

GMLight: Lighting Estimation via Geometric Distribution Approximation

Fangneng Zhan, Yingchen Yu, Changgong Zhang, Rongliang Wu, Wenbo Hu, Shijian Lu*, Feiying Ma, Xuansong Xie, Ling Shao, *Fellow, IEEE*

Abstract—Inferring the scene illumination from a single image is an essential yet challenging task in computer vision and computer graphics. Existing works estimate lighting by regressing representative illumination parameters or generating illumination maps directly. However, these methods often suffer from poor accuracy and generalization. This paper presents Geometric Mover’s Light (GMLight), a lighting estimation framework that employs a regression network and a generative projector for effective illumination estimation. We parameterize illumination scenes in terms of the geometric light distribution, light intensity, ambient term, and auxiliary depth, which can be estimated by a regression network. Inspired by the earth mover’s distance, we design a novel geometric mover’s loss to guide the accurate regression of light distribution parameters. With the estimated light parameters, the generative projector synthesizes panoramic illumination maps with realistic appearance and high-frequency details. Extensive experiments show that GMLight achieves accurate illumination estimation and superior fidelity in relighting for 3D object insertion. The codes are available at <https://github.com/fnzhan/Illumination-Estimation>.

Index Terms—Lighting estimation, image synthesis, generative adversarial networks.

I. INTRODUCTION

ESTIMATING scene illumination from a single image has attracted increasing attention thanks to its wide spectrum of applications in image composition [1], object relighting [2] in mixed reality, etc. As the formation of image is influenced by several factors including reflectance, scene geometry and illumination, recovering illumination from a single image is a typical under-constrained problem. The problem becomes more challenging as high-dynamic-range (HDR) and full-view illumination map is expected to be inferred from low-dynamic-range (LDR) image with limited field of view.

Early works tackle this problem by utilizing auxiliary information such as scene geometry [4] and user annotation [5]. With the advancement of deep learning, recent studies [6], [7], [2] attempted to estimate illumination through direct generating illumination maps or regressing the representation parameters without utilizing auxiliary information. In particular, [2], [8], [9] aim to generate or render the illumination maps directly through neural networks. On the other hand, [10], [6] propose to recover the scene illumination by regressing the

F. Zhan, Y. Yu, R. Wu, and S. Lu are with the School of Computer Science and Engineering, Nanyang Technological University, Singapore. C. Zhang, F. Ma, and X. Xie are with the Alibaba DAMO Academy, China. W. Hu is with the Chinese University of Hong Kong. L. Shao is with the Inception Institute of Artificial Intelligence, UAE.

* indicates the corresponding author. Email: shijian.lu@ntu.edu.sg




Fig. 1. Illustration of scene geometry in EMLight [3] and the proposed GMLight. GMLight regresses the illumination distribution in real geometry spaces as defined by scene depth, while EMLight simplifies the scene geometry to be a spherical surface. Inspired by earth mover’s distance, we design a geometric mover’s loss to regress the illumination distribution by minimizing the total cost required to move the predicted lighting to ground truth (GT) lighting.

Spherical Harmonic (SH) parameters, while Spherical Gaussian (SG) representation is adopted for regression in [7], [11]. However, the regression-based methods [6], [12] often lack realistic lighting details, while the generation-based methods [2], [13] may incur inaccurate prediction of light properties (e.g., light directions) and suffer from poor generalization ability.

In our previous work [3], we designed EMLight which combines regression-based method and generation-based method for the prediction of environmental lighting. Specially, the scene illumination is decomposed into a discrete distribution defined on anchor points of a spherical surface (namely spherical distribution). Then a spherical mover’s loss (SML) is designed based on earth mover’s distance (EMD) [14] to regress the spherical distribution by treating the distance between anchor points as the moving cost of EMD. However, EMLight employs a simplified spherical surface to represent the illumination scene and compute the distance between anchor points without considering scene depths. As the formation of images are jointly determined by several intrinsic factors including the scene illumination, scene geometry, etc., oversimplification of the scene geometry will cause the lighting estimation from a single image to be unreliable. Besides, a simplified scene geometry also incapacitates the network in handling *spatially-varying illuminations* with the method

described in [7]. In this work, we propose Geometric Mover’s Light (GMLight) to model light distributions in a geometric space (namely geometric distributions) as indicated by scene depth, which is closer to real-world light distributions compared with the simplified spherical distribution in EMLight as shown in Fig. 1.

The proposed GMLight consists of a regression network for light parameters prediction and a generative projector for illumination map synthesis. For light parameters estimation, we propose a geometric distribution representation method to decompose illumination scene into four components: *depth value*, *light distribution*, *light intensity*, *ambient term*. Note that the last two are scalars and can be directly regressed with a naive L2 loss. Light distributions and scene depths, in contrast, are spatially-localized in scenes, and thus are not suitable to be regressed by a naive L2 loss which does not capture geometry information or SML in EMLight [3] which simplifies the scene geometry. In this work, we design a Geometric Mover’s Loss (GML), that regresses light distributions and scene depths with an ‘earth mover’ in a geometric space as indicated by the scene depth. GML aims to search for an optimal plan to move one distribution to another with the minimal total cost which is defined by the moving distance in the scene. With known scene depth, the depth value and light distribution can be effectively regressed by GML with consideration of the scene geometry.

With the estimated illumination scene parameters, the generative projector generates illumination maps with realistic appearance and details in an adversarial manner. In detail, Spherical Gaussian function [7] is adopted to map the estimated light parameters into a Gaussian map (panoramic image), which serves as the guidance in illumination map generation. The Gaussian map can be reconstructed at each position in a scene with the knowledge of scene depth, thus achieving the estimation of *spatially-varying* illuminations. Since illumination maps are panoramic images that usually suffer from spherical distortions at different latitudes, we adopt spherical convolutions [15] to accurately generate panoramic illumination maps. To provide progressive guidance in the illumination generation process, an adaptive radius strategy is designed to generate progressive Gaussian maps in a coarse-to-fine manner. More details will be provided in Section III-C.

Compared with our previous work EMLight [3], the contributions of this work can be summarized in three aspects. First, we propose a geometric distribution representation to parameterize the illumination scene for lighting estimation and enable effective estimation of *spatially-varying* illumination. Second, we introduce a novel geometric mover’s loss that leverages the real scene geometry through depth values in the regression of light distribution. Third, we design a generative projector with progressive guidance in illumination generation process.

The rest of this paper is organized as follows. Section II presents related works. The proposed method is then described in detail in Section III. Experimental results are further presented and discussed in Section IV. Finally, concluding remarks are drawn in Section V.

II. RELATED WORKS

Lighting estimation targets to predict HDR illumination from a single image, which has been widely applied in relighting for object insertion [17], [18], [19], [20], [21] and image composition [22], [1]. The earlier works in this field heavily rely on auxiliary information to tackle this problem. In particular, the scene is typically decomposed into geometry, reflectance, and illumination, and then the lighting is estimated with known scene geometry or reflectance. For instance, Karsch *et al.* [5] acquire scene geometry through user annotations for lighting estimation. Maier *et al.* [23] employ additional depth information to estimate Spherical Harmonics (SH) representation of the illumination. Moreover, Zhang *et al.* [4] utilize a multi-view 3D reconstruction of scenes to recover illumination, while Lombardi and Nishino [24] achieve illumination estimation by incorporating a low-dimensional reflectance model based on the objects with known shapes. Sengupta *et al.* [25] jointly estimate the albedo, normals, and lighting of an indoor scene from a single image and Barron *et al.* [26] estimate shape, lighting, and material but rely on data-driven priors to compensate for the lack of geometry information.

Thanks to the tremendous success of deep learning, data-driven approaches have become prevalent in recent years and demonstrated the feasibility of estimating lighting without auxiliary information. They roughly fall into two main categories: 1) regression-based methods that regress the lighting representation parameters [10], [7], [11]; 2) generation-based methods [2], [8] that directly generate illumination maps with neural networks [27], [28], [29]. For the regression-based methods, Cheng *et al.* [10] and Garon *et al.* [6] represent the illumination with spherical harmonic (SH) parameters and predict the scene illumination by regressing the corresponding parameters. On the other hand, Gardner *et al.* [7] parameterize the scene illumination into the light directions, light intensities, and light color with spherical Gaussian functions to enable the explicit regression of individual light source in the scene. In the aspect of the generation-based method, Gardner *et al.* [2] employ a two-step training strategy to generate illumination maps. Song *et al.* [8] utilize a convolutional network to predict unobserved content in the environment map with predicted per-pixel 3D geometry. Legendre *et al.* [30] generate the illumination maps by matching the LDR ground-truth sphere images to those rendered with the predicted illumination using image-based relighting. Besides, Srinivasan *et al.* [9] utilize volume rendering to generate incident illuminations according to a 3D volumetric RGB model of the scene. Instead of predicting the illumination maps, Sun *et al.* [31] propose a framework to achieve relighting on RGB portrait images given any illumination maps. Moreover, several works [32], [1] adopt generative adversarial network to generate shadows on RGB images directly without explicitly estimating the illumination map.

The previous lighting estimation works only adopt either the regression-based method or generation-based method alone, which tends to lose realistic lighting details or predict inaccurate light properties (e.g., light directions) respectively. In

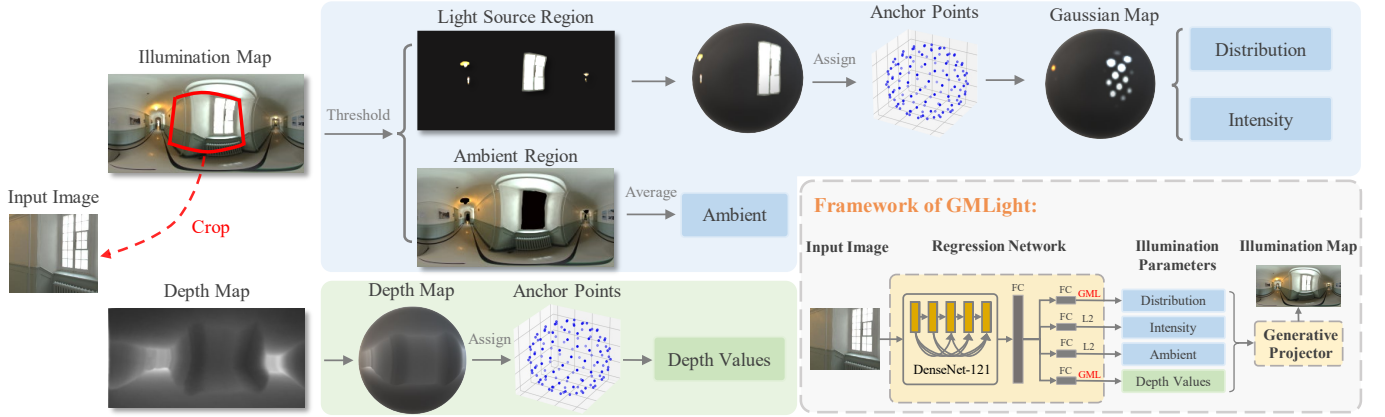


Fig. 2. Illustration of the geometric distribution representation to acquire training data and the overall framework of GMLight. Given an *Illumination Map*, we first determine a *Light Source Region* and *Ambient Region* via thresholding. The pixels in ambient region are averaged to obtain the ambient term and pixels in light source region are assigned to N anchor points to obtain a light intensity and a light distribution as visualized by the *Gaussian Map*. Scene depths (dataset provided) are also assigned to the N anchor points as ground truth. The illumination map (scene) is thus decomposed to representative parameters including light distribution, light intensity, ambient term, and depth. With the image cropped from the illumination map as input, the regression network in GMLight aims to regress the representative parameters, followed by a generative projector to generate the final illumination map. FC denotes fully-connected layer, GML denotes the proposed geometric mover's loss.

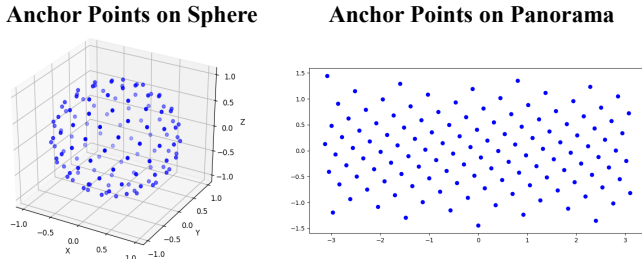


Fig. 3. Visualization of 128 anchor points on a unit sphere and panorama. (pre-defined with the method described in [16])

contrast, the proposed GMLight combines the merits of the regression-based and generation-based methods for accurate yet realistic illumination estimation.

III. PROPOSED METHOD

The proposed GMLight consists of a regression network and a generative projector that are jointly optimized. A geometric distribution representation is proposed to parameterize illumination scenes with a set of parameters which are to be estimated by the regression network. The predicted illumination parameters serve as guidance for the generative projector to synthesize realistic yet accurate illumination maps.

A. Geometric Distribution for Illumination Representation

As illumination maps are high-dynamic-range (HDR) images, the light intensity of different illumination maps may vary drastically, which is deleterious for the parameter regression. We thus propose to normalize the light intensity of illumination maps and design a novel geometric distribution method for illumination representation.

As illustrated in Fig. 2, the geometric distribution illumination representation decomposes the scene illumination into four parameters including light distribution P , light intensity

I , ambient term A , and depth D . With a HDR illumination map, we first separate the light sources from the full image since the light sources in a scene play the most critical role for relighting. Consistent with [7], the light sources are selected as top 5% of pixels with the highest values in the HDR illumination map. The light intensity I and the ambient term A can then be determined by the sum of all light-source pixels and the average of remaining pixels, respectively. To formulate discrete distributions, we generate N anchor points (using the method described in [16]) that are uniformly distributed on a unit sphere as illustrated in Fig. 3. Light-source pixels are assigned to their corresponding anchor point based on the minimum radian distance, and further determine the light-source value of the anchor point by summing all its affiliated pixels. Afterward, the value of light-source pixels is normalized by the intensity I so that the N anchor points on the unit sphere form a discrete distribution (i.e. the light distribution P). Since Laval Indoor HDR dataset [2] provides pixel-level depth annotations, the depth D at each anchor point can be determined by averaging depth values of all the pixels that are assigned to the anchor point (similarly based on the minimum radian distance).

With a limited-view image cropped from the illumination map to serve as the network input, the illumination parameters and the original illumination map serve as the ground truth for the regression network and generative projector, respectively.

B. Regression of Illumination Parameters

In the regression network, four branches with shared backbone network are adopted to regress the four sets of parameters, respectively. The light intensity I and ambient term A are scalars which can be regressed with a naive L2 loss. However, the light distribution P and depth values D are vectors of N dimension that are spatially localized on a sphere, and the naive L2 loss cannot effectively utilize the proximity of the geometric distribution and the property of the standard

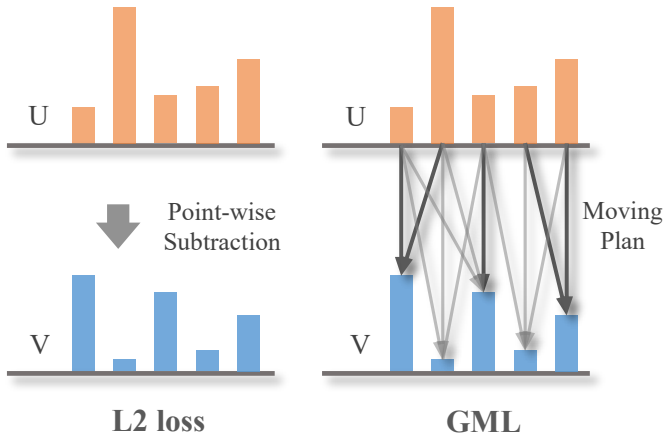


Fig. 4. Comparison between L2 loss and the proposed geometric mover's loss (GML) for computing the distance between distributions U and V . L2 loss is computed through a point-wise subtraction which cannot leverage the proximity of geometric distribution. Instead, the geometric mover's loss aims to find the moving plan of minimal cost with consideration of the scene geometry which is correlated with the moving cost. The thickness of connecting lines denotes the amount of probability moved between two points.

distribution (for light distribution P , the summation of all anchor point values is equal to one). Therefore, we propose a geometric mover's loss based on earth mover's distance [14] to effectively measure the discrepancy between distributions with consideration of the scene geometry as shown in Fig. 4.

Geometric Mover's Loss: Our previous work EMLight proposes a spherical earth mover's loss (SML) to achieve geometry-aware regression. However, the SML over-simplifies the scene geometry to be a spherical surface which ignores the complex geometry of real scenes. As an image is produced as a result of the scene illumination, scene geometry, etc., ignoring the scene geometry will deteriorate the inverse process of inferring illumination from a single image. We thus propose a geometric mover's loss (GML) which effectively leverages the real scene geometry through depth values. To derive the proposed GML, two discrete distributions U and V with N points are defined as shown in Fig. 4. According to earth mover's distance, GML can be defined as the minimum total cost to convert distribution U to distribution V , where the cost is measured by the product of the amounts of value (or 'earth') to be moved and the distance to be moved. Thus, a moving plan matrix T and a cost matrix C both of size (N, N) are defined, where entries T_{ij} and C_{ij} denote the amounts of moved 'earth' between point U_i and point V_j and the unit cost of moving U_i to V_j , respectively. In our previous spherical mover's distance, the unit cost between point U_i and point V_j is measured by their radian distance along the unit sphere as shown in Fig. 5, which is unable to take advantage of the real scene geometry in regression. We thus propose to facilitate the scene geometry through a geometric distance as determined by the scene depth. As shown in Fig. 5, the distance of a GMLight anchor point to the center is measured by a depth value D instead of the radius, thus effectively reflecting the real geometry of the scene. The geometric distance (or unit cost) between anchor points U_i and V_j can be computed

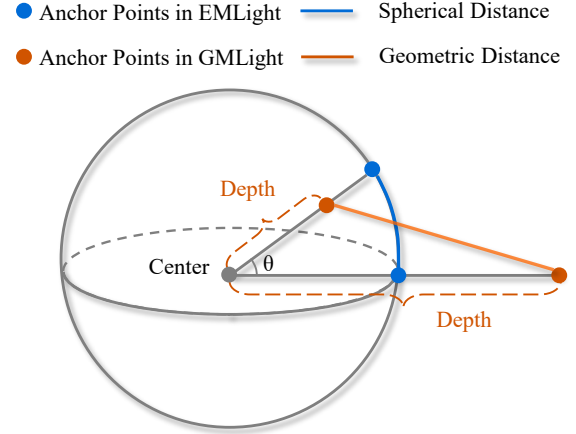


Fig. 5. Geometric distances in GMLight and spherical distances in EMLight. Spherical distances in EMLight assume that anchor points are distributed on a spherical surface which neglects the complex geometry of real scenes. Geometric distances in GMLight instead capture real geometries of anchor points by using scene depths, i.e., the distance from the anchor point to the spherical center. Since the spherical angle θ between anchor points is known [16], geometric distances between anchor points can be computed.

according to their depth values D_i , D_j and their spherical angle θ , as follows:

$$C_{ij} = D_i^2 + D_j^2 - 2 D_i D_j \cos(\theta). \quad (1)$$

With the defined transportation plan matrix T and cost matrix C , GML can be formulated as the minimum total cost for the transport between U and V :

$$GML = \min_T \left(\sum_{i=1}^N \sum_{j=1}^N C_{ij} T_{ij} \right) = \min_T \langle C, T \rangle \quad (2)$$

subject to $T \cdot \vec{1} = U, \quad T^\top \cdot \vec{1} = V.$

On the other hand, different from the illumination distributions, the depth values of anchor points don't form a standard distribution (the sum of all depth values is not constant). Thus, we introduce an unbalanced setting of GML for the regression of depth values. Specifically, it is handled by introducing a relaxed version of classical earth mover's distance, namely unbalanced earth mover's distance [33]. It aims to determine an optimal transport plan between measures (depth values) of different total masses. We formulate unbalanced GML by replacing the 'Hard' conservation of masses in (2) by a 'Soft' penalty with a divergence metric. An unbalanced GML as denoted by GML^u can thus be formulated as follows:

$$GML^u = \min_T \left[\langle C, T \rangle + \text{KL}(T\vec{1}|U) + \text{KL}(T^\top\vec{1}|V) \right] \quad (3)$$

where KL is the Kullback-Leibler divergence which is defined as:

$$\text{KL}(a||b) = \sum_{i=1}^n a_i \log\left(\frac{a_i}{b_i}\right) - a_i + b_i. \quad (4)$$

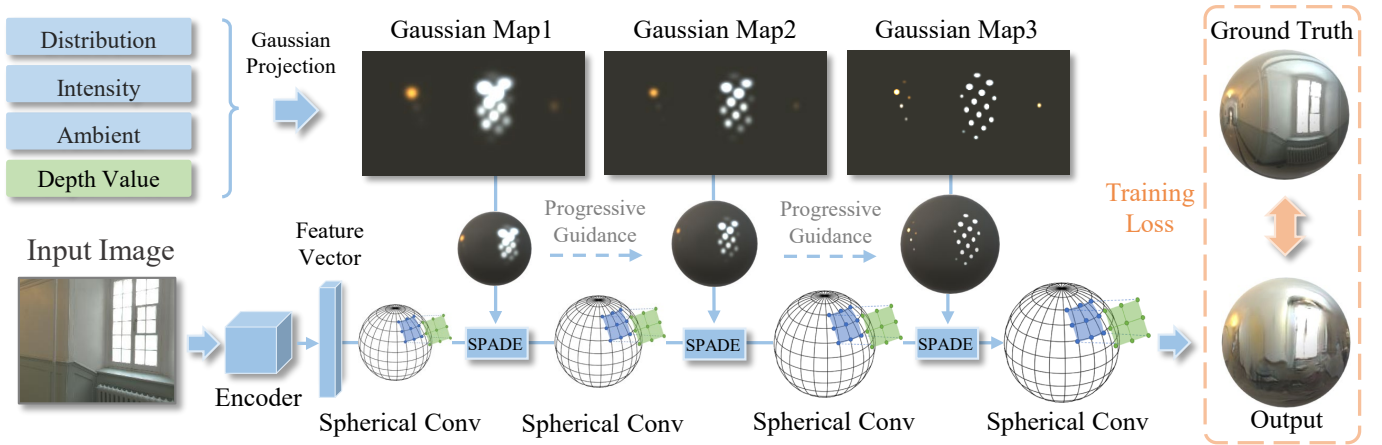


Fig. 6. Structure of the generative projector. *Input Image* is fed to an *Encoder* to produce a *Feature Vector* for the ensuing multi-stage spherical convolution. Progressive guidance (e.g., *Gaussian Map1*, *Gaussian Map2*, *Gaussian Map3*) are acquired from the regressed parameters (distribution, intensity, ambient, depth) via Gaussian projection with different radius parameters, and are injected into the multi-stage generation process to synthesize the *Output* illumination map. SPADE denotes spatially adaptive normalization [27] for feature injection.

Grounded in well-studied entropic regularization [34] for differentiable optimization, the entropic version of GML (2) can be formulated as below:

$$\begin{aligned} GML &= \min_T \langle C, T \rangle - \epsilon H(T) \\ &\text{subject to } T \cdot \vec{1} = U, \quad T^\top \cdot \vec{1} = V, \end{aligned} \quad (5)$$

where $H(T)$ is the entropic term as defined by $H(T) = -\sum_{i=1}^N \sum_{j=1}^N T_{ij} \log T_{ij}$, ϵ is the regularization coefficients denoting the smoothness of the transportation plan matrix T . In our model, ϵ is empirically set to 0.0001. The unbalanced GML (3) can be regularized similarly by adding the entropic term $H(T)$. Then the entropic version of the problem (2) and (3) can be solved in a differentiable yet efficient way through Sinkhorn iteration [34] in training.

The proposed geometric mover's loss makes the regression sensitive to the global geometry which effectively penalizes the spatial discrepancy between the predicted distribution and ground-truth distribution. Besides, the geometric mover's loss is smoother than L2 loss which enables more stable optimization in training process.

C. Illumination Generation

The regressed light parameters indicate accurate light properties including the light intensity, ambient and light distributions. To project these parameters into illumination maps, we propose a novel generative projector based on generative adversarial networks (GANs) [35] as illustrated in Fig. 6. Specifically, the light distribution P , light intensity I , and ambient term A are firstly projected to Gaussian map M through spherical Gaussian function [7] as follows:

$$M = \sum_{i=1}^N v_i * \exp \frac{o_i * u - 1}{s} + A, \quad (6)$$

where N is the number of anchor points, $v_i = P_i * I$ and o_i denote the RGB value and the direction of an anchor point respectively, u is a unit vector on the sphere and s is the angular size which is set to 0.0025 empirically. Then the

illumination generation can be formulated as an image-to-image translation task conditioned on the Gaussian map.

Fig. 6 illustrates the image translation process, where the input image is encoded into a feature vector to serve as the network input. To provide effective guidance in the generation process, the Gaussian map is injected into generation process through a spatially adaptive normalization (SPADE) [27] as shown in Fig. 6. As the illumination map is a panoramic image where pixels are stretched at different latitudes, the vanilla convolution suffers from heavy distortions around the polar regions of the illumination map. To address this, Benjamin *et al.* [15] propose to encode the invariance against latitude distortions into convolutional neural networks by adjusting the locations of convolution operations, as known as spherical convolution. We thus employ the spherical convolution (Spherical Conv) in the generative projector to synthesize panoramic illumination map which effectively mitigates the distortions of different latitudes on the illumination map.

In EMLight [3], the same conditional Gaussian map is injected into the generation process at different levels. Ideally, coarse Gaussian map guidance is expected in low generation layers to learn the overall illumination condition, while fine Gaussian map is expected in the high layers to indicate the accurate illumination distribution. Thus, we design an adaptive radius strategy in the Gaussian function to generate coarse-to-fine Gaussian maps for different generation layers, as shown in 'Gaussian Map1', 'Gaussian Map2', and 'Gaussian Map3' in Fig. 6.

Spatially-Varying Projection: Spatially-varying illumination prediction aims to recover the illumination of different positions in a scene from a single image. As there are no annotations for spatially-varying illumination in the Laval Indoor HDR dataset, we are unable to train a spatially varying model directly. Previous research [7] proposed to incorporate depth values into the projection to approximate the effect of spatially-varying illumination. We follow a similar idea with [7] to achieve the estimation of spatially varying illumination,

as described below.

The Gaussian map is constructed through a spherical Gaussian function, as shown in Eq. (6). When we move the insertion position by ∇o , the new direction of the anchor point i can be denoted by $o_i + \nabla o$. The depth of the original insertion position and the new position are l_i and $l_i + \nabla l$, which can be obtained from the predicted depth value of N anchor points. The light intensity at the new insertion position can thus be approximated by $v_i * \frac{l_i}{l_i + \nabla l}$, and the Gaussian map M_{∇} of the new insertion position can be constructed as follows:

$$M_{\nabla} = \sum_{i=1}^N v_i \left(\frac{l_i}{l_i + \nabla l} \right) * \exp \frac{(o_i + \nabla o) * u - 1}{s} + A \quad (7)$$

The Gaussian map is then fed into the generative projector to synthesize the final illumination map. Fig. 8 illustrates several samples of predicted Gaussian maps, generated illumination maps, visualized intensity maps and the corresponding ground truth. Fig. 11 illustrates the generated spatially-varying illumination maps at different insertion positions.

Loss Functions: Several losses are employed in the generative projector to yield realistic yet accurate illumination maps. For clarity, the input Gaussian map, ground-truth illumination map, and generated illumination map are denoted by x , y , and x' , respectively.

To synthesize realistic illumination maps, a discriminator with Patch-GAN [36] structure is included to impose an adversarial loss as denoted by \mathcal{L}_{adv} . Different from the vanilla GAN discriminator which discriminates real from faked at image level, Patch-GAN discriminator has a fully convolutional architecture and only conduct the discrimination at the scale of local patches. Image discrimination results across all the patches are averaged to provide the ultimate output of discriminator. The Patch-GAN achieves better synthesis quality and detail generation in image translation tasks [36] compared with vanilla GAN discriminator.

Then a feature matching loss \mathcal{L}_{feat} is introduced to stabilize the training by matching the intermediate features of the discriminator between the generated illumination map and the ground truth:

$$\mathcal{L}_{feat} = \sum_l \lambda_l \|D_l(x, x') - D_l(x, y)\|_1, \quad (8)$$

where D_l represents the activation of layer l in the discriminator and λ_l denotes the balanced coefficients. Targeting to yield similar illumination distribution regardless of the absolute intensity, a cosine similarity is employed to measure the distance between the generated illumination map and ground truth as follows:

$$\mathcal{L}_{cos} = 1 - \text{Cos}(x', y), \quad (9)$$

Then, the generative projector is optimized under the following objective:

$$\mathcal{L} = \min_G \max_D (\lambda_1 \mathcal{L}_{feat} + \lambda_2 \mathcal{L}_{cos} + \lambda_3 \mathcal{L}_{adv}). \quad (10)$$

where λ balances the loss terms. As the regression network and generative projector are both differentiable, the whole framework can be optimized end-to-end.

IV. EXPERIMENTS

A. Datasets and Experimental Settings

We benchmark GMLight with other SOTA lighting estimation models on the Laval Indoor HDR dataset [2]. To acquire paired data for model training, we crop eight images with limited fields of view from each panorama in the Laval Indoor HDR dataset, which finally produces 19,556 training pairs for our experiments. Specifically, the image warping operation as described in [2] is applied to each panorama to mimic the light locality for indoor scenes. In our experiments, 200 images are randomly selected as the testing set and the rest images are used for training. In addition to the Laval Indoor HDR dataset, we also qualitatively evaluate GMLight on the dataset¹ introduced in [6].

Following [7] and [6], we use DenseNet-121 [37] as backbone in regression network. The default size N of the discrete light distribution is 128. The sizes of the input image and reconstructed illumination map are 192 and 128×256 . The detailed network structure of the generative projector is provided in the supplementary material. We implemented GMLight in PyTorch and adopted the Adam algorithm [38] as the optimizer that employs a learning rate decay mechanism (the initial learning rate is 0.001). The network is trained on two NVIDIA Tesla P100 GPUs with a batch size of 4 for 100 epochs.



Fig. 7. The scenes in evaluations are spheres of three types of materials including diffuse gray, matte silver, and mirror silver.

B. Evaluation Method and Metrics

For accurate and comprehensive evaluation of the model performance, we create a virtual scene consisting of three spheres made of gray diffuse, matte silver and mirror silver as illustrated in Fig. 7. By comparing the sphere images rendered (by Blender [40]) with the ground-truth illumination and the predicted illuminations, the illumination estimation performance can be effectively measured. Several metrics are employed to compare the rendered images including root mean square error (RMSE) which mainly evaluates the accuracy of light intensity, scale-invariant RMSE (si-RMSE) and RGB angular error [30] which mainly evaluate the accuracy of light directions. Besides, we also perform crowdsourcing user study through Amazon Mechanical Turk (AMT) to assess the realism of images rendered with illumination maps predicted by different methods. Two images rendered by the ground truth and each compared method are shown to 20 users who

¹<https://lvsn.github.io/fastindoorlight/>

TABLE I

COMPARISON OF GMLIGHT WITH SEVERAL STATE-OF-THE-ART LIGHTING ESTIMATION METHODS. THE EVALUATION METRICS INCLUDE THE WIDELY USED RMSE, SI-RMSE, ANGULAR ERROR, USER STUDY AND GMD. D, S, AND M DENOTE DIFFUSE, MATTE SILVER, AND MIRROR MATERIALS OF THE RENDERED OBJECTS, RESPECTIVELY.

Metrics	RMSE ↓			si-RMSE ↓			Angular Error ↓			User Study ↑			GMD ↓
	D	S	M	D	S	M	D	S	M	D	S	M	N/A
Gardner <i>et al.</i> [2]	0.146	0.173	0.202	0.142	0.151	0.174	8.12°	8.37°	8.81°	28.0%	23.0%	20.5%	6.842
Gardner <i>et al.</i> [7]	0.084	0.112	0.147	0.073	0.093	0.119	6.82°	7.15°	7.22°	33.5%	28.0%	24.5%	5.524
Li <i>et al.</i> [12]	0.203	0.218	0.257	0.193	0.212	0.243	9.37°	9.51°	9.81°	25.0%	21.5%	17.5%	7.013
Garon <i>et al.</i> [6]	0.181	0.207	0.249	0.177	0.196	0.221	9.12°	9.32°	9.49°	27.0%	22.5%	19.0%	7.137
EMLight [3]	0.062	0.071	0.089	0.043	0.054	0.078	6.43°	6.61°	6.95°	40.0%	35.0%	25.0%	5.131
NeedleLight [39]	0.072	0.074	0.091	0.051	0.062	0.084	6.61°	6.78°	7.04°	43.0%	33.5%	29.0%	4.213
GMLight	0.051	0.064	0.078	0.037	0.049	0.074	6.21°	6.50°	6.77°	42.0%	35.5%	31.0%	4.892



Fig. 8. Illustration of GMLight illumination estimation. For the *Input Images* in column 1, columns 2, 3, and 4 show the estimated *Gaussian Maps* (based on the regressed illumination parameters), *Lighting Maps* and *Intensity Maps*, respectively. Columns 5 and 6 show the ground truth of the corresponding intensity maps and illumination maps (i.e. *Lighting Maps (GT)* and *Intensity Maps (GT)*), respectively.

are asked to pick a more realistic image. The score is the percentage of rendered images that are deemed more realistic than the ground-truth rendering. The testing scene of three spheres is mainly used for *quantitative evaluation*. We also design 25 virtual 3D scenes on the testing set for object insertion to evaluate the *qualitative performance* in various scenes.

Besides, we introduce a Geometric Mover’s distance (GMD) based on the geometric mover’s loss as described in (2) to measure the discrepancy between light distributions of illumination maps as below:

$$GMD = \min_T [\langle C, T \rangle], \quad T \cdot \vec{1} = U, \quad T^\top \cdot \vec{1} = V, \quad (11)$$

where U and V are the normalized illumination maps, and the pixels in the maps form geometric distributions. The GMD metric is sensitive to the scene geometry with the cost matrix C , thus achieving a more accurate evaluation of illumination distribution (or directions) compared with si-RMSE.

C. Quantitative Evaluation

We conduct quantitative comparison between GMLight and other SOTA lighting estimation methods as shown in Table I. Specifically, each compared method predicts 200 illumination maps from the testing set to render the testing scene (three spheres of different materials). The experimental results is

tabulated in Table I, where D , S and M denote diffuse, matte silver and mirror material objects, respectively. As can be observed in Table I, GMLight consistently outperforms all compared methods under different evaluation metrics and materials. EMLight [3] simplifies the light distribution of scenes to be spherical, ignoring the complex scene geometry. GMLight introduces depth to model scene geometry which leads to more accurate illumination estimation. Gardner *et al.* [2] generate illumination maps directly, but it tends to overfit training data due to the unconstrained nature of illumination estimation from a single image. Gardner *et al.* [7] regress the spherical Gaussian parameters of light sources, but it often loses useful frequency information and generates inaccurate shading and shadows. Li *et al.* [12] adopt spherical Gaussian functions to reconstruct the illumination maps in the spatial domain but it often loses high-frequency illumination. Garon *et al.* [6] recover lighting by regressing spherical harmonic coefficients, but their model struggles to regress light directions and recover high-frequency information. Although Garon *et al.* [6] adopt a masked L2 loss to preserve high-frequency information, it does not fully solve the problem as illustrated in Fig. 9. NeedleLight [39] achieves lighting estimation in both spatial and frequency domains, while it is unable to recover illumination maps with fine details as it doesn’t employ a generative projector. In contrast, GMLight firstly regresses the

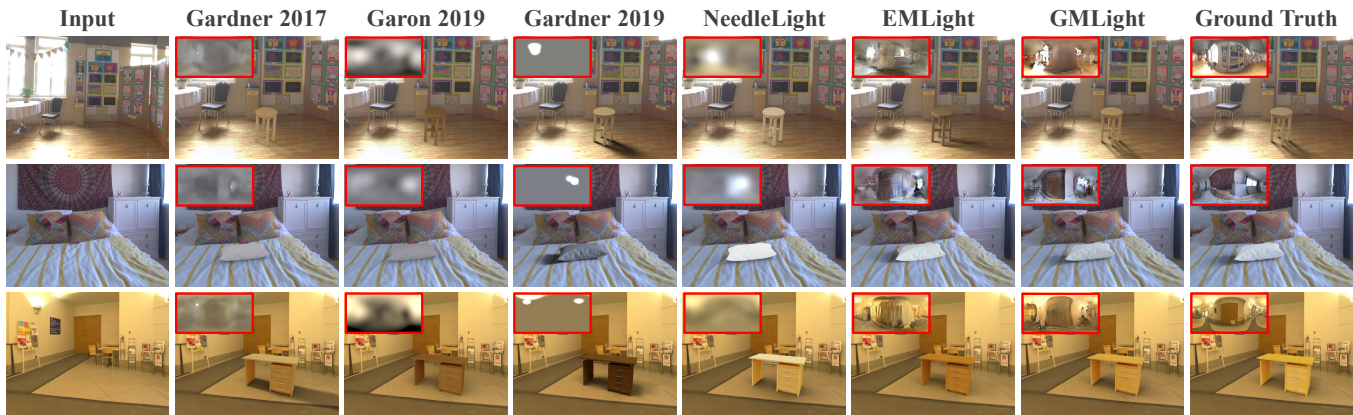


Fig. 9. Qualitative comparison of GMLight with state-of-the-art methods: With illumination maps predicted by different methods (shown at the top-left corner of each rendered image), the rendered objects demonstrate different light intensities, colors, shadows, and shades.



Fig. 10. Object relighting over sample images from the Internet. For sample images in Column 1, our GMLight estimates illumination maps (highlighted by red boxes) automatically and applies them to relight virtual objects realistically as shown in Column 2.

accurate light parameters with a regression network, followed by a generative projector to synthesize realistic yet accurate illumination maps under guidance of the regressed parameters.

D. Qualitative Evaluation

To demonstrate that the regressed parameters provide accurate guidance for the generative projector, we visualize the Gaussian maps, lighting maps, and intensity maps in GMLight and the corresponding ground truth in Fig. 8. As shown in Fig. 8, the predicted Gaussian maps indicate the light distributions in the scenes accurately, which enables the generative projector to synthesize illumination maps with

accurate light directions. We also visualize the intensity maps of the predicted lighting maps, and the plausible intensity map compared with the ground truth demonstrate that the generative projector synthesizes accurate HDR lighting maps.

We conduct qualitative comparison between GMLight and four SOTA lighting estimation methods on object insertion task as shown in Fig. 9. Specifically, illumination maps are predicted from the input images and the inserted objects are rendered with the predicted illumination maps. As illustrated in Fig. 9, the illumination maps predicted by GMLight present realistic and accurate light sources with fine details, thus enabling to render object with plausible shading and shadow effects. On the other hand, Gardner *et al.* [2] generate the HDR illumination maps directly which make it hard to synthesize accurate light sources; Gardner *et al.* [7] and Li *et al.* [12] regress the representative Gaussian parameters to reconstruct the illumination maps which however causes unrealistic illumination maps with losing details especially frequency information. Garon *et al.* [6] and Zhan *et al.* [39] regress the representation parameters to recover the scene illumination but are often constrained by the limited order of representation basis, which may incur low-frequency illumination that produces weak shading and shadows effects.

Besides the testing set, we also conduct object insertion on wild images collected from the Internet as shown in Fig. 10. The accurate illumination maps and realistic rendering results demonstrate the exceptional generalization capability of the proposed method.

E. Spatially-varying Illumination

Spatially-varying illumination prediction aims to recover the illumination at different positions of a scene. Fig. 11 shows the spatially-varying illumination maps that are predicted at different insertion positions (center, left, right, up, and down) by GMLight. It can be seen that GMLight estimates illumination maps of different insertion positions nicely, largely due to the auxiliary depth branch that estimates scene depths and recovers the scene geometry accurately. We also evaluate spatially-varying illumination estimation quantitatively and Table II




Fig. 11. GMLight is capable of estimating spatially-varying illumination from different insertion positions (e.g. center, left, right, up, and down as illustrated). The samples are from the Laval Indoor HDR dataset [2].



Fig. 12. Illustration of spatially-varying illumination estimation by GMLight: Images in column 1 are background images [6] for 3D object insertion, where the silver sphere in each image shows real illuminations at one specific position. Images in column 2 show object relighting at different insertion positions. With GMLight-estimated illuminations, the relighting at different positions is well aligned with that of the silver spheres.

shows experimental results. We can see that GMLight outperforms other methods consistently in all insertion positions. The superior performance is largely attributed to the accurate geometry modeling of light distributions with scene depths.

Fig. 12 illustrates the 3D insertion results with the estimated spatially-varying illuminations. The sample images are from [6], where a silver sphere is employed to indicate spatially-varying illuminations at different scene positions which serve as references for evaluating the realism of 3D insertion. As Fig. 12 shows, the inserted objects (clock) at different positions present consistent shading and shadow effect with the silver sphere, demonstrating the high-quality estimation of spatially-

TABLE II
QUANTITATIVE COMPARISON OF SPATIALLY-VARYING ILLUMINATION ESTIMATED BY GMLIGHT AND STATE-OF-THE-ART METHODS: EVALUATIONS WERE PERFORMED FOR THREE INSERTION POSITIONS (LEFT, CENTER, AND RIGHT) OVER IMAGES FROM THE LAVAL INDOOR HDR DATASET [2]. THE SCORES OF THE THREE SPHERES IN FIG. 7 ARE AVERAGED TO OBTAIN THE FINAL SCORE.

Methods	RMSE ↓			si-RMSE ↓		
	Left	Center	Right	Left	Center	Right
Gardner <i>et al.</i> [2]	0.168	0.176	0.171	0.148	0.159	0.152
Gardner <i>et al.</i> [7]	0.102	0.114	0.104	0.085	0.097	0.087
Garon <i>et al.</i> [6]	0.186	0.199	0.182	0.174	0.184	0.173
GMLight	0.059	0.066	0.058	0.043	0.051	0.041

varying illuminations by GMLight.

F. Ablation Study

We developed several GMLight variants as listed in Table III to evaluate the effectiveness of the proposed designs. The baseline model is selected as SG+L2 which regresses spherical Gaussian parameters with L2 loss. GD+L2, GD+SML, and GD+GML denote regress the geometric distribution of illumination with L2 loss, SML in EMLight [3], and GML in GMLight, respectively. GD+GML+GP denotes the standard GMLight with all proposed designs. All variant models are employed to render 200 images of the testing scene (3 spheres) and the rendering results are evaluated by various metrics including the RMSE, si-RMSE, Angular Error, User Study, and the proposed GMD as shown in Table III. It can be observed that GD+L2 outperforms SG+L2 clearly which demonstrates the effectiveness of geometric distributions for lighting representation. The superiority of the proposed GML is also verified as GD+GML produces better estimation than GD+L2 and GD+SML. With the including of the generative projector (GP), the performance across all metrics is improved by a large margin which demonstrates that generative projector improves illumination realism and accuracy significantly.

We also ablate the effect of different number of anchor points as shown in Table IV. It should be noted that the experimental results in RMSE and si-RMSE are averaged

TABLE III

ABLATION STUDY OF THE PROPOSED GMLIGHT. SG AND GD DENOTE SPHERICAL GAUSSIAN REPRESENTATION AND OUR PROPOSED GEOMETRIC DISTRIBUTION REPRESENTATION OF ILLUMINATION MAPS. L2 AND GML DENOTE L2 LOSS AND GEOMETRIC MOVER’S LOSS THAT ARE USED IN THE REGRESSION OF LIGHT PARAMETERS. GP DENOTES OUR PROPOSED GENERATIVE PROJECTOR. ‘GD+GML+GP’ DENOTES THE STANDARD GMLIGHT.

Models	RMSE ↓			si-RMSE ↓			Angular Error ↓			User Study ↑			GMD ↓
	D	S	M	D	S	M	D	S	M	D	S	M	
SG+L2	0.204	0.213	0.238	0.188	0.203	0.229	9.18°	9.42°	9.73°	26.0%	22.5%	18.0%	5.631
GD+L2	0.133	0.161	0.178	0.117	0.132	0.161	7.60°	7.88°	8.12°	30.5%	25.5%	22.0%	5.303
GD+SML	0.080	0.103	0.117	0.072	0.087	0.106	6.78°	6.98°	7.12°	34.0%	31.5%	26.0%	5.163
GD+GML	0.073	0.091	0.102	0.062	0.069	0.092	6.61°	6.85°	7.04°	35.5%	32.0%	25.5%	5.031
GD+GML+GP	0.051	0.064	0.078	0.037	0.049	0.074	6.21°	6.50°	6.77°	42.0%	35.5%	31.0%	4.892

TABLE IV

ABLATION STUDIES OVER ANCHOR POINTS, LOSS FUNCTIONS, AND CONVOLUTION OPERATORS: GMLIGHT DENOTES THE STANDARD SETTING WITH 128 ANCHOR POINTS, GEOMETRIC MOVER’S LOSS (GML), AND SPHERICAL CONVOLUTION. WE CREATE FIVE GMLIGHT VARIANTS THAT USE 64 AND 196 ANCHOR POINTS, REPLACE GML WITH CROSS-ENTROPY LOSS (WITH CEL), REPLACE SPHERICAL CONVOLUTION WITH VANILLA CONVOLUTION (WITH VCONV), AND EMPLOY A FIXED RADIUS IN THE GAUSSIAN FUNCTION (WITH FIXED RADIUS), RESPECTIVELY (OTHER CONDITIONS UNCHANGED).

Methods	RMSE ↓			si-RMSE ↓		
	D	S	M	D	S	M
Anchor=64	0.071	0.085	0.102	0.064	0.071	0.093
Anchor=196	0.053	0.062	0.081	0.036	0.050	0.075
With CEL	0.062	0.074	0.094	0.055	0.059	0.078
With VConv	0.056	0.069	0.082	0.044	0.054	0.083
With Fixed Radius	0.063	0.071	0.085	0.047	0.056	0.081
GMLight	0.051	0.064	0.078	0.037	0.049	0.074

on three materials. Compared with the standard 128 anchor points, the prediction performance with 64 anchor points drops slightly and increasing anchor points to 196 also doesn’t bring clear performance gain. We conjecture that the larger number of parameters with 196 anchor points affects the regression accuracy negatively. In addition, we also compare the geometric mover’s loss and cross-entropy loss for distribution regression, compare the spherical convolution and vanilla convolution for panorama generation, and study the effect of adaptive radius in Gaussian projection. We can see that GML outperforms cross-entropy loss (CEL) clearly as GML captures spatial information of geometric distributions effectively. In addition, spherical convolution performs better than vanilla convolution consistently in panoramic image generation and adaptive radius strategy in Gaussian projection also brings notably improvement to the performance.

V. CONCLUSION

We present a lighting estimation framework GMLight which combines the merits of regression-based method and generation-based method. We formulate the illumination prediction as a distribution regression problem within a geometric space and design a geometric mover’s loss to achieve accurate regression of illumination distribution by leveraging the real scene geometry. With the including of depth branch, the proposed method also enables effective estimation of spatially-varying illumination. To synthesize panoramic illumination

maps with fine details from the light distribution, a novel generative projector with progressive guidance is designed to adversarially generate illumination maps. Extensive experiments demonstrate that the proposed GMLight significantly outperforms previous methods in terms of relighting in object insertion.

VI. ACKNOWLEDGEMENT

This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

REFERENCES

- [1] F. Zhan, S. Lu, C. Zhang, F. Ma, and X. Xie, “Adversarial image composition with auxiliary illumination,” in *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [2] M.-A. Gardner, K. Sunkavalli, E. Yumer, X. Shen, E. Gambaretto, C. Gagné, and J.-F. Lalonde, “Learning to predict indoor illumination from a single image,” in *SIGGRAPH Asia*, 2017.
- [3] F. Zhan, C. Zhang, Y. Yu, Y. Chang, S. Lu, F. Ma, and X. Xie, “Emlight: Lighting estimation via spherical distribution approximation,” *arXiv preprint arXiv:2012.11116*, 2020.
- [4] E. Zhang, M. F. Cohen, and B. Curless, “Emptying, refurbishing, and relighting indoor spaces,” *TOG*, 2016.
- [5] K. Karsch, V. Hedau, D. Forsyth, and D. Hoiem, “Rendering synthetic objects into legacy photographs,” *TOG*, 2011.
- [6] M. Garon, K. Sunkavalli, S. Hadap, N. Carr, and J.-F. Lalonde, “Fast spatially-varying indoor lighting estimation,” in *CVPR*, 2019.
- [7] M.-A. Gardner, Y. Hold-Geoffroy, K. Sunkavalli, C. Gagné, and J.-F. Lalonde, “Deep parametric indoor lighting estimation,” in *ICCV*, 2019.
- [8] S. Song and T. Funkhouser, “Neural illumination: Lighting prediction for indoor environments,” in *CVPR*, 2019.
- [9] P. P. Srinivasan, B. Mildenhall, M. Tancik, J. T. Barron, R. Tucker, and N. Snavely, “Lighthouse: Predicting lighting volumes for spatially-coherent illumination,” in *CVPR*, 2020.
- [10] D. Cheng, J. Shi, Y. Chen, X. Deng, and X. Zhang, “Learning scene illumination by pairwise photos from rear and front mobile cameras,” *Computer Graphics Forum*, 2018.
- [11] 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.
- [12] M. Li, J. Guo, X. Cui, R. Pan, Y. Guo, C. Wang, P. Yu, and F. Pan, “Deep spherical gaussian illumination estimation for indoor scene,” in *MM Asia*, 2019.
- [13] Z. Chen, A. Chen, G. Zhang, C. Wang, Y. Ji, K. N. Kutulakos, and J. Yu, “A neural rendering framework for free-viewpoint relighting,” *arXiv:1911.11530*, 2019.
- [14] Y. Rubner, C. Tomasi, and L. J. Guibas, “The earth mover’s distance as a metric for image retrieval,” *IJCV*, 2000.
- [15] B. Coors, A. P. Condurache, and A. Geiger, “Spherenet: Learning spherical representations for detection and classification in omnidirectional images,” in *ECCV*, 2018.

- [16] H. Vogel, "A better way to construct the sunflower head," *Mathematical biosciences*, 1979.
- [17] J. F. Lalonde, A. A. Efros, and S. G. Narasimhan, "Estimating the natural illumination conditions from a single outdoor image," *IJCV*, 2012.
- [18] Y. Hold-Geoffroy, K. Sunkavalli, S. Hadap, E. Gambaretto, and J.-F. Lalonde, "Deep outdoor illumination estimation," in *CVPR*, 2017.
- [19] L. Murmann, M. Gharbi, M. Aittala, and F. Durand, "A dataset of multi-illumination images in the wild," in *ICCV*, 2019.
- [20] M. Boss, V. Jampani, K. Kim, H. P. Lensch, and J. Kautz, "Two-shot spatially-varying brdf and shape estimation," in *CVPR*, 2020.
- [21] Z. Liao, K. Karsch, H. Zhang, and D. Forsyth, "An approximate shading model with detail decomposition for object relighting," *International Journal of Computer Vision*, vol. 127, no. 1, pp. 22–37, 2019.
- [22] F. Zhan and C. Zhang, "Spatial-aware gan for unsupervised person re-identification," *Proceedings of the International Conference on Pattern Recognition*, 2020.
- [23] R. Maier, K. Kim, D. Cremers, J. Kautz, and M. Nießner, "Intrinsic3d: High-quality 3d reconstruction by joint appearance and geometry optimization with spatially-varying lighting," in *ICCV*, 2017.
- [24] S. Lombardi and K. Nishino, "Reflectance and illumination recovery in the wild," *TPAMI*, 2016.
- [25] S. Sengupta, J. Gu, K. Kim, G. Liu, D. W. Jacobs, and J. Kautz, "Neural inverse rendering of an indoor scene from a single image," in *ICCV*, 2019.
- [26] J. T. Barron and J. Malik, "Intrinsic scene properties from a single rgb-d image," *TPAMI*, 2015.
- [27] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic image synthesis with spatially-adaptive normalization," in *CVPR*, 2019.
- [28] F. Zhan, Y. Yu, R. Wu, K. Cui, A. Xiao, S. Lu, and L. Shao, "Bi-level feature alignment for semantic image translation & manipulation," *arXiv preprint*, 2021.
- [29] F. Zhan, Y. Yu, K. Cui, G. Zhang, S. Lu, J. Pan, C. Zhang, F. Ma, X. Xie, and C. Miao, "Unbalanced feature transport for exemplar-based image translation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [30] C. LeGendre, W.-C. Ma, G. Fyffe, J. Flynn, L. Charbonnel, J. Busch, and P. Debevec, "Deeplight: Learning illumination for unconstrained mobile mixed reality," in *CVPR*, 2019.
- [31] T. Sun, J. T. Barron, Y.-T. Tsai, Z. Xu, X. Yu, G. Fyffe, C. Rhemann, J. Busch, P. Debevec, and R. Ramamoorthi, "Single image portrait relighting," in *TOG*, 2019.
- [32] D. Liu, C. Long, H. Zhang, H. Yu, X. Dong, and C. Xiao, "Arshadowgan: Shadow generative adversarial network for augmented reality in single light scenes," in *CVPR*, 2020.
- [33] L. Chizat, G. Peyré, B. Schmitzer, and F.-X. Vialard, "Scaling algorithms for unbalanced transport problems," in *arXiv:1607.05816*, 2016.
- [34] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," in *NIPS*, 2013.
- [35] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," in *NIPS*, 2014, pp. 2672–2680.
- [36] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *CVPR*, 2017.
- [37] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.
- [38] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [39] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 12 830–12 839.
- [40] R. Hess, *Blender Foundations: The Essential Guide to Learning Blender 2.6*. Focal Press, 2010.

Fangneng Zhan received the B.E. degree in Communication Engineering and Ph.D. degree in Computer Science & Engineering from University of Electronic Science and Technology of China and Nanyang Technological University, respectively. His research interests include deep generative models and image synthesis & manipulation. He contributed to the research field by publishing more than 10 articles in prestigious conferences. He also served as a reviewer or program committee member for top journals and conferences including TPAMI, TIP, ICLR, NeurIPS, CVPR, ICCV, and ECCV.

Yingchen Yu obtained the B.E. degree in Electrical & Electronic Engineering at Nanyang Technological University, and M.S. degree in Computer Science at National University of Singapore. He is currently pursuing the Ph.D. degree at School of Computer Science and Engineering, Nanyang Technological University under Alibaba Talent Programme. His research interests include computer vision and machine learning, specifically for image synthesis and manipulation.

Changgong Zhang is currently an algorithm engineer of Alibaba DAMO Academy. He received the PhD degree in Computer Graphics and Visualization from Delft University of Technology in 2017. His research interests include 3D vision, inverse rendering, and scientific visualization.

Rongliang Wu received the B.E. degree in Information Engineering from South China University of Technology, and M.S. degree in Electrical and Computer Engineering from National University of Singapore. He is currently pursuing the Ph.D. degree at School of Computer Science and Engineering, Nanyang Technological University. His research interests include computer vision and deep learning, specifically for facial expression analysis and generation.

Wenbo Hu is currently a Ph.D. candidate in the Department of Computer Science and Engineering, The Chinese University of Hong Kong. He received his B.Sc. degree in computer science and technology from Dalian University of Technology, China, in 2018. His research interests include computer vision, computer graphics and deep learning.

Shijian Lu is an Assistant Professor in the School of Computer Science and Engineering, Nanyang Technological University. He received his PhD in Electrical and Computer Engineering from the National University of Singapore. His research interests include computer vision and deep learning. He has published more than 100 internationally refereed journal and conference papers. Dr Lu is currently an Associate Editor for the journals of Pattern Recognition and Neurocomputing.

Feiyang Ma is a senior algorithm engineer at Alibaba Damo Academy and currently is in charge of openAI platform on Alibaba Cloud. Her research interests include image processing, virtual reality, augmented reality, real-time rendering, and 3D reconstruction.

Xuansong Xie is a senior staff engineer and technical director of Damo Academy, joined Alibaba, in 2012, and currently is in charge of the Alibaba Design Intelligence and Health Intelligence team, focusing on vision generation and enhancement, image retrieval, medical image & language intelligence, and other AI technology R&D.

Ling Shao is the CEO and the Chief Scientist of the Inception Institute of Artificial Intelligence (IIAI), Abu Dhabi, United Arab Emirates. He was also the initiator and the Founding Provost and Executive Vice President of the Mohamed bin Zayed University of Artificial Intelligence (the world's first AI University), UAE. His research interests include computer vision, deep learning, medical imaging and vision and language. He is a fellow of the IEEE, the IAPR, the IET, and the BCS.