

ReasonX: MLLM-Guided Intrinsic Image Decomposition

Alara Dirik^{1*}, Tuanfeng Wang², Duygu Ceylan², Stefanos Zafeiriou¹, Anna Frühstück²
¹Imperial College London ²Adobe Research

*Work done during internship at Adobe.




Figure 1. We propose ReasonX, a novel framework for MLLM-guided improvement of intrinsic decomposition models via relative intrinsic judgments on RGB input images.

Abstract

Intrinsic image decomposition aims to separate images into physical components such as albedo, depth, normals, and illumination. While recent diffusion- and transformer-based models benefit from paired supervision from synthetic datasets, their generalization to diverse, real-world scenarios remains challenging. We propose ReasonX, a novel framework that leverages a multimodal large language model (MLLM) as a perceptual judge providing relative intrinsic comparisons, and uses these comparisons as GRPO rewards for fine-tuning intrinsic decomposition models on unlabeled, in-the-wild images. Unlike RL methods for generative models, our framework aligns conditional intrinsic predictors by rewarding agreement between the judge’s relational assessments and analytically derived relations from the model’s outputs. ReasonX is model-agnostic and can be applied to different intrinsic predictors. Across multiple base architectures and modalities, ReasonX yields significant improvements, including 9-25% WHDR reduction on IIW albedo and up to 46% depth accuracy gains on ETH3D, highlighting the promise of MLLM-guided comparative supervision to bridge low- and high-level vision reasoning.

1. Introduction

Recovering intrinsic scene properties such as albedo, surface normals, depth, and illumination from a single RGB image is a fundamental inverse problem in computer vision. Remarkable progress has been made on this task with the adoption of diffusion models and vision transformers as

backbone architectures [6, 12, 14, 18, 19]. Recent methods have demonstrated impressive capabilities in intrinsic scene understanding, achieving high-quality decompositions on complex indoor scenes, and enabling a variety of downstream applications such as relighting and material editing. These advances have been driven not only by sophisticated model architectures, but also by the development of large-scale training frameworks [49] that leverage multiple heterogeneous datasets [32, 53].

However, a fundamental bottleneck remains: the reliance on paired synthetic datasets with ground-truth intrinsic labels. State-of-the-art methods depend heavily on synthetic data rendered from physically-based simulation engines [21, 22, 32, 53], which, while photorealistic, are inherently limited in scope and diversity. These datasets typically focus on specific domains (e.g., indoor scenes) and fail to capture the full complexity of in-the-wild imagery. Creating comprehensive ground truth for intrinsic properties across diverse real-world scenarios remains prohibitively expensive and technically challenging. Consequently, models trained primarily on synthetic data often struggle to generalize to images outside their training distribution, limiting their practical applicability.

Meanwhile, multimodal large language models (MLLMs) have emerged as powerful foundational models with extensive visual and semantic understanding capabilities [2, 54]. Recent work has demonstrated that MLLMs achieve impressive performance on diverse visual reasoning tasks, even approaching or surpassing specialized models in certain domains [7, 23]. This raises a compelling

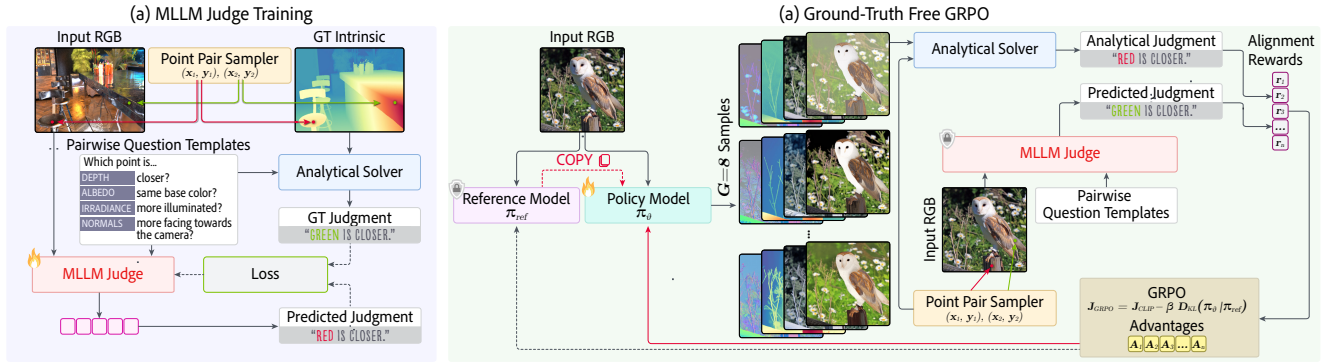


Figure 2. **Overview of our ReasonX framework.** (a) We fine-tune an MLLM to judge relative intrinsic properties from RGB images using sampled point pairs. (b) The frozen judge then provides rewards within a GRPO loop to refine an intrinsic decomposition model π : for each RGB image, we generate a group of $G = 8$ samples, query the judge across point pairs and modalities, and compute group-relative rewards to update π without ground-truth intrinsics.

question: can we leverage the priors from pretrained MLLMs to overcome the generalization bottleneck in intrinsic image decomposition? We take inspiration from recent findings that MLLMs excel at *relative* spatial reasoning [41, 44, 50]. This observation aligns with principles of human perception as people are naturally better at making comparative judgments (e.g., ‘which point is closer?’, ‘do these regions share the same material?’) than absolute measurements [4]. Building on this insight, we propose to fine-tune pre-trained MLLMs as *judges* that evaluate the quality of intrinsic decompositions through pairwise comparisons.

Our key insight is to treat these relative intrinsic judgments as a supervision signal. Rather than learning intrinsics from absolute labels or heuristic losses, we train an MLLM to reason about relative intrinsic comparisons by sampling point pairs on images and asking questions about their intrinsic properties (e.g., relative depth, material similarity, lighting differences). We then employ this judge as a reward function within a Group Relative Policy Optimization (GRPO) framework [35]. Because intrinsic decomposition is tightly conditioned on the input RGB image and therefore largely deterministic, we introduce exploration by rewarding consistency between the judge’s comparative assessments and analytical relations computed from the model’s predicted intrinsics. During fine-tuning of the base intrinsic decomposition model, we sample multiple point pairs per image and compute alignment rewards by comparing the judge’s assessment with predicted intrinsics. This approach requires no ground-truth intrinsic images during the base model fine-tuning, we only need RGB images and an MLLM judge trained to reason about intrinsic properties through relative comparisons.

We evaluate the effectiveness of our ReasonX¹ framework by coupling our MLLM-based judge with different base intrinsic methods, including flow matching models

(PRISM [12]) and diffusion-based models (Marigold IID Lighting v1.1 [18]). The supervision from the MLLM judge enables training of the base model on diverse in-the-wild images, achieving 9–25% WHDR reduction on the IIW albedo benchmark and up to 46% depth accuracy improvements on ETH3D. We also establish an out-of-distribution dataset and show that our method significantly improves the base model’s generalization to challenging real-world scenarios. Our main contributions include:

- An MLLM-based judge that reasons about intrinsic image decompositions through pairwise comparisons,
- A GRPO-based approach that uses the MLLM judge as a reward signal to fine-tune base intrinsic models on diverse and unlabeled images,
- Improved generalization to in-the-wild scenes across different intrinsic modalities (albedo, depth, normals, and irradiance) and multiple base models.

2. Related Work

Intrinsic Image Decomposition Models. Early intrinsic decomposition methods differ mainly in how they model the interaction between reflectance, geometry, and illumination, ranging from explicit physical priors [20] to data-driven reflectance–illumination separation [3, 29]. Among recent advances, Ordinal Shading [5] presents a dual-stage convolutional framework that enforces scale- and shift-invariant ordering within shading predictions. By focusing on relative rather than absolute intensity, it preserves global coherence and fine detail, demonstrating the benefits of ordinal constraints for intrinsic decomposition. Our approach extends this principle beyond shading, using a multimodal large language model (MLLMs) as a perceptual judge to evaluate relative relationships across arbitrary intrinsic modalities—such as depth, reflectance, or illumination. Through this generalized ordinal reasoning and GRPO-based optimization, our method enables cross-modal fine-tuning of existing models without paired supervision.

Complementary to these advances, vision transformers

¹Code and model for the MLLM judge will be publicly released at <https://github.com/alaradirik/reasonx>.

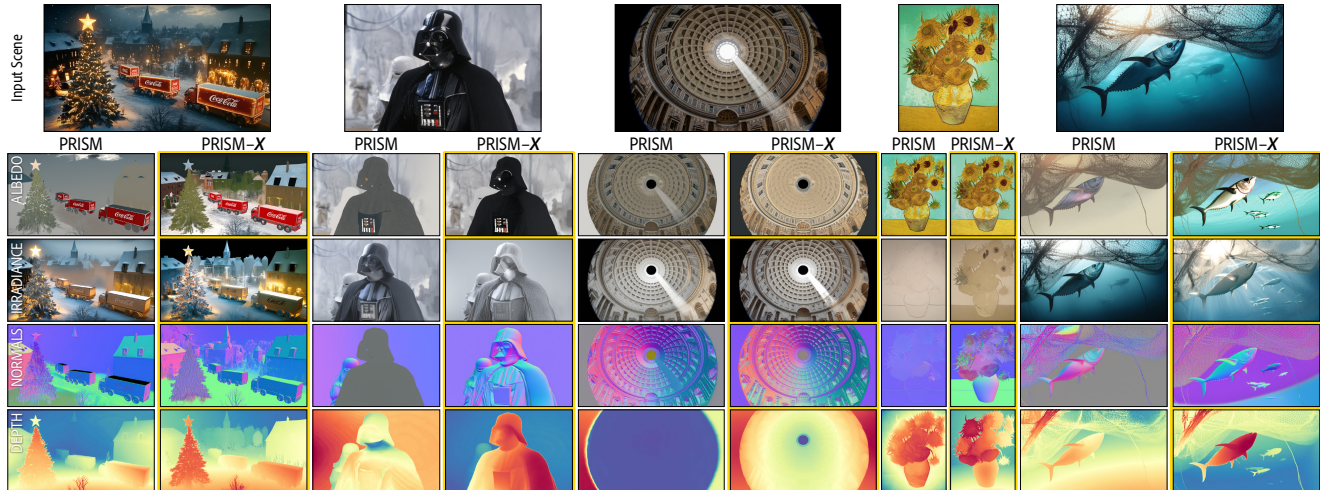


Figure 3. Intrinsic decomposition samples on challenging out-of-distribution images. Our PRISM-X significantly improves its base model PRISM across all intrinsic channels with respect to decomposition quality and in-the-wild generalization performance.

and diffusion-based architectures have enabled multi-modal intrinsic prediction with improved consistency [6, 12, 14, 18]. Frameworks such as PRISM [12] and Marigold [18] jointly predict intrinsic properties, demonstrating strong performance on synthetic datasets, while our method provides an orthogonal route toward generalization through MLLM-guided, order-aware refinement. For a broader overview of intrinsic decomposition methods, we refer readers to recent surveys [15].

Synthetic Data and Generalization. Due to the scarcity of real-world intrinsic labels, many methods rely on physically-based rendering datasets such as HyperSim [32], InteriorVerse [53], OpenRooms [22], and MatrixCity [21]. While photorealistic, these datasets are typically biased toward indoor environments, limiting generalization to natural scenes. To address this, prior work has explored cross-domain learning, semi-supervised approaches, and auxiliary modalities such as RGBX supervision [21, 49].

Multimodal Large Language Models for Visual Reasoning. Multimodal large language models (MLLMs) [2, 54] combine vision encoders with large-scale language backbones, enabling reasoning about spatial and semantic relationships in images. Initially applied to high-level tasks such as captioning and visual question answering, MLLMs have recently been adapted for geometric reasoning in 3D scenes [8, 41, 43, 44, 50]. These models excel at relative and comparative judgments, although absolute metric estimation remains challenging.

MLLM-Based Feedback and Reward-Guided Learning. Recent works investigate using MLLMs as verifiers or reward models for feedback-guided generation [26, 33, 38, 39, 47, 52], extending reinforcement learning from human feedback [11] to multimodal tasks. Techniques like Group Relative Policy Optimization (GRPO) [35] have been used

to align generative models via learned critics [25].

Relative Reasoning and Comparative Learning. Relative supervision has been recognized as a robust alternative to absolute labels in vision tasks. Works on ordinal depth estimation [9, 55] and pairwise reflectance comparisons [4] illustrate how comparative cues provide stronger structural constraints and improved generalization. Recent MLLM research demonstrates strong spatial and relational judgment capabilities [8, 41, 50], suggesting a natural connection between comparative supervision and multimodal reasoning.

3. Method

Our goal is to fine-tune a base intrinsic image decomposition model with real, unlabeled images. To achieve this, we train a multimodal large language model (MLLM) to act as a perceptual judge that makes *relative intrinsic comparisons*, and then fine-tune the base model to satisfy these comparisons through Group Relative Policy Optimization (GRPO). As illustrated in Figure 2, our ReasonX framework consists of two components: (i) a pretrained intrinsic model that predicts multiple physical modalities from an RGB image, and (ii) an MLLM judge that evaluates these predictions through point-pair comparisons.

Base Models. ReasonX is flexible and can be applied to various intrinsic decomposition models. We demonstrate this using two complementary architectures: PRISM [12], a diffusion transformer trained under the rectified flow formulation for text- and image-conditioned RGBX generation, and Marigold [18] IID Lighting v1.1, a diffusion-based model for joint albedo and irradiance estimation.

3.1. MLLM Judge: Relative Intrinsic Reasoning

We adopt a multimodal language model (MLLM) with image understanding capability to perform pairwise intrinsic



Figure 4. Our MLLM judge is trained on synthetic data to make relative intrinsic judgments from point-pair annotated RGB images and modality-specific questions. Ground-truth answers are derived from the corresponding intrinsic maps.

judgments on RGB images. More specifically, we fine-tune the open source InternVL2.5-4B [10] MLLM on synthetic RGB images from the base model’s training set, where ground-truth intrinsics are available. For each training image, we sample point pairs $(x_1, y_1), (x_2, y_2)$, overlay colored visual markers, and formulate modality-specific relative questions (see Figure 4) such as "which point is closer?" or "do these points share the same material?". This focuses the learning task on *comparative* reasoning, which we find transfers far more reliably to real images than absolute predictions (see supplementary for analysis). Ground-truth comparisons are derived analytically from corresponding synthetic intrinsic maps as follows:

- **Depth:** direct comparison of scalar depth values; pairs with minimal difference are excluded to avoid ambiguity.
- **Normals:** assuming the camera is looking at the $+z$ axis; comparison of front-facingness via the normal’s z -component.
- **Irradiance:** comparison of luminance-albedo ratios under Lambertian assumptions; the point with the higher ratio is labeled as more illuminated.
- **Albedo:** thresholded perceptual color differences.

Once trained, the MLLM judge remains fixed. For a predicted intrinsic map \mathbf{I}_m , we sample N point pairs on the RGB input, overlay visual markers, and query the frozen judge with a modality-specific question to get a relative judgment. We then derive analytic predictions from \mathbf{I}_m , and compute their agreement as our alignment reward:

$$r(\mathbf{I}_{\text{RGB}}, \mathbf{I}_m) = \frac{1}{N} \sum_{i=1}^N \left[\text{MLLM}(\mathbf{I}_{\text{RGB}}^{(i)}, q_m) = g_m(\mathbf{I}_m, \mathbf{p}_i) \right], \quad (1)$$

where $\mathbf{I}_{\text{RGB}}^{(i)}$ is the RGB image with markers at point pair \mathbf{p}_i , q_m is the modality-specific query, and g_m denotes the deterministic relation computed from the predicted intrinsic map. This produces a reward signal grounded in both the model’s physical structure and the judge’s comparative perception. We next describe how this reward is used to fine-tune the base intrinsic model.

3.2. Intrinsic-GRPO: Learning from Relative Rewards

Fine-tuning an intrinsic decomposition model on real images is challenging as the task is *RGB-conditioned and largely deterministic*, leaving little room for the exploration required by policy-gradient methods. We therefore formulate the fine-tuning of the base intrinsic decomposition model as a reinforcement learning problem. Specifically, we perform such a fine-tuning on a set of real images where no ground truth annotations exist. We obtain intrinsic predictions from the current base model, compute rewards for these predictions using the MLLM-based judge, and back-propagate this reward as a supervision signal.

Injecting Exploration. Both PRISM and Marigold generate intrinsic predictions through a deterministic, image-conditioned denoising trajectory. In both cases, the model applies a learned update $f_\theta(\mathbf{x}_t, t, \mathbf{c})$ to the latent state \mathbf{x}_t at each timestep t given the condition \mathbf{c} . To enable exploration, we follow Flow-GRPO [25] and introduce a small stochastic term during sampling, converting the deterministic trajectory into a lightly stochastic one via the Euler–Maruyama update:

$$\mathbf{x}_{t+\Delta t} = \mathbf{x}_t + f_\theta(\mathbf{x}_t, t, \mathbf{c}) \Delta t + \sigma_t \sqrt{\Delta t} \epsilon, \quad (2)$$

where $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and σ_t controls the noise level. We note that independent noise vectors yield multiple plausible intrinsic predictions for the same RGB image, which is crucial for comparative, group-based optimization in an otherwise deterministic setting.

Group-based reward optimization. During policy learning on real images, we keep the judge MLLM and the VAE of the base model frozen. For each real RGB image, we randomly choose a target modality m , generate a group of G intrinsic predictions using independent noise, compute a reward for each sample using the MLLM judge, and update the model using GRPO. Given group rewards $\{r_i\}_{i=1}^G$, we compute normalized advantages as

$$\hat{A}_i = \frac{r_i - \mu_G}{\sigma_G},$$

where μ_G and σ_G are the group mean and standard deviation. GRPO optimizes a clipped PPO-style objective [35] with a KL regularizer that constrains the updated policy to remain close to a frozen reference copy of the pretrained model:

$$\mathcal{J}(\theta) = \mathbb{E}_{\pi_\theta} \left[\min \left(\rho_t \hat{A}, \text{clip}(\rho_t, 1 - \epsilon, 1 + \epsilon) \hat{A} \right) - \beta D_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right], \quad (3)$$

where $\rho_t = \pi_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)/\pi_{\theta_{\text{old}}}(\mathbf{x}_{t-1}|\mathbf{x}_t)$ is the importance ratio. We note that the KL regularization helps prevent reward hacking [25], such as collapsing to near-constant intrinsics, and ensures that improvements reflect genuine alignment with comparative judgments. Under the stochastic update above, the transition distribution is Gaussian, allowing the KL divergence to be computed in closed form by comparing the model and reference velocity fields.

Our formulation adapts GRPO to the intrinsic decomposition setting by combining RGB-conditioned exploration with structured, relative rewards from the MLLM judge. Using our ReasonX framework, the base intrinsic decomposition model is refined to satisfy physics-guided relational constraints directly on real images, without requiring any ground-truth intrinsic maps. Unlike prior uses of GRPO in image generative settings [25], our method targets deterministic, RGB-conditioned intrinsic prediction, where the reward is derived from modality-specific relational comparisons rather than preference scores.

3.3. Implementation Details

We fine-tune each base model on a dataset of 10,000 real RGB images from the COCO training set [24]. At each iteration, we sample $N = 40$ point pairs and use $T = 15$ denoising steps (inference uses $T = 50$). PRISM is conditioned with an empty text prompt. We optimize both base models using AdamW with a learning rate of 10^{-5} , cosine annealing schedule, and gradient clipping at norm 1.0. The group size is set to $G = 8$, the SDE noise level to $a = 0.7$. We train our models on 6 H100 GPUs for 3 epochs.

4. Evaluation

ReasonX is a generic framework that can be used in conjunction with different base intrinsic decomposition models. Hence, we center our evaluations to showcase the relative performance gains achieved through our ReasonX, both qualitatively and quantitatively, across multiple modalities and datasets. In particular, we apply ReasonX on top of two representative base models - PRISM and Marigold IID Lighting v1.1 [18] (a Marigold variant that jointly estimates albedo, irradiance and residual images) - resulting in enhanced variants denoted as PRISM-X and Marigold-X, respectively. To further contextualize these improvements,




Figure 5. Depth estimation comparisons of our PRISM-X with its base model PRISM and SOTA baseline method Marigold Depth v1.0 on samples from the NYUv2, DIODE and ETH3D datasets. PRISM-X performs significantly better on challenging images.

we include an additional MLLM baseline where we fine-tune OmniGen2 [40], a versatile MLLM with image generation and editing capabilities to generate intrinsic modalities (e.g. depth, albedo) given an RGB image and a text prompt cue denoting the task (e.g. “convert this image to a depth map.”). We use the same synthetic dataset that we train our MLLM judge for this fine-tuning. While this baseline provides competitive performance, it often is inferior to ReasonX variants. In addition, we report results from task-specific state-of-the-art methods to position our models within the broader literature, noting that many of these SOTA approaches are trained on substantially larger or domain-specific datasets. While this may limit direct comparability, we observe that our method is often able to match the performance of specialized models. We refer to the supplementary material for a more comprehensive evaluation including additional baseline methods.

Evaluation Metrics and Datasets. We use various modality specific datasets and metrics for a thorough comparison. For albedo decomposition, we perform our comparisons on the real IIW [4] and MAW [42] datasets. We follow the experiment protocol in Careaga and Aksoy [5, 6] and report the mean weighted human disagreement rate (WHDR), intensity and chromaticity metrics.

For depth estimation, we perform zero-shot evaluation



Figure 6. Normals estimation comparisons of our PRISM-X with its base model PRISM and SOTA baseline method Marigold Normals v1.1 on NYUv2 and DIODE indoor and outdoor samples.

on the NYU-v2 [36] (indoor) and ETH3D [34] (mixed indoor and outdoor) datasets and report Absolute Relative Error (AbsRel) and Threshold Accuracy (δ_1) metrics following previous work. For surface normals estimation, we report the mean angular error and the percentage of pixels with angular error lower than 11.5° on the NYU-v2 and DIODE [37] datasets. Unless specified otherwise, we perform all experiments on the designated test set of each dataset.

4.1. Intrinsic Decomposition Performance

Albedo Estimation. We start by evaluating our method on two real datasets with respect to the albedo estimation quality and report the numbers in Tables 1 and 2, respectively. IIW [4] provides pairwise human annotations of albedo brightness between sparsely sampled pixels. The MAW dataset [42] consists of ~ 850 indoor images and measures albedo within specific masked regions of the image. For MAW, we report the intensity and chromaticity of the albedo estimates.

As shown in Tables 1 and 2, both ReasonX variants (denoted Marigold-X and PRISM-X) significantly outperform their respective base models, with PRISM-X yielding SOTA zero-shot results on the IIW dataset that are comparable to CRefNet [27], a non-competing method trained on the IIW dataset. Similarly, ReasonX variants achieve

Table 1. WHDR error of albedo estimates on the IIW dataset [4] for our base models, their ReasonX variants and select SOTA baselines. CRefNet [27], a non-competing method trained on IIW, is denoted with gray. Our ReasonX variants significantly outperform their respective base models (9–25% improvement in WHDR 10%) and achieve SOTA results compared to highest performing zero-shot methods.

Method	WHDR 10% ↓	WHDR 20% ↓
CRefNet [27]	12.8	10.8
Zhu et al. [53]	34.7	24.1
Careaga and Aksoy [5]	24.8	19.2
Kocsis et al. [19]	26.1	20.7
RGB \leftrightarrow X [49]	23.6	21.1
Careaga and Aksoy [6]	16.8	15.6
OmniGen2 [40]	17.8	14.3
Marigold (base variant) [18]	16.7	14.8
Marigold-X	15.2	14.0
<i>Improvement</i>	+9.0%	+5.4%
PRISM [12]	17.2	15.9
PRISM-X	12.9	11.9
<i>Improvement</i>	+25.0%	+25.2%

Table 2. Intensity and chromaticity of albedo estimates on the MAW dataset [42] for our method and baselines. ReasonX variants achieve 16.3–39.4% improvement in albedo estimation intensity, denoting highly improved robustness to over-exposure and generalization capability.

Method	Intensity ($\times 100$) ↓	Chromaticity ↓
Zhu et al. [53]	1.44	4.94
Careaga and Aksoy [5]	0.57	6.56
Kocsis et al. [19]	1.13	5.35
RGB \leftrightarrow X [49]	0.82	3.96
Careaga and Aksoy [6]	0.54	3.37
OmniGen2 [40]	0.41	4.40
Marigold (base variant) [18]	0.49	4.10
Marigold-X	0.41	4.01
<i>Improvement</i>	+16.3%	+2.2%
PRISM [12]	0.71	3.92
PRISM-X	0.43	3.78
<i>Improvement</i>	+39.4%	+3.6%

16–39% improvement over their base model on MAW intensity and achieve SOTA and comparable results, highlighting our method’s effectiveness in improving decomposition robustness.

Depth Estimation. We separately evaluate the depth estimation quality of our method against a number of state-of-the-art monocular relative and metric depth estimation methods including Marigold Depth v1.0 [17] and Depth Anything V2 [45, 46]. We perform the comparisons on the NYUv2 [36] and ETH3D [34] datasets - two datasets neither our base model or ReasonX variant have seen during training. To evaluate the depth estimation quality, we report Absolute Relative Error (AbsRel) and Threshold Accuracy (δ_1) metrics. As PRISM is trained to estimate disparity, we first convert PRISM and our PRISM-X estimated

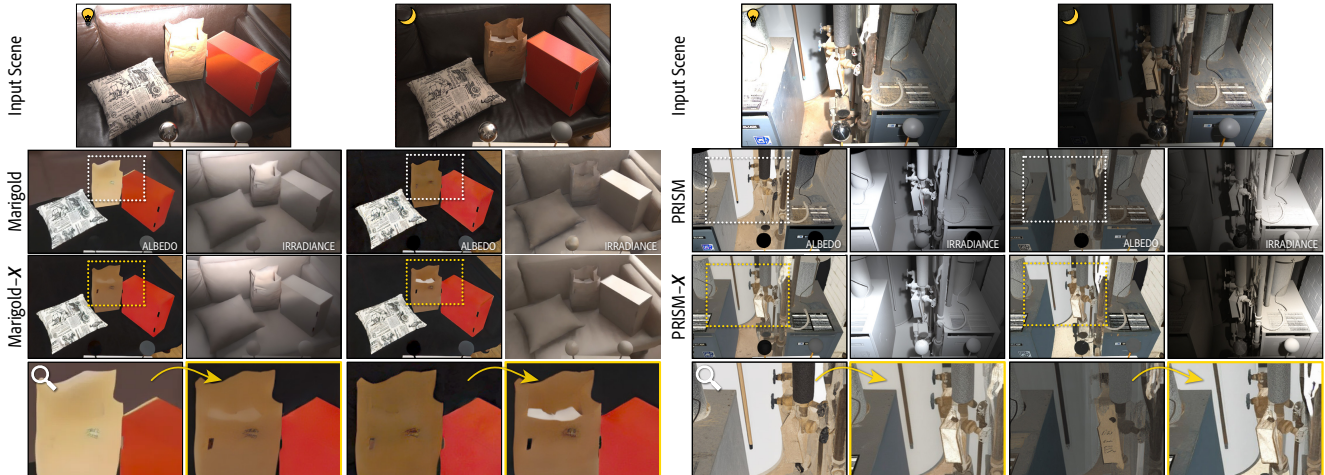


Figure 7. Effect of ReasonX on albedo and irradiance decomposition under varying lighting conditions. Note how both models are able to significantly improve their stability under very difficult lighting settings, such as these over/underexposed images.

disparity maps to relative depth maps. As shown in Table 3, our PRISM-X model achieves SOTA or comparable results, as well as large improvements over its base PRISM model despite having been trained on real RGB images with no ground truth depth maps. We note that the improvement from the base model is more prominent on the ETH3D dataset, which is a predominantly outdoor dataset, whereas NYUv2 consists of indoor images.

Table 3. Zero-shot relative depth estimation on NYUv2 and ETH3D datasets. Methods that require more than 2M samples for fine-tuning are denoted with gray. Our method achieves depth estimation results comparable to current SOTA methods despite having been trained on a fraction of the data, with our PRISM-X variant achieving 45.8% AbsRel improvement over its base model on the mixed indoor / outdoor ETH3D dataset.

Method	NYUv2		ETH3D	
	AbsRel ↓	δ_1 ↑	AbsRel ↓	δ_1 ↑
Omnidata [13]	0.074	0.945	0.166	0.778
Depth Anything V2 [46]	0.045	0.979	0.131	0.865
MiDaS [30]	0.111	0.885	0.184	0.752
DPT [31]	0.098	0.903	0.078	0.946
HDN [51]	0.069	0.948	0.121	0.833
OmniGen2 [40]	0.111	0.886	0.123	0.897
Marigold Depth v1.0 [17]	0.055	0.964	0.065	0.960
PRISM [12]	0.061	0.922	0.142	0.836
PRISM-X	0.053	0.958	0.077	0.950
<i>Improvement</i>	+13.1%	+3.9%	+45.8%	+13.6%

Surface Normals Estimation. We evaluate surface normal estimation performance of our PRISM-X model on the NYUv2 and DIODE [37] datasets and present quantitative comparisons with state-of-the-art methods in Table 4. Our PRISM-X achieves the best mean angular error of 15.7° on NYUv2 and 14.5° on DIODE, demonstrating consistent improvements over the base PRISM model (16.1° and 14.6° respectively). While the improvements are

modest compared to gains in other modalities, we argue that they demonstrate that our MLLM-guided approach can effectively refine geometric predictions even when the base model already performs competitively. Notably, our method outperforms recent specialized normal estimators including DSINE [1], GeoWizard [14], and StableNormal [48], despite being trained without ground-truth normal supervision. We present qualitative depth and surface normals estimation comparisons with Marigold Depth v1.0, Marigold Normals v1.1 and the base PRISM model in Figures 5 and 6, where PRISM-X produces more accurate surface orientations, particularly in regions with complex geometry.

Table 4. Zero-shot quantitative comparison of surface normals estimators on NYUv2 and DIODE datasets. Mean metric is in absolute angles, 11.25° metric in percentage. Our PRISM-X variant achieves consistently reports lower mean angular error compared to its base model PRISM, indicating increased robustness to edge cases.

Method	NYUv2		DIODE	
	Mean ↓	11.25° ↑	Mean ↓	11.25° ↑
DSINE [1]	16.4	59.6	19.9	41.8
GeoWizard [14]	19.0	50.0	24.7	30.1
StableNormal [48]	17.8	54.2	19.3	53.8
Lotus-G [16]	16.9	59.1	21.2	39.7
OmniGen2 [40]	17.7	54.6	20.5	40.5
Marigold-Normals v1.1 [18]	16.1	60.5	18.8	45.5
PRISM [12]	16.1	56.8	14.6	42.6
PRISM-X	15.7	57.0	14.5	41.3
<i>Improvement</i>	+2.5%	+0.4%	+0.7%	-3.1%

Effect on Cross Modal Alignment. A common issue when estimating each intrinsic channel separately is the alignment across different modalities, e.g., the white balance ambiguity between albedo and irradiance, alignment between depth and surface normals. We evaluate the effect of ReasonX on alignment across modalities. In partic-

Table 5. Given depth predictions on the ETH3D and COCO test sets, we compute the depth gradients and compare to the estimated normals to evaluate the alignment between the two channels.

Dataset	Method	RMSE ↓	PSNR ↑	SSIM ↑	LPIPS ↓
ETH3D	PRISM [12]	0.146	17.3	0.582	0.281
	PRISM-X	0.099 (+32%)	18.0 (+4.0%)	0.640 (+10.0%)	0.239 (+14.9%)
COCO	PRISM [12]	0.202	16.5	0.495	0.330
	PRISM-X	0.137 (+32.2%)	17.5 (+6.1%)	0.579 (+17.0%)	0.280 (+15.0%)

ular, we evaluate the alignment between depth and surface normal predictions by computing normals from estimated depth gradients and comparing to the estimated normals. We evaluate base model PRISM and our ReasonX variant, PRISM-X, on the ETH3D [34] test set and 2000 samples from the COCO [24] test set, and report the reconstruction metrics in Table 5. We exclude our Marigold-X variant from this experiment as both this variant and its base model (Marigold IID Lighting v.1.1) are trained to estimate albedo and irradiance only. As shown in Table 5, our PRISM-X achieves significant improvement over its base PRISM on both datasets, highlighting the effectiveness of our method in improving geometric consistency.

Qualitative Results. We showcase ReasonX qualitatively on a variety of challenging indoor and outdoor images in Figure 3. To evaluate the robustness under varying illumination, we test on the MIT Multi-Illumination Dataset (MIIW) [28], which contains multiple images of each scene captured under different lighting conditions. Figure 7 shows decomposition results for the same scenes under two different illuminations using both base models (Marigold, PRISM) and their ReasonX variants (Marigold-X, PRISM-X). The ReasonX variants demonstrate significantly improved consistency across challenging lighting changes while base models produce noticeably different albedo estimates for the same scene under different illumination. We note that this improved disentanglement is particularly evident in regions with strong directional lighting or shadows, where PRISM-X and Marigold-X better preserve true material properties in the albedo channel, hence improving the albedo-irradiance alignment.

4.2. Judge Accuracy and Validation

To assess the reliability of the intrinsic judge before applying it for policy fine-tuning, we evaluate its performance on the held-out test sets of InteriorVerse and HyperSim. Judge evaluation is inherently challenging, particularly for albedo and irradiance, where spatially extended markers and within-material color variation introduce visual ambiguity. Given an RGB image and its ground-truth intrinsic maps, we sample 20 point pairs and generate point-pair-annotated RGB images. For each {RGB image, intrinsic image, point pair} triplet, the judge answers modality-specific questions, and we analytically derive the corresponding ground-truth judgment from the intrinsic image to compute the judge’s accuracy and F1 score (Table 6).

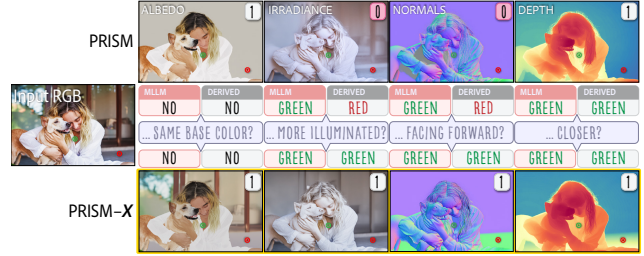


Figure 8. Alignment of MLLM judgment and estimated intrinsics pre- and post-ReasonX training of PRISM.

Because our visual markers cover small regions (not single pixels), the model must reason over textured neighborhoods rather than isolated points, making the notion of "same material color" or "same illumination" inherently ambiguous. Nevertheless, as shown in qualitative examples (Fig. 8), the judge generally provides semantically correct feedback that improves global consistency in the fine-tuned model. Further implementation details and ablations on relative vs. absolute judgments are provided in the supplementary.

Table 6. Our judge VLM performance on the held-out test set.

Metric	Depth	Normal	Albedo	Irradiance
Accuracy	0.962	0.935	0.894	0.876
Macro F1 score	0.962	0.933	0.889	0.878

5. Conclusion

We introduced ReasonX, a framework that uses relative intrinsic comparisons from an MLLM judge as rewards for GRPO-based fine-tuning of intrinsic decomposition models on unlabeled, in-the-wild images. By replacing absolute supervision with structured comparative feedback, ReasonX enables effective fine-tuning of intrinsic decomposition models without requiring ground-truth intrinsic maps. Unlike prior GRPO uses with internal critics, our method leverages an external, MLLM-based reward model capable of cross-modal comparative reasoning. This enables supervision across intrinsic modalities without paired data. Our experiments show that ReasonX consistently improves multiple base architectures and modalities, yielding substantial gains on real-world benchmarks—for example, a 25% WHDR reduction on IIW and up to a 46% depth accuracy improvement on ETH3D - despite the absence of annotated data during fine-tuning.

Limitations include the reliance on point-pair sampling and modality-wise rewards. Hence, future work could explore joint multi-modal or reconstruction-based signals and extensions to broader inverse rendering tasks. Overall, ReasonX demonstrates that MLLM-guided comparative supervision offers an effective pathway for bridging physical intrinsic reasoning with high-level perceptual understanding, providing a general paradigm for learning intrinsic decomposition models directly from real images.

References

- [1] Gwangbin Bae and Andrew J. Davison. Rethinking inductive biases for surface normal estimation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9535–9545, 2024. 7
- [2] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-VL: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. 1, 3
- [3] Anil S. Baslamisli, Thomas T. Groenestegge, Partha Das, Hoang-An Le, Sezer Karaoglu, and Theo Gevers. Joint learning of intrinsic images and semantic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [4] Sean Bell, Kavita Bala, and Noah Snavely. Intrinsic images in the wild. *ACM SIGGRAPH Conference Proceedings*, 33(4), 2014. 2, 3, 5, 6
- [5] Chris Careaga and Yağız Aksoy. Intrinsic image decomposition via ordinal shading. *ACM Transactions on Graphics (ToG)*, 2023. 2, 5, 6
- [6] Chris Careaga and Yağız Aksoy. Colorful diffuse intrinsic image decomposition in the wild. *ACM Transactions on Graphics (ToG)*, 43(6), 2024. 1, 3, 5, 6
- [7] Gongwei Chen, Leyang Shen, Rui Shao, Xiang Deng, and Liqiang Nie. Lion: Empowering multimodal large language model with dual-level visual knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26540–26550, 2024. 1
- [8] Guizhen Chen, Weiwen Xu, Hao Zhang, Hou Pong Chan, Deli Zhao, Anh Tuan Luu, and Yu Rong. GeoPQA: Bridging the visual perception gap in MLLMs for geometric reasoning. *arXiv preprint arXiv:2509.17437*, 2025. 3
- [9] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. In *Advances in Neural Information Processing Systems (NeurIPS)*, page 730–738, 2016. 3
- [10] Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, abs/2412.05271, 2024. 4
- [11] Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 3
- [12] 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. 1, 2, 3, 6, 7, 8
- [13] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10786–10796, 2021. 7
- [14] Xiao Fu, Wei Yin, Mu Hu, Kaixuan Wang, Yuexin Ma, Ping Tan, Shaojie Shen, Dahua Lin, and Xiaoxiao Long. GeoWizard: Unleashing the diffusion priors for 3d geometry estimation from a single image. *arXiv preprint arXiv:2403.12013*, 2024. 1, 3, 7
- [15] Elena Garces, Carlos Rodriguez-Pardo, Dan Casas, and Jorge López-Moreno. A survey on intrinsic images: Delving deep into lambert and beyond. *International Journal of Computer Vision*, 130:836 – 868, 2021. 3
- [16] Jing He, Haodong Li, Wei Yin, Yixun Liang, Leheng Li, Kaiqiang Zhou, Hongbo Liu, Bingbing Liu, and Yingcong Chen. Lotus: Diffusion-based visual foundation model for high-quality dense prediction. *arXiv preprint arXiv:2409.18124*, 2024. 7
- [17] 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. 6, 7
- [18] Bingxin Ke, Kevin Qu, Tianfu Wang, Nando Metzger, Shengyu Huang, Bo Li, Anton Obukhov, and Konrad Schindler. Marigold: Affordable adaptation of diffusion-based image generators for image analysis. *IEEE transactions on pattern analysis and machine intelligence*, PP, 2025. 1, 2, 3, 5, 6, 7
- [19] Peter Kocsis, Vincent Sitzmann, and Matthias Nießner. Intrinsic image diffusion for indoor single-view material estimation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5198–5208, 2024. 1, 6
- [20] Edwin Herbert Land and John J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, 61 1:1–11, 1971. 2
- [21] Yixuan Li, Lihan Jiang, Linning Xu, Yuanbo Xiangli, Zhenzhi Wang, Dahua Lin, and Bo Dai. MatrixCity: A large-scale city dataset for city-scale neural rendering and beyond. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3182–3192, 2023. 1, 3
- [22] Zhengqin Li, Ting-Wei Yu, Shen Sang, Sarah Wang, Meng Song, Yuhan Liu, Yu-Ying Yeh, Rui Zhu, Nitesh Gundavarapu, Jia Shi, Sai Bi, Hong-Xing Yu, Zexiang Xu, Kalyan Sunkavalli, Miloš Hašan, Ravi Ramamoorthi, and Manmohan Chandraker. OpenRooms: An open framework for photorealistic indoor scene datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 3
- [23] Zhiyuan Li, Dongnan Liu, Chaoyi Zhang, Heng Wang, Tengfei Xue, and Weidong Cai. Enhancing advanced visual reasoning ability of large language models. *arXiv preprint arXiv:2409.13980*, 2024. 1
- [24] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision*, 2014. 5, 8
- [25] Jie Liu, Gongye Liu, Jiajun Liang, Yangguang Li, Jiaheng Liu, Xintao Wang, Pengfei Wan, Di Zhang, and Wanli

- Ouyang. Flow-GRPO: Training flow matching models via online rl. *arXiv preprint arXiv:2505.05470*, 2025. 3, 4, 5
- [26] Grace Luo, Jonathan Granskog, Aleksander Holynski, and Trevor Darrell. Dual-process image generation. *arXiv preprint arXiv:2506.01955*, 2025. 3
- [27] Jundan Luo, Nanxuan Zhao, Wenbin Li, and Christian Richardt. CRefNet: Learning consistent reflectance estimation with a decoder-sharing transformer. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 2023. 6
- [28] Lukas Murmann, Michaël Gharbi, Miika Aittala, and Frédo Durand. A dataset of multi-illumination images in the wild. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4079–4088, 2019. 8
- [29] Takuya Narihira, Michael Maire, and Stella X Yu. Learning lightness from human judgement on relative reflectance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2965–2973, 2015. 2
- [30] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 44:1623–1637, 2019. 7
- [31] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12179–12188, 2021. 7
- [32] 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 (ICCV)*, pages 10912–10922, 2021. 1, 3
- [33] Juan Rocamonde, Victoriano Montesinos, Elvis Nava, Ethan Perez, and David Lindner. Vision-language models are zero-shot reward models for reinforcement learning. *arXiv preprint arXiv:2310.12921*, 2023. 3
- [34] Thomas Schöps, Johannes L. Schönberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2538–2547, 2017. 6, 8
- [35] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Jun-Mei Song, Mingchuan Zhang, Y. K. Li, Yu Wu, and Daya Guo. DeepSeekMath: Pushing the limits of mathematical reasoning in open language models. *ArXiv*, abs/2402.03300, 2024. 2, 3, 5
- [36] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2012. 6
- [37] Igor Vasiljevic, Nick Kolkin, Shanyi Zhang, Ruotian Luo, Haochen Wang, Falcon Z. Dai, Andrea F. Daniele, Mohammadreza Mostajabi, Steven Basart, Matthew R. Walter, and Gregory Shakhnarovich. DIODE: A Dense Indoor and Outdoor DEpth Dataset. *arXiv preprint arXiv:1908.00463*, 2019. 6, 7
- [38] David Venuto, Sami Nur Islam, Martin Klissarov, Doina Precup, Sherry Yang, and Ankit Anand. Code as reward: Empowering reinforcement learning with vision-language models. *arXiv preprint arXiv:2402.04764*, 2024. 3
- [39] Yufei Wang, Zhanyi Sun, Jesse Zhang, Zhou Xian, Erdem Biyik, David Held, and Zackory Erickson. RL-VLM-F: reinforcement learning from vision language foundation model feedback. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2024. 3
- [40] Chenyuan Wu, Pengfei Zheng, Ruiran Yan, Shitao Xiao, Xin Luo, Yuezhe Wang, Wanli Li, Xiyang Jiang, Yexin Liu, Junjie Zhou, et al. OmniGen2: Exploration to advanced multimodal generation. *arXiv preprint arXiv:2506.18871*, 2025. 5, 6, 7
- [41] Diankun Wu, Fangfu Liu, Yi-Hsin Hung, and Yueqi Duan. Spatial-MLLM: Boosting MLLM capabilities in visual-based spatial intelligence. *arXiv preprint arXiv:2505.23747*, 2025. 2, 3
- [42] Jiaye Wu, Sanjoy Chowdhury, Hariharmano Shanmugaraja, David Jacobs, and Soumyadip Sengupta. Measured albedo in the wild: Filling the gap in intrinsics evaluation. *2023 IEEE International Conference on Computational Photography (ICCP)*, pages 1–12, 2023. 5, 6
- [43] Penghao Wu, Yushan Zhang, Haiwen Diao, Bo Li, Lewei Lu, and Ziwei Liu. Visual jigsaw post-training improves MLLMs. *arXiv preprint arXiv:2509.25190*, 2025. 3
- [44] Runsen Xu, Weiyao Wang, Hao Tang, Xingyu Chen, Xiaodong Wang, Fu-Jen Chu, Dahua Lin, Matt Feiszli, and Kevin J. Liang. Multi-SpatialMLLM: Multi-frame spatial understanding with multi-modal large language models. *arXiv preprint arXiv:2505.17015*, 2025. 2, 3
- [45] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 6
- [46] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024. 6, 7
- [47] Michihiro Yasunaga, Luke Zettlemoyer, and Marjan Ghazvininejad. Multimodal RewardBench: Holistic evaluation of reward models for vision language models. *arXiv preprint arXiv:2502.14191*, 2025. 3
- [48] 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)*, 43:1 – 18, 2024. 7
- [49] 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*, 2024. 1, 3, 6
- [50] Weichen Zhan, Zile Zhou, Zhiheng Zheng, Chen Gao, Jinqiang Cui, Yong Li, Xinlei Chen, and Xiao-Ping Zhang. Open3DVQA: A benchmark for comprehensive spatial reasoning with multimodal large language model in open space. *arXiv preprint arXiv:2503.11094*, 2025. 2, 3

- [51] Chi Zhang, Wei Yin, Zhibin Wang, Gang Yu, Bin Fu, and Chunhua Shen. Hierarchical normalization for robust monocular depth estimation. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 7
- [52] Ke Zhang, Cihan Xiao, Yiqun Mei, Jiacong Xu, and Vishal M. Patel. Think before you diffuse: Llm-guided physics-aware video generation. *arXiv preprint arXiv:2505.21653*, 2025. 3
- [53] Jingsen Zhu, Fujun Luan, Yuchi Huo, Zihao Lin, Zhihua Zhong, Dianbing Xi, Rui Wang, Hujun Bao, Jiayang Zheng, and Rui Tang. Learning-based inverse rendering of complex indoor scenes with differentiable monte carlo raytracing. In *ACM SIGGRAPH Asia Conference Proceedings*. ACM, 2022. 1, 3, 6
- [54] Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, et al. InternVL3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025. 1, 3
- [55] Daniel Zoran, Phillip Isola, Dilip Krishnan, and William T Freeman. Learning ordinal relationships for mid-level vision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2015. 3