

Uni MS-PS: a Multi-Scale Encoder Decoder Transformer for Universal Photometric Stereo

Clément Hardy¹, Yvain Quéau¹, and David Tschumperlé¹

¹Normandie Univ, UNICAEN, CNRS, ENSICAEN, GREYC laboratory, Caen, France

Abstract

Photometric Stereo (PS) addresses the challenge of reconstructing a three-dimensional (3D) representation of an object by estimating the 3D normals at all points on the object’s surface. This is achieved through the analysis of at least three photographs, all taken from the same viewpoint but with distinct lighting conditions. This paper introduces a novel approach for Universal PS, i.e., when both the active lighting conditions and the ambient illumination are unknown. Our method employs a multi-scale encoder-decoder architecture based on Transformers that allows to accommodate images of any resolutions as well as varying number of input images. We are able to scale up to very high resolution images like 6000 pixels by 8000 pixels without losing performance and maintaining a decent memory footprint. Moreover, experiments on publicly available datasets establish that our proposed architecture improves the accuracy of the estimated normal field by a significant factor compared to state-of-the-art methods. Code and dataset available at: <https://clement-hardy.github.io/Uni-MS-PS/index.html>

1 Introduction

Photometric stereo (PS) is a technique for recovering surface normals of an object by capturing multiple images of it from the same perspective but under varying light conditions. For decades, traditional image processing methods have focused on the ideal Lambertian case with a controlled and parallel light beam as well as no ambient light [42]. However in practice most light effects on real-world objects deviate from Lambert’s law, exhibiting complex effects such as specular components or translucency (e.g., transparent materials). On the other hand, the emergence of deep learning approaches has enabled significant advancements in managing more complex geometries and challenging objects that do not adhere to Lambert’s law.

Three types of approaches are considered in the literature to adress the PS problem: calibrated, uncalibrated, and Universal methods. The difference between calibrated and uncalibrated methods lies in whether we know the light parameters (positions, intensities,...). Additionally, most of these methods (uncalibrated or calibrated) assume the ideal case of perfect directional lighting in a dark environment with no external light. Obtaining this ideal case in real life is challenging, requiring special equipment to capture images under such conditions. Universal methods overcome this limitation by reconstructing objects in any lighting conditions, thus largely simplifying the process from the end-user perspective. They simultaneously address two major challenges:

- reconstructing the normal map for non-Lambertian materials like specular ones;
- handling complex illumination conditions, including ambient illumination.

In our conference paper [9], we introduced a multi-scale approach to improve the performance of *calibrated* PS on challenging materials. In the present article, we extend this multi-scale approach to solve the Universal PS problem. We define a multi-scale architecture combined with an encoder-decoder architecture. The multi-scale architecture can process input images of any size without loss of performance, even when considering very high-resolution images, as presented in Figure 1. In this example, our algorithm takes 11 images of size 6000×8000 as input, and we observe that our method recovers all details of the scene. Our method is also able to manage more difficult materials as illustrated in Figure 2. Our Universal method achieves state-of-the-art results in any environment or lighting conditions, including directional or non-parallel lighting beams, as demonstrated on real benchmarking datasets.




Figure 1: Reconstruction of a normal map of size 6000×8000 of the Marsoulas cave by photometric stereo with the proposed method. Our method is able to both achieve accurate 3D-reconstruction and manage very high image resolution. Indeed, all details in the rock are reconstructed and the resolution of the RGB images is preserved. PS image credits: A. Laurent 2023 (INPT, UMR 5505 IRIT), C. Fritz and G. Tosello team (CREAP-E.Cartailhac), MSHS-T (UAR 3414).



Figure 2: Reconstruction of a normal map of a snail from [38] and a coin from [39] by photometric stereo with our proposed method. We can see that our method manages very difficult materials such as metal.

The main advantage of our proposed extension of [9] is the ability to solve the PS problem *in any light conditions*, without any prior knowledge of the environment (our previous method was only capable of handling calibrated PS in a controlled dark environment). The extension includes new two major features:

1. A Transformer-based approach is chosen for its effectiveness in Universal PS, instead of the CNN approach in [9].
2. A new training dataset is designed and synthesised to better cope with the Universal context.

The rest of this work is organised as follow. In Section 2, we present an overview of deep learning methods for photometric stereo. In Section 3, we describe our multi-scale Transformer method and our new training dataset. Finally, in Section 4, we present qualitative and quantitative results on many benchmark datasets and compare the performance of our method with state-of-the-art PS methods.

2 Related Work

Deep learning methods for PS proposed in the literature initially considered specific lighting conditions. Most of them, such as [6],[5], [12, 14],[43] and [44], consider the light beam as parallel, which is not realistic in practice. [24] introduced methods for non-parallel light beams, but they still require images captured in a very controlled environment. All of these constraints limit their real-world applicability.

Calibrated PS — To solve the problem of calibrated photometric stereo, multiple learning-based methods have been introduced in the literature.

[34] propose a first approach based entirely on a fully connected network. This approach assumes that the light directions are fixed and identical between the training and testing phases, which is its major drawback. Photometric stereo methods should be able to handle a different number of input images. Indeed, to make PS methods more suitable for experimentation, the neural network architecture should allow the use of an arbitrary number of images to avoid training a different neural network for each possible number of input images. Two main alternatives have been proposed in the literature to solve this problem:

- the first approach is to aggregate information from different images using a pooling layer, such as in the work of [6, 5], [9], [16, 17],[25] and [41].
- the second approach is to project all observations corresponding to the same pixel location under different illuminations onto a fixed-size space, such as in the work of [12, 13],[21],[27] and [44].

More recently, [14] show the relevance of using Transformers and attention mechanisms in the context of calibrated PS with a small number of different lights.

Uncalibrated PS — Uncalibrated PS is a category of PS where the prior light information, such as its direction and intensity, is unknown. In the context of parallel light beams, a common practice is to use a first neural network to infer the missing light information for a standard calibrated PS neural network as presented in [5]. Then, a second neural network handles the problem as in the calibrated case. This approach has also been successfully applied to non-parallel light beams in [24]. However, it is difficult to apply this approach to Universal PS with natural light/ambient light because the physics of natural light/ambient light is complex to model.

Another practice to solve uncalibrated PS is to use an inverse rendering based method[23],[22],[19]. Those methods optimize an image reconstruction loss (between the reconstructed images and the input images) to get the normal map, albedo...

Universal PS — Recently, Ikehata introduced the notion of Universal PS with the UNI-PS[11] and SDM-UniPS methods[15]. These new methods solve the PS problem under unknown and arbitrary lighting conditions using a pure data-driven approach without complex prior light assumptions. This method is based on an encoder-decoder model, where the encoder extracts a global lighting context from a fixed 'canonical' resolution image. Images are resized if their size does not match the canonical resolution. The decoder takes as input the original resolution images and the output of the encoder, i.e., the global lighting context interpolated to the original resolution. This allows for inference on very high-resolution images, as the decoder processes the images pixel by pixel.

The idea behind using a global lighting context, rather than a global lighting model, is due to the spatially-varying light direction. Indeed, intensity could not be encoded by a few global values. Even though this method allows for processing high-resolution images by downsampling and inferring




Figure 3: Global architecture of the proposed multi-scale method. The network can be either our calibrated Transformer, or our Universal Transformer. The first scale is independent from the other scale, then the others scales share parameters.

the normal map pixel by pixel, this technique decreases the performance of the reconstruction due to the loss of information during downsampling and the lack of spatial information during pixel-by-pixel inference.

To address this problem, [15] introduced a way to use all available information in a non-local way, even on very high-resolution images. The method is based on a scale-invariant spatial-light feature encoder, which allows for a fixed input size without resizing the images. The encoder splits an image into P^2 sub-images, where P is the size of the input of the model, by taking a single pixel every $P \times P$ pixels. It then extracts feature maps from these sub-images, which are eventually merged back to reconstruct one image.

During the encoding phase, spatial information is extracted using ConvNeXt layers [26] and information over the light axis is extracted using Self-Attention Blocks [20]. Afterwards, another pixel sampling strategy is used and several Transformer layers are applied in both the spatial and light dimensions. Finally, the normal map is inferred using two linear layers.

UNI-PS [11] tends to infer blurry and inaccurate normal maps with a lack of detail, especially with high-resolution images. SDM-UniPS [15] produces better normal maps, but its performance decreases as object materials become more complex. In [15], Ikehata explains the importance of objects and materials diversity in the training dataset and so introduce a new dataset for training UNI-PS [11].

Moreover, we have showed in [9] the more diverse and representative the training dataset is, the better the results are. To this end, we generate in the present paper a way more diverse, complete training dataset for Universal PS. Additionally, processing objects with complex materials, such as highly specular materials, requires both global and local information. We also showed in [9] that a multi-scale network architecture can extract both local and global information and gives thus gives excellent performance on non-Lambertian reflectance materials.

While UniPS allows for the processing of high-resolution images, it may be difficult to scale to very high resolutions without losing performance. Indeed, keeping only one pixel every P pixels is a problem if P is large (for instance if $P \geq 100$). For example, the geometry of a small detail in the object would be completely invisible in each of the sub-images.

To be able to infer on very high resolution images, we propose here to use a multi-scale approach as introduced in [9] and to extend it to the Universal problem. In [9], we use an architecture based on the convolution network proposed by [6]. It consists of an encoder and a decoder. Each input image is processed independently by the encoder. Then, the extracted feature maps are synthesized using max pooling to create a single feature map for all images. The decoder takes this feature map to generate an estimation of the normal map. In the next section, we show how to enhance this architecture in a multi-scale manner, and extend it to the universal setup by relying on Transformers.

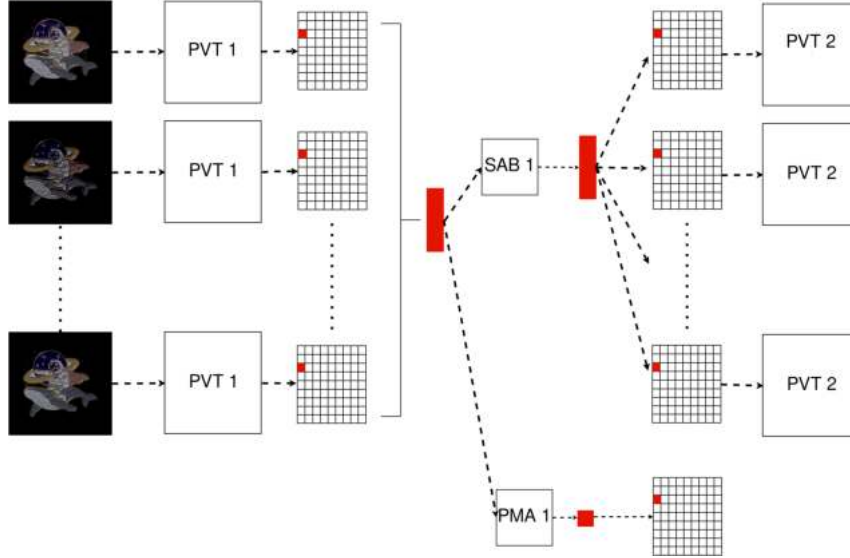


Figure 4: Detailed architecture of one scale in the encoder part of our proposed method. We process images at pixel level in order to catch all geometric details.

3 Proposed Multi-Scale Method

In this section, we first present our multi-scale architecture, and then introduce our new synthetic training dataset.

We propose a multi-scale framework to perform equally well on both low-frequency geometry and high-frequency details. Our multi-scale approach progressively refines the estimated normal map as the spatial scale increases. Our model starts by focusing on the global aspect of the object and then progressively refines details such as holes, cracks, or slight bumps. While a model should be indeed able to handle high-resolution images and different resolutions, a model with a fixed number of convolution layers may still lack sufficient convolutions to effectively synthesize information across an entire arbitrary large image.

To address this, we separate the first scale of our network from the other scales. The first network, used for the first scale, takes as input the downsampled images at the resolution 32×32 pixels and estimates a first normal map at the same resolution. Actually, the first network is used to predict a normal map which can be seen as a initialization prediction for the second network. The latter is used for the remaining scales of our multi-scale approach. It iteratively refines the normal map estimation each time by a factor of 2. Thus, it takes as input an upsampled version of the normal map estimated at the lower resolution as well as images downsampled to the same resolution. A new normal map estimation is predicted at the same resolution as the input images. This process is repeated until we get a normal map at the same resolution of the original images. The key point is that both networks (i.e., for first scale and other scales) have exactly the same architecture, but with different weights. It is necessary to have two independent architectures because the first scale takes as input only the images, while the other scales take as input the concatenation of the images and an estimation of the normal maps. The weights of the second network are identical for every scale.

Our overall architecture is presented in Figure 3.

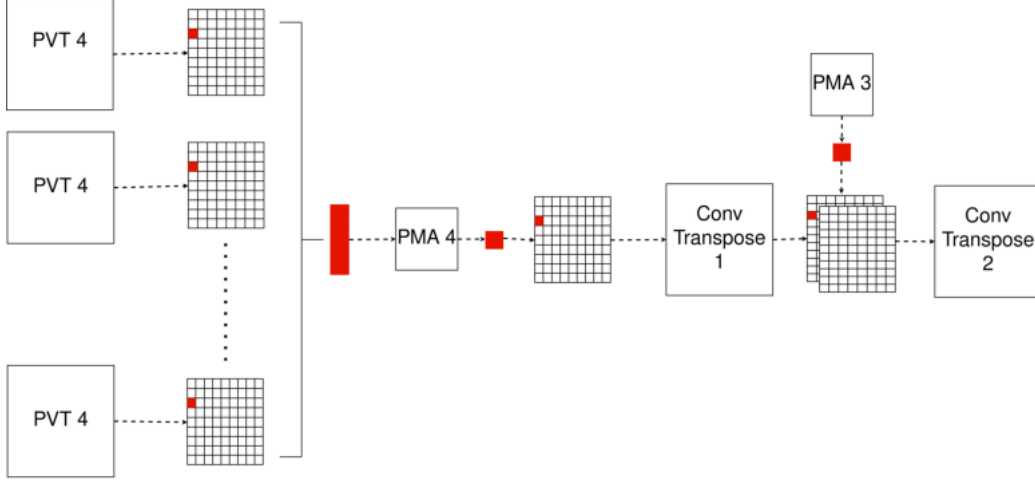


Figure 5: Architecture of the decoder for one scale of the proposed method. The encoder part is detailed in Figure 4 and the decoder part is composed of transposed convolution.

3.1 Transformers-based Multi-Scale Architecture for Universal PS

Our proposed architecture is a multi-scale encoder-decoder backbone composed of Pyramid Vision Transformer (PVT) blocks [40], Self Attention Blocks (SAB) [20] and Pooling by Multihead Attention (PMA) blocks [20].

Additionally, to test the robustness and performance of our architecture on different PS problems, we propose a variant for solving the calibrated problem. The only difference is the first convolutional embedding layer, which is modified to take either only the images for Universal PS or the images concatenated with the lighting directions for calibrated PS.

Scales Architecture of the Encoder Part — Here we detail the structure of one scale of our proposed method. Each scale is composed of the same encoder and the same decoder. The encoder part combines three modules: the first one extracts the spatial information for each image independently, the second one extracts the lighting information for all images at each pixel location, and the third one pools the information for the skip connections. The decoder part is mainly composed of regression modules.

The spatial extractor module is based on the PVT (Pyramid Vision Transformer) [40]. Indeed, this kind of architecture generates high-resolution features and also features at different scales, allowing us to consider problems at the pixel level. The main advantage is the ability to take in input images of different sizes while keeping moderate computation times. This last point is very important for the photometric stereo problem, because it is necessary to consider the full size of the images to get a better reconstruction of the normal map.

Then, the lighting extractor module extracts light information at the pixel level. To do so, we use a Self Attention Block (SAB) module. Indeed, at a fixed pixel location, we concatenate the pixel value of each image in order to merge the information at this location. Therefore, we can apply the attention block at each pixel location independently.

Finally, we use a Pooling by Multihead Attention (PMA) module in parallel with the SAB module to aggregate the information given by the PVT block. The aim is to create a feature map and use it for the skip connection in the decoder part.

Figure 4 presents the architecture of the encoder for one scale of our proposed method.

Scales Architecture of the Decoder Part — Once the four encoding blocks are processed, the normal maps are reconstructed with the decoder. For the decoder, we consider three transposed convolutions with skip connections to the PMA map. Indeed, at each step, we concatenate the PMA map obtained in the encoder with the output of the transposed convolution, and so on until we have the desired resolution. The final step consists of a 3×3 convolution to fuse the first PMA map with the output of the last transposed convolution without changing the shape of the feature map, and to create the final normal map.

The decoder architecture is presented in Figure 5.

3.2 Training Dataset

To obtain the best normal map reconstruction possible, a proper dataset needs to be used for the training stage. Most available training datasets are built for photometric stereo in dark environments with parallel light beams [6], [12]. For Universal PS, Ikehata introduced the PS-Wild training dataset [11]. Unfortunately, this dataset has some issues, such as a lack of diversity in geometry, materials, and environments (see Table 1). This appears to be not enough to calibrate a neural network properly to be able to handle all possible materials and geometries.

Training database	samples	shapes	materials	ambient environments
PS-Wild	10 099	410	926	31
Our training database	100 000	11 000	200 000	1 100

Table 1: Comparison between our training dataset and PS-Wild. Our training dataset proposes more objects with a larger variety of geometries, shapes and environment than PS-Wild [11].

To solve these issues, we create a new training dataset. To do so, we render 11,000 diverse objects from [35], [37], and [31] using the Blender software [3]. To complete the lack of smooth surfaces that can exist on these types of objects, we also generate 3,000 distinct objects using the sum of random Gaussian potentials and the Marching Cubes algorithm [29] to extract isosurfaces.

Each time we render a scene, we apply a random material to the object. For materials, we use more than 1,100 ‘real’ materials taken from [2] and around 200,000 materials from Deep-materials [8]. Moreover, we generate random materials to complete all possible materials. These new materials are created with the BRDF layer of the Cycle render engine [7] in Blender by choosing random values as inputs of this layer.

Finally, for the ambient lighting environment, we use 1,100 360° HDR (High Dynamic Range) images from diverse sources, such as [32], [2] and [1].

For each sample, we render 50 images with random light distribution over the hemisphere. To ensure that our model can handle different types of lights, we use directional lights and non-parallel lights. Each time the non-parallel light type is chosen for the scene, the size of the bulb and other light parameters are also chosen randomly. On the contrary, the power of the light varies between each image, regardless of the type of lamp chosen. In total, we generate 100,000 samples. A comparison between our training dataset and PS-Wild is shown in Table 1.

Examples of training images generated by our pipeline are given in Figure 6. As we can see, materials, shapes, geometries and environments are diverse.




Figure 6: Examples of training images generated by our pipeline. We can see the variety of geometries, shapes and environments.

3.3 Training Process

To train our multi-scale network, we use images with a resolution of 128 by 128 pixels. To reach the resolution of 128 pixels, 3 stages are necessary: 32 by 32 pixels, 64 by 64 pixels, and 128 by 128 pixels. Because of the small resolution of our training images, we are able to give 23 images per view during the training process. A batch size of 2 is enough for the training. Therefore, we are able to use a single A100 to train our method. The Adam optimizer is used with a learning rate of 10^{-4} . The three stages are trained together, and the cosine similarity loss is used. This loss measures the angular difference between the estimated 3D normals and the ground truth normals. It is defined as:

$$l_{normal} = \sum_{ij} \left(1 - N_{ij} \cdot \hat{N}_{ij} \right), \quad (1)$$

where \hat{N}_{ij} is the predicted normal at pixel (i, j) , and N_{ij} is the ground truth normal. Everything was implemented with the Pytorch framework.

3.4 Inference on Very High Resolution Image

As mentioned previously, our method can be used for any size of input image. However, performing inference on very high-resolution images is challenging, because even with a batch size of 1, the image may not fit on a single graphics card.

Therefore, to run our network on very high-resolution images, we embed our multi-scale approach in a patch-based heuristic. For images up to 256x256 pixels (i.e., 32 by 32, 64 by 64, 128 by 128, and 256 by 256), we use the full resolution. For larger images, we cut each image and its corresponding predicted normal map into 256x256 patches with an overlap of 64 pixels. We then process each patch independently. Finally, we merge all patches together using a spatial weighted average, with Gaussian weights as defined below:

$$w(x, y) = e^{-\frac{\| (x, y) - (x_c, y_c) \|^2}{2\sigma^2}}, \quad (2)$$

where (x_c, y_c) is the center of the patch and $\sigma = 25$. This value has been chosen empirically.

This method allows us to avoid computing the attention map on the full resolution image. Indeed, computing the attention map can significantly increase the memory requirement in the PVT module when the image size increases.

However, the results in Table 2 show that performance can degrade if the patch size is too small. Therein, we tested our network combined with the proposed patch-based inference on DiLiGenT10² [33] at the full image resolution with 30 images per object.

Finally, we choose to take a patch size equal to 256 as it seems to be a reasonable compromise between memory cost and accuracy.

Patch size (px)	Overlap (px)	Memory usage (Go)	MAE (°)
128	32	3.5	15.52
256	64	21	13.19
512	128	130	12.96

Table 2: Memory usage and mean angular error of the proposed patch inference method on the full image resolution of DiLiGenT10². A patch size of 256 px seem to be a good compromise between performance and memory cost.

4 Experiments

From a quantitative point of view, we compare our approach with all state-of-the-art methods. Indeed, we consider calibrated [12],[4],[10],[27],[18],[24],[28],[9], uncalibrated [5],[22],[24] and universal methods [11],[15].

4.1 Description of the testing datasets

We evaluate quantitatively our method on the three publicly available dataset with directional light DiLiGenT [36], DiLiGenT10² [33] and DiLiGenT-Pi [39]. We also test the generalization capacity of our method on a dataset with non-parallel light directions named Lucas [30]. Examples of images of each dataset are presented in Figure 7.




Figure 7: Examples of images of DiLiGenT [36], DiLiGenT10² [33], DiLiGenT-Pi [39] and Lucas [30] datasets.

- The DiLiGenT dataset [36] is a real-world image dataset containing 96 images from the same viewpoint captured under known light directions and light intensities. It contains 10 objects with ground truth normal maps, obtained by scanning objects with a 3D scanner.

- DiLiGenT10² [33] contains 10 objects, each manufactured in 10 different materials. The diversity of materials in this dataset is quite large. Indeed, diffuse, moderately specular, metallic materials with anisotropic reflectance, and translucent materials are all present in this dataset. This diversity offers the possibility to test the performance of our algorithms on challenging materials. The ground truth is also available, but was obtained using 3D digital models and not by a 3D scanner as in the DiLiGenT dataset.

- The last dataset with directional light is DiLiGenT-Pi [39]. This dataset was created to test photometric stereo methods on near-planar surfaces with rich details, such as coins and badges. It contains four groups of materials: metallic, specular, rough, and translucent surfaces. In addition, the dataset contains 30 objects, each with 100 photographs provided. As with the DiLiGenT dataset, the ground truth normals were obtained using a scanner.

- Finally, Lucas [30] is a dataset with non-parallel and near-lighting. It contains 14 objects and 52 images per object. The light directions, intensities, and camera calibration are known. As with the other datasets, the ground truth normal maps are available and were obtained using a scanner.

All of these datasets offer the opportunity to test our method on a wide variety of object shapes, materials, contexts, and so on. However, to complete our evaluation and test the performance of our method on different types of light, environments, and cameras, we also visually test the performance on publicly available datasets where no ground truth is given, such as Skoltech3D [38], Shape and Material [25], UNI-PS, [11], and SDM-UniPS [15]. The image acquisition setup varies from dataset to dataset. In

Skoltech3D [38], an industrial camera is used and images are captured in the dark with directional light. In UNI-PS [11], an 8-bit smartphone camera is used with near light (within 30 cm of the object) to have spatially varying lighting effects. In SDM-UniPS [15] a digital camera is used, and in Shape and Material [25] an iPhone is used. The variety in the acquisition setup is important for testing the generalization capability of our method to any setup or environment.

	type	Ball	Bear	Buddha	Cat	Cow	Goblet	Harvest	Pot1	Pot2	Reading	average
PS-FCN [4]	C	2.67	7.72	7.52	4.75	6.72	7.84	12.39	6.17	7.15	10.92	7.39
CNN-PS [12]	C	2.2	4.6	7.9	4.1	8.0	7.3	14.0	5.4	6.0	12.6	7.2
OB-Cnn [10]	C	2.49	3.59	7.23	4.69	4.89	6.89	12.79	5.10	4.98	11.08	6.37
PX-NET [27]	C	2.03	3.58	7.61	4.39	4.69	6.90	13.10	5.08	5.10	10.26	6.28
NormAttention-PSN [18]	C	2.93	5.48	7.12	4.65	5.99	7.49	12.28	5.96	6.42	9.93	6.83
Our previous method, MS-PS [9]	C	2.05	4.24	7.03	3.9	4.00	7.57	11.01	4.94	5.22	8.47	5.84
SDPS-Net [5]	UC	2.8	6.9	9.0	8.1	8.5	11.9	17.4	8.1	7.5	14.9	9.5
SCPS-NIR [22]	UC	1.24	3.82	9.28	4.72	5.53	7.12	14.96	6.73	6.50	10.54	7.05
UNI-PS [11]	UC/Uni	4.9	9.1	19.4	13.0	11.6	24.2	25.2	10.8	9.9	18.8	14.7
SDM-UniPS [15]	UC/Uni	1.5	3.6	7.5	5.4	4.5	8.5	10.2	4.7	4.1	8.2	5.8
Our (K=30)	C	1.93	2.64	5.88	3.05	3.76	6.40	10.44	3.85	4.32	7.31	4.96
Our (K=96, all images)	UC/Uni	1.92	3.14	6.16	3.60	4.04	6.35	8.84	4.08	4.88	7.09	5.01
Our (K=30)	UC/Uni	1.84	3.14	6.04	3.45	3.99	6.49	8.9	4.12	4.7	7.0	4.97
Our (K=15)	UC/Uni	1.93	3.05	6.31	3.97	4.06	7.0	9.27	4.25	4.9	7.41	5.22
Our (K=6)	UC/Uni	2.4	3.7	7.14	4.52	4.7	8.06	12.43	5.32	5.84	9.4	6.35
Our (K=3)	UC/Uni	3.58	4.83	11.46	7.13	6.68	17.8	18.05	8.79	7.75	15.65	10.17

Table 3: Mean angular error (in degrees) on the DiLiGenT benchmark [36]. The type C means calibrated PS, UC is uncalibrated PS and Uni is Universal PS as defined in [11]. The best result is indicated in bold, and the second best one is underlined. The proposed method gives best state-of-the-art results.

4.2 Quantitative comparison

We first evaluate our methods on DiLiGenT in Table 3. We compare the performance of both our Universal and calibrated Transformer methods with calibrated, uncalibrated, and Universal state-of-the-art methods for PS. Our Universal method clearly outperforms all other methods by at least 16% on all objects. Our calibrated method also achieves state-of-the-art performance compared with calibrated methods. Interestingly though our Universal method reaches results comparable to the calibrated one.

In addition, we test our proposed Universal method with different numbers of images (K=3, 6, 15, 30, and 96). We can see that with only 6 images, our method obtains results that are close to the state-of-the-art using all the available images. Moreover, the results are already the best with only 30 images.

Then, we compare our methods on a more challenging dataset, DiLiGenT10² [33], in Table 4. On this dataset, we can see that our Universal method still achieves state-of-the-art results. Our Universal method improves the state-of-the-art results of Universal PS by 13%, from 14.96° to 13.19°. On difficult geometries, like Turbine, the improvement is significant compared to other methods, even compared to calibrated ones. Also we still obtain good results on specular material like AL (aluminium), CU (brass) or STEEL, contrary to other Universal PS methods.

Our transformer calibrated method is the best performer on this dataset (see Table 5). We note that all multi-scale methods get much better results than non-multi-scale methods like CNN-PS [12] which is the best performer of the non-multi-scale methods.

The second challenging dataset is DiLiGenT-Pi [39]. Again, our Universal method outperforms all other Universal and uncalibrated methods, see Table 6. Compared to the calibrated methods, our Universal method tends to have slightly lower performance on near-flat objects, but it still achieves competitive results. We note that the average is not necessarily the best metric to compare the performance of calibrated and uncalibrated methods. Indeed, for some objects all uncalibrated or Universal state-of-the-art methods predict an inverted normal map compared to the ground truth (for example, see Figure 9). This is likely because uncalibrated methods are unable to determine the direction of incoming light and tend to assume that it is coming from the opposite direction to the actual direction. As shown in Figure 9a, it is difficult to tell if the light is coming from above or below. Both possibilities are equally plausible, but would result in opposite normal maps.

Our methods are way more robust to this problem than other uncalibrated and Universal methods, as we only have 2 objects inverted compared to 11 for SDM-UniPS [15] and 8 for SDPS-NET [5]. Indeed, as shown in Figure 10, our uncalibrated methods is able to predict correctly the normal map contrary to SDM-UniPS [15] and Uni-PS [11].

In this dataset, our calibrated Transformer gives also very good results. Overall, our calibrated Transformer method gives a significant improvement of 12% compare to the second best. And again, all

mean: 14.96

	POM	PP	NYLON	PVC	ABS	PAKELITE	AI	CU	STEEL	ACRYLIC
BALL	1.7	1.2	2.5	2.8	2.7	2.4	2.8	5.0	6.9	3.6
GOLF	12.0	6.2	13.0	5.0	12.0	7.1	6.7	6.3	7.5	9.2
SPIKE	12.0	6.8	11.0	6.5	8.7	7.6	8.4	5.7	9.2	13.0
NUT	16.0	5.1	18.0	4.7	8.4	4.8	18.0	13.0	7.5	20.0
SQUARE	23.0	5.4	25.0	5.3	12.0	7.3	19.0	5.5	20.0	32.0
PENTAGON	24.0	8.1	29.0	8.4	13.0	10.0	25.0	26.0	25.0	29.0
HEXAGON	18.0	6.5	20.0	4.8	11.0	7.0	20.0	14.0	19.0	34.0
PROPELLER	28.0	8.3	44.0	6.1	24.0	7.2	22.0	12.0	19.0	28.0
TURBINE	46.0	10.0	51.0	9.4	36.0	11.0	31.0	25.0	25.0	31.0
BUNNY	36.0	8.9	44.0	6.3	19.0	8.1	27.0	7.8	11.0	28.0

(a) SDM-UniPS [15] (Universal)

mean: 13.19

	POM	PP	NYLON	PVC	ABS	PAKELITE	AI	CU	STEEL	ACRYLIC
BALL	7.4	6.3	7.5	8.4	9.9	7.2	9.0	9.8	13.0	43.0
GOLF	14.0	7.6	14.0	5.3	11.0	7.6	8.0	8.2	9.6	38.0
SPIKE	9.3	7.1	10.0	6.7	7.4	6.7	11.0	7.7	13.0	29.0
NUT	16.0	6.5	20.0	5.8	8.8	8.1	9.9	10.0	9.1	35.0
SQUARE	18.0	6.0	21.0	7.4	9.1	6.5	6.7	5.5	7.4	35.0
PENTAGON	20.0	8.8	24.0	9.9	12.0	9.7	13.0	13.0	12.0	40.0
HEXAGON	17.0	8.2	18.0	5.5	11.0	6.2	9.4	8.0	7.9	41.0
PROPELLER	15.0	8.1	20.0	6.2	8.4	7.3	16.0	9.2	12.0	21.0
TURBINE	30.0	9.9	33.0	9.6	15.0	10.0	22.0	13.0	17.0	33.0
BUNNY	15.0	8.0	23.0	6.0	7.2	8.1	10.0	7.8	8.2	12.0

(b) Our Universal transformer

Table 4: Mean angular error (in degrees) on the DiLiGenT10² benchmark, with the results of SDM-UniPS [15] indicated for comparison. The lower the values, the better the results are. Our Universal method gives best state-of-the-art results.

our multi-scale architectures obtain the best results in their categories.

Finally, our Universal method is able to also manage non-parallel light beam as showned in Table 7. We obtain the best performance over other Universal PS methods and uncalibrated methods especially built for non parallel scenario. Compared to the best calibrated results our method is very close in term of performances. Our method reaches comparable result with [24], yet these methods use an explicit lighting model whose parameters must be calibrated (location, intensities, dispersion,...).

Inference with no mask — One advantage of SDM-UniPS of [15] compared to the other PS methods is its ability to solve the PS problem without using any object mask. This is a novel feature for deep learning-based methods, as common deep learning PS methods require to mask the background to work properly. To test this type of inference only the quantitative datasets DiLiGenT [36] and Lucas [30] are suitable for use, as the backgrounds in the other datasets are completely dark or cropped.

We infer the normal map without masking the background and then compute the normal error only on the object part. This technique allows us to test the impact of masking the background for our Universal method. The performance decreases without masking the background (see Table 8 and Table 9), but our method still maintains really good results and outperforms SDM-UniPS [15], the only one so far able to manage inference with no mask. In Figure 8, we show an example on two objects: the Reading object from DiLiGenT [36] and the Alligator from SDM-UniPS [15]. We can see that our Universal method not only generates a proper normal map for the desired object, but also reconstructs the background correctly, which is not necessarily the case for SDM-UniPS [15]. For example, the carpet is reconstructed much better with our Universal method.

mean: 15.78

	POM	PP	NYLON	PVC	ABS	PAKELITE	AI	CU	STEEL	ACRYLIC
BALL	5.1	6.4	4.2	4.5	6.9	7.3	16.0	14.0	16.0	19.0
GOLF	14.0	8.0	12.0	6.8	14.0	9.4	12.0	9.2	13.0	22.0
SPIKE	11.0	9.4	11.0	11.0	12.0	9.5	14.0	8.3	16.0	28.0
NUT	20.0	8.8	19.0	6.9	17.0	8.0	16.0	13.0	14.0	22.0
SQUARE	21.0	8.1	22.0	6.7	19.0	8.1	13.0	4.9	7.9	18.0
PENTAGON	26.0	9.5	26.0	9.8	22.0	9.6	15.0	13.0	15.0	23.0
HEXAGON	18.0	7.5	19.0	7.2	17.0	28.0	18.0	10.0	17.0	21.0
PROPELLER	28.0	12.0	35.0	8.4	23.0	11.0	16.0	9.6	9.8	17.0
TURBINE	54.0	20.0	51.0	16.0	39.0	21.0	25.0	22.0	21.0	32.0
BUNNY	24.0	11.0	27.0	7.8	21.0	9.1	12.0	7.7	12.0	14.0

(a) CNN-PS [12] (calibrated)

mean: 11.33

	POM	PP	NYLON	PVC	ABS	PAKELITE	AI	CU	STEEL	ACRYLIC
BALL	9.3	3.4	8.7	5.2	8.4	4.3	8.5	12.0	14.0	8.6
GOLF	10.0	7.3	9.8	5.8	10.0	6.87	7.9	7.7	9.8	12.0
SPIKE	12.0	8.8	9.9	6.3	8.5	7.9	12.0	7.6	12.0	17.0
NUT	14.0	8.9	15.0	5.8	10.0	5.8	9.2	7.6	8.2	16.0
SQUARE	18.0	11.0	17.0	8.2	14.0	5.5	12.0	7.2	7.9	11.0
PENTAGON	18.0	8.4	17.0	8.0	17.0	9.4	11.0	9.4	13.0	20.0
HEXAGON	16.0	7.5	15.0	6.1	13.0	7.1	11.0	8.1	11.0	20.0
PROPELLER	13.0	8.9	11.0	7.9	16.0	9.7	11.0	8.4	8.2	19.0
TURBINE	21.0	12.0	24.0	11.0	18.0	15.0	23.0	16.0	18.0	22.0
BUNNY	17.0	8.2	16.0	6.5	12.0	8.3	8.6	7.3	8.0	18.0

(b) Our previous method, MS-PS [9] (calibrated)

mean: 11.01

	POM	PP	NYLON	PVC	ABS	PAKELITE	AI	CU	STEEL	ACRYLIC
BALL	4.5	3.3	5.0	3.6	5.8	4.1	3.6	7.8	8.8	6.4
GOLF	13.0	6.4	14.0	5.1	11.0	6.8	7.0	6.4	8.1	9.3
SPIKE	10.0	7.3	11.0	7.5	9.1	8.3	8.5	8.3	9.0	11.0
NUT	11.0	5.0	19.0	4.5	8.1	5.0	6.6	6.8	6.4	24.0
SQUARE	18.0	8.5	23.0	7.7	13.0	7.1	7.9	5.0	7.6	19.0
PENTAGON	15.0	8.3	22.0	8.9	13.0	8.4	11.0	9.5	9.5	21.0
HEXAGON	16.0	5.8	20.0	5.9	12.0	6.4	7.1	5.2	6.7	20.0
PROPELLER	16.0	9.2	32.0	8.2	9.6	6.9	15.0	7.8	10.0	17.0
TURBINE	30.0	9.4	30.0	10.0	19.0	9.8	22.0	15.0	16.0	23.0
BUNNY	12.0	8.0	29.0	6.0	8.0	8.3	9.4	8.5	8.4	13.0

(c) Our calibrated transformer

Table 5: Mean angular error (in degrees) on the DiLiGenT10² benchmark, with the results of CNN-PS [12] and our previous multi-scale CNN [9] indicated for comparison. The lower the values, the better the results are. Our calibrated Transformer method gives best state-of-the-art results.

	Type	Astro Lung	Bagua-R Ocean	Bagua-T Panda-R	Bear Panda-T	Bird Para	Cloud-R Queen	Cloud-T Rhino	Crab Sail	Fish Ship	Flower Sun	Lion-R TV	Lion-T Taichi	Lions Tree	Lotus-R Wave	Lotus-T Whale	average
NormAttention-PSN [18]	C	7.2	7.8	13.9	16.6	4.2	4.9	<u>5.1</u>	<u>5.2</u>	4.9	5.6	7.6	9.7	9.6	6.1	8.7	9.2
PS-FCN [6]	C	7.2	7.8	13.0	16.8	7.4	7.2	14.3	17.8	5.3	4.6	4.6	18.4	21.2	<u>4.5</u>	11.8	13.6
	C	9.7	5.8	14.8	17.2	4.7	4.7	5.3	<u>5.1</u>	6.1	6.7	<u>8.0</u>	10.2	10.6	6.8	12.2	9.85
CNN-PS [12]	C	<u>6.0</u>	12.2	16.4	7.4	<u>5.8</u>	14.6	17.2	<u>4.5</u>	<u>4.2</u>	<u>4.7</u>	15.8	20.3	4.7	10.9	13.5	9.16
	C	<u>5.7</u>	4.6	14.2	16.6	3.9	5.4	4.9	<u>5.2</u>	4.9	<u>5.3</u>	8.3	<u>7.8</u>	11.3	<u>5.3</u>	11.6	
Our previous method, MS-PS [9]	C	5.96	11.32	15.1	<u>6.9</u>	7.69	13.28	14.74	4.58	4.68	5.43	14.37	15.71	5.5	11.92	12.8	<u>8.78</u>
	C	7.51	<u>4.97</u>	14.75	14.72	<u>4.09</u>	6.37	5.18	5.26	5.14	6.46	8.63	9.91	<u>8.22</u>	<u>5.29</u>	<u>7.09</u>	
SDPS-Net [5]	UC	37.7	22.5	28.9	30.7	17.6	27.4	27.5	20.5	23.6	12.8	20.8	23.6	19.6	21.7	26.5	25.93
	UC	40.2	31.4	21.8	23.7	19.8	16.5	24.9	16.7	19.0	31.5	26.9	34.1	41.1	39.1	29.8	
SDM-UniPS [15]	UC/Uni	37.8	14.6	17.1	23.8	26.5	17.1	19.2	25.4	24.5	15.2	15.9	16.2	9.2	11.8	13.6	23.34
	UC/Uni	46.6	34.6	17.1	17.6	23.2	10.6	17.0	10.5	22.0	26.2	36.6	47.2	34.4	34.9	33.8	
Our (k=30)	C	6.03	<u>9.57</u>	11.75	<u>6.72</u>	<u>6.55</u>	<u>12.61</u>	<u>11.01</u>	5.75	<u>4.11</u>	4.85	13.12	11.43	5.37	<u>10.17</u>	<u>8.09</u>	7.75
	C	5.41	5.44	<u>12.98</u>	<u>11.39</u>	4.73	5.69	5.22	6.66	6.25	5.9	10.24	<u>7.26</u>	<u>6.98</u>	<u>5.48</u>	<u>6.71</u>	
Our (k=100, all images)	UC/Uni	7.58	10.19	11.12	<u>12.49</u>	8.14	<u>12.45</u>	<u>11.63</u>	6.0	8.32	5.88	<u>12.66</u>	<u>11.24</u>	6.63	11.29	<u>10.38</u>	11.35
	UC/Uni	42.1	6.35	13.5	<u>11.9</u>	7.2	7.43	6.69	7.2	5.35	6.54	10.39	8.54	47.27	6.11	7.84	
Our (k=30)	UC/Uni	7.14	10.43	<u>11.69</u>	14.09	7.35	13.08	11.92	5.32	5.96	5.14	12.73	11.2	6.16	11.51	10.39	11.38
	UC/Uni	41.98	5.91	<u>13.28</u>	12.22	7.13	9.54	6.68	6.62	5.65	6.05	11.5	8.95	47.15	5.93	8.77	
Our (k=15)	UC/Uni	10.93	10.46	13.44	12.16	8.21	12.71	14.04	8.23	8.76	8.02	14.19	12.2	7.31	11.79	11.19	12.54
	UC/Uni	43.73	10.46	13.81	12.65	7.34	7.68	6.83	8.2	5.99	8.05	11.37	10.45	48.9	7.49	9.49	

Table 6: Mean angular error (in degrees) on the DiLiGenT-Pi benchmark [39]. Best results are in bold, and the second best ones are underlined. The type C means calibrated PS, UC is uncalibrated PS and Uni is Universal PS as defined in [11]. The best result is indicated in bold, and the second best one is underlined. The proposed method gives best state-of-the-art results.

		Ball	Bell	Bowl	Buddha	Bunny	Cup	Die	Hippo	House	Jar	Owl	Queen	Squirrel	Tool	average
Fast-PS (v1) [24]	C	<u>8.55</u>	6.20	7.0	12.69	8.63	17.28	5.16	<u>8.01</u>	29.00	5.32	12.32	12.90	13.00	12.33	11.32
L22 [28]	C	8.84	7.51	5.95	11.59	7.06	15.35	<u>5.19</u>	5.60	22.97	6.19	8.89	9.97	11.77	11.64	9.90
Fast-PS (v2) [24]	UC	6.59	<u>7.17</u>	10.17	14.50	11.75	18.98	8.63	10.64	31.00	9.14	15.92	18.39	15.97	18.61	14.11
UNL-PS [11]	UC/Uni	11.012	24.12	23.84	27.90	23.51	28.64	16.24	21.41	35.93	14.53	32.87	28.36	25.36	19.03	23.77
SDM-UniPS [15]	UC/Uni	13.30	12.76	8.44	18.58	<u>8.53</u>	19.67	7.25	8.86	26.07	8.30	12.67	15.97	16.01	12.54	13.50
Our (K=52, all images)	UC/Uni	10.20	10.52	6.98	12.83	9.60	13.68	6.19	8.33	25.29	6.30	11.47	12.45	11.36	11.79	11.21
Our (K=30)	UC/Uni	10.29	10.51	<u>6.79</u>	<u>12.57</u>	9.6	13.35	6.27	8.44	25.46	<u>6.10</u>	11.38	15.97	<u>11.37</u>	12.22	<u>11.10</u>
Our (K=15)	UC/Uni	10.47	10.8	7.91	13.14	9.90	13.96	6.52	8.54	25.30	6.49	11.82	<u>12.49</u>	11.64	11.89	11.10
Our (K=6)	UC/Uni	10.94	11.40	9.38	13.75	11.029	15.38	7.80	9.41	26.68	7.37	12.62	12.85	12.79	12.47	12.42
Our (K=3)	UC/Uni	10.93	15.95	12.07	16.78	14.53	16.09	9.09	11.06	31.61	10.49	15.73	14.99	15.67	15.69	15.05

Table 7: Mean angular error (in degrees) on the Lucas benchmark [30]. Best results are in bold, and the second best ones are underlined. The type C means calibrated PS, UC is uncalibrated PS and Uni is Universal PS as defined in [11]. The best result is indicated in bold, and the second best one is underlined. The proposed method gives best state-of-the-art results.

		Ball	Bear	Buddha	Cat	Cow	Goblet	Harvest	Pot1	Pot2	Reading	average
Without mask	SDM-UniPS [15]	4.42	4.21	8.54	5.59	7.24	10.37	14.92	5.44	6.72	12.97	8.04
	Our Universal	11.46	4.64	7.46	4.11	7.80	7.14	10.34	5.27	5.59	7.93	7.17
With mask	SDM-UniPS [15]	1.5	3.6	7.5	5.4	4.5	8.5	10.2	4.7	4.1	8.2	5.8
	Our Universal	1.84	3.14	6.04	3.45	3.99	6.49	8.9	4.12	4.7	7.0	4.97

Table 8: Mean angular error (in degrees) on the DiLiGenT benchmark [36] without masking the background before processing. For comparison the results with background is also shown.

		Ball	Bell	Bowl	Buddha	Bunny	Cup	Die	Hippo	House	Jar	Owl	Queen	Squirrel	Tool	average
Without mask	SDM-UniPS [15]	10.45	14.27	10.94	21.29	11.91	10.69	7.56	9.34	27.47	7.33	13.69	16.23	17.16	15.37	13.84
	Our Universal	14.6	13.95	8.29	12.5	8.81	9.19	8.54	9.78	25.52	9.61	12.49	12.26	12.44	16.95	12.49
With mask	SDM-UniPS [15]	13.30	12.76	8.44	18.58	8.53	19.67	7.25	8.86	26.07	8.30	12.67	15.97	16.01	12.54	13.50
	Our Universal	10.20	10.52	6.98	12.83	9.60	13.68	6.19	8.33	25.29	6.30	11.47	12.45	11.36	11.79	11.21

Table 9: Mean angular error (in degrees) on the Lucas benchmark [30] without masking the background before processing. For comparison the results with background is also show.

4.3 Qualitative evaluation

Next, to test the robustness of our Universal method in the most diverse contexts and environments, we use several available qualitative datasets. We compare our Universal method only to SDM-UniPS [15] as all other methods are not Universal. Indeed, Uni-PS [11] is a Universal PS method, but its performance is below SDM-UniPS on all quantitative datasets, so it is not of interest to compare to it. Note that in this section, we only focus on inference with masked backgrounds.

Overall, the results seem good for both methods, but our Universal method outperforms SDM-UniPS on surface details (see Figures 11 and 12). In the Owl object in Figure 11, the results seem similar, but when zooming on the talons, artifacts appear on the prediction of SDM-UniPS [15] which is not the case with our Universal method. Finally, with our multi-scale method, the results remain good regardless of the resolution. For example, in Figure 13 and 14, the images resolution are really high and our Universal method obtains excellent results. Our method performs way better than SDM-UniPS [15] especially on nearly 'flat' objects like the ceramic (Figure 13).




Figure 8: Comparison on the Alligator object of SDM [15] and the Reading object of DiLiGenT [36] without masking background. We can see that considering the background does not degrade reconstruction of normals. We reconstruct more accurate details than SDM [15].

4.4 Limitations

The main limitation of Universal/Uncalibrated methods is the normal map reconstruction on translucent material like acrylic. None of the state-of-the-art methods give accurate reconstruction. Indeed, as the material is translucent, it is very difficult to know from which side the light is coming. For example, in some objects like acrylic balls, the light passes through the ball. So it is actually really hard to determine if the light source is located on the left or the right of the ball. Another example is shown in Figure 9a. Without any prior knowledge of the object shape, it is difficult to find out precisely where the light is coming from. This greatly impacts methods for uncalibrated PS, as the two opposite incoming light directions would lead to the perfectly opposite normals. So, our Universal method can be improved on this type of material.


5 Conclusion

To conclude, we propose a new multi-scale approach based on Transformers with encoder and decoder for each scale. Our method gives excellent results over a large panel of benchmark datasets with a large diversity of acquisition setup and environments which show its robustness. Our method also shows its capacity to manage very high resolution image to get the smallest details of the geometry and to keep very high normal reconstruction performance.


Acknowledgment This work was granted access to the HPC resources of IDRIS under the allocation 2022-AD010613775 made by GENCI.




(a) Gamma corrected image




(b) UNI-PS [[11]]



(c) SDM-UniPS [[15]]



(d) Our calibrated Transformer




(e) Our Universal Transformer


Figure 9: Normal prediction of our methods, SDM-UniPS [15] and Uni-PS [11] on the Lung object of [39]. We can see that this material is challenging for uncalibrated and Universal approaches because of the light reflection. Indeed, normal maps are inverted.




(a) Gamma corrected image




(b) UNI-PS [[11]]



(c) SDM-UniPS [[15]]




(d) Our calibrated Transformer



(e) Our Universal Transformer

Figure 10: Normal prediction of our methods, SDM-UniPS [15] and Uni-PS [11] on the Whale object of [39]. Our Universal method gives the correct normal orientation. The other uncalibrated and Universal approaches fail, inverting the normals orientations.



SDM-UniPS [[15]]

SDM-UniPS [[15]]

Our Universal Transformer

Figure 11: Comparison on objects without ground truth from [11, 15]. The first column is the RGB images, the second one is the SDM method [11] and the last one is our method. Then, for all object, we present the full image and a zoom part. For all objects, we reconstruct more accurate details than SDM [11].






Figure 12: Comparison on objects without ground truth from [38, 25]. The first column is the RGB images, the second one is the SDM method [11] and the last one is our method. Then, for all object, we present the full image and a zoom part on the next line. For all objects, we reconstruct more accurate details than SDM [11].




SDM-UniPS [[15]]




Our

Figure 13: Visual comparison between SDM-UniPS [15] method and our Universal method on the Seasons mosaic object. The image resolution is 5 500 by 8 200 pixels. We can see that our method can manage very high images and outperforms SDM-UniPS [15] in terms of normal map reconstruction. PS image credits: A. Laurent (INRT, UMR 5055 IRIT)/MAN.



SDM-UniPS [[15]]



Our

Figure 14: Visual comparison between SDM-UniPS [15] method and our Universal method on the Sweet object of [15]. The image resolution is 4 000 by 4 000 pixels. We can see that our method can manage very high resolution images and outperforms SDM-UniPS [15] in terms of normal map reconstruction.

References

- [1] Alexandre Duret-Lutz, 2023.
- [2] AmbientCG. <https://ambientcg.com/>.
- [3] Blender-Foundation. *Blender - a 3D modelling and rendering package*. Stichting Blender Foundation, Amsterdam.
- [4] G. Chen, K. Han, B. Shi, Y. Matsushita, and K. K. Wong. Deep photometric stereo for non-lambertian surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(01):129–142, jan 2022.
- [5] Guanying Chen, Kai Han, Boxin Shi, Yasuyuki Matsushita, and Kwan-Yee K. K. Wong. Self-calibrating deep photometric stereo networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8731–8739, 2019.
- [6] Guanying Chen, Kai Han, and Kwan-Yee K. Wong. Ps-fcn: A flexible learning framework for photometric stereo. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Proceedings of the IEEE/CVF European Conference on Computer Vision (ECCV)*, pages 3–19, Cham, 2018.
- [7] Cycles-developper. *Cycles is a physically based production renderer developed by the Blender project*.
- [8] Valentin Deschaintre, Miika Aittala, Frédo Durand, George Drettakis, and Adrien Bousseau. Single-image svbrdf capture with a rendering-aware deep network. *ACM Transactions on Graphics (SIGGRAPH)*, 37(128), 2018.
- [9] Clément Hardy, Yvain Quéau, and David Tschumperlé. MS-PS: A Multi-Scale Network for Photometric Stereo With a New Comprehensive Training Dataset. In *International Conference on Computer Graphics, Visualization and Computer Vision (WSCG)*, volume 3301 of *Computer Science Research Notes*, pages 194–203, 2023.
- [10] D. Honzatko, E. Turetken, P. Fua, and L. Dunbar. Leveraging spatial and photometric context for calibrated non-lambertian photometric stereo. In *Proceedings of the International Conference on 3D Vision (3DV)*, pages 394–402, 2021.
- [11] S. Ikehata. Universal photometric stereo network using global lighting contexts. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12581–12590, 2022.
- [12] Satoshi Ikehata. Cnn-ps: Cnn-based photometric stereo for general non-convex surfaces. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Proceedings of the IEEE/CVF European Conference on Computer Vision (ECCV)*, 2018.
- [13] Satoshi Ikehata. Does physical interpretability of observation map improve photometric stereo networks? In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 291–295, 2022.
- [14] Satoshi Ikehata. Ps-transformer: Learning sparse photometric stereo network using self-attention mechanism. In *British Machine Vision Conference (BMVC)*, 2022.
- [15] Satoshi Ikehata. Scalable, detailed and mask-free universal photometric stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13198–13207, 2023.
- [16] Yakun Ju, Junyu Dong, and Sheng Chen. Recovering surface normal and arbitrary images: A dual regression network for photometric stereo. *IEEE Transactions on Image Processing (TIP)*, pages 3676–3690, 2021.
- [17] Yakun Ju, Kin-Man Lam, Yang Chen, Lin Qi, and Junyu Dong. Pay attention to devils: A photometric stereo network for better details. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 694–700, 2020.

- [18] Yakun Ju, Boxin Shi, Muwei Jian, Lin Qi, Junyu Dong, and Kin-Man Lam. Normattention-psn: A high-frequency region enhanced photometric stereo network with normalized attention. *International Journal of Computer Vision (IJCV)*, 130:3014–3034, Dec 2022.
- [19] B. Kaya, S. Kumar, C. Oliveira, V. Ferrari, and Van G. Uncalibrated neural inverse rendering for photometric stereo of general surfaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3804–3814, 2021.
- [20] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosior, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 3744–3753, 2019.
- [21] J. Li, A. Robles-Kelly, S. You, and Y. Matsushita. Learning to minify photometric stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7560–7568, 2019.
- [22] Junxuan Li and Hongdong Li. Self-calibrating photometric stereo by neural inverse rendering. In *Proceedings of the IEEE/CVF European Conference on Computer Vision (ECCV)*, pages 166–183, 2022.
- [23] Zongrui Li, Qian Zheng, Boxin Shi, Gang Pan, and Xudong Jiang. Dani-net: Uncalibrated photometric stereo by differentiable shadow handling, anisotropic reflectance modeling, and neural inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8381–8391, 2023.
- [24] D. Lichy, S. Sengupta, and D. W. Jacobs. Fast light-weight near-field photometric stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12602–12611, jun 2022.
- [25] D. Lichy, J. Wu, S. Sengupta, and D. W. Jacobs. Shape and material capture at home. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6119–6129, 2021.
- [26] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11966–11976, 2022.
- [27] F. Logothetis, I. Budvytis, R. Mecca, and R. Cipolla. PX-net: Simple and efficient pixel-wise training of photometric stereo networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12737–12746, 2021.
- [28] Fotios Logothetis, Roberto Mecca, Ignas Budvytis, and Roberto Cipolla. A cnn based approach for the point-light photometric stereo problem. *International Journal of Computer Vision (IJCV)*, 131:101–120, Jan 2022.
- [29] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, volume 21 of *SIGGRAPH '87*, page 163–169. Association for Computing Machinery, 1987.
- [30] Roberto Mecca, Fotios Logothetis, Ignas Budvytis, and Roberto Cipolla. Lucas: A dataset for near-field point light source photometric stereo. In *British Machine Vision Conference (BMVC)*, 2021.
- [31] MyMiniFactory. <https://www.myminifactory.com/fr>.
- [32] Poly-Haven. *Poly Haven is a curated public asset library for visual effects artists and game designers, providing useful high quality 3D assets in an easily obtainable manner.*, 2023.
- [33] Jieji Ren, Feishi Wang, Jiahao Zhang, Qian Zheng, Mingjun Ren, and Boxin Shi. Diligent102: A photometric stereo benchmark dataset with controlled shape and material variation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12571–12580, 2022.

- [34] Hiroaki Santo, Masaki Samejima, Yusuke Sugano, Boxin Shi, and Yasuyuki Matsushita. Deep photometric stereo network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 501–509, 2017.
- [35] Scan the world, 2023.
- [36] B. Shi, Z. Wu, Z. Mo, D. Duan, S. Yeung, and P. Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3707–3716, 2016.
- [37] Sketchfab. <https://sketchfab.com>.
- [38] Oleg Voynov, Gleb Bobrovskikh, Pavel Karpyshev, Saveliy Galochkin, Andrei-Timotei Ardelean, Arseniy Bozhenko, Ekaterina Karmanova, Pavel Kopanov, Yaroslav Labutin-Rymsho, Ruslan Rakhimov, Aleksandr Safin, Valerii Serpiva, Alexey Artemov, Evgeny Burnaev, Dzmitry Tsetserukou, and Denis Zorin. Multi-sensor large-scale dataset for multi-view 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21392–21403, June 2023.
- [39] Feishi Wang, Jieji Ren, Heng Guo, Mingjun Ren, and Boxin Shi. Diligent-pi: Photometric stereo for planar surfaces with rich details - benchmark dataset and beyond. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12571–12580, 2023.
- [40] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)*, pages 548–558, 2021.
- [41] Xi Wang, Zhenxiong Jian, and Mingjun Ren. Non-lambertian photometric stereo network based on inverse reflectance model with collocated light. *IEEE Transactions on Image Processing (TIP)*, 29:6032–6042, 2020.
- [42] Robert J. Woodham. Photometric Method For Determining Surface Orientation From Multiple Images. *Optical Engineering*, page 513–531, 1980.
- [43] Zhuokun Yao, Kun Li, Ying Fu, Haofeng Hu, and Boxin Shi. Gps-net: Graph-based photometric stereo network. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 10306–10316. Curran Associates, Inc., 2020.
- [44] Q. Zheng, Y. Jia, B. Shi, X. Jiang, L. Duan, and A. Kot. Spline-net: Sparse photometric stereo through lighting interpolation and normal estimation networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8548–8557, 2019.