh Tenenbaum. Self-supervised intrinsic image decomposition. *Advances in neural information processing systems*, 30, 2017. 2
- [20] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [21] Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *Advances in neural information processing systems*, 36:36652–36663, 2023. 5, 6
- [22] Peter Kocsis, Vincent Sitzmann, and Matthias Niešner. Intrinsic image diffusion for indoor single-view material estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5198–5208, 2024. 2, 7
- [23] Peter Kocsis, Lukas Höllein, and Matthias Niešner. Intrinsic-six: High-quality pbr generation using image priors. *arXiv preprint arXiv:2504.01008*, 2025. 2, 3, 4, 5, 6
- [24] Dawid Jan Kopiczko, Tijmen Blankevoort, and Yuki M Asano. VeRA: Vector-based random matrix adaptation. In *The Twelfth International Conference on Learning Representations*, 2024. 3
- [25] Nupur Kumari, Bingliang Zhang, Han Zhang, Jianfeng Zhang, Shijie Li, Yu Zhu, and Jia-Bin Chen. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13481–13490, 2023. 3
- [26] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2023. 3, 5
- [27] Edwin H Land and John J McCann. Lightness and retinex theory. *Journal of the Optical society of America*, 61(1):1–11, 1971. 2
- [28] Louis Lettry, Kenneth Vanhoey, and Luc Van Gool. Unsupervised deep single-image intrinsic decomposition using

- illumination-varying image sequences. In *Computer graphics forum*, pages 409–419. Wiley Online Library, 2018. 7
- [29] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023. 5, 6
- [30] Zhengqi Li and Noah Snavely. Learning intrinsic image decomposition from watching the world. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9039–9048, 2018. 1, 2
- [31] Ruofan Liang, Zan Gojcic, Huan Ling, Jacob Munkberg, Jon Hasselgren, Zhi-Hao Lin, Jun Gao, Alexander Keller, Nandita Vijaykumar, Sanja Fidler, and Zian Wang. Diffusionrenderer: Neural inverse and forward rendering with video diffusion models. *arXiv preprint arXiv:2501.18590*, 2025. 2
- [32] Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. Dora: Weight-decomposed low-rank adaptation. In *Forty-first International Conference on Machine Learning*, 2024. 3
- [33] Ivan Lopes, Fabio Pizzati, and Raoul de Charette. Material palette: Extraction of materials from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 3
- [34] Jundan Luo, Zhaoyang Huang, Yijin Li, Xiaowei Zhou, Guofeng Zhang, and Hujun Bao. Niid-net: Adapting surface normal knowledge for intrinsic image decomposition in indoor scenes. *IEEE Transactions on Visualization and Computer Graphics*, 26(12):3434–3445, 2020. 7
- [35] Jundan Luo, Duygu Ceylan, Jae Shin Yoon, Nanxuan Zhao, Julien Philip, Anna Fr’uhst’uck, Wenbin Li, Christian Richardt, and Tuanfeng Wang. Intrinsicdiffusion: Joint intrinsic layers from latent diffusion models. In *ACM SIGGRAPH Conference Proceedings*, page 74, 2024. 2
- [36] Konstantin Mishchenko and Aaron Defazio. Prodigy: An expeditiously adaptive parameter-free learner. *arXiv preprint arXiv:2306.06101*, 2023. 5
- [37] Ivan V Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011. 3, 5
- [38] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10912–10922, 2021. 5
- [39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 3
- [40] Soumyadip Sengupta, Jinwei Gu, Kihwan Kim, Guilin Liu, David W Jacobs, and Jan Kautz. Neural inverse rendering of an indoor scene from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8598–8607, 2019. 2
- [41] Shanlin Sun, Yifan Wang, Hanwen Zhang, Yifeng Xiong, Qin Ren, Ruogu Fang, Xiaohui Xie, and Chenyu You. Ouroboros: Single-step diffusion models for cycle-consistent forward and inverse rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10386–10397, 2025. 2
- [42] Shimon Vainer, Mark Boss, Mathias Parger, Konstantin Kutsy, Dante De Nigris, Ciara Rowles, Nicolas Perony, and Simon Donn’e. Collaborative control for geometry-conditioned pbr image generation. *CoRR*, abs/2402.05919, 2024. 3
- [43] Zian Wang, Jonah Philion, Sanja Fidler, and Jan Kautz. Learning indoor inverse rendering with 3d spatially-varying lighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12538–12547, 2021. 2
- [44] Yuxiang Wei, Yushun Yang, Jianmin Zhang, Jiayu Hu, Qingming Hu, and Xiaoming Wang. Elite: Encoding visual concepts into textual embeddings for customized text-to-image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20218–20228, 2023. 3
- [45] Zhiyu Xie, Yuqing Zhang, Xiangjun Tang, Yiqian Wu, Dehan Chen, Gongsheng Li, and Xiaogang Jin. Styletex: Style image-guided texture generation for 3d models. *ACM Transactions on Graphics (TOG)*, 43(6):1–14, 2024. 3
- [46] Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36:15903–15935, 2023. 5, 6
- [47] Bowen Xue, Giuseppe Claudio Guarnera, Shuang Zhao, and Zahra Montazeri. Diffusion-based g-buffer generation and rendering. *arXiv preprint arXiv:2503.15147*, 2025. 3
- [48] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. 6
- [49] Chongjie Ye, Lingteng Qiu, Xiaodong Gu, Qi Zuo, Yushuang Wu, Zilong Dong, Liefeng Bo, Yuliang Xiu, and Xiaoguang Han. Stablenormal: Reducing diffusion variance for stable and sharp normal. *ACM Transactions on Graphics (ToG)*, 2024. 2
- [50] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapt: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023. 3
- [51] Zheng Zeng, Valentin Deschaintre, Iliyan Georgiev, Yannick Hold-Geoffroy, Yiwei Hu, Fujun Luan, Ling-Qi Yan, and Miloš Hašan. Rgb $\leftrightarrow$ x: Image decomposition and synthesis using material- and lighting-aware diffusion models. In *ACM SIGGRAPH Conference Proceedings*, page 75, 2024. 2, 5, 7, 8, 4
- [52] Jingyang Zhang, Jian Li, Li He, Zhaomeng Fan, Jingyu Zhang, Shengfeng Zhang, Wei Yang, and Fan Liu. Jointnet: Extending text-to-image diffusion for dense distribution modeling. *arXiv preprint arXiv:2310.06347*, 2023. 3
- [53] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5219–5229, 2023. 3, 8

- [54] Qingru Zhang, Minshuo Chen, Alexander Bukharin, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. Adaptive budget allocation for parameter-efficient fine-tuning. In *The Eleventh International Conference on Learning Representations*, 2023. [3](#)
- [55] Canyu Zhao, Mingyu Liu, Huanyi Zheng, Muzhi Zhu, Zhiyue Zhao, Hao Chen, Tong He, and Chunhua Shen. Diception: A generalist diffusion model for visual perceptual tasks. *ArXiv abs/2502.17157*, 2025. [2](#)
- [56] Qibin Zhao, Guoxu Zhou, Shengli Xie, Liqing Zhang, and Andrzej Cichocki. Tensor ring decomposition. *arXiv preprint arXiv:1606.05535*, 2016. [3](#)
- [57] Rongjia Zheng, Qing Zhang, Chengjiang Long, and Wei-Shi Zheng. Dnf-intrinsic: Deterministic noise-free diffusion for indoor inverse rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10342–10352, 2025. [2](#)
- [58] Yang Zhou, Zesong Liu, Chunping Chen, Zhiheng Chen, Feng Deng, Wei Lv, and Cheng Luo. Attention distillation: A unified approach to visual characteristics transfer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025. [3](#), [4](#)
- [59] Jingsen Zhu, Fujun Luan, Yuchi Huo, Zihao Lin, Zhihua Zhong, Dianbing Xi, Rui Wang, Hujun Bao, Jiaxiang Zheng, and Rui Tang. Learning-based inverse rendering of complex indoor scenes with differentiable monte carlo raytracing. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–8, 2022. [2](#), [7](#)
- [60] Rui Zhu, Zhengqin Li, Janarбек Matai, Fatih Porikli, and Manmohan Chandraker. Irisformer: Dense vision transformers for single-image inverse rendering in indoor scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2822–2831, 2022. [2](#)

LumiX: Structured and Coherent Text-to-Intrinsic Generation

Supplementary Material

A. Method Details

A.1. Additional Details on Query Broadcast Attention

Pseudocode We provide minimal pseudocode for the Query Broadcast Attention forward pass in Algorithm 1.

Algorithm 1 Forward pass of Query Broadcast Attention

Require: $x \in \mathbb{R}^{M \times L \times D}$ $\triangleright M$ intrinsic properties, L tokens per property, feature dimension D
Require: attn (attention block with H heads)
Require: TensorLoRA^K, TensorLoRA^V $\triangleright$ Tensor LoRA applied to K/V projections
Require: pe $\triangleright$ positional encoding

```

1:  $q, k, v \leftarrow \text{attn.qkv}(x)$   $\triangleright [M, H, L, D_{\text{head}}]$ 
2:  $q_{\text{shared}} \leftarrow q[\text{color}]$   $\triangleright [1, H, L, D_{\text{head}}]$ 
3:  $q_{\text{shared}} \leftarrow \text{expand}(q_{\text{shared}}, M)$   $\triangleright [M, H, L, D_{\text{head}}]$ 
4:  $\Delta k \leftarrow \text{TensorLoRA}^K(x)$   $\triangleright [M, L, D]$ 
5:  $\Delta v \leftarrow \text{TensorLoRA}^V(x)$   $\triangleright [M, L, D]$ 
6:  $\Delta k \leftarrow \text{reshape\_to\_heads}(\Delta k)$   $\triangleright [M, H, L, D_{\text{head}}]$ 
7:  $\Delta v \leftarrow \text{reshape\_to\_heads}(\Delta v)$   $\triangleright [M, H, L, D_{\text{head}}]$ 
8:  $k_{\text{mod}} \leftarrow k + \Delta k$ 
9:  $v_{\text{mod}} \leftarrow v + \Delta v$ 
10:  $y \leftarrow \text{attention}(q_{\text{shared}}, k_{\text{mod}}, v_{\text{mod}}, \text{pe})$ 
11: return  $y$ 

```

A.2. Additional Details on Tensor LoRA

Fused Contraction The main paper presents Tensor LoRA using three tensor contractions:

```

# N d R1 R2, M d -> N M R1 R2
Ch = einsum('ndrs, md -> nmrs', C, h)
# N M R2, N M R1 R2 -> N R1
BCh = einsum('nms, nmrs -> nr', B, Ch)
# N d R1, N R1 -> N d
ABCh = einsum('ndr, nr -> nd', A, BCh)

```

These operations follow directly from the element-wise form in Eq. 14, where the rank indices (α_1, α_2) are contracted across the three components A , B , and C .

Single-Step Equivalent Because all intermediate indices are internal and do not appear in the final output, the entire computation can be expressed as a single contraction:

$$\Delta[i, o] = \sum_{j, l, \alpha_1, \alpha_2} C[i, l, \alpha_1, \alpha_2] h[j, l] B[i, j, \alpha_2] A[i, o, \alpha_1]. \quad (15)$$

This form can be implemented using a single einsum call:

```
ABCh = einsum('ndrs, md, nms, ndr -> nd',
               C, h, B, A)
```

which avoids constructing the intermediate tensors in the three-step version and improves memory efficiency.

B. Implementation Details

Additional Training Details All intrinsic properties are processed together by concatenating them along the batch dimension. All weights of FLUX.1-dev remain frozen during training. The only trainable parameters are the Tensor LoRA updates inserted into the key and value projections of each attention block, while all original query, key, and value weights are kept fixed. The flow-matching loss is computed independently for each intrinsic property and then averaged. For data preparation, we follow the official Hypersim pipeline and apply their tone mapping and gamma correction to convert HDR color, irradiance, and albedo images into LDR inputs.

Evaluation Metrics Evaluating intrinsic images such as albedo is inherently challenging due to their ill-posed nature. To assess perceptual quality, we use two state-of-the-art human preference models, ImageReward and PickScore, which we find to correlate well with visual realism in our setting. To measure cross-property consistency, we rely on Qwen3-VL, a strong vision-language model capable of fine-grained visual reasoning. For each pair of intrinsic maps (e.g., color and albedo), the model is asked to evaluate their alignment. It assigns five global scores in $[0, 1]$ for scene identity, structural layout, object arrangement, material-region consistency, and reflectance correspondence, summed to a global score in $[0, 5]$. If the global score is at least 2.0, it further assigns five local scores in $[0, 1]$ for edge precision, geometric detail, surface boundaries, shadow/illumination separation, and small-object/material consistency, yielding a local score in $[0, 5]$. The final alignment score is the total in the interval $[0, 10]$.

C. Additional Results

We include additional examples from our quantitative evaluation in Figs. 7 to 12.

<https://github.com/black-forest-labs/flux>
<https://github.com/apple/ml-hypersim>
<https://github.com/zai-org/ImageReward>
<https://github.com/yuvalkirstain/PickScore>
<https://huggingface.co/Qwen/Qwen3-VL-4B-Instruct>








Input	“an astronaut inside a bedroom with soft fabric curtains and neon reflections” “an astronaut sitting on a couch in a cozy living room with warm lighting” “a cat sleeping on a stack of ancient scrolls inside a sunlit library” “a glass greenhouse where trees grow through futuristic metal furniture”				
IntrinsiX					Rendering
					Albedo
					Normal
LumiX (Ours)					Color
					Albedo
					Normal

Figure 7. **More Text-to-Intrinsic Generation Results.** Compared with IntrinsiX [23], our method yields higher-quality images and consistent intrinsic maps.




Figure 8. **More Intrinsic Decomposition Results.** Our method produces cleaner albedo and more accurate, consistent normal maps compared with prior baselines.




Figure 9. **More Intrinsic Decomposition Results on Indoor Scenes.** We show the albedo, irradiance, and normal predictions for RGB $\leftrightarrow$ X [51] and our method. Both methods produce high-quality albedo and irradiance, while ours yields more accurate and consistent normal maps.



Figure 10. **More Intrinsic Decomposition Results on Indoor Scenes.** We visualize albedo, irradiance, and normal maps from RGB $\leftrightarrow$ X [51] and our method. Both achieve strong albedo and irradiance performance, but ours attains higher overall quality.




Figure 11. **More Intrinsic Decomposition Results on Out-of-Domain Data.** Our method produces high-quality and consistent intrinsic maps even outside indoor scenes, with cleaner albedo (less baked-in lighting) and more accurate normals.




Figure 12. **More Intrinsic Decomposition Results on Out-of-Domain Data.** We visualize albedo, irradiance, and normal maps on clearly out-of-domain images. Despite being trained on only 3K examples (far fewer than the 190K images used by RGB $\leftrightarrow$ X [51]), our method still produces high-quality and consistent intrinsic maps, with cleaner albedo and more accurate normal maps.