

Tailored photometric stereo: Optimization of light source positions for various materials

Christian Kapeller^{1,2}; Doris Antensteiner¹; Svorad Štolc¹

¹ AIT Austrian Institute of Technology GmbH, High Performance Vision, 1210 Vienna, Austria

² Vienna University of Technology, Institute of Visual Computing, 1040 Vienna, Austria

Abstract

Industrial machine vision applications frequently employ Photometric Stereo (PS) methods to detect fine surface defects on objects with challenging surface properties. To achieve highly precise results, acquisition setups with a vast amount of strobed illumination angles are required. The time-consuming nature of such an undertaking renders it inapt for most industrial applications. We overcome these limitations by carefully tailoring the required light setup to specific applications. Our novel approach facilitates the design of optimized acquisition setups for inline PS inspection systems. The optimal positions of light sources are derived from only a few representative material samples without the need for extensive amounts of training data. We formulate an energy function that constructs the illumination setup which generates the highest PS accuracy. The setup can be tailored for fast acquisition speed or cost efficiency. A thorough evaluation of the performance of our approach will be given on a public data set, evaluated by the mean angular error (MAE) for surface normals and root mean square (RMS) error for albedos. Our results show, that the obtained optimized PS setups can deliver a reconstruction performance close to the ground truth, while requiring only a few acquisitions.

Introduction

Photometric Stereo (PS) [10] is a reliable method to reconstruct fine surface details of objects by illuminating them from different directions. To reconstruct the surface of a perfectly Lambertian material, 3 illumination sources are needed. More complex reflections require a higher amount of strobed illumination directions, resulting in a slow acquisition process, unfit for industrial applications. This paper focuses on generating tailored designs for PS setups by defining an optimization model. We tailor optimal setups, which are apt for industrial challenges in terms of speed and reconstruction quality. Such PS setups are especially important in context of real-time inline machine vision applications, where PS methods can be used to detect fine defects on objects with various surface properties. When designing a PS system for an industrial application, the prime factors to be considered are the achieved reconstruction quality, acquisition speed and cost effectiveness. Accurate reconstruction of surface normals from objects with challenging material properties (e.g. specular, glossy, sub-surface scattered, etc.), requires acquisition setups that are carefully tailored to the specific application. Furthermore, the number of employed illumination sources in a setup influences speed and accuracy.

For PS surface reconstruction, light sources are placed in a

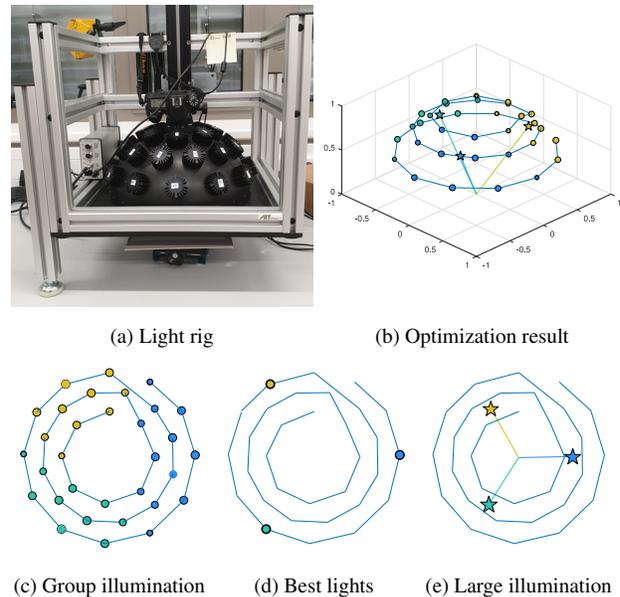


Figure 1: Illustration of a light rig with 32 illumination sources (1a) and an optimized group arrangement with 3 groups indicated by node colors (1b). We derive three kinds of setups: Joint group illumination (1c), best light illumination for existing fixtures (1d) and optimized light positions for new light rigs (1e).

rig structure around a scene and strobed separately for each image during the acquisition process. Such light rigs consist of light sources placed on top of an inspected material or a light source moved over it dynamically. Specifically, we are utilizing a light rig in a dome structure illustrated in Fig. 1a previously described in [8]. Surface reconstructions using light rigs can yield accurate results due to high sampling of available illumination directions. Since industrial setups typically require fast acquisition and processing interval, using a multitude of light sources is not viable. Under temporal limitations, grouping several light sources for each PS capture can lead to superior results. Our method provides a systematic approach for optimizing the PS light groups and positions for specific material samples.

We explore three applications of our model: (1) *Group illumination setup* (Fig. 1c): A small number of groups of jointly activated light sources allows for fast acquisitions, compared to single light illumination, while keeping a comparable reconstruction error. (2) *Best light source setup* (Fig. 1d): Cost-effective and fast setups can be created by using only the most relevant

light sources of a light rig for surface reconstruction. We show that this configuration can lead to satisfactory accuracy in terms of the normal and albedo quality while the existing rig can be utilized. (3) *Large illumination setup* (Fig. 1e): A new optimal setup with a small number of highly diffused light sources placed on the mean position of the inferred optimal illumination. This case is intended for designing new, optimized acquisition setups for industrial inspection systems.

State of the art

Finding good light source positions is a fundamental issue for Photometric Stereo. From reasoning by pure theory based on the Lambertian PS formulation Drbholav and Chantler [5] derive, that for 3 light sources, lights should be placed equally spaced, 120° apart, in a circle around the object at a slant angle of 54.74°¹. Moreover, any orthogonal light triplet is found to be optimal, while a slant angle of 54.74° is even optimal independent of the number of light sources. While these results give solid advice for general illumination placement, they fail to regard the actual surface properties of the material to inspect. Contrary, our proposed model finds optimal solutions by taking the material into account.

Empirical studies, such as [1, 2, 3], use manual selection of lights of an existing setup to find optimal light source configurations for specific materials. This method requires significant experimental effort and poses the risk of missing optimal configurations due to possibly insufficient number of available lights. Our method, on the other hand, optimizes a smooth, quadratic energy function sampled from fully sampled light rig acquisitions and can determine optimal positions in between existing light source positions.

Recently, machine learning-based models, like [7, 4, 6], have been used for determining optimal light source configurations. They are, however, usually supervised, and often require excessive training before enabling the generation of bearing solutions. The proposed method achieves optimal results for specific material samples in a computationally efficient way without the need of pre-trained models.

Method

Our algorithm determines the optimal light positions relative to given ground truth (GT) surface normals and albedo for PS under the Lambertian assumption. The GT either stems from a public data set (e.g. [11]) or is calculated from full light rig acquisitions, as described in the following section. In our approach, we determine a solution, which approximates the GT result with a limited need of strobing, allowing for industrial applications due to a high acquisition speed. Given ground truth normals and albedo, we parameterize our model with the chosen number of light groups and the optimization parameter weights λ_i with $i \in \{1, 2, 3\}$ as described in the next section. We obtain a solution that assigns membership of each light source to a certain group, as well as the amount of activation intensity within the assigned group. From the resulting light group occupation we are able to derive optimal setups for PS inspection tasks.

In the following section we describe our model formulation and the optimization process in detail.

¹Measured as angular distance from vertical orientation

Photometric Stereo Reconstruction

Consider an object captured from the same viewpoint under different illumination angles, where each acquisition is expressed as a discretized surface of size $p \times q$ pixels. Note that, we denote given data structures with a hat symbol, e.g. $\hat{\mathbf{N}}$. PS methods are well-studied and reconstruct the surface normals $\mathbf{N}_{i,j} \in \mathbb{R}^3$ and the albedo $\rho_{i,j} \in \mathbb{R}$, for all pixel locations (i, j) for $1 \leq i \leq p$ and $1 \leq j \leq q$. In case of Lambertian reflectance behaviour of a surface, normals are reconstructed under known illumination sources $\hat{\mathbf{L}} \in \mathbb{R}^{n \times 3}$. The observed n intensities are defined as $\hat{\mathbf{I}}_{i,j} \in \mathbb{R}^n$. The following tensors hold vectors in each pixel location and are denoted with bold characters:

$$\mathbf{M}_{i,j} = \rho_{i,j} \mathbf{N}_{i,j}, \quad (1)$$

$$\mathbf{M}_{i,j} = (M_{i,j,x}, M_{i,j,y}, M_{i,j,z}), \quad (2)$$

$$\mathbf{N}_{i,j} = (N_{i,j,x}, N_{i,j,y}, N_{i,j,z}), \quad (3)$$

$$\hat{\mathbf{I}}_{i,j} = (\hat{I}_{i,j,1}, \dots, \hat{I}_{i,j,n}). \quad (4)$$

Surface normals and albedo can be recovered from the observed pixel intensities using the following least squares (LS) formulation:

$$\min_{\mathbf{M}_{i,j}} \frac{1}{2} \|\hat{\mathbf{L}} \cdot \mathbf{M}_{i,j} - \hat{\mathbf{I}}_{i,j}\|^2. \quad (5)$$

As per definition, a surface normal is a unit vector, the length of the vector $\mathbf{M}_{i,j}$ is defined by the albedo value $\rho_{i,j}$:

$$\sqrt{N_{i,j,x}^2 + N_{i,j,y}^2 + N_{i,j,z}^2} = 1, \quad (6)$$

$$\rho_{i,j} = \sqrt{M_{i,j,x}^2 + M_{i,j,y}^2 + M_{i,j,z}^2}. \quad (7)$$

Using this photometric formulation, we calculate our GT surface normals and albedo using all available n illumination sources.

Light Source Position Optimization

To optimize the light source position for a specific object, we formulate a optimization model. Our ground truth tensor $\hat{\mathbf{M}}_{i,j} \in \mathbb{R}^3$ was calculated as defined in Eq. 5. Note that this tensor was variable in the previous equation but is considered given in the following equations and hence is denoted with a hat hereinafter. We have given intensities $\hat{\mathbf{I}}_k \in \mathbb{R}^{p \times q}$, where $k \in \{1, \dots, \hat{n}\}$ defines the indices of the observed light sources and a ground truth surface normal $\hat{\mathbf{N}} \in \mathbb{R}^{p \times q \times 3}$. The activation of each light is defined in the grouping matrix $G \in \mathbb{R}^{\hat{n} \times \hat{g}}$, where \hat{g} defines the number of light groups we optimize for. Our standard setup is a light rig with $\hat{n} = 32$ illuminations with a final target of $\hat{g} = 3$ optimized light groups. While our given intensities $\hat{\mathbf{I}}_k$ hold an image-sized matrix at each light source location, the grouping matrix G_k defines the membership ratio for each light source to a group and holds a vector of size \hat{g} :

$$\hat{\mathbf{I}}_k = (\hat{I}_{1,1}, \dots, \hat{I}_{p,q}), \quad (8)$$

$$G_k = (G_1, \dots, G_{\hat{g}}). \quad (9)$$

Our energy term E is optimizing for the grouping matrix, hence, finds the optimal light vectors L for the defined number of groups. It consists of three constraining energy terms and a

boundary constraint on the grouping matrix G , which is projecting it to a value between 0 and 1. The first constraining term is a *Photometric Stereo constraint*. It restricts the final solution to be close the ground truth surface orientation. Second, an *orthogonality constraint* on the grouping matrix enforces orthogonality between the groups. High orthogonality results in light groups having a low overlap to each other in terms of the occupied light sources. Third, a *sparsity constraint* enforces sparsity within each group by favoring overall weights close to 1.

$$\min_G E(G) \quad s.t. \quad 0 \leq G_k \leq 1 \quad \forall k \in \{1, \dots, \hat{g}\} \quad (10)$$

$$E(G) = E_{PS}(G) + E_{Ortho}(G) + E_{Sparse}(G) \quad (11)$$

In the following, we describe the energy constraining terms guiding the light position optimization in more detail. We optimize the energy function by iterative gradient descent steps.

Photometric Stereo constraint

The Photometric Stereo constraint ensures the final result of the albedo scaled surface orientation to be close to the ground truth tensor $\hat{\mathbf{M}}$. For this, we use Eq. 5 and activate the illumination elements according to their membership ratio in the grouping matrix G . The membership ratio for each light to the light groups is illustrated in Fig. 2b, where for each light the strength of influence to each group (shown with $\hat{g} = 3$) is indicated.

$$E_{PS}(G) = \frac{\lambda_1}{2} \left\| \left(\frac{G}{\hat{\mathbf{L}}} \cdot \hat{\mathbf{I}} \right) - \hat{\mathbf{M}} \right\|^2 \quad (12)$$

This constraint is weighted by the scalar λ_1 . The approximation of the matrix division $\frac{G}{\hat{\mathbf{L}}}$ at each iteration is a time-consuming step. Hence, we accelerate the computation by the following energy approximation.

$$E_{PS}(G) = \frac{\lambda_1}{2} \left\| G \cdot \frac{\hat{\mathbf{I}}}{\hat{\mathbf{M}}} \cdot \hat{\mathbf{M}} - G \right\|^2 \quad (13)$$

We formulate $E_{PS}(G)$ by utilizing pseudo-inverse and matrix transposition (denoted with $^+$ and T respectively) of matrices $\hat{\mathbf{M}}$ and $\hat{\mathbf{I}}$:

$$E_{PS}(G) = \frac{\lambda_1}{2} \left\| G \cdot (\hat{\mathbf{M}}^T \cdot \hat{\mathbf{I}}^+)^+ \cdot \hat{\mathbf{M}}^T \cdot \hat{\mathbf{I}}^+ - G \right\|^2. \quad (14)$$

This formulation of the Photometric Stereo constraint allows a stable and fast solution, and binds the grouping matrix to converge to a solution close to the ground truth formulation.

Orthogonality constraint

Between groups we enforce orthogonality by a scalar factor λ_2 as follows:

$$E_{Ortho}(G) = \frac{\lambda_2}{2} \left\| - (G \cdot G^T) \cdot \left(\mathbb{1} - \frac{1}{\hat{g}} \right) \right\|, \quad (15)$$

where $\mathbb{1} \in \mathbb{R}^{\hat{g} \times \hat{g}}$ symbolizes a unit matrix. This term enforces the group matrix to converge to a result where the light groups have a high orthogonality to each other. The constraint ensures different light sources to be activated in each group.

Sparsity constraint

The sparsity term constrains the size of each group.

$$E_{Sparse}(G) = \frac{\lambda_3}{2} \left\| \sum_{k=1}^{\hat{g}} \left(\left(\sum_{l=1}^{\hat{n}} G_{l,k} \right) - 1 \right) \right\|^2 \quad (16)$$

For a large λ_3 more light sources are allowed in each group, than for a small λ_3 value.

Experiments

In this section we evaluate our model qualitatively as well as quantitatively in terms of the normal and albedo error measurements. The accuracy of our optimized surface normals is measured by the mean angular error (MAE). Reconstructed albedos are evaluated as root mean squared (RMS) error:

$$EN_{MAE} = \frac{1}{pq} \sum_i^p \sum_j^q \cos^{-1} \left(\frac{\mathbf{N}_{i,j} \cdot \hat{\mathbf{N}}_{i,j}}{|\mathbf{N}_{i,j}| |\hat{\mathbf{N}}_{i,j}|} \right) \quad (17)$$

$$EA_{RMS} = \sqrt{\frac{1}{pq} \sum_i^p \sum_j^q (\rho_{i,j} - \hat{\rho}_{i,j})^2}. \quad (18)$$

We compare results for two optimized configurations, (1) group illumination and (2) best light source illumination configuration. For group illumination (1) we utilize all light sources in a grouped manner according to their activation value in G . Best light source configurations (2) use only the highest activated light source of each group. Additionally, the configuration performance is analyzed by measuring the effect of varying the amount of light groups. The run time of the optimization procedure is evaluated in terms of the required optimization iterations, as well run time measured in seconds. The first experiment is carried out on self-acquired data sets of varying materials. In the second experiment we validate the generality of our method on a publicly available data set.

Evaluation on a real-world data set

In these experiments, we assess our model on self-acquired data sets of materials with varying surface characteristics (coin, leather, fabric), which are relevant for industrial applications.

We use a light rig with 32 white LEDs, as depicted in Fig. 1a. To prevent interference effects, such as speckles, a diffuser is placed between the light sources and the object. Lights are arranged evenly spaced within each of the three rings that are located at slant angles of 54° , 38° and 25° . Objects are acquired with a Nikon D7200 DSLR and a 28 mm lens. Cropped image sizes range from 1300×1300 to 4300×3100 pixels.

We generate ground truth surface normals $\hat{\mathbf{N}}$ and albedo $\hat{\rho}$ from 32 single illuminant acquisitions $\hat{\mathbf{I}}$ as described above. We obtain light directions $\hat{\mathbf{L}}$ using specular spheres [12]

For our experiments, we experimentally determine parameter values $\lambda_1 = 0.0005$, $\lambda_2 = 0.00005$, and $\lambda_3 = 0.01$, and perform 100000 optimization iterations.

Fig. 3 shows a qualitative comparison of the optimization result of the *coin* data set optimized for $\hat{g} = 3$ light groups. The top row illustrates our albedo- and the lower row our surface normal results. The first column shows the GT albedo on top, and GT normals at the bottom. Columns 2 and 3 illustrate the albedo and surface normals for the the optimized group configuration and

their corresponding error maps. The last two columns show the results of the configuration consisting only of the strongest light source in any group. An example of obtained results for the coin data set are shown for group illumination (Fig. 1e) and best light configurations (Fig. 1d). The rig is partitioned into three approximately equally sized sections with average group light vectors spaced evenly approximately 120° apart, a result that is consistent with existing literature [9]. The same holds for best light configurations, where the light sources are located at a 54° slant angle, similarly evenly spaced in the azimuth direction.

The group illumination result for 3 groups approximates our GT with $EN_{MAE} \approx 2.3^\circ$ for normals, and $EA_{RMS} = 0.008$ for albedo. In the best light configuration, $EN_{MAE} \approx 6.9^\circ$ for normals, and $EA_{RMS} = 0.058$ for albedo. Here, group illumination improves the best lights result by 40% with respect to normal error and 12.9% for albedo error.

Next, we analyze accuracy for a varying number of groups $\hat{g} \in \{3, 4, 5\}$ on our data set. The results are summarized in Table 1. For $\hat{g} = 3$ group illumination approximates the GT with $EN_{MAE} \approx 2.9^\circ$ for normals. Increasing the number of groups to 4 and 5 improves normal accuracy to a mean of 1.19° and 1.21° respectively as shown in Fig. 4 (left). Accuracy improvement for more than 4 groups is statistically negligible. In this experiment, we find, that a setup with four light groups is a choice that allows for a low number of acquisitions while achieving good accuracy.

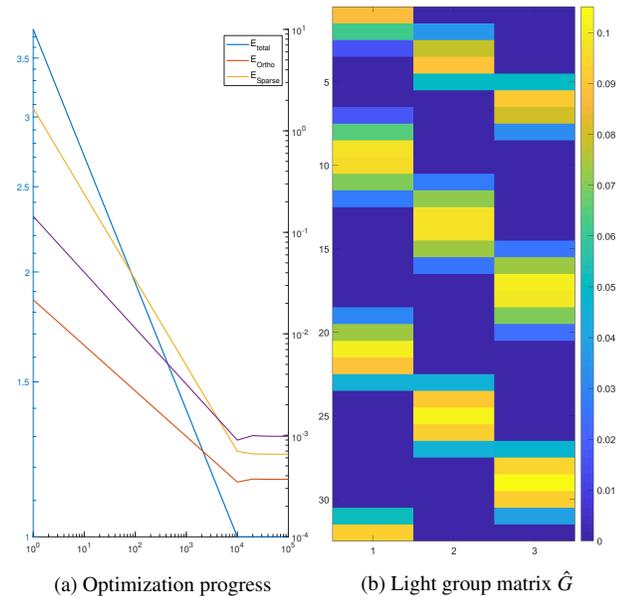


Figure 2: Illustration of (a) the optimization progress for the *coin* data set with three light groups ($\hat{g} = 3$) and (b) the resulting light source membership ratio.

Evaluation on a public data set

We evaluate our algorithm on the public Harvard data set [11]. It provides both ground truth normals and albedo of 7 real-world objects illuminated from 20 directions determined using a laser scanner. Contrary to our light rig, light directions are irregularly spaced around the object, as shown in Fig. 5 (2nd row, left). We use the same model parameterization as in the first experiment, but only compute errors in areas with valid data.

Table 1: Optimization Results for self-acquired data sets.

Data set	Normal Error (MAE)			Albedo Error (RMS)		
	3	4	5	3	4	5
	Optimized light groups					
coin	2.9008	1.5110	1.5225	0.0076	0.0060	0.0060
sponge	1.1713	0.5292	0.5322	0.0028	0.0026	0.0026
leather	3.2496	1.5457	1.5800	0.0019	0.0014	0.0014
Average	2.4405	1.1953	1.2116	0.0041	0.0033	0.0033
	Best light sources only					
coin	6.8907	5.5638	5.4300	0.0576	0.0555	0.0554
sponge	3.0645	2.3958	2.2440	0.0232	0.0230	0.0229
leather	8.0611	7.1012	7.0603	0.0162	0.0152	0.0149
Average	6.0055	5.0203	4.9119	0.0323	0.0312	0.0311

Qualitative results of this experiments are shown in Fig. 5. Despite the irregular light placement our result partitions the available light sources into approximately equal groups, with lights of highest activation located at extreme locations. The results appear similar as in our first experiment.

Our quantitative results for a varying number of groups $\hat{g} \in \{3, 4, 5\}$ are summarized in Table 2. Group configurations for $\hat{g} = 3$ approximate the GT with a normal error of $EN_{MAE} = 5.44^\circ$ for normals on average, and GT albedo with $EA_{RMS} \approx 0.12$, while the best lights configuration exhibits normal error of $EN_{MAE} \approx 8.8^\circ$ and $EA_{RMS} \approx 0.18$. Here, group illumination improves the best light configuration by 62% in terms of normal error and 66% with respect to albedo.

Increasing the number of groups \hat{g} from 3 to 4 yields and an accuracy improvement of 4% in terms of normal error. Using more than 4 groups does not result in any further accuracy gain. This fact can be explained with the data set's irregular light positions contrary to our regularly shaped light rig.

Table 2: Accuracy results for Harvard data set.

Data set	Normal Error (MAE)			Albedo Error (RMS)		
	3	4	5	3	4	5
	Optimized light groups					
cat	3.8127	3.3002	3.7513	0.1136	0.1129	0.1133
hippo	3.9599	3.1553	3.2699	0.1003	0.0994	0.0998
frog	7.8819	8.2342	8.2896	0.1281	0.1279	0.1280
lizard	4.4301	4.5306	4.2319	0.0950	0.0937	0.0923
scholar	8.8715	8.6916	8.6067	0.1691	0.1667	0.1677
turtle	3.4088	3.3556	3.4581	0.1065	0.1061	0.1054
pig	5.7151	5.3215	5.3407	0.1243	0.1234	0.1233
Average	5.4400	5.2270	5.2783	0.1196	0.1186	0.1185
	Best light sources					
cat	5.0618	4.5757	5.6830	0.1696	0.1473	0.1589
hippo	8.1835	5.2484	6.9989	0.1225	0.1358	0.1103
frog	10.5156	12.1304	10.5156	0.1961	0.2473	0.1961
lizard	7.5467	7.5467	6.9220	0.1695	0.1695	0.1537
scholar	16.9706	14.5729	13.5128	0.2390	0.2458	0.1926
turtle	5.6868	5.9360	5.7776	0.1976	0.1861	0.1747
pig	7.6755	6.0151	6.8109	0.1668	0.1707	0.1730
Average	8.8058	8.0036	8.0315	0.1801	0.1861	0.1656

Optimization runtime

Solving for G can be achieved by computing the terms $(\hat{M}^T \cdot \hat{I}^+)^+$ and $\hat{M}^T \cdot \hat{I}^+$ in Eq. 13 before starting the gradient descent. Therefore, the computational effort of subsequent optimization iterations scales with the size of G , namely $\hat{n}\hat{g}$.

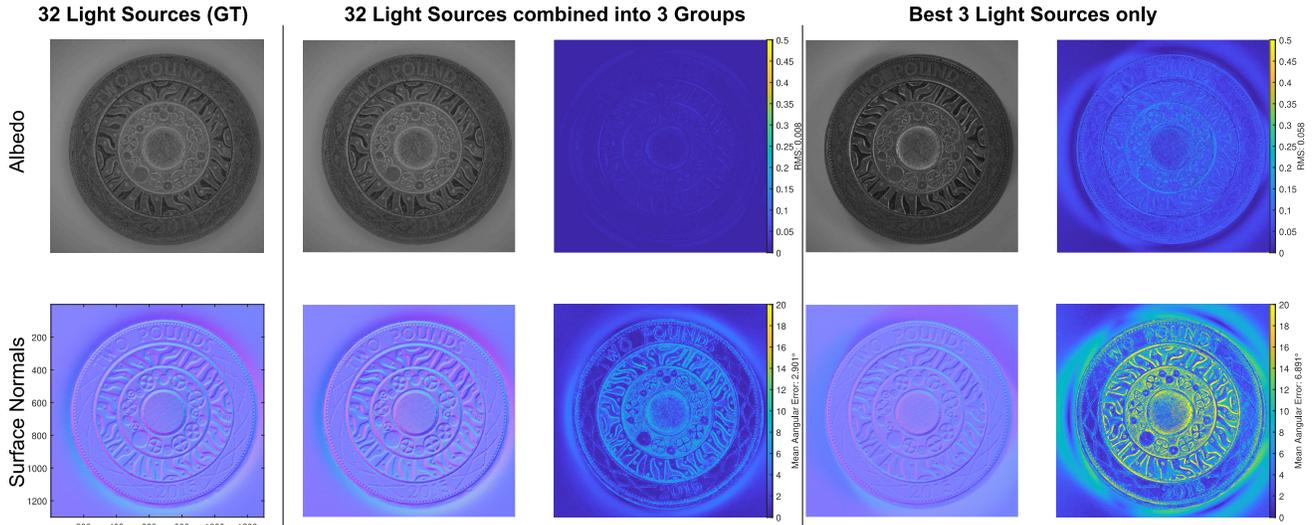


Figure 3: Illustration example of qualitative optimization results for our *coin* data set.

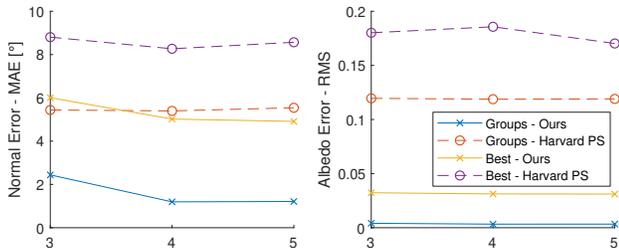


Figure 4: Illustration of the optimized accuracy with respect to the number of optimization groups \hat{g} .

Table 3: Optimization runtimes in seconds.

Number of groups	3	4	5	Size [MPix]
Own	11.4	12.8	19.8	6.620
HarvardPS	4.2	4.8	7.3	0.219

Table 3 summarizes mean runtimes² in seconds needed to obtain results of the presented experiments. Optimization for 5 groups on our own data set (32 images with 6.62 MPix size), takes approximately 20 seconds. Solving for 5 groups on the Harvard data sets (20 images with 0.219 MPix size), converges after roughly 7.3 seconds.

As shown in Fig. 2a, our model converges already at 10^4 iterations, a tenth of the total number of iterations used in our experiments.

Conclusion and Future Work

In this paper, we have presented a novel method for designing light configuration setups for industrial inspection of materials using photometric stereo. We have shown how to generate light source setups optimized for the concurring requirements of high acquisition speed, high accuracy and cost effectiveness.

In industrial setups, long light strobing sequences are often

²Using our CPU-based single threaded MATLAB implementation on a Intel Xeon W-2145 CPU @ 3.7 GHz.

not feasible due to time limitations or moving objects. We limit the amount of required strobe times and deter weak results by using an optimal light source positioning as well as light groups. Our approach determines light configurations, allowing for a reconstruction closest to ground truth. Dealing with real-world acquisitions, this enables us to optimize for specific material types and topographies. We undertake this research by formulating an energy function, which consists of three constraints. A photometric stereo term keeps the solution close to ground truth normals. An orthogonality term guides light groups to occupy different light sources in the occupancy (group) matrix. A sparsity term rewards light groups which minimize the overlap of light source activations in specific groups. Our energy is optimized using a standard gradient descent approach. In our experiments on both self-acquired and public data we have validated the adequacy of our approach to generate highly accurate, as well as fast acquisition setups. We have found that accurate and fast PS surface reconstruction can be achieved with a low number of 3 to 4 lights.

Future research will target an experimental validation on tailored optimized PS setups by placing highly diffuse, large light sources at the position of the computed mean angular vectors. Moreover, while in this paper, we described optimal light placement for PS under the Lambertian assumption, we plan to extend our model to more generalized PS formulations.

References

- [1] V. Argyriou, S. Zafeiriou, B. Villarini, and M. Petrou. A sparse representation method for determining the optimal illumination directions in photometric stereo. *Signal Processing*, 93(11):3027–3038, Nov 2013.
- [2] F. Azhar, K. Emrith, S. Pollard, M. Smith, G. Adams, and S. Simske. Testing the validity of lamberts law for micro-scale photometric stereo applied to paper substrates. In *VISAPP*, pages 246–253, 2015.
- [3] S. Brenner, S. Zambanini, and R. Sablatnig. An investigation of optimal light source setups for photometric stereo reconstruction of historical coins. In *Eurographics Workshop on Graphics and Cul-*

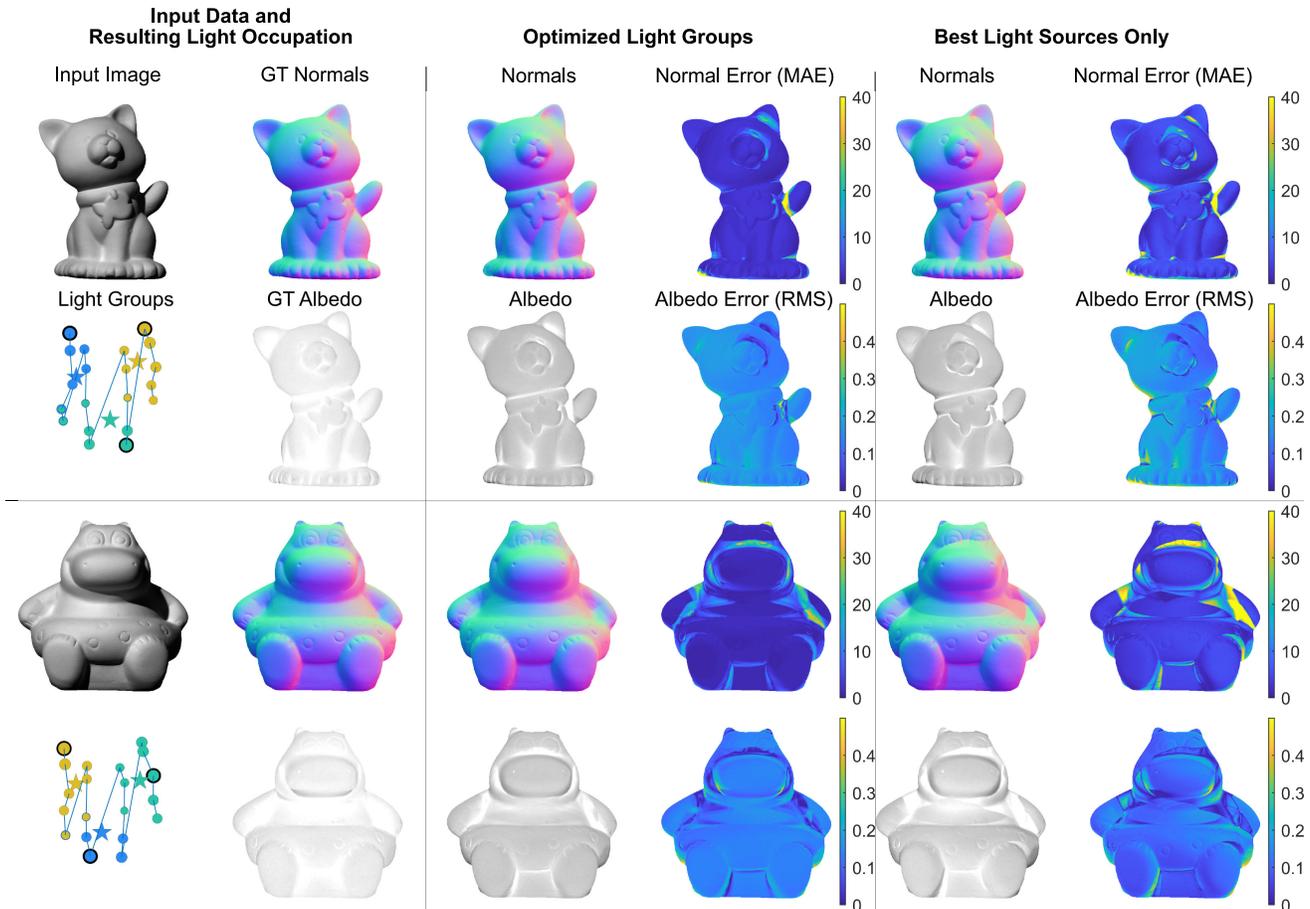


Figure 5: Illustration of qualitative optimization results for the Harvard photometric stereo data set.

tural Heritage, pages 2312–6124, 2018.

- [4] G. Chen, K. Han, and K.-Y. K. Wong. PS-FCN: A flexible learning framework for photometric stereo. In *ECCV*, pages 3–19, 2018.
- [5] O. Drbohlav and M. Chantler. On optimal light configurations in photometric stereo, 2005.
- [6] S. Ikehata. CNN-PS: CNN-based photometric stereo for general non-convex surfaces. In *ECCV*, pages 3–19, 2018.
- [7] J. Li, A. Robles-kelly, and S. You. Learning to minify photometric stereo. In *IEEE CVPR*, pages 7568–7576, 2019.
- [8] D. Soukup, S. Štolc, and R. Huber-Mörk. Analysis of optically variable devices using a photometric light-field approach. In *Proc. SPIE 9409, Media Watermarking, Security, and Forensics*, pages 250–258, 2015.
- [9] A. Spence and M. Chantler. Optimal illumination for three-image photometric stereo acquisition of texture. In *Proc. Int. Workshop on Texture Analysis and Synthesis*, pages 89–94, 2003.
- [10] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
- [11] Y. Xiong, A. Chakrabarti, R. Basri, S. J. Gortler, D. W. Jacobs, and T. Zickler. From shading to local shape. *IEEE TPAMI*, 37(1):67–79, 2015.
- [12] W. Zhou and C. Kambhampettu. Estimation of illuminant direction and intensity of multiple light sources. In *ECCV*, pages 206–220, 2002.

Author Biography

Christian Kapeller received his master’s degree in computer science from the Vienna University of Technology in 2018. Since then he has worked at the High Performance Vision Department of AIT Austrian Institute of Technology, Vienna as junior scientist since 2019, as well as the Institute of Visual Computing & Human-Centered Technology at the Vienna University of Technology. His main research areas are computational imaging and computer vision.

Doris Antensteiner received her master’s degree with distinction at the Technical University of Vienna in 2011 in the field of computer science. After that she worked at Kapsch, Vienna, at the R&D unit “Vision and Sensor” in the field of computer vision until she joined the AIT in 2015. In 2018, she earned her PhD in Computer Science from Graz University of Technology, where her main research focus was the fusion of light field and photometric stereo data. Since then she has worked as a scientist in the Center for Vision, Automation & Control of the AIT Austrian Institute of Technology GmbH, Vienna.

Svorad Štolc is a senior scientist in the Center for Vision, Automation & Control of the AIT Austrian Institute of Technology GmbH, Vienna. He graduated in Computer Science from Comenius University, Bratislava and earned his PhD degree in Bionics and Biomechanics from Technical University of Košice. He is (co)author of more than 50 peer-reviewed scientific papers and holds a number of patents in machine vision. Since 2016, he coordinates the computational imaging research activities at AIT with the focus on industrial inspection and document security.

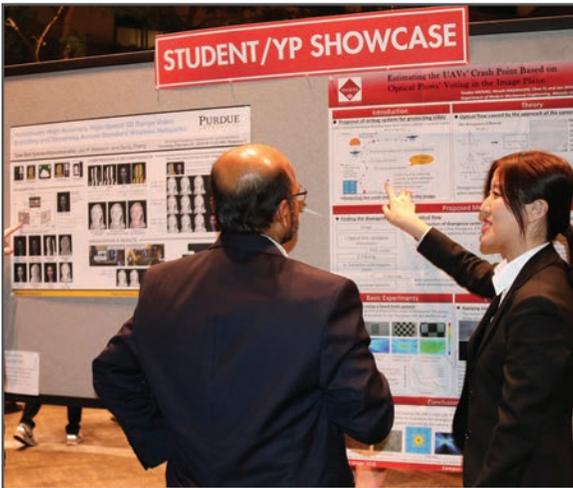
JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

www.electronicimaging.org

