

# Physically-Based Photometric Bundle Adjustment in Non-Lambertian Environments

Lei Cheng\*, Junpeng Hu\*, Haodong Yan, Mariia Gladkova, Tianyu Huang,  
Yun-Hui Liu, Daniel Cremers, and Haoang Li<sup>2,†</sup>

**Abstract**—Photometric bundle adjustment (PBA) is widely used in estimating the camera pose and 3D geometry by assuming a Lambertian world. However, the assumption of photometric consistency is often violated since the non-diffuse reflection is common in real-world environments. The photometric inconsistency significantly affects the reliability of existing PBA methods. To solve this problem, we propose a novel physically-based PBA method. Specifically, we introduce the physically-based weights regarding material, illumination, and light path. These weights distinguish the pixel pairs with different levels of photometric inconsistency. We also design corresponding models for material estimation based on sequential images and illumination estimation based on point clouds. In addition, we establish the first SLAM-related dataset of non-Lambertian scenes with complete ground truth of illumination and material. Extensive experiments demonstrated that our PBA method outperforms existing approaches in accuracy.

## I. INTRODUCTION

Bundle adjustment is a key component of SLAM [1], [2], 3D reconstruction [3], and structure from motion [4]. By directly aligning the pixels, photometric bundle adjustment (PBA) has demonstrated notable performance in estimating camera poses and reconstructing scene geometry, even in low-texture environments [5].

The classic photometric error of PBA is based on the assumption of photometric consistency, implying that corresponding pixels in different views receive identical radiance. However, this assumption holds only for Lambertian surfaces, while non-Lambertian objects are common in real-world situations. These surfaces frequently result in extensive glossy areas, particularly under complex illumination. As illustrated in Fig. 1, such glossy regions pose significant challenges to the robustness of existing PBA methods.

Existing strategies of handling photometric inconsistencies can be classified into three categories: invariance-based methods [6]–[9], affine model-based methods [5], [10]–[12], and statistical methods [13]–[17]. Invariance-based methods,

Lei Cheng and Junpeng Hu contributed equally to this work. Haoang Li is corresponding author (haoangli@hkust-gz.edu.cn).

This work was supported by the ERC Advanced Grant SIMULACRON, the Munich Center for Machine Learning, the EPSRC Programme Grant VisualAI EP/T028572/1, the Shenzhen Portion of Shenzhen-Hong Kong Science and Technology Innovation Cooperation Zone under HZQB-KCZYB-20200089, and the InnoHK of the Government of Hong Kong via the Hong Kong Centre for Logistics Robotics.

L. Cheng, J. Hu, M. Gladkova, and D. Cremers are with Technical University of Munich and Munich Center for Machine Learning, Munich, Germany. H. Yan and H. Li are with The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China. T. Huang and Y.-H. Liu are with The Chinese University of Hong Kong, Hong Kong, China.




Fig. 1. Effectiveness of our PBA method in a non-Lambertian environment. DSM-PBA [17] improves the accuracy of ORB-SLAM2 [18], and our PBA method is more accurate than DSM-PBA. The color bar shows the magnitude of the absolute trajectory error.

using relative pixel intensity ordering or intensity changes between neighbors, struggle with spatially varying illumination and are sensitive to noise. Affine model-based approaches often assume a linear relationship between illumination conditions. This simplification is not always applicable in environments with highly reflective surfaces or a wide dynamic range of illumination. Statistical methods like [16], [17] can lead to inaccuracies in complex environments where illumination conditions vary locally, inadequately representing outlier distribution caused by photometric inconsistencies. These limitations become pronounced under intensive illumination changes, reducing their effectiveness.

To overcome the above limitations, we propose a novel physically-based photometric bundle adjustment method. We integrate the physically-based reflection model into the classical PBA framework. We design a weighted photometric error to eliminate the impact of pixels corresponding to the non-Lambertian surface points. However, it is challenging to model the weight scalar directly. Based on the physically-based rendering (PBR) model, radiance values are determined by a combination of factors including surface physical properties, scene geometry, and illumination. Therefore, the weight should encode the 1) illumination, 2) object material, and 3) light paths to a point in 3D space. To achieve this, we necessitate material and illumination as prior information.

Object material can be estimated from images. However, a single image may not be sufficient under challenging illumination conditions. Intuitively, sequential images from multiple viewpoints can resolve potential ambiguities and improve the estimation. Existing works either focus on single

frame input [19]–[23], or necessitate individual optimization on every scene [24]–[26]. Therefore, we develop our new method of material estimation, which processes sequential images with a transformer and outputs the per-pixel roughness, without per-scene optimization. More specifically, we use a transformer to process sequential inputs and capture spatial–temporal correlation over frames.

Illumination of various positions is required as a prior in our PBA. Moreover, the estimated illumination should be spatially varying and consistent without necessitating per-scene optimization. Existing works either lack a consistent representation [22], [23] or have a limited field of view [27]. Therefore, we design our new pipeline of illumination estimation, utilizing a point cloud collected from multiple frames of the scene as input. The output panorama image, termed an “environment map”, can be used as illumination information for the queried 3D point.

Material and illumination information, serving as priors, will not undergo further optimization in our PBA. However, the light paths are determined by the camera poses and estimated point clouds. Therefore, our weight can be treated as a function of camera poses and point depths. We integrate this term into the classic photometric loss and minimize it to obtain optimal light paths.

Existing SLAM-related datasets [28], [29] do not simultaneously contain sufficient non-Lambertian surfaces and information regarding material and illumination. To overcome their limitations in evaluating non-Lambertian SLAM-related work, we establish our SpecularRooms dataset.

In summary, we make the following contributions:

- We propose a novel weighting scheme for photometric error function based on illumination and material information to satisfy the photometric consistency in challenging environments with non-Lambertian surfaces.
- We propose new pipelines for material estimation based on sequential images and illumination estimation using point cloud data, without per-scene optimization.
- We establish the first PBR dataset of sequential images in non-Lambertian environments with ground truth material and illumination.

Our SpecularRooms dataset and supplementary document are available on the project website<sup>1</sup>.

## II. RELATED WORKS

### A. Methods for Photometric Consistency

Depending on the way photometric inconsistencies are handled, we classify existing methods into three primary categories, i.e., illumination invariance-based methods, affine model-based methods, and statistical methods.

**Invariance-Based Methods.** Census Transform [6] utilizes relative ordering of intensity, making it invariant concerning monotonic illumination variations. Alismail et al. [9] transform images into 8-channel, illumination-invariant binary descriptor images for minimizing photometric error, demonstrating enhanced robustness in subterranean environments.

In contrast to the alignment of raw pixels, the approaches presented by Crivellaro et al. [7] and Dai et al. [8], align gradient magnitudes. This method results in invariance only to local bias changes, making it sensitive to noise and less effective against complex spatially varying illumination.

**Affine Model-Based Methods.** Works by Klose et al. [11] and Engel et al. [12] focus on optimizing the relative pose along with global linear affine brightness transfer parameters between images. DSO [5] further incorporates photometric calibration to address the camera’s response, vignetting, and exposure variations. Additionally, Jin et al. [10] employ affine models for local patches to handle subtle illumination variations, yet this approach falls short under drastic spatial illumination changes.

**Statistical Methods.** Some statistics in pixel space are employed to handle photometric inconsistency. Scandaroli et al. [13] employ zero-mean normalized cross correlation in planar template tracking. Median bias normalization, as explored by Meilland et al. [14] and Gonçalves et al. [15], compensates for intensity biases on both global and patch levels by utilizing the median of residuals. This approach thus achieves partial invariance to variations in intensity bias. To enhance the robustness of PBA, [17] applies a t-distribution to down-weight outliers and achieve state-of-the-art, aiming to mitigate photometric errors. However, this approach falls short of accurately modeling photometric errors under complex illumination conditions, suggesting a need for more sophisticated error modeling techniques.

In contrast, our PBA method considers the physically-based reflection model of each 3D point. By leveraging illumination and material prior information, we can calculate an adaptive weight for every photometric error specifically. Therefore, our method achieves superior performance in challenging scenes than previous works.

### B. Material and Illumination Estimation

As mentioned above, our method relies on the estimation of material and illumination. In this part, we shall review related methods.

**Material Estimation.** Traditional methods for material estimation, as the works of Barron et al. [19], [20], predominantly rely on some strong heuristic priors within an optimization framework. However, these heuristic priors may not be reliable, particularly in indoor scenes characterized by complex geometry and challenging illumination conditions. Representation-based methods [24]–[26] usually take multi-view input and conduct per-scene optimization for their representation. These approaches usually focus on the simultaneous reconstruction of material properties, illumination, and geometry, but rely on per-scene optimization. Other deep learning-based methods [21], [23], without per-scene optimization, predominantly focus on single-frame or stereo inputs. However, information from a single viewpoint may lead to ambiguity caused by complicated illumination and material. In contrast, our method, without per-scene optimization, uses sequential images as input to overcome their limitations.

<sup>1</sup><https://sites.google.com/view/haoangli/projects/nlb-pba>




Fig. 2. Physically-based reflection model based on point light source. The incident direction  $\alpha_k$  is associated with the radiance  $I_k$ . The radiance received by a 3D point  $\mathbf{P}$  is then reflected as the radiance  $R_k$  along the reflective direction  $\beta$ . The pixel  $\mathbf{p}$  receives the sum of reflective radiances along the direction  $\beta$ .

**Illumination Estimation.** Early works [30]–[32] estimate a single environment map from a single image. However, a single environment map is not enough for complex indoor scenes. Other works [23], [33] consider the spatial variance of illumination, but with a simplified representation. This approximation cannot fully preserve fine-grained illumination information in the observed part of the scene. Srinivasan et al. [27] propose a spatial consistent estimation based on multi-plane images and multi-scale volumes. Nonetheless, it takes only stereo inputs, with a limited field of view. Another work [34] handles spatial and temporal coherence of video input. However, their RNN-based structure limits the capability of representing illumination in long image sequences. In our method, we feed point cloud into our model to utilize multi-view information in a consistent way.

### III. BACKGROUND

#### A. Photometric Bundle Adjustment

The classical photometric error is formulated as a pixel-wise alignment of two images. As shown in Fig. 2, we back-project a pixel  $\mathbf{p} \in \mathbb{R}^2$  from the reference frame  $\mathcal{F}$  into a 3D point  $\mathbf{P} \in \mathbb{R}^3$  using the depth  $d$  of  $\mathbf{p}$ . Then we project the 3D point  $\mathbf{P}$  into a pixel  $\mathbf{p}'$  on the  $i_{th}$  visible frame  $\mathcal{F}'$  using the unknown rotation  $\mathbf{R}$  and translation  $\mathbf{t}$  between two frames. Accordingly, the pixel  $\mathbf{p}'$  is with respect to the pose  $\mathbf{R}, \mathbf{t}$  and the depth  $d$ . With using the pinhole model and given camera intrinsics in  $\mathbf{p}'(\mathbf{R}, \mathbf{t}, d)$  implicitly, the classic photometric error is defined by

$$E_{\mathbf{p}i} = \sum_{\mathbf{p} \in \mathcal{N}_{\mathbf{p}}} \|\mathcal{F}(\mathbf{p}) - \mathcal{F}'(\mathbf{p}'(\mathbf{R}, \mathbf{t}, d))\|, \quad (1)$$

where  $\mathcal{N}_{\mathbf{p}}$  denote a set of pixels around the pixel  $\mathbf{p}$ . This photometric error is predicted on the assumption that the 3D point  $\mathbf{P}$  resides on a Lambertian surface, where the radiance, power received by pixels  $\mathbf{p}$  and  $\mathbf{p}'$ , is uniform across view-points. However, for non-Lambertian objects, their surface reflection varies with viewing direction intensively.

Fundamentally, PBA optimizes scene geometry and camera poses jointly by minimizing the photometric error of projecting all estimated 3D points on their visible frames. The classical PBA loss is given by

$$\mathcal{L}_{\text{photo}} = \sum_{j \in \mathcal{F}} \sum_{\mathbf{p} \in \mathbf{P}_j} \sum_{i \in \text{obs}(\mathbf{p})} E_{\mathbf{p}i}, \quad (2)$$

where  $j$  traverses through all frames  $\mathcal{F}$ , and within each frame  $j$ ,  $\mathbf{p}$  iterates over all points  $\mathbf{P}_j$ . The  $\text{obs}(\mathbf{p})$  denotes all visible frames for  $\mathbf{p}$ .

#### B. Physically-Based Reflection Model

To formulate the PBA with the physically-based reflection model, we start with a conceptual exposition using direct lighting. As illustrated in Fig. 2, at a 3D point  $\mathbf{P}$ , an incident light direction  $\alpha_1$  is associated with the radiance  $I_1$ . The “valid” radiance received by  $\mathbf{P}$  is the component along the surface normal  $\mathbf{n}$ , i.e.,  $I_1 \cdot \cos \theta_{\alpha_1, \mathbf{n}}$ . For the surface around the point  $\mathbf{P}$ , we denote its normal by  $\mathbf{n}$ . After reflection at point  $\mathbf{P}$ , the valid radiance is distributed to different reflective directions such as  $\beta$  and  $\beta'$ . Bidirectional reflectance distribution function (BRDF) [35]  $f_{\mathbf{P}}$  determines the distributed radiances  $r_1$  and  $r'_1$  along the directions  $\beta$  and  $\beta'$ , i.e.,

$$r_1 = I_1 \cdot \cos \theta_{\alpha_1, \mathbf{n}} \cdot f_{\mathbf{P}}(\alpha_1, \mathbf{n}, \beta), \quad (3a)$$

$$r'_1 = I_1 \cdot \cos \theta_{\alpha_1, \mathbf{n}} \cdot f_{\mathbf{P}}(\alpha_1, \mathbf{n}, \beta'). \quad (3b)$$

### IV. OUR METHOD

Given an initial estimate of the scene points  $\{\mathbf{P}_j\}_{j=1}^N$  and viewing parameters  $\{\mathbf{T}_i\}_{i=1}^M$  per frame in the world, we first estimate each pixel’s roughness and normal in every frame. Meanwhile, we get the colored point cloud of the scene and select control points. Then the illumination network predicts environment maps of control points, with the point cloud as input. Leveraging environment maps, surface roughness, and surface normals, we optimize the trajectory and the scene geometry in our PBA method.

#### A. PBA with a Physically-Based Reflection Model

In this section, we introduce a weighted photometric error in PBA loss, which is defined as

$$\mathcal{L} = \sum_{j \in \mathcal{F}} \sum_{\mathbf{p} \in \mathbf{P}_j} \sum_{i \in \text{obs}(\mathbf{p})} \delta_{\mathbf{p}i} \cdot E_{\mathbf{p}i}, \quad (4)$$

where we define the weight as

$$\delta_{\mathbf{p}i} = e^{-\theta|r-r'|}. \quad (5)$$

Here,  $r$  and  $r'$  correspond to radiance reflected by 3D point  $\mathbf{p}$  in two view directions. In static scenes, unoccluded points with significant variations in reflected radiance result in large photometric errors. To decrease the impact of these points, we propose a straightforward inversely proportional function to determine the weight  $\delta$  in Eq. (5). We empirically found that  $\theta = 14.6$  leads to superior PBA performance. Ablation studies are available in the supplementary document.

In practice, the calculation of radiance  $r$  in Eq. (5) should consider both direct and indirect lighting. To define our radiance model, we first introduce the global radiance received by the pixel  $\mathbf{p}$ . As shown in Fig. 2, each incident direction  $\alpha_k$  is associated with the radiance  $I_k$ . The unit hemisphere  $\mathcal{H}$  covers all the incident directions. Analogous with Eq. (3a), the reflective radiance along with the reflective direction  $\beta$  becomes  $I_k \cdot \cos \theta_{\mathbf{n}, \alpha_k} \cdot f_{\mathbf{P}}(\alpha_k, \mathbf{n}, \beta)$ . By integrating the

radiance function over the hemisphere  $\mathcal{H}$ , we compute the total radiance  $r$  received by the pixel  $\mathbf{p}$  as

$$r = \int_{\mathcal{H}} I_k \cdot \cos \theta_{\mathbf{n}, \alpha_k} \cdot f_{\mathbf{P}}(\alpha_k, \mathbf{n}, \beta) d\alpha_k. \quad (6)$$

To represent the physical reflection process, i.e., the light path, we detail the formulation of the vectors  $\mathbf{n}$ ,  $\beta$ , and  $\beta'$  from Eq. (6) in the following subsections.

**Normal of Facet  $\mathbf{n}$ .** We back-project the pixel  $\mathbf{p}$  and its neighbor pixels using the measured depths  $d$  and its neighbouring depth  $\{d_i\}$  from the depth map, obtaining 3D points  $\mathbf{P}(d)$  and its surrounding points within a 2cm vicinity  $\{\mathbf{N}_i(d_i)\}$ . Each 3D point triplet, e.g.,  $\{\mathbf{P}(d), \mathbf{N}_1(d_1), \mathbf{N}_2(d_2)\}$  leads to a normal, e.g.,  $\mathbf{n}(d, d_1, d_2) = (\mathbf{P}(d) - \mathbf{N}_1(d_1)) \times (\mathbf{P}(d) - \mathbf{N}_2(d_2))$ . Then we use the moving least square method [36] to smooth the estimated normal.

**Reflective Direction  $\beta$  and  $\beta'$ .** With the relative rotation  $\mathbf{R}$  and translation  $\mathbf{t}$ , the reflective directions  $\beta$  and  $\beta'$  can be expressed as

$$\beta(d) = \mathbf{P}(d), \quad (7)$$

$$\beta'(\mathbf{R}, \mathbf{t}, d) = \mathbf{P}(d) + \mathbf{R}^\top \mathbf{t}. \quad (8)$$

Accordingly, we re-formulate the photometric loss  $\mathcal{L}_{\text{photo}}$  (Eq. (2)) as

$$\mathcal{L} = \sum_{j \in \mathcal{F}} \sum_{\mathbf{p} \in \mathbf{P}_j} \sum_{i \in \text{obs}(\mathbf{p})} \delta_{\mathbf{p}_i}(\mathbf{R}, \mathbf{t}, d, \{d_i\}) \cdot E_{\mathbf{p}_i}, \quad (9)$$

which additionally incorporates a physically-based reflection model. Nevertheless, the proposed loss  $\mathcal{L}$  cannot be directly minimized using a gradient descent algorithm since the partial derivatives of several integrals in the loss over the camera parameters  $\{\mathbf{R}, \mathbf{t}, d, \{d_i\}\}$  are difficult to compute [37]. To avoid computing these integrals in Eq. (6), we leverage the estimated environment lighting and roughness in BRDF [35].

Given the surface material, reflective direction  $\beta$  and surface normal  $\mathbf{n}$ , we can determine the solid angle area  $\mathcal{A}_{\mathbf{P}}(\mathbf{n}, \beta)$  in BRDF. Hence, we split the irradiance  $r$  in Eq. (6) as

$$r \approx \underbrace{\frac{\int_{\mathcal{A}_{\mathbf{P}}(\mathbf{n}, \beta)} I_k d\alpha_k}{\int_{\mathcal{A}_{\mathbf{P}}(\mathbf{n}, \beta)} d\alpha_k}}_{T_{\text{Light}}} \cdot \underbrace{\int_{\mathcal{H}} \cos \theta_{\mathbf{n}, \alpha_k} \cdot f_{\mathbf{P}}(\alpha_k, \mathbf{n}, \beta) d\alpha_k}_{T_{\text{BRDF}}}. \quad (10)$$

- **Term  $T_{\text{Light}}$ .** Given the initial structure and estimated cameras, we use a viewing direction  $\beta$  and  $\mathbf{n}$  to compute the area  $\mathcal{A}_{\mathbf{P}}$  and to further index the pre-computed integral of the environment map.
- **Term  $T_{\text{BRDF}}$ .** Term  $T_{\text{BRDF}}$  is predominantly influenced by the surface roughness  $r_s$  of the surface material, which ranges from 0 to 1.

**Illumination at Control Points.** According to Eq. (6), we need environment lighting  $I$  for every point in PBA. However, it is impractical and unnecessary to predict all of them, as points in the vicinity are under similar illumination conditions. Therefore, we select representative ‘‘control points’’ on the surfaces by employing SLIC [38] clustering on




Fig. 3. Material estimation pipeline. Our approach uses the current frame and  $K - 1$  previous frames to predict the roughness of the current frame. A CNN encoder extracts features from these  $K$  frames. A transformer and a decoder process these features to predict the roughness of the current frame.

the normal maps. Subsequently, we estimate the environment maps for these control points as introduced in Section IV-C. For each 3D point, we employ the predicted environment map of its nearest ‘‘control point’’ in Eq. (6).

Utilizing the estimated roughness  $r_s$ , as detailed in Section IV-B, along with environment lighting for each point, we compute the specular radiance  $r$  of 3D points across all views. Subsequently, the radiance is applied in Eq. (5) to weigh each photometric error present in the loss function  $\mathcal{L}$ . To optimize Eq. (9), initial camera poses and scene geometry can be obtained within SLAM frameworks. During optimization, the refined camera poses and geometry informs the light path determination in subsequent optimizations. We follow [39] to use the Levenberg-Marquardt algorithm to minimize the weighted PBA loss.

## B. Material Estimation

The appearance of surface material varies significantly under different illumination conditions, leading to ambiguities in material estimation from single-frame inputs. To handle these ambiguities, we use the multi-view information from sequential images. Our material estimation network leverages previous  $K - 1$  frames to enhance the estimation of the current frame. As shown in Fig. 3, the architecture of our material estimation model consists of three principal components: (1) a CNN encoder to extract features from individual frames, (2) a transformer to capture and model the spatial-temporal relationships across frames, and (3) a CNN decoder to produce the per-pixel roughness.

**Single-Frame Encoder.** Our model adopts the CNN encoder structure as described in IRCIS [23]. Initially, frame-level features are extracted for each frame using this encoder. These extracted features from  $K$  frames are then concatenated to construct a multi-frame feature map, denoted as  $f \in \mathbb{R}^{K \times C \times H' \times W'}$ .

**Spatial-temporal Correlation.** We use a transformer to capture the spatial-temporal correlations among different locations across frames. We flatten the feature maps of  $K$  frames into  $f \in \mathbb{R}^{d \times (K \cdot H' \cdot W')}$ , which are  $K \cdot H' \cdot W'$  tokens of dimension  $d$ . These tokens are fed into the transformer, with temporal and spatial positional information encoded by 3D position embedding as in [40].

**Output Decoder.** We adopt the CNN decoder structure from IRCIS [23] to output per-pixel roughness.




Fig. 4. Illumination estimation pipeline. The colored point cloud is fed into the point encoder, which extracts per-point features. These features are then projected and used as inputs for the image decoder. The image decoder then produces a coarse output. The coarse output is combined with a color-depth map in the RefineNet, yielding a refined output.

**Loss Function.** During training, the model estimates roughness maps of  $i - K$  to  $i$  frames. A mean squared error (MSE) loss term  $L_{\text{MSE}}$  measures pixel-level error and a structural similarity index measure (SSIM) [41] term  $L_{\text{SSIM}}$  assesses the similarity of the overall structure. The final loss function  $L_{\text{mat}}$  is a weighted sum of the two terms:

$$L_{\text{mat}} = \lambda_{\text{MSE}} \cdot L_{\text{MSE}} + \lambda_{\text{SSIM}} \cdot L_{\text{SSIM}}, \quad (11)$$

where  $\lambda_{\text{MSE}}$  and  $\lambda_{\text{SSIM}}$  are the corresponding weights of MSE loss and SSIM loss.

### C. Illumination Estimation

For every control point in the scene, its illumination varies with its position, especially in indoor scenes with complex occlusion and illumination. Therefore we need to estimate a spatially varying illumination of the whole scene. Compared to single-image input, multi-view images of the scene brings a larger field of view, which benefits this task. However, individually processing these images might lead to spatial inconsistency, since they are from different viewpoints. To avoid this inconsistency, we need to integrate those images into a unified 3D representation. In our PBA framework, a colored point cloud of the scene is available and we utilize it as input for our illumination prediction pipeline. Our illumination estimation pipeline is shown in Fig. 4. Given the scene’s point cloud, the model outputs an environment map as the environment lighting of the queried position.

**Point Encoder.** To extract multi-level geometry information of the scene, we adopt the point encoder developed based on the PointNet++ [43], following Synpt [44]. It is designed to extract information from points and their proximal surroundings. This model employs four layers of set abstraction and four levels of propagation, enabling the concurrent capture of localized and global features. The identified features are then interpolated within the feature propagation layers.

**Image Decoder.** Following [44], point features are projected onto various levels of feature map spheres aligned with subpointset sizes. The image decoder employs these projected feature maps for generating the environment map.

**RefineNet.** We use a RefineNet module to enhance coarse environment maps from the image decoder as [44]. In this module, we feed the color and depth information of points into the pipeline once again and introduce a 2D self-attention block to capture the long-range interactions.

**Loss Function.** We follow the setup of [27], where the loss function comprises an adversarial loss term  $L_{\text{adv}}$  and

a perceptual loss term. The perceptual loss quantifies the high-level discrepancies between images. Following [27], we use a pre-trained VGG-19 network for feature extraction and compute the  $L_1$  loss of features across layers as  $L_{\text{vgg}}$ . The total loss function  $L_{\text{light}}$  is defined as a weighted sum:

$$L_{\text{light}} = \lambda_{\text{vgg}} L_{\text{vgg}} + \lambda_{\text{adv}} L_{\text{adv}}, \quad (12)$$

where  $\lambda_{\text{vgg}}$  and  $\lambda_{\text{adv}}$  are the weights of perceptual loss and adversarial loss.

## V. EXPERIMENTS

We first introduce our SpecularRooms dataset. Then we compare our PBA, material estimation, and illumination methods with corresponding state-of-the-art works. Ablation studies, additional results, and more implementation details are available in our supplementary document. All the following tests were conducted on a computer equipped with NVIDIA RTX A4000 GPU.

### A. Our SpecularRooms Dataset

Existing SLAM-related datasets, such as TUM-RGBD [29] and EuRoC [28], can not fulfill the requirements for evaluating non-Lambertian SLAM methods. They have at least one of the following limitations: 1) Incomplete ground truth for materials and illumination. 2) Lack of non-Lambertian surfaces. To address these limitations, we establish a new dataset that encompasses all the aforementioned criteria. We name this dataset “SpecularRooms”. It contains more than 17000 images from 85 sequences of different lengths and complete ground truth trajectories, material, and illumination. The experiments of the following sections are conducted on our SpecularRooms dataset. In the experiments of material estimation and illumination estimation, we use 80% of images for training and 20% for testing. To validate our PBA method, we choose 6 longest trajectories.

### B. Photometric Bundle Adjustment

**Experimental Setup.** Our pixel selection strategy adheres to the approach outlined in [5]. The input trajectory is derived by estimating relative poses pairwise, which can be accomplished by any visual odometry or SLAM system, such as ORB-SLAM2 [18].

**Methods for Comparison.** Here we refer to the classical PBA as LBT-PBA since it assumes a Lambertian world. We compare our PB-PBA with LBT-PBA [42] and the state-of-the-art t-distribution based PBA work DSM-PBA [17].

- LBT-PBA is implemented based on the photometric consistency assumption, without considering non-Lambertian surfaces.
- DSM-PBA models the photometric error distribution through a selected t-distribution, thereby down-weighting outlier points.

**Metrics.** We follow [45], [46] to evaluate the trajectory accuracy based on the absolute trajectory error (ATE).

**Results.** As shown in Fig. 5, DSM-PBA enhances the accuracy compared to the classical DSM approach. Nonetheless,

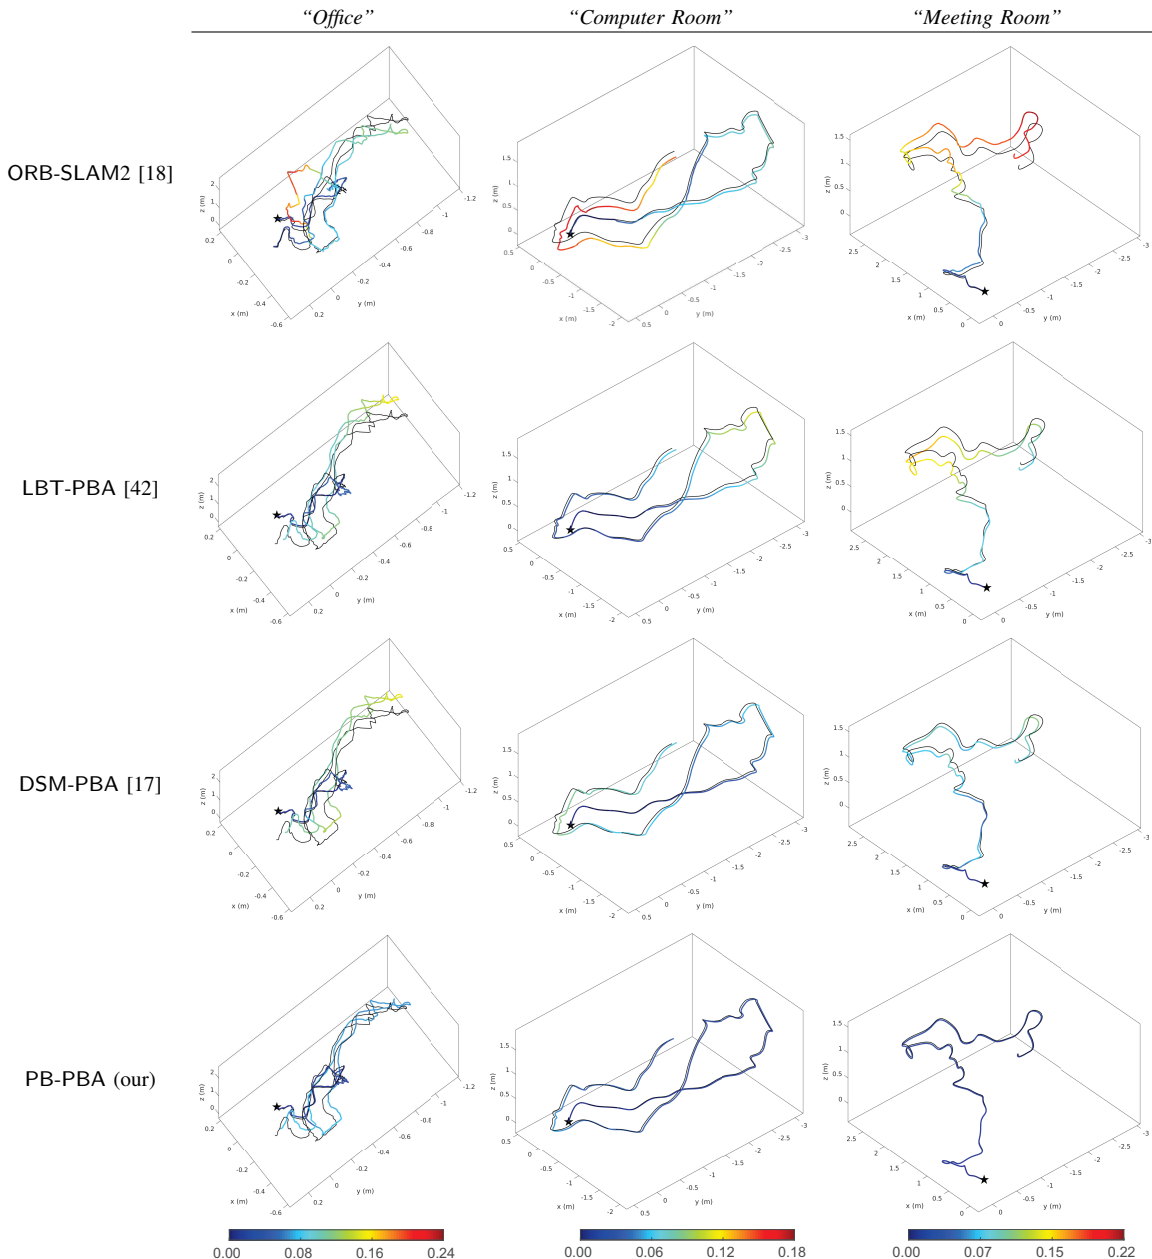


Fig. 5. Trajectories estimated by ORB-SLAM2 [18] and various PBA methods on three image sequences of our SpecularRooms dataset. The colored and black lines denote the estimated and ground truth trajectories, respectively. The color bar indicates the magnitude of the absolute trajectory error.

the precision of DSM-PBA is somewhat constrained due to its inability to adjust photometric error in a comprehensible manner. In contrast, our method attains superior accuracy, primarily because it incorporates a reflection model into the weight calculation process.

### C. Material Estimation

**Methods for Comparison.** We compare our method with IRCIS [23], a state-of-the-art work with single-frame input. To our knowledge, there is no open-source scene-level method with multi-view input, without requiring per-scene optimization. We re-train IRCIS on our SpecularRooms dataset using its recommended parameters for a fair comparison.

**Metrics.** Following [23], we use mean squared error (MSE) as our evaluation metric.

TABLE I  
ABSOLUTE TRAJECTORY ERRORS OF VARIOUS PBA METHODS ON OUR SPECULARROOMS DATASET (UNIT: METERS)

Sequence	LBT-PBA [42]	DSM-PBA [17]	PB-PBA (our)
<i>Club</i>	0.080	0.078	<b>0.074</b>
<i>Office</i>	0.093	0.088	<b>0.052</b>
<i>Computer Room</i>	0.060	0.058	<b>0.032</b>
<i>Meeting Room</i>	0.105	0.069	<b>0.018</b>
<i>Study Room</i>	0.107	0.070	<b>0.056</b>
<i>Restroom</i>	0.115	0.062	<b>0.012</b>
Average	0.106	0.067	<b>0.031</b>

**Results.** In Table II, we list the mean and median of per-frame MSE of our method with 5 frames (Ours@5) and




Fig. 6. Error maps of roughness estimation results on our SpecularRooms dataset. Here we list input images (in the first row) and the corresponding error maps of roughness predictions. The second row corresponds to IRCIS and the third row corresponds to our method. The color bar indicates the magnitude of the difference between roughness predictions and ground truth. Lighter pixels correspond to larger errors. Note that our method works better on glossy and featureless surfaces.

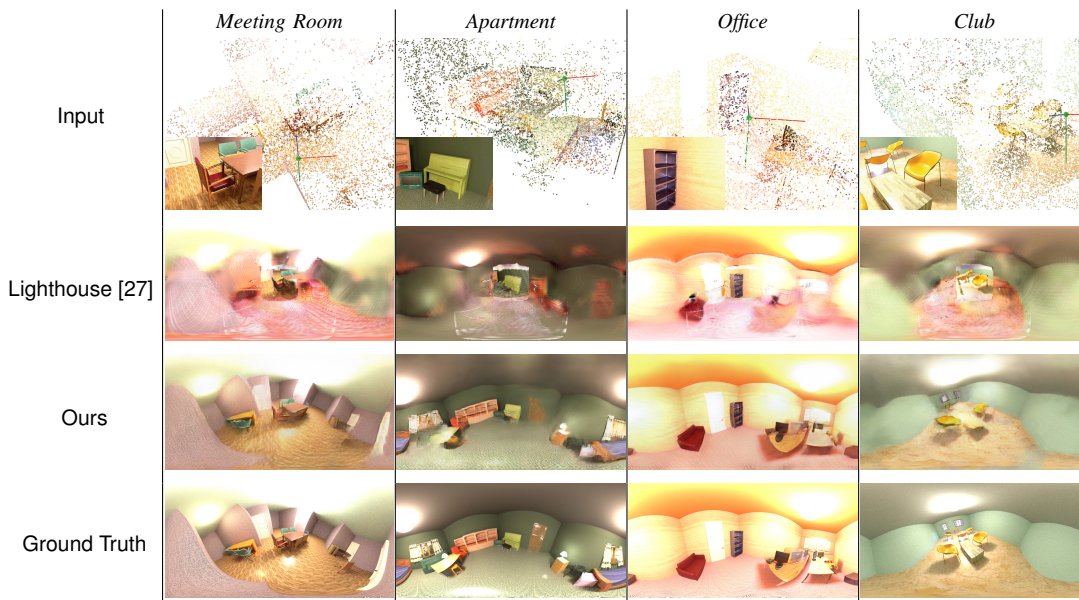


Fig. 7. Illumination estimation results on our SpecularRooms dataset. We list the input reference frames of Lighthouse and the input point cloud of our method in the first row. Our method retains more detailed information of a larger part of the scene.

TABLE II  
QUANTITATIVE RESULTS OF ROUGHNESS ESTIMATION ON OUR  
SPECULARROOMS DATASET

	IRCIS [23]	IRCIS+	Ours@5	Ours@10
Mean ( $\times 10^{-2}$ )	2.00	1.68	1.46	<b>1.26</b>
Median ( $\times 10^{-2}$ )	0.72	0.56	0.55	<b>0.50</b>

10 frames (Ours@10). IRCIS+ denotes IRCIS with extra supervision information, e.g., depth, normal, albedo, etc. Some representative results are available in Fig. 6. Compared to IRCIS, our method with multi-view inputs achieves better results, since material’s intrinsic properties are better demonstrated under varying viewpoints.

#### D. Illumination Estimation

**Methods for Comparison.** We compare our model with Lighthouse [27]. It uses a multiscale volumetric scene illumination representation, which can also produce spatial-consistent illumination results. For a fair comparison, we finetune the pre-trained model of Lighthouse on our SpecularRooms dataset with its recommended parameters.

**Metrics.** Following [27] and [32], we report color angular error and peak signal-to-noise ratio (PSNR) for the predictions versus ground truth environment maps.

**Results.** Quantitative results are shown in Table III. We also list some representative results in Fig. 7. Note that Lighthouse has a limited field of view, which means a large part of its prediction is from unseen areas and thus is unreliable. Our method performs better in most cases thanks

TABLE III  
QUANTITATIVE RESULTS OF ILLUMINATION ESTIMATION ON OUR  
SPECULARROOMS DATASET

	Angular Error ( $^{\circ}$ ) $\downarrow$		PSNR (dB) $\uparrow$	
	Lighthouse [27]	Ours	Lighthouse	Ours
<i>Meeting Room</i>	6.62	<b>2.44</b>	13.00	<b>19.93</b>
<i>Computer Room</i>	1.55	<b>1.09</b>	17.03	<b>19.22</b>
<i>Office</i>	8.58	<b>2.85</b>	13.54	<b>19.26</b>
<i>Club</i>	4.98	<b>1.67</b>	15.63	<b>21.84</b>
<i>Apartment Room</i>	6.58	<b>2.30</b>	11.27	<b>19.22</b>
Average	5.66	<b>2.07</b>	14.09	<b>19.89</b>

to the larger visibility of the scene from the point cloud.

## VI. CONCLUSIONS

In this paper, we introduce a novel PBA method for non-Lambertian environments. Our PBA method utilizes the material and illumination information to overcome the photometric inconsistency. The information is from our material and illumination estimation pipeline, which can effectively utilize multi-view input. Furthermore, we establish a new dataset to provide an evaluation framework for PBA methods in non-Lambertian environments. Extensive experiments demonstrated that our PBA method outperforms existing approaches in accuracy.

## REFERENCES

- [1] H. Li, J. Zhao, J.-C. Bazin, P. Kim, K. Joo, Z. Zhao, and Y.-H. Liu, "Hong Kong World: Leveraging structural regularity for line-based SLAM," *TPAMI*, 2023.
- [2] G. Wang, X. Wu, S. Jiang, Z. Liu, and H. Wang, "Efficient 3D deep LiDAR odometry," *TPAMI*, 2022.
- [3] T. Huang, H. Li, K. He, C. Sui, B. Li, and Y.-H. Liu, "Learning accurate 3D shape based on stereo polarimetric imaging," in *CVPR*, 2023.
- [4] H. Li, J. Zhao, J.-C. Bazin, and Y.-H. Liu, "Robust estimation of absolute camera pose via intersection constraint and flow consensus," *TIP*, 2020.
- [5] J. Engel, V. Koltun, and D. Cremers, "Direct sparse odometry," *TPAMI*, 2017.
- [6] R. Zabih and J. Woodfill, "Non-parametric local transforms for computing visual correspondence," in *ECCV*, 1994.
- [7] A. Crivellaro and V. Lepetit, "Robust 3D tracking with descriptor fields," in *CVPR*, 2014.
- [8] A. Dai, M. Nießner, M. Zollhöfer, S. Izadi, and C. Theobalt, "Bundlefusion: Real-time globally consistent 3D reconstruction using on-the-fly surface reintegration," *TOG*, 2017.
- [9] H. Alismail, M. Kaess, B. Browning, and S. Lucey, "Direct visual odometry in low light using binary descriptors," *RAL*, 2016.
- [10] P. Favaro, Hailin J., and S. Soatto, "A semi-direct approach to structure from motion," in *ICIAP*, 2001.
- [11] S. Klose, P. Heise, and A. Knoll, "Efficient compositional approaches for real-time robust direct visual odometry from RGB-D data," in *IROS*, 2013.
- [12] J. Engel, J. Stückler, and D. Cremers, "Large-scale direct SLAM with stereo cameras," in *IROS*, 2015.
- [13] G. Scandaroli, M. Meilland, and R. Richa, "Improving ncc-based direct visual tracking," in *ECCV*, 2012.
- [14] M. Meilland, A. Comport, and P. Rives, "Real-time dense visual tracking under large lighting variations," in *BMVC*, 2011.
- [15] T. Gonçalves and A. Comport, "Real-time direct tracking of color images in the presence of illumination variation," in *ICRA*, 2011.
- [16] C. Kerl, J. Sturm, and D. Cremers, "Robust odometry estimation for RGB-D cameras," in *ICRA*, 2013.
- [17] J. Zubizarreta, I. Aguinaga, and J. Montiel, "Direct sparse mapping," *TRO*, 2020.
- [18] R. Mur-Artal and J. Tardós, "ORB-SLAM2: An open-source SLAM system for monocular, stereo, and RGB-D cameras," *TRO*, 2017.
- [19] J. Barron and J. Malik, "Intrinsic scene properties from a single RGB-D image," in *CVPR*, 2013.
- [20] J. Barron and J. Malik, "Shape, illumination, and reflectance from shading," *TPAMI*, 2015.
- [21] C. Sun, G. Cai, Z. Li, K. Yan, C. Zhang, C. Marshall, J. Huang, S. Zhao, and Z. Dong, "Neural-PBIR reconstruction of shape, material, and illumination," in *ICCV*, 2023.
- [22] R. Zhu, Z. Li, J. Matai, F. Porikli, and M. Chandraker, "Irisformer: Dense vision transformers for single-image inverse rendering in indoor scenes," in *CVPR*, 2022.
- [23] Z. Li, M. Shafiei, R. Ramamoorthi, K. Sunkavalli, and M. Chandraker, "Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and SVBRDF from a single image," in *CVPR*, 2020.
- [24] H. Jin, I. Liu, P. Xu, X. Zhang, S. Han, S. Bi, X. Zhou, Z. Xu, and H. Su, "TensorIR: Tensorial inverse rendering," in *CVPR*, 2023.
- [25] Y. Zhang, J. Sun, X. He, H. Fu, R. Jia, and X. Zhou, "Modeling indirect illumination for inverse rendering," in *CVPR*, 2022.
- [26] Z. Li, L. Wang, M. Cheng, C. Pan, and J. Yang, "Multi-view inverse rendering for large-scale real-world indoor scenes," in *CVPR*, 2023.
- [27] P. Srinivasan, B. Mildenhall, M. Tancik, J. Barron, R. Tucker, and N. Snavely, "Lighthouse: Predicting lighting volumes for spatially-coherent illumination," in *CVPR*, 2020.
- [28] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, "The euroc micro aerial vehicle datasets," *IJRR*, 2016.
- [29] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of RGB-D SLAM systems," in *IROS*, Oct. 2012.
- [30] M. Gardner, K. Sunkavalli, E. Yumer, X. Shen, E. Gambaretto, C. Gagné, and J. Lalonde, "Learning to predict indoor illumination from a single image," *arXiv*, 2017.
- [31] S. Sengupta, J. Gu, K. Kim, G. Liu, D. Jacobs, and J. Kautz, "Neural inverse rendering of an indoor scene from a single image," in *ICCV*, 2019.
- [32] C. LeGendre, W. Ma, G. Fyffe, J. Flynn, L. Charbonnel, J. Busch, and P. Debevec, "Deeplight: Learning illumination for unconstrained mobile mixed reality," in *CVPR*, 2019.
- [33] F. Zhan, C. Zhang, W. Hu, S. Lu, F. Ma, X. Xie, and L. Shao, "Sparse needlets for lighting estimation with spherical transport loss," in *ICCV*, 2021.
- [34] Z. Li, L. Yu, M. Okunev, M. Chandraker, and Z. Dong, "Spatiotemporally consistent HDR indoor lighting estimation," *TOG*, 2023.
- [35] R. Cook and K. Torrance, "A reflectance model for computer graphics," *TOG*, 1982.
- [36] D. Levin, "The approximation power of moving least-squares," *Mathematics of computation*, 1998.
- [37] H. Li, J. Yao, J.-C. Bazin, X. Lu, Y. Xing, and K. Liu, "A monocular SLAM system leveraging structural regularity in Manhattan world," in *ICRA*, 2018.
- [38] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "SLIC superpixels compared to state-of-the-art superpixel methods," *TPAMI*, 2012.
- [39] H. Li, Y. Xing, J. Zhao, J.-C. Bazin, Z. Liu, and Y.-H. Liu, "Leveraging structural regularity of Atlanta world for monocular SLAM," in *ICRA*, 2019.
- [40] Y. Wang, Z. Xu, X. Wang, C. Shen, B. Cheng, H. Shen, and H. Xia, "End-to-end video instance segmentation with transformers," in *CVPR*, 2021.
- [41] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *TIP*, 2004.
- [42] H. Alismail, B. Browning, and S. Lucey, "Photometric bundle adjustment for vision-based SLAM," in *ACCV*, 2017.
- [43] C. Qi, L. Yi, H. Su, and L. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," *NeurIPS*, 2017.
- [44] Z. Song, W. Chen, D. Campbell, and H. Li, "Deep novel view synthesis from colored 3D point clouds," in *ECCV*, 2020.
- [45] H. Li, W. Chen, J. Zhao, J.-C. Bazin, L. Luo, Z. Liu, and Y.-H. Liu, "Robust and efficient estimation of absolute camera pose for monocular visual odometry," in *ICRA*, 2020.
- [46] G. Wang, C. Zhang, H. Wang, J. Wang, Y. Wang, and X. Wang, "Unsupervised learning of depth, optical flow and pose with occlusion from 3D geometry," *TITS*, 2020.