

Universal Photometric Stereo Network using Global Lighting Contexts

Satoshi Ikehata

National Institute of Informatics (NII)

sikehata@nii.ac.jp

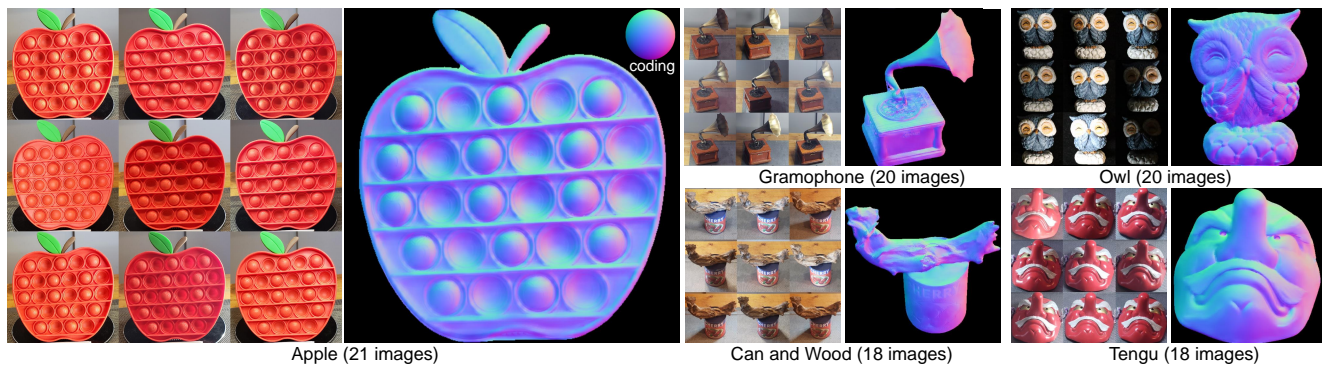


Figure 1. Given images of an object under different lighting conditions, our method can recover the detailed surface normal map. Images were captured by a smartphone camera with moving handheld lights under the indoor natural illumination. Please see details in Section 5.3.

Abstract

This paper tackles a new photometric stereo task, named universal photometric stereo. Unlike existing tasks that assumed specific physical lighting models; hence, drastically limited their usability, a solution algorithm of this task is supposed to work for objects with diverse shapes and materials under arbitrary lighting variations without assuming any specific models. To solve this extremely challenging task, we present a purely data-driven method, which eliminates the prior assumption of lighting by replacing the recovery of physical lighting parameters with the extraction of the generic lighting representation, named global lighting contexts. We use them like lighting parameters in a calibrated photometric stereo network to recover surface normal vectors pixelwisely. To adapt our network to a wide variety of shapes, materials and lightings, it is trained on a new synthetic dataset which simulates the appearance of objects in the wild. Our method is compared with other state-of-the-art uncalibrated photometric stereo methods on our test data to demonstrate the significance of our method.

1. Introduction

Photometric stereo is a problem of recovering the surface normal map from appearances of an object under varying lighting conditions. For decades, a broad spectrum of techniques have been proposed to expand the scope of target geometry, material and acquisition setup [15, 23, 24, 42] in

the framework of the physics-based inverse rendering. Recently, advances in deep learning have been eliminating the dependence on physics-based modeling from photometric stereo methods, which contributes to handle complex optical phenomena which are hardly described in a mathematically tractable form [12, 21, 39, 44].

However, despite the long journey in this research field, each photometric stereo algorithm is still limited to a specific physical lighting model, which severely compromises its usability. In reality, most recent (semi-)calibrated [11, 21, 39, 40, 44] and uncalibrated [12, 27] photometric stereo methods still assume the single lighting in a dark environment. Others address natural lighting conditions, however both calibrated and uncalibrated methods still assume convex Lambertian surfaces and their lighting models are limited to spherical harmonics lighting [8, 17], dominant sun lighting [6, 20] and equivalent directional lighting [16, 36] which cannot represent the complex illumination.

To address this limitation, this paper proposes the “third” task in photometric stereo problem following calibrated and uncalibrated tasks. We name it *universal* photometric stereo (UniPS) which denotes the setup without prior assumption of physical lighting models; hence, arbitrary lighting conditions should be considered unlike calibrated and uncalibrated tasks that consider specific ones as in Table 1.¹

¹It would be ideal if the task were universal on materials as well. However, since some objects such as mirrors and transparent objects must be excluded, we consider only lighting conditions to be universal in this task.

Table 1. Comparison of different photometric stereo tasks.

	Lighting Calibration	Lighting Model	Lighting Condition
Calibrated	Required	Required	Specific
Uncalibrated	Free	Required	Specific
Universal	Free	Free	Arbitrary

In this paper, we present a first viable method for UniPS based on key insights as follows. Conventional uncalibrated photometric stereo algorithms recovered physical lighting parameters and surface normals sequentially or simultaneously; thus, constrained by specific lighting models. However, we show that the recovery of physical lighting parameters is *not* essential in a UniPS network, and can be replaced with the extraction of *global lighting contexts* from individual images by their interaction with others. In order to keep the receptive field of our network constant, contexts are extracted at the predefined *canonical resolution* which is independent of the input image resolution. By using global lighting contexts like calibrated lighting parameters, surface normal vectors can be recovered in similar to existing pixel-wise calibrated photometric stereo networks (*e.g.* [21, 22]) which easily scale to high-resolution images. While our network drops a prior assumption of lighting, its adaptation to diverse shapes, materials and lightings has to be ensured by training data. Since existing training datasets for photometric stereo networks are limited to the single, directional lighting setup, we create a dataset for our task by physically rendering the appearance of objects with more than 10,000 combinations of shape, material and lighting using high quality commercial 3-D assets. We also create an evaluation dataset with 50 sets of attributes using different assets to compare our method with state-of-the-art uncalibrated photometric stereo algorithms specifically designed for directional lighting [12] and natural lighting [17, 36]. Finally, the qualitative evaluation demonstrates that our method even works for real objects under the challenging spatially-varying lighting conditions that were conventionally considered to be intractable (See Fig. 1).

2. Related Work

Here we briefly review uncalibrated photometric stereo methods under the directional or natural lighting. Due to the space limit, other categories (*e.g.* calibrated setup [21–25], near-light setup [26, 40] or multi-view setup [29, 38]) are not included in this survey. However, we emphasize that all the methods more or less assume any of physical lighting models; hence, are limited to specific lighting setups.

Uncalibrated, Directional Lighting: Since Woodham [46] presented the first Lambertian photometric stereo algorithm in 1980, methods following its setup have been called calibrated photometric stereo which assumes an orthographic camera and a known single directional lighting. The uncal-

ibrated task is almost identical to the calibrated one except that lighting parameters are unknown. Until very recently, most uncalibrated photometric stereo algorithms assumed Lambertian integrable surfaces and the goal was to resolve the General Bas-Relief ambiguity [18] between geometry and light. Various cues were employed for resolving this ambiguity, which include inter-reflections [13], entropy of surface albedo [7], color profile [41], diffuse maxima [14], reflectance symmetry [47] and perspective geometry [37]. Though there were very few methods for non-Lambertian surfaces due to their ill-posed nature, Lu *et al.* [35] utilized the statistical distribution of the intensity profile to recover isotropic non-Lambertian surfaces.

In 2018, the first deep neural network for uncalibrated photometric stereo was presented [11]. However, it simply dropped the lighting channel from its calibrated variant presented at the same time and the performance was quite limited. Hence, authors extended their work to the two-step approach [10] where only lighting information was firstly recovered, and then it was used as input of the calibrated photometric stereo network. Later, authors further updated their work by feeding the surface normal estimation result back to the lighting prediction to improve its accuracy [12]. Building upon this work, Kaya *et al.* [27] have recently utilized the lighting prediction result from [12] to recover surface normals in the neural inverse rendering framework.

Uncalibrated, Natural Lighting: Although most photometric stereo methods assume the single directional lighting condition, some literature exists which addresses natural lighting conditions. However, inversely decomposing the natural light reflected on non-Lambertian surfaces is intractable in the inverse rendering framework, so it is common to assume convex Lambertian surfaces and approximated lighting models [8, 9, 16, 17, 36].

The first uncalibrated photometric stereo algorithm under natural lighting was proposed by Barsri *et al.* [8] which approximated natural lighting as global first-order spherical harmonics. Though a global concave-convex ambiguity exists in estimated surface normals due to the lighting approximation with an orthographic camera, Brahimi *et al.* [9] recently proved that the perspective integrability constraint makes the problem well-posed. Mo *et al.* [36] presented another uncalibrated photometric stereo algorithm under natural lighting. They proposed the equivalent directional lighting model which decomposed the entire task into the patch-wise directional uncalibrated photometric stereo problem. This work was later extended by authors where MRF-based global optimization and rotation averaging were introduced for a better normal patch integration [16]. These methods worked under an orthographic camera, however when solving the global orthogonal ambiguity with integrability, there was a binary ambiguity left which requires to be solved manually and a non-integrable surface can not be recov-

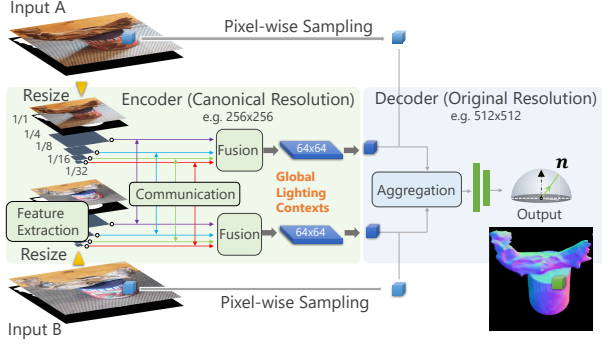


Figure 2. Illustration of the proposed method. Given arbitrary number of image crops (*e.g.* two in this figure) of an object under different lighting conditions, they are resized to the *canonical resolution* and passed to the encoder to extract global lighting contexts. Then, for each pixel in the *original image*, raw image values and an interpolated global lighting context are concatenated to be fed to the decoder that aggregates features on the lighting axis and recovers the surface normal vector at the location.

ered in theory. To address this issue, Haefner *et al.* [17] presented a variational method which directly recovered a depth map rather than a surface normal map to handle non-integrable surfaces. However, this method requires initial geometry which is basically recovered from the object outline and its performance drastically affected by its quality.

3. Universal Photometric Stereo

In this section, we preliminarily define the UniPS problem setup. Given a set of q RGB images $I_{\{1, \dots, q\}}$ captured under unknown, arbitrarily varying lighting conditions with a fixed camera, and a binary mask M to specify the target object in the image, our goal is to recover the unit surface normal vector \mathbf{n} at pixels which belong to the object. The object mask is friendly provided using any existing foreground extraction methods such as [43]. UniPS doesn't put any particular constraints on the camera model, however an orthographic, linear camera is practically assumed in this work as training images are rendered by the camera.

The main difficulty in UniPS is due to the lack of prior knowledge of lighting which disables most of existing photometric stereo algorithms designed under specific physical lighting models. In UniPS, lighting conditions could include near or distant, directional, point, area, natural or even mixture of them (*e.g.* putting active near area lights under passive natural light as demonstrated in Fig. 1). It is highly possible that the lighting is *spatially-varying* which cannot be represented by a global lighting model. This may be the case, for example, in most indoor natural scenes [30]. In a general context, even inter-reflections and cast shadows can also be a part of spatially-varying lightings.

Since the ultimate goal of UniPS is to realize a truly prac-

tical photometric stereo method, the available geometries and materials (except for extreme cases such as transparent and perfect mirror objects) should be diverse as well and the number of input images and their resolution should be arbitrary. Although accepting arbitrary number of input images was considered as an important requirement in recent works (*e.g.* [11, 21, 44]), there are still some methods that do not meet this requirement (*e.g.* [31]). Scalability is basically not a problem in pixel-by-pixel algorithms (*e.g.* [21, 24]), but it often becomes a major issue in methods that use the entire image information such as convolutional neural networks and global optimization (*e.g.* [12, 36]).

4. Method

4.1. UniPS Network with Global Lighting Contexts

This paper presents a first viable universal photometric stereo network. As illustrated in Fig. 2, our network consists of an encoder and decoder. The encoder extracts *global lighting contexts* from images and an object mask, which is a generic lighting representation that corresponds to physical lighting parameters (*e.g.* light direction) in deep uncalibrated photometric stereo networks [10, 12]. The decoder takes all raw image values and the interpolated global lighting context at each pixel and predict its surface normal.

Our network architecture has two major differences from basic encoder-decoder architectures that take a single image as input. First, our network takes multiple images as input; hence, features must be embedded in the latent space considering the interaction of them. Therefore, we perform feature communication in the encoder and aggregation in the decoder, respectively. Second, unlike typical architectures where encoded features are directly passed to the decoder, we use different working resolutions for the encoder and decoder. The working resolution for the decoder is same as the original image resolution, but the encoder takes as input images that have been resized to the pre-defined *canonical resolution*, which is basically smaller than the original resolution and its output is passed to the decoder after inversely converted for the decoder's working resolution.

There are two major advantages to use different working resolutions. First, the scalability to the image size is ensured because memory requirements for the encoder depend only on the canonical resolution, not on the original one, while the decoder processes each pixel one by one. The second and more important reason is to keep the receptive field of the encoder invariant to the input image size. Without this, the networks' receptive field may not cover the entire object in extremely high-resolution test images.

4.2. Framework Components

Preprocessing: The pixel value range of individual images can significantly vary under different lightings. There-

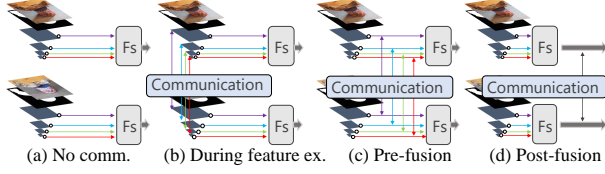


Figure 3. The illustration of four different variants of the encoder with regard to the timing of the inter-image communication.

fore, we divide each image by its mean for normalization. Note that the common zero-mean normalization is not used to avoid emphasizing the area of low signal-to-noise ratio. We then crop the rectangular object bounding region with a small margin (*e.g.* four pixels) based on the object mask to confirm that the object is placed in the middle of the crop and covered from edge to edge.

Encoder: Preprocessed crops of the images and mask are bilinearly resized to the predefined $s \times s$ canonical resolution. They are then concatenated $C_{\{1, \dots, q\}} \in \mathbb{R}^{s \times s \times (3+1)}$ and passed to the encoder. The output of the encoder is embedded features for individual images whose size is a quarter of the canonical resolution $\mathcal{G}_{\{1, \dots, q\}} \in \mathbb{R}^{\frac{s}{4} \times \frac{s}{4} \times d_e}$ where d_e is the embed dimension. We call them as global lighting contexts because the difference of features among images should only be attributed to the difference of lighting conditions. Global lighting contexts are also an analogy of physical lighting parameters similarly recovered by the conventional uncalibrated photometric stereo methods. However, unlike physical parameters were basically assigned to each image globally (*e.g.* light direction [10], spherical harmonics [17]), a unique global lighting context is assigned at each location (therefore we name context's); hence has the capacity to represent spatially-varying phenomena such as near-lighting, inter-reflections and cast shadows.

The encoder mutually embeds global lighting contexts via the image-wise feature extraction and inter-image feature communication. The former extracts multi-scale feature maps using a backbone (*e.g.* SwinTransformer [32]) followed by the basic multi-scale feature fusion [48]. The latter propagates features across images without changing the feature dimension by pixelwisely employing a single Transformer layer [45] as with the recent calibrated photometric stereo network [22]. As illustrated in Fig. 3, there are multiple feasible designs of the encoder depending on how information propagate via inter-/intra-image interactions. The best design of them will be discussed in the ablation study in Section 5.1. As it will turn out, the feature communication should be conducted just before the multi-scale feature fusion (*i.e.* pre-fusion). Please refer to the supplementary for more details about individual implementations of the feature extraction and communication.

Decoder: For each coordinates x of the *original* resolution, the decoder inputs a set of image values $I_{\{1, \dots, q\}}(x)$ and

a global lighting context $\mathcal{G}_{\{1, \dots, q\}}(\mathcal{S}(x))$, and outputs the unit normal vector at the location. Here, \mathcal{S} is the sampling operation to fetch the value of the global lighting context corresponding to x using bilinear interpolation. Note that avoiding the resize of global lighting contexts to the original resolution also contributes to the scalability.

The decoder is composed of the feature aggregation and the surface normal recovery. The former further propagates information across different lighting conditions with extra Transformer layers, and then squeezes the lighting channel with the pooling by multi-head attention (PMA) [28]. The output of PMA is a pixelwise feature vector independent of number of input images. The latter feeds the feature to the multi-layer perceptron with one hidden layer to output the surface normal vector. We should note that when no communication is conducted in the encoder as in Fig. 3, it can be deemed that global lighting contexts have been extracted at the aggregation step in the decoder.

Training Loss: The mean squared error between ground truth and predicted surface normals is used as the training loss. Because of the efficiency, we compute the loss only on samples from specific locations, not from the entire image. For each object, we uniformly sample pixels at the original resolution whose centers can be projected exactly onto centers of pixels at the canonical resolution. Then, we further sample pixels from a limited number of random locations (*e.g.* 2500) at the original resolution which allows for the normal estimation with a sub-pixel global lighting context.

Network Analysis: Our method is data-driven; hence satisfies the requirements of UniPS. This is achieved by global lighting contexts, which eliminate the need for the physical lighting parameter recovery and makes our method feasible to handle complex spatially-varying lighting effects.

Our framework is scalable by introducing the canonical resolution for encoding global lighting contexts and pixelwisely applying the decoder. The network's receptive field is independent of the input image size as well as there is no upper limit to the test image size. By utilizing the self-attention mechanism for the feature interaction across images, the network is applicable to an arbitrary number of images as long as the computational resource permits. Given this capacity, the availability of varying shapes, materials and lightings are secured by training data.

4.3. PS-Wild Training Dataset

We need sufficient training examples that properly characterize the universal photometric stereo task. Though two major synthetic photometric stereo datasets (*i.e.* Blobby and Sculpture [11], CyclesPS [21]) have been presented to date, they were for the methods under the single directional lighting and not applicable to UniPS. Therefore, we present a new photometric stereo dataset, named *PS-Wild*, which sim-

ulates the appearance of objects with diverse geometries and materials captured under various lighting conditions in the wild. The idea is simply to ask a physically-based renderer (*i.e.* Blender Cycles [2]) to synthesize a large number of images under general lightings, taking full advantage of high quality commercial 3-D assets.

To assemble an appropriate collection of object appearances for our PS-Wild dataset, we browsed through online marketplaces looking for 3-D assets that satisfied three main desiderata. First, we should use 3-D models that have sufficient complexity. We don’t want the objects whose surface normal distribution deviates from ones in the wild (*e.g.* too many planar objects or low-poly models are not adequate). Second, we want 3-D materials to be as diverse as possible. In our problem, we don’t assume specific materials as well as lighting models, so the training data needs to be chosen to cover diverse materials in the wild. In addition, textures need to be realistic. In the existing synthetic datasets, the entire image was rendered with a single BRDF [11], or pixelwisely different random BRDFs [21, 33], but the actual surface texture has some realistic rules as mentioned in previous studies [7, 15]. Hence, it is worth using texture maps that were designed by professionals. Third, we want to render scenes under diverse lighting conditions. In this regard, the most practical method we found is to use the HDR image-based (HDRI) lighting that covers various indoor and outdoor scenes. To add more details, this method uses an HDR environment map to place an omnidirectional light source. Each ray from the light is self-shielded before it hits the surface, so a variety of spatially-varying lighting effects (*e.g.* cast shadows, near-lighting) are baked.

Based on these desiderata, we choose Adobe Stock [1] as our assets to create the training dataset. This collection consists of over 17,000 3-D assets including 3-D models, materials (texture maps) and lightings (environment maps). From all available assets, we downloaded 410 3-D models, 926 materials and 31 lightings. Most of data in Adobe Stock assets are 3-D models and we actually used all the materials and lightings available at the time we accessed there (Oct. 2021). The number of assets itself may not seem very large, but in reality, a countless number of images can be synthesized by various augmentation techniques such as rotation and color transformation.

After acquiring 3-D assets, we apply our computational pipeline to generate images with ground truth surface normal maps. For each 3-D object in assets, we randomly assign one material and one environment map. Then we apply several random rotations of the object until the Shannon entropy of its surface normal distribution becomes more than a threshold (*i.e.* 4.0 in our method) and if the maximum entropy is less than the threshold, the 3-D model is discarded from assets. We then scale the object to make sure that the object’s outline reaches the edge of the image. We render 10

of 512×512 , 16-bit images per object. For each rendering, we randomly rotate the environment map on the spherical axis to make the variation of lighting conditions and automatically adjust the exposure of the camera to make the dynamic range of rendered images consistent. We turn on the ray tracer to render the cast shadow and inter-reflection for adding the spatially-varying lighting effects. Finally, we got 10,099 objects with a different pose, material and lighting.

4.4. PS-Wild Test Dataset

We also create a test dataset for the evaluation purpose. The computational pipeline of generating images is same as one for the training dataset but different 3-D assets are used for the fair evaluation; 25 objects from CGTrader [3], 50 materials from ShareTextures [4] and 50 environment maps from sIBL Archive [5]. For each 3-D model, we assign two sets of a material and environment map resulting in 50 different sets of the object, material and environment map. Unlike training dataset, we carefully pick six textures per each texture category; Concrete, Fabric, Floor, Ground, Wood and Metal as categorized in ShareTextures. In order to properly evaluate the performance of a method under various lighting conditions, we render images for the same set of the object and material using three different lighting methods; (a) single directional lighting (uniformly sampled), (b) HDRI lighting (same as training) and (c) mixture of (a) and (b). The image resolution is also 512×512 but the number of images is 32 to evaluate the performance with varying number of input images.

Dataset Analysis: Our test dataset characterizes the universal photometric stereo task. Objects with both convex and non-convex geometries and a variety of spatially-varying materials involving diffuse, specular and *metallic* are rendered with three different lighting methods including the HDRI lighting which exhibits challenging spatially-varying lighting effects. Therefore, methods that assume a specific physical lighting model are not suitable for this dataset.

5. Results

We conduct experiments on synthetic and real test data. We first ablate the important design elements of our architecture and then we compare the proposed universal photometric stereo network with the previous state-of-the-arts on the uncalibrated photometric stereo task [10, 17, 36]. For the convenience, we henceforth refer to the proposed method as UniPS-GLC (Universal Photoemtric Stereo network using Global Lighting Contexts) as necessary.

Training Details: Our network was trained from scratch on a NVIDIA Quadro RTX 8000 machine with AdamW [34] optimizer for 20 epochs using a step decay learning rate schedule ($\times 0.8$ every three epochs). A batch size of 3, an initial learning rate of 0.0001, and a weight decay of 0.05

Table 2. Comparison with different image feature extractors.

	SwinT [32]	ViL [49]	ResNet101	ResNet50
Directional	19.7	29.5	48.6	50.6
HDRI	16.8	23.2	45.8	44.0
Dir.+HDRI	16.1	24.0	45.2	45.0

Table 4. Comparison with different canonical resolutions.

	128 × 128	256 × 256	512 × 512
Directional	19.9	19.7	N.A
HDRI	18.1	16.8	N.A
Dir.+HDRI	17.0	16.1	N.A

Table 6. Comparison with different feature aggregation methods.

	Max-pool	TF+PMA (3L)	TF+PMA (6L)
Directional	59.3	19.7	37.5
HDRI	39.3	16.8	29.4
Dir.+HDRI	39.0	16.1	30.6

were used. The number of random samples from the original resolution was fixed by 2500. In total, it took roughly 48 hours to train the network for each configuration.

Inference Time: The inference time of our method depends on the number and resolution of input images. In the case of 32 of 512×512 images as input, it takes less than a few seconds excluding IO on GPU. This is slightly slower than the simpler deep photometric stereo networks (*e.g.* 0.5 sec in [10]), however more than hundred times efficient than inverse rendering based methods [17, 27, 36, 44].

Evaluation metric: Evaluation is based on the mean angular errors (MAE) between predicted and true surface normal maps measured in degrees (0 to 180). In our evaluation, we apply each algorithm to PS-Wild test dataset with three different lighting methods and discuss the result mainly based on the averaged MAE over 50 different objects.

5.1. Ablation study

Base architecture: As needed, we define the base architecture for the ablation study with following design elements: SwinTransformer [32] for feature extraction; No communication in the encoder; 256×256 canonical resolution; Three stacks of transformer layers followed by PMA for the feature aggregation in the decoder; In ablation, only the target property was changed from this base architecture.

Encoder: Table 2 shows the comparison of four different image feature extractors: SwinTransformer [32], Vision-Longformer [49] and ResNet-50/101 [19]. We observe that Transformer-based encoders, especially SwinTransformer outperformed ResNets probably due to the larger receptive field of the Transformer model. Because the main purpose of this work is to present a viable method to demonstrate our ideas, the further discussion remains future work.

Table 3 compared four different variants of the feature

Table 3. Comparison with different encoder designs.

	Base	Dur-Ext.	Pre-Fus.	Post-Fus.
Directional	19.7	22.5	17.0	25.6
HDRI	16.8	20.1	14.5	22.5
Dir.+HDRI	16.1	19.4	13.8	21.1

Table 5. Spatially-varying vs uniform global lighting contexts.

	Spatially-varying	Spatially-uniform
Directional	19.7	40.5 (trivial)
HDRI	16.8	40.5 (trivial)
Dir.+HDRI	16.1	40.5 (trivial)

communication; no communication in the encoder (Base) during the image feature extraction (Dur-Ext), before the multi-scale fusion (Pre-Fus) and after the multi-scale fusion (Post-Fus) as illustrated in Fig. 3. As have been mentioned already, the best performance was obtained when the communication was done just before the multi-scale fusion. Interestingly, the second best was obtained when no communication among images was done in the encoder, but only in the decoder. We analysed this result and found that the communication during the image feature extraction simply broke the feature extraction process and the post-fusion communication made the optimization unstable since it is equivalent to increasing the transformer layers in the decoding process which is consistent with later results.

Table 4 compared different canonical resolutions either from 128×128 or 256×256 and we observe that 256×256 worked slightly better. This result indicates that though the lower resolution is helpful to see the entire object, important information for details could be discarded at very low resolution. On the other hand, it is surprising to see that 32×32 global lighting contexts from 128×128 canonical resolution still provides a reasonable reconstruction for 512×512 output resolution. Unfortunately, we couldn't get the result of 512×512 canonical resolution due to the memory limit.

Finally, we compared our spatially-varying global lighting contexts with the spatially-uniform ones by applying the global average pooling to shrink $\{\mathcal{G}\} \in \mathbb{R}^{s \times s \times d_e}$ to $\{\mathcal{G}_{uni}\} \in \mathbb{R}^{1 \times 1 \times d_e}$ and feeding the spatially same vector to the decoder. In reality, this is a similar procedure in existing uncalibrated photometric stereo networks [10, 27] where the single lighting parameter is firstly recovered for each image and used as the input of the surface normal predictor. However, Table 5 shows that the network with the uniform context always provided the trivial solution since it couldn't capture spatially-varying lighting effects.

Decoder: We compared three different variants of the pixel-wise feature aggregation strategy from a max pooling, three Transformer layers with PMA and six Transformer layers with PMA; The further to the right we go, the more complex interaction can theoretically take place. Note that the latent vector size in the Transformer layer was fixed by 384 and

Table 7. Different number of input images (Ours, pre-fusion).

	1	4	8	16	32
Directional	35.4	22.6	18.8	17.5	17.0
HDRI	27.2	19.5	16.5	14.7	14.5
Dir.+HDRI	28.5	20.6	17.4	15.6	13.8

Table 8. Comparison with other methods (pre-fusion).

	Ours	GCNet [10]	MPM [36]	Var. [17]
Directional	17.0	17.7	32.9	33.0
HDRI	14.5	24.8	30.4	37.5
Dir.+HDRI	13.8	31.5	30.8	32.5

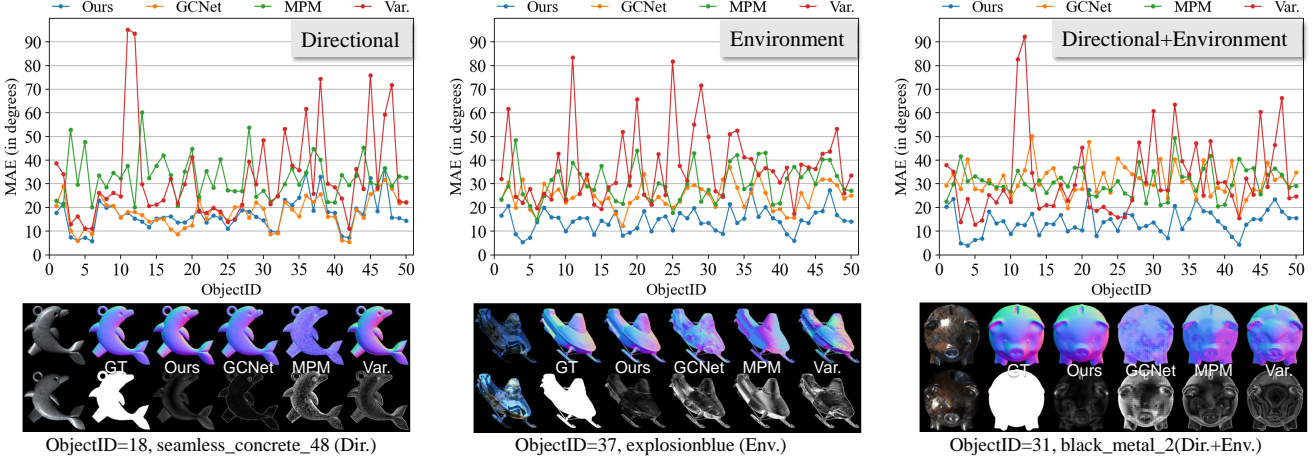


Figure 4. We compared our method with three uncalibrated photometric stereo methods specifically designed for single directional lighting (*i.e.* GCNet [10]) and natural lighting (*i.e.* MPM [36] and Variational [17]) on our synthetic PS-Wild test dataset. There were three different methods of lighting; directional, HDRI (and the mixture of them. The performance was evaluated based on the mean angular errors in degrees. We showed the examples in the dataset and the recovered surface normal maps/error maps (80 degrees at maximum).

the hidden dimension of the feedforward layer was 1024. The result is illustrated in Table 6. Though there is no surprise to see the simple max-pooling didn’t work at all, the reason for the worse results for the deeper Transformer networks seemed to be due to the over-fitting. It should be noted that the Transformer network is known to be a difficult network to train, and careful adjustment of the hyper-parameters may yield different results.

Different number of input images: In Table 7, we evaluated our network (base+pre-fusion) with different number of input images. As expected, the accuracy dropped as the number of input images decreased. However, unlike existing calibrated and uncalibrated photometric stereo methods which had been shown that they didn’t work when the number of images was small (*e.g.* 10) [10, 21], the degree of decline with our method was not quite significant.

5.2. Quantitative Evaluation on PS-Wild Test Data

We compared our method to state-of-the-arts on the uncalibrated photometric stereo problem on our synthetic test dataset. The first comparison point is GCNet [12] which is a learning-based method under the directional lighting condition that alternately estimates physical lighting parameters and surface normals. The second point is the matrix-based patch merging (MPM) [36] which is the Lambertian photometric stereo method for uncalibrated natural lighting based on the equivalent directional lighting model. The

third point is the variational uncalibrated photometric stereo method (Variational) [17] which approximates Lambertian reflectance model through a spherical harmonic expansion. We used authors’ official codes for the evaluation and our GLC-UniPS architecture was “base + pre-fusion” configuration that has yielded the best result so far.

The results are illustrated in Fig. 4 and Table 8². We observe that our UniPS-GLC reasonably worked independent of the lighting methods. As expected, GCNet [10] worked fine for objects under the directional lighting condition (*e.g.* ID 48), however problematic with non-uniform materials and non-directional lightings. Overall, MPM [36] and Variational [17] were inferior to GCNet for diverse objects due to their assumptions of convex, Lambertian surfaces. Especially, MPM had a problem in handling spatially varying non-Lambertian textures or non-convex geometries (*e.g.* ID 37) and variational also had a problem in handling non-Lambertian reflections especially for metallic surfaces (*e.g.* ID 31). Interestingly, better performance of our method was observed for objects of the same shapes and materials under the environment lighting than the directional one. This result may be attributed to the fact that in an image under the directional lighting condition, the shadow or inter-reflection behaves negatively, whereas under the environment light, they positively yields spatially-varying lighting effects as a source of information.

²The input images and recovered normal map for individual objects are presented in the supplementary.

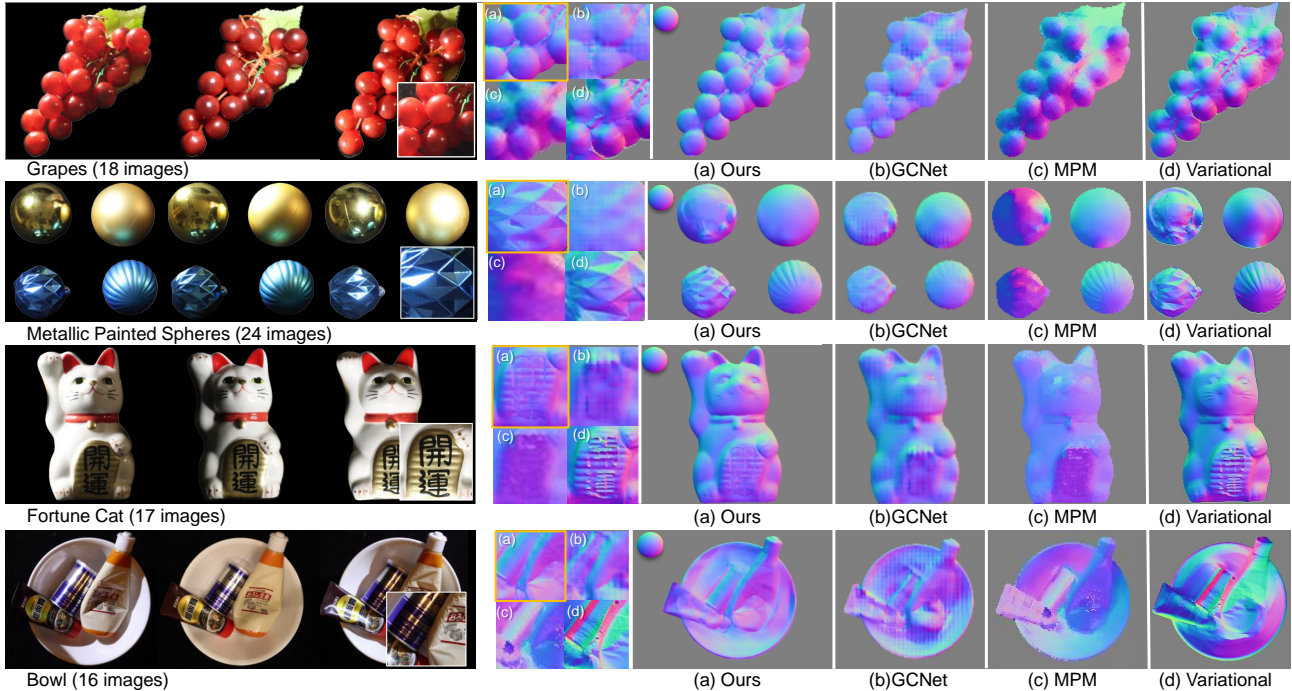


Figure 5. The quantitative comparison on real images under challenging spatially-varying lighting conditions. All the images were captured with a moving near area light source and static indoor natural illumination. Foreground objects have been cropped off using object masks.

5.3. Qualitative Evaluation on Real Data

Fig. 1 validates our method (base+pre-fusion) on real images. Objects of varying materials such as ceramic, metallic and clear coat were captured by a 8-bit smartphone camera under a few hand-held area lights and/or the indoor natural illumination. The area light sources were intentionally put within 30 cm of an object so that strong spatially-varying lighting effects were observed. Although quantitative evaluation is not possible due to the absence of the ground truth, we observe that overall reasonable normal maps were recovered even under these extremely challenging setups. We also observe that the very low resolution of GLiCo (*i.e.* 64×64) is sufficient to recover surface details.

For a more objective evaluation, Fig. 5 provides a comparison with GCNet, MPM and Variational. Images were captured by putting a single area light near an object under the static indoor natural illumination; therefore, could not be represented by physical lighting models such as the directional lighting nor spherical harmonics. We picked four objects with different difficulties and included two “spherical” objects (*i.e.* *Grapes* and *Metallic Painted Spheres*) so that the quality of normal maps could be evaluated based on their spherical parts. *Grapes* exhibit complex cast shadows because of its non-convex geometry. *Metallic Painted Spheres* consists of four metallic-painted balls with different surface roughness and structures which exhibit strong inter-reflections with each other. *Fortune Cat* contains large “black” regions, which are known to be difficult to distin-

guish from shadows. *Bowl* includes multiple challenging objects put in a bowl which is hard to resolve the convex-concave ambiguity. We emphasize that these are extremely difficult objects that existing photometric stereo methods have not even attempted to handle them. Though our results are still far from perfect (*i.e.*, we clearly observe errors in our results due to inter-reflection in *Metallic Painted Spheres* and convex geometries in *Bowl*), our method obviously recovered most reasonable surface normal maps.

6. Conclusion

This paper tackled a new photometric stereo task, named universal photometric stereo (UniPS) which drops prior assumptions of physical lighting models. To this end, we presented the first viable UniPS network based on the generic lighting representation named global lighting contexts. We also presented synthetic training and evaluation datasets for our UniPS task and our extensive evaluation demonstrated the performance of our method.

Despite the significant progress toward a practical photometric stereo method, there are some limitations remained. First, the deeper analysis of global lighting contexts is missing. To explore its potential for applications other than the surface normal recovery is a future work. Second, our training method specifies a camera model (*e.g.*, orthographic, linear camera) and we may also want to remove this dependence. Finally, we need a real dataset with the ground truth for evaluating UniPS methods in a more quantitative manner.

References

- [1] Adobe Stock. <https://stock.adobe.com/>. 5
- [2] Blender. <https://www.blender.org/>. 5
- [3] CGTrader. <https://www.cgtrader.com/>. 5
- [4] ShareTextures. <https://www.sharetextures.com/>. 5
- [5] sIBL Archive. <http://www.hdrlabs.com/sibl/archive.html>. 5
- [6] J. Ackermann, F. Langguth, S. Fuhrmann, and M. Goesele. Photometric stereo for outdoor webcams. In *CVPR*, 2012. 1
- [7] N. Alldrin, S. Mallick, and D. Kriegman. Resolving the generalized bas-relief ambiguity by entropy minimization. *CVPR*, 2007. 2, 5
- [8] Ronen Basri, David Jacobs, and Ira Kemelmacher. Photometric stereo with general, unknown lighting. *International Journal of computer vision*, 72(3):239–257, 2007. 1, 2
- [9] Mohammed Brahimi, Yvain Quéau, Bjoern Haefner, and Daniel Cremers. On the well-posedness of uncalibrated photometric stereo under general lighting. In *Advances in Photometric 3D-Reconstruction*, pages 147–176. Springer, 2020. 2
- [10] G. Chen, K. Han, B. Shi, Y. Matsushita, and K. K. K. Wong. Self-calibrating deep photometric stereo networks. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8731–8739, 2019. 2, 3, 4, 5, 6, 7
- [11] G. Chen, K. Han, and K-Y. K. Wong. Ps-fcn: A flexible learning framework for photometric stereo. *ECCV*, 2018. 1, 2, 3, 4, 5
- [12] Guanying Chen, Michael Waechter, Boxin Shi, Kwan-Yee K Wong, and Yasuyuki Matsushita. What is learned in deep uncalibrated photometric stereo? In *European Conference on Computer Vision*, pages 745–762. Springer, 2020. 1, 2, 3, 7
- [13] Ondřej Drbohlav and Radim Šára. Specularities reduce ambiguity of uncalibrated photometric stereo. In *ECCV*, pages 46–60, 2002. 2
- [14] P. Favaro and T. Papadhimetri. A closed-form solution to uncalibrated photometric stereo via diffuse maxima. In *CVPR*, 2012. 2
- [15] D. Goldman, B. Curless, A. Hertzmann, and S. Seitz. Shape and spatially-varying brdfs from photometric stereo. In *ICCV*, October 2005. 1, 5
- [16] Heng Guo, Zhipeng Mo, Boxin Shi, Feng Lu, Sai Kit Yeung, Ping Tan, and Yasuyuki Matsushita. Patch-based uncalibrated photometric stereo under natural illumination. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2021. 1, 2
- [17] Bjoern Haefner, Zhenzhang Ye, Maolin Gao, Tao Wu, Yvain Quéau, and Daniel Cremers. Variational uncalibrated photometric stereo under general lighting. pages 8539–8548, 2019. 1, 2, 3, 4, 5, 6, 7
- [18] H. Hayakawa. Photometric stereo under a light source with arbitrary motion. *JOSA*, 11(11):3079–3089, 1994. 2
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [20] Yannick Hold-Geoffroy, Paulo Gotardo, and Jean-François Lalonde. Single day outdoor photometric stereo. *IEEE transactions on pattern analysis and machine intelligence*, 43(6):2062–2074, 2019. 1
- [21] S. Ikehata. Cnn-ps: Cnn-based photometric stereo for general non-convex surfaces. In *ECCV*, 2018. 1, 2, 3, 4, 5, 7
- [22] S. Ikehata. Ps-transformer: Learning sparse photometric stereo network using self-attention mechanism. In *BMVC*, 2021. 2, 4
- [23] S. Ikehata and K. Aizawa. Photometric stereo using constrained bivariate regression for general isotropic surfaces. In *CVPR*, 2014. 1, 2
- [24] S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa. Robust photometric stereo using sparse regression. In *CVPR*, 2012. 1, 2, 3
- [25] S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa. Photometric stereo using sparse bayesian regression for general diffuse surfaces. *IEEE TPAMI*, 36(9):1816–1831, 2014. 2
- [26] Yuji Iwahori, Hidezumi Sugie, and Naohiro Ishii. Reconstructing shape from shading images under point light source illumination. In *ICPR*, volume 1, pages 83–87. IEEE, 1990. 2
- [27] Berk Kaya, Suryansh Kumar, Carlos Oliveira, Vittorio Ferrari, and Luc Van Gool. Uncalibrated neural inverse rendering for photometric stereo of general surfaces. pages 3804–3814, 2021. 1, 2, 6
- [28] J. Lee, Y. Lee, J. Kim, A. Kosior, S. Choi, and Y. W. Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *ICML*, pages 3744–3753, 2019. 4
- [29] Min Li, Zhenglong Zhou, Zhe Wu, Boxin Shi, Changyu Diao, and Ping Tan. Multi-view photometric stereo: A robust solution and benchmark dataset for spatially varying isotropic materials. *IEEE Transactions on Image Processing*, 29:4159–4173, 2020. 2
- [30] Zhengqin Li, Mohammad Shafiei, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and svbrdf from a single image. In *CVPR*, pages 2475–2484, 2020. 3
- [31] Daniel Lichy, Jiaye Wu, Soumyadip Sengupta, and David W Jacobs. Shape and material capture at home. In *CVPR*, pages 6123–6133, 2021. 3
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 4, 6
- [33] Fotios Logothetis, Ignas Budvytis, Roberto Mecca, and Roberto Cipolla. Px-net: Simple and efficient pixel-wise training of photometric stereo networks. In *CVPR*, pages 12757–12766, 2021. 5
- [34] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. 2019. 5

- [35] Feng Lu, Yasuyuki Matsushita, Imari Sato, Takahiro Okabe, and Yoichi Sato. Uncalibrated photometric stereo for unknown isotropic reflectances. In *CVPR*, pages 1490–1497, 2013. 2
- [36] Zhipeng Mo, Boxin Shi, Feng Lu, Sai-Kit Yeung, and Yasuyuki Matsushita. Uncalibrated photometric stereo under natural illumination. pages 2936–2945. IEEE Computer Society, 2018. 1, 2, 3, 5, 6, 7
- [37] T. Papadhimetri and P. Favaro. A new perspective on uncalibrated photometric stereo. In *CVPR*, 2013. 2
- [38] J. Park, S. Sinha, Y. Matsushita, Y. Tai, and I. Kweon. Multi-view photometric stereo using planar mesh parameterization. *ICCV*, 2013. 2
- [39] H. Santo, M. Samejima, Y. Sugano, B. Shi, and Y. Matsushita. Deep photometric stereo network. In *International Workshop on Physics Based Vision meets Deep Learning (PBDL) in Conjunction with IEEE International Conference on Computer Vision (ICCV)*, 2017. 1
- [40] Hiroaki Santo, Michael Waechter, and Yasuyuki Matsushita. Deep near-light photometric stereo for spatially varying reflectances. In *ECCV*, pages 137–152, 2020. 1, 2
- [41] B. Shi, Y. Matsushita, Y. Wei, C. Xu, and P. Tan. Self-calibrating photometric stereo. In *CVPR*, 2010. 2
- [42] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. Bi-polynomial modeling of low-frequency reflectances. *IEEE TPAMI*, 36(6):1078–1091, 2014. 1
- [43] Konstantin Sofiiuk, Ilia Petrov, and Anton Konushin. Reviving iterative training with mask guidance for interactive segmentation. *arXiv preprint arXiv:2102.06583*, 2021. 3
- [44] T. Taniai and T. Maehara. Neural Inverse Rendering for General Reflectance Photometric Stereo. In *ICML*, 2018. 1, 3, 6
- [45] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 4
- [46] P. Woodham. Photometric method for determining surface orientation from multiple images. *Opt. Engg.*, 19(1):139–144, 1980. 2
- [47] Z. Wu and P. Tan. Calibrating photometric stereo by holistic reflectance symmetry analysis. *CVPR*, 2013. 2
- [48] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *ECCV*, pages 418–434, 2018. 4
- [49] Pengchuan Zhang, Xiyang Dai, Jianwei Yang, Bin Xiao, Lu Yuan, Lei Zhang, and Jianfeng Gao. Multi-scale vision longformer: A new vision transformer for high-resolution image encoding. *ICCV*, 2021. 6