



HAL
open science

Colour Dynamic Photometric Stereo for Textured Surfaces

Zsolt Janko, Amael Delaunoy, Emmanuel Prados

► **To cite this version:**

Zsolt Janko, Amael Delaunoy, Emmanuel Prados. Colour Dynamic Photometric Stereo for Textured Surfaces. ACCV 2010 - 10th Asian Conference on Computer Vision, Nov 2010, Queenstown, New Zealand. pp.55-66, 10.1007/978-3-642-19309-5_5. inria-00525869

HAL Id: inria-00525869

<https://inria.hal.science/inria-00525869v1>

Submitted on 14 Oct 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Colour Dynamic Photometric Stereo for Textured Surfaces

Zsolt Jankó^{1,2*}, Amaël Delaunoy¹, and Emmanuel Prados¹

¹ INRIA Rhône-Alpes, Montbonnot, France

² Computer and Automation Research Institute, Budapest, Hungary

Abstract. In this paper we present a novel method to apply photometric stereo on textured dynamic surfaces. We aim at exploiting the high accuracy of photometric stereo and reconstruct local surface orientation from illumination changes. The main difficulty derives from the fact that photometric stereo requires varying illumination while the object remains still, which makes it quite impractical to use for dynamic surfaces. Using coloured lights gives a clear solution to this problem; however, the system of equations is still ill-posed and it is ambiguous whether the change of an observed surface colour is due to the change of the surface gradient or of the surface reflectance.

In order to separate surface orientation from reflectance, our method tracks texture changes over time and exploits surface reflectance’s temporal constancy. This additional constraint allows us to reformulate the problem as an energy functional minimisation, solved by a standard quasi-Newton method. Our method is tested both on real and synthetic data, quantitatively evaluated and compared to a state-of-the-art method.

1 Introduction and related works

1.1 Motivation

Reconstructing dynamic surfaces from images has become a popular research topic of computer vision during the last few years. Extending the time-honoured methods of *static* surface reconstruction to the *dynamic* case is reasonable. However, *active* techniques (laser scanner, structured light projection system, *etc.*) that project time-varying textures on the surface are rather difficult to apply on deformable surfaces. *Passive* techniques, such as conventional stereo or multi-camera systems, are not biased by temporal variation, and they are more frequently applied for dynamic surfaces [1–7]. However, their accuracy is not sufficient to reconstruct small details, such as wrinkles and creases of garments.

Since the seminal work of Woodham [8], **photometric stereo** has become a powerful technique to recover local surface orientation from brightness variations generated by the shading. In contrast to multi-camera systems, photometric stereo uses image sequence acquired from the same viewpoint, under differing

* Corresponding author: janko@sztaki.hu

known directional illumination conditions, and it is capable of reconstructing small details, such as surface bumpiness, as well. However, when one wishes to apply photometric stereo on dynamic surfaces, he meets a technical difficulty: how to alter lights while the object remains still. It is clear that for intensively varying surfaces switching the lights is rather inconvenient and impractical.

The work of Smith and Smith [9] distinguishes between temporal and spectral multiplexing approaches based on how they isolate illumination conditions from one another. *Temporal multiplexing* means adaptation of static to dynamic photometric stereo, that is alternating lights. In works [10, 11] light sources are rapidly pulsed in synchronisation with the camera using a timing controller. This solution is technically difficult to attain and requires slow object motion. Additionally, for a human actor vibration of the light can be very annoying. *Spectral multiplexing* approaches can avoid light pulsation by using light frequency to separate illumination conditions. Here all the light sources are in play simultaneously. This category is also termed as **colour photometric stereo**, as coloured lights are used to recover surface shape [9, 12–14].

Our method falls within the latter category: we use three light sources with colours red, green and blue, respectively. These three colours appear almost uniformly in the visible light spectrum (400-700nm), hence they are easy to separate from one another. Furthermore, RGB is a standard colour model for sensing, representation and display of images and applied in many conventional cameras. The application of coloured lights to photometric stereo is described in more details below.

1.2 Problem formulation

A general model of image formation is

$$\begin{aligned} I(\mathbf{u}) &= \int_{\lambda=400}^{\lambda=700} \left(\int_{\Omega} (\mathbf{i}^T \mathbf{n}(\mathbf{u})) E(\mathbf{i}, \lambda) R(\mathbf{i}, \mathbf{u}, \lambda) d\mathbf{i} \right) S(\lambda) d\lambda \\ &= \int_{\Omega} (\mathbf{i}^T \mathbf{n}(\mathbf{u})) \left(\int_{\lambda=400}^{\lambda=700} E(\mathbf{i}, \lambda) R(\mathbf{i}, \mathbf{u}, \lambda) S(\lambda) d\lambda \right) d\mathbf{i}, \quad (1) \end{aligned}$$

where $I(\mathbf{u})$ is the response of a camera sensor at image point \mathbf{u} , $\mathbf{n}(\mathbf{u})$ the corresponding surface normal, and $\Omega \subset \mathbb{S}^2$ is the set of all possible directions of incoming lights \mathbf{i} (\mathbb{S}^2 being the unit sphere of \mathbb{R}^3). Denoting the wavelength by λ , $E(\mathbf{i}, \lambda)$ is the light’s spectral power distribution, $R(\mathbf{i}, \mathbf{u}, \lambda)$ is the reflectance of the surface point projected to the image point \mathbf{u} , and $S(\lambda)$ is the sensor’s spectral sensitivity.

In this work we assume that the scene is illuminated by three light sources (a Red one, $\lambda_r = 650\text{nm}$, a Green one, $\lambda_g = 510\text{nm}$, and a Blue one, $\lambda_b = 475\text{nm}$) and that it is viewed by a single camera with three sensors which mostly respond to the same wavelengths of the lights (λ_r , λ_g and λ_b for the red, green and blue captors, respectively). (See fig. 1.) Sensors’ sensitivity can then be well-approximated by Dirac delta functions as $S_k(\lambda) = s_k \delta(\lambda - \lambda_k)$, $k = r, g, b$, with suitable constants s_k .




Fig. 1. The spectral sensitivities of a typical camera. From [15].

Since the sensors are sensitive to different, well-separable wavelengths, using three light sources with three different colours (according to the above, red, green and blue, respectively) means that one image is enough to separate the three lights that are necessary for photometric stereo.

Based on these assumptions and approximations, the response of the three sensors, *i.e.* eq. (1) simplifies to

$$I_k(\mathbf{u}) = \mathbf{l}_k^T \mathbf{n}(\mathbf{u}) \alpha_k(\mathbf{u}), \quad k = r, g, b. \quad (2)$$

Here $\mathbf{l}_k = \mathbf{i}_k E(\mathbf{i}_k, \lambda_k)$ is the unit vector \mathbf{i}_k towards the light source k multiplied by its intensity³. $\alpha_k(\mathbf{u}) = R(\mathbf{i}_k, \mathbf{u}, \lambda_k) s_k$ is proportional to the *albedo*, *i.e.* the fraction of incident energy reflected by the surface. As in previous works that consider the same problem, here we assume the scene to be Lambertian. Also, we will call the triplet $\boldsymbol{\alpha} = (\alpha_r, \alpha_g, \alpha_b)$ albedo further on.

Carefully examining eq. (2) one can see that the reconstruction problem (*i.e.* recovering normals and albedos from the colours observed) is ill-posed. Beyond the image colour values $I_k(\mathbf{u})$, vectors \mathbf{l}_k are also given, assuming calibrated light sources. The rest of the unknowns are to be recovered, namely the normal \mathbf{n} (2 unknowns) and the albedo $\boldsymbol{\alpha}$ (3 unknowns). On the contrary, we have only 3 equations, for $k = r, g, b$. (Note that in conventional photometric stereo, when surface is illuminated by three light sources *separately* at three different time instants, we have three colour images instead of one, and 6 additional equations. Although the normals can then easily be calculated, altering lights is not practical for dynamic surfaces.)

In order to solve eq. (2) and to recover surface normals and albedos, one should add further constraints. Such a constraint can be that $\alpha_r = \alpha_g = \alpha_b$, *i.e.* the surface is white or gray; however, this assumption is rather restrictive. Another possibility is to capture an additional image under a broad spectrum white light and to initially measure reflectance properties. The light source should be

³ For simplicity we assume directional light sources at infinite distance.

located close to the camera, and acquisition must be performed at a separate time instant to avoid coupling with the RGB lights. The latter condition makes it impractical to apply for time-varying surfaces.

The method of Hernandez et al. [13] works on uniformly coloured surfaces, and a calibration tool is employed to estimate this colour. In contrast to them, we do not want to restrict our method to uniform albedos, either not to apply an initial estimation of reflectance properties. The work of Smith and Smith [9] meets these requirements, as they use narrow band colour photometric stereo in the infra-red region of the spectrum. In the infra-red region the reflected light is less sensitive to surface colour changes, *i.e.* different colours become metameric to one another. Due to this, surface orientation and colour can be separated. However, their system is rather complicated and requires special and expensive tools.

In this work we present a novel method that extends the approach of Hernandez et al. [13] to textured surfaces. The principal idea is to track surface motion and exploit the fact that surface colour, *i.e.* the albedo is constant through time. Based on this assumption we formulate an energy functional according to eq. (2) that penalises albedos' temporal change and reinforces spatial and temporal smoothness.

Contribution. The principal contribution of this work is a novel photometric stereo method for textured dynamic surfaces. To the best of our knowledge, our method is the first one in the category of dynamic photometric stereo which successfully works on **textured** surfaces as well, using only a single standard camera with a simple and practical light configuration. The problem is formulated by an energy functional and its optimisation. The method is able to reconstruct local surface orientations and to separate them from reflectance properties.

2 Proposed solution

In this paper, we focus on a representation of dynamic surfaces as *animated normal maps*. An animated normal map is a sequence of normal maps where temporal correspondence is given.

As discussed above, we consider a setup of one colour camera and three light sources filtered with red, green and blue filters, respectively. We denote the image sequence by I_t , the image domain by D_t and the image points by $\mathbf{u}_t \in D_t$, where $t \in T$ means a discrete time instant, $T = \{0, 1, \dots\}$. Here $\mathbf{u}_0, \mathbf{u}_1, \mathbf{u}_2, \dots$ denote the corresponding image points, that belong to the same surface part, but not the same co-ordinates. That is, $\mathbf{u}_{t+1} = \mathbf{u}_t + \mathbf{d}(\mathbf{u}_t)$, where $\mathbf{d}(\mathbf{u}_t)$ is a 2D displacement vector. The dynamic surface \mathcal{S} to be reconstructed can then be represented by the normals, the albedos and the displacement vectors:

$$\mathcal{S} = \{\mathbf{n}(\mathbf{u}_t), \boldsymbol{\alpha}(\mathbf{u}_t), \mathbf{d}(\mathbf{u}_t) : \mathbf{u}_t \in D_t, t \in T\}. \quad (3)$$

Although the normal map representation does not reflect the complete structure of the surface, it can easily be integrated to obtain a depth map and a surface

mesh. Note that recovering depth map from normals gives bad low frequency details; however, if a rough surface mesh is already given, a normal map can be used to represent fine details, small variations of the surface (*e.g.* creases on garments) as a *bump map*.

2.1 Energy Functional

The optimality of surface \mathcal{S} is quantified by an energy functional composed of two terms. The first term forces the model to suit the photometric constraint while enforcing albedo temporal constancy and unit length normals (data term); the second term guarantees spatial and temporal smoothness (regularisation term):

$$E(\mathcal{S}) = E_D(\mathcal{S}) + E_R(\mathcal{S}). \quad (4)$$

The first part of the **data term** is local, it does not consider correspondences between consecutive frames. In accordance with eq. (2) it measures the photometric error:

$$E_{D_{ph}}(\mathcal{S}) = \frac{1}{|T|} \sum_{t \in T} \frac{1}{|D_t|} \sum_{\mathbf{u}_t \in D_t} \sum_{k=r,g,b} |I_{t,k}(\mathbf{u}_t) - \alpha_k(\mathbf{u}_t) \max(\mathbf{l}_k^T \mathbf{n}(\mathbf{u}_t), 0)|^2. \quad (5)$$

To guarantee unit length of normals and albedo temporal constancy we define the following terms:

$$E_{D_n}(\mathcal{S}) = \frac{1}{|T|} \sum_{t \in T} \frac{1}{|D_t|} \sum_{\mathbf{u}_t \in D_t} \|\mathbf{n}(\mathbf{u}_t)^T \mathbf{n}(\mathbf{u}_t) - 1\|^2, \quad (6)$$

$$E_{D_a}(\mathcal{S}) = \frac{1}{|T|} \sum_{t \in T} \frac{1}{|D_t|} \sum_{\mathbf{u}_t \in D_t} \|\alpha(\mathbf{u}_{t+1}) - \alpha(\mathbf{u}_t)\|^2. \quad (7)$$

The data term is then the weighted sum of the three terms above:

$$E_D(\mathcal{S}) = \mu_{ph} E_{D_{ph}}(\mathcal{S}) + \mu_n E_{D_n}(\mathcal{S}) + \mu_a E_{D_a}(\mathcal{S}). \quad (8)$$

The **regularisation term** is destined to measure the smoothness of the model. It is composed of a weighted sum of three terms:

$$E_R(\mathcal{S}) = \mu_s E_{R_s}(\mathcal{S}) + \mu_r E_{R_r}(\mathcal{S}) + \mu_t E_{R_t}(\mathcal{S}). \quad (9)$$

The first term quantifies the *spatial smoothness*, *i.e.* the deviation of neighbouring normals:

$$E_{R_s}(\mathcal{S}) = \frac{1}{|T|} \sum_{t \in T} \frac{1}{|D_t|} \sum_{\mathbf{u}_t \in D_t} \frac{1}{|\mathcal{N}|} \sum_{\mathbf{v}_t \in \mathcal{N}(\mathbf{u}_t)} \|\mathbf{n}(\mathbf{u}_t) - \mathbf{n}(\mathbf{v}_t)\|^2, \quad (10)$$

where \mathcal{N} means the set of neighbouring image points. The second term measures the *rigidity* of the model. We require displacements of nearby points to be as similar as possible:

$$E_{R_r}(\mathcal{S}) = \frac{1}{|T|} \sum_{t \in T} \frac{1}{|D_t|} \sum_{\mathbf{u}_t \in D_t} \frac{1}{|\mathcal{N}|} \sum_{\mathbf{v}_t \in \mathcal{N}(\mathbf{u}_t)} \|\mathbf{d}(\mathbf{u}_t) - \mathbf{d}(\mathbf{v}_t)\|^2. \quad (11)$$

Third, normals and displacement vectors are reinforced to change smoothly in time (*temporal smoothness*):

$$E_{R_t}(\mathcal{S}) = \frac{1}{|T|} \sum_{t \in T} \frac{1}{|D_t|} \sum_{\mathbf{u}_t \in D_t} \|\mathbf{n}(\mathbf{u}_{t+2}) - 2\mathbf{n}(\mathbf{u}_{t+1}) + \mathbf{n}(\mathbf{u}_t)\|^2 + \|\mathbf{d}(\mathbf{u}_{t+2}) - 2\mathbf{d}(\mathbf{u}_{t+1}) + \mathbf{d}(\mathbf{u}_t)\|^2, \quad (12)$$

where $\mathbf{u}_{t+1} = \mathbf{u}_t + \mathbf{d}(\mathbf{u}_t)$ and $\mathbf{u}_{t+2} = \mathbf{u}_t + \mathbf{d}(\mathbf{u}_t) + \mathbf{d}(\mathbf{u}_t + \mathbf{d}(\mathbf{u}_t))$.

2.2 Optimisation

In optimisation, quasi-Newton methods are widely used to minimise energy functionals. We have selected one of the most popular members of this class, namely the Broyden-Fletcher-Goldfarb-Shanno (**BFGS**) method to optimise our energy functional (4). In our experiments we used the limited memory variant (**L-BFGS**) of the algorithm, included in the ALGLIB library implemented by Sergey Bochkhanov [16].

For an iterative optimisation method, such as the L-BFGS, it is essential to choose a good starting point. Our **initialisation** consists of three phases. First, motion is tracked to estimate the displacement of the points, second, a rough guess is given to the albedos based on the variation of the colour values, and finally the normals are computed using the albedos.

Due to the fact that the scene is illuminated by three coloured light sources, image colour is not sufficient to follow surface motion, as surface points change their colour significantly over time. Consider fig. 2a where the same surface particles change their colour from red to green and blue, depending on whether they look towards the red, the green or the blue light source. Although from a global aspect the colours vary, local changes of the texture remain the same, thus they are traceable. According to this, we apply a feature detector method, namely the Sobel edge detector, in order to emphasise local variations in the images. (See fig. 2b.) To track the motion and estimate displacement vectors $\mathbf{d}(\mathbf{u}_t)$ we applied the high-accuracy optical flow method of Brox et al. [17] and its variation [18], implemented by Chari [19], on the edge maps.

Note that the related work of Hernandez et al. [13] also uses optical flow to track the deformation of the surface. However, they use optical flow *after* normal and depth maps are already computed and their goal is to put them in correspondence. In contrast, we use optical flow *before* obtaining the normals, in order to constrain the ill-posed problem by assuming temporally constant albedos. Some other methods [11, 20, 21] also use optical flow to improve photometric stereo for dynamic objects, but since they do not use coloured light sources (no spectral multiplexing), we leave discussing them in details.

With the estimated displacement vectors in hand, we can give a rough approximation to the albedos. On closer inspection of eq. (2) one can see that $I_k(\mathbf{u}_t)$ reaches its maximum at a time instant t if $\mathbf{n}(\mathbf{u}_t)$ looks towards light source k , and at that time $\mathbf{l}_k^T \mathbf{n}(\mathbf{u}_t) = \|\mathbf{l}_k\|$, *i.e.* $\alpha_k(\mathbf{u}_t) = I_k(\mathbf{u}_t) / \|\mathbf{l}_k\|$. Consequently, if we




Fig. 2. a) Due to coloured illumination, deforming the surface can change its colour from red to green and blue; b) in contrast, edge detection emphasises local colour changes.

assume the surface *dynamic enough* and the sequence *long enough*, so that all normals turn towards each light sources at least once, the following estimation is a good approximation to the albedos:

$$\alpha_k^*(\mathbf{u}) = \max_{t \in T} \frac{I_k(\mathbf{u}_t)}{\|\mathbf{l}_k\|}, \quad k = r, g, b. \quad (13)$$

In the final step of the initialisation normals are computed solving a system of linear equations based on eq. (2) using the already estimated albedos.

3 Experiments and discussion

3.1 Datasets

Three real test sequences have been created in our studio using three colour light sources and one camera. The datasets contain sequences of about 150 images acquired from the same viewpoint. The image resolution is 1624×1224 .

Since our method requires calibrated light sources, we applied a simple calibration method. For this we used some additional calibrated cameras and created a short sequence of a white balance card that provides a precise uniform surface that is spectrally neutral under all lighting conditions. (See fig. 3.) Detecting the card’s corners their 3D co-ordinates and hence the fitting plane can be determined using the multiple camera setup. The light sources’ properties (direction and intensity) are computed from the plane’s normal and the sensor response observed.

In order to quantitatively assess the results, we have also created a synthetic dataset (fig. 6, first column). The dataset contain a 10 frames sequence of 800×800 images about textured waves. Ground truth normals enable to measure the reconstruction error and to compare our result to another state-of-the-art method.

3.2 Results

Fig. 4 shows the results on real datasets: the original input images, the albedo maps and the colour normal images. Our method is capable of separating surface



Fig. 3. Calibration sequence: from the corners of the white balance card the plane’s varying normal can be computed and used to calibrate light sources.

reflectance (albedo) from surface direction (normal). This is clearly demonstrated by the results: the albedo map is constant over time, the red-green-blue colour changes (visible in the input images) disappear; at the same time, the texture, for instance the stripes of the shirt are not present in the normal image. If we assumed uniform albedo, as in [13], stripes would appear in the normal image, not in the albedo map.

To visualise the quality of the results we have integrated the normal maps to obtain 3D surfaces. For this we have used the method of Kovese [22] and his implementation that is available on his web-site [23]. Fig. 5 shows examples for the surfaces obtained for each real datasets. Our supplemental material contains the complete videos of the reconstructed normals and surfaces.

We have tested our method on synthetic data as well and compared the result to the state-of-the-art method of Hernandez et al. [13]. Since their method assumes uniform albedo, texture changes incorrectly appear in the normal maps. (See figs. 6 and 7.) On the contrary, our method clearly separates albedo and normal maps.

To quantitatively assess the results, we have compared the normals obtained to the ground truth data. Considering the angle (in degrees) between the result and the ground truth normal, the mean and the confidence interval were calculated: $3.45^\circ \pm 0.0035$ with our results and $15.32^\circ \pm 0.0106$ for the method of [13]. Fig. 8 shows the histogram of the error image, which illustrates the error distribution. Note that the reconstruction’s accuracy decreases towards the edge of the image, where tracking is less precise (fig. 8, left).

4 Conclusion

In this paper we have discussed the problem of applying photometric stereo on dynamic surfaces. We have presented a novel colour photometric stereo method that successfully recovers normals for textured surfaces as well. Our method is an extension of [13], which works only for surfaces with uniform albedo. Assuming the albedos’ temporal constancy and tracking the surface’s local changes initially, we have formulated the problem as an optimisation that forces the model to fit the photometric constraint and encourage spatial and temporal smoothness. Tests on real and synthetic datasets demonstrate the efficiency of our method.




Fig. 4. Results on real datasets: input images (left), albedos (centre), normals (right).




Fig. 5. Examples for reconstructed surfaces for each real datasets.




Fig. 6. Results on a synthetic dataset: input images (left), albedos (centre), normals (right). First two rows: our method. Last two rows: assuming uniform albedo [13].

Acknowledgments. This work was supported by the French National Agency for Research (ANR) under grant ANR-06-MDCA-007 and partially by the NKTH-OTKA grant CK78409 and by the HUNOROB project (HU0045, 0045/NA/2006-2/P-9), a grant from Iceland, Liechtenstein and Norway through the EEA Financial Mechanism and the Hungarian National Development Agency.

The authors thank Visesh Chari for his optical flow code and Thomas Dupeux for his help at data acquisition.




Fig. 7. Examples for reconstructed surfaces for a synthetic dataset. Left: our method. Right: assuming uniform albedo [13].

References

1. Aganj, E., Pons, J.P., Ségonne, F., Keriven, R.: Spatio-temporal shape from silhouette using four-dimensional Delaunay meshing. In: Proc. ICCV'07. (2007)
2. Courchay, J., Pons, J.P., Monasse, P., Keriven, R.: Dense and accurate spatio-temporal multi-view stereovision. In: Proc. ACCV. (2009) 11–22
3. Furukawa, Y., Ponce, J.: Dense 3D motion capture from synchronized video streams. In: Proc. CVPR'08. (2008)
4. Pons, J.P., Boissonnat, J.D.: Delaunay deformable models: Topology-adaptive meshes based on the restricted Delaunay triangulation. In: Proc. CVPR'07. (2007)
5. Starck, J., Hilton, A.: Surface capture for performance based animation. *IEEE Computer Graphics and Applications* **27** (2007) 21–31
6. Vedula, S., Baker, S., Kanade, T.: Image-based spatio-temporal modeling and view interpolation of dynamic events. *ACM Trans. on Graphics* **24** (2005) 240–261
7. Vlastic, D., Baran, I., Matusik, W., Popović, J.: Articulated mesh animation from multi-view silhouettes. *ACM Transactions on Graphics* **27** (2008)
8. Woodham, R.J.: Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In: *Image Understanding Systems and Industrial Applications*, Proc. SPIE. Volume 155. (1978) 136–143
9. Smith, M.L., Smith, L.N.: Dynamic photometric stereo—a new technique for moving surface analysis. *Image and Vision Computing* **23** (2005) 841–852
10. Timo, P., Detlef, P., Pertti, K. (inventors): Arrangement and method for inspection of surface quality. European Patent Application, EP 1 030 173 A1 (2000)
11. Vlastic, D., Peers, P., Baran, I., Debevec, P., Popović, J., Rusinkiewicz, S., Matusik, W.: Dynamic shape capture using multi-view photometric stereo. In: Proc. SIGGRAPH Asia '09. (2009) 1–11
12. Drew, M.S.: Shape from color. Technical report, School of Computing Science, Simon Fraser University, Vancouver, B.C. (1992)

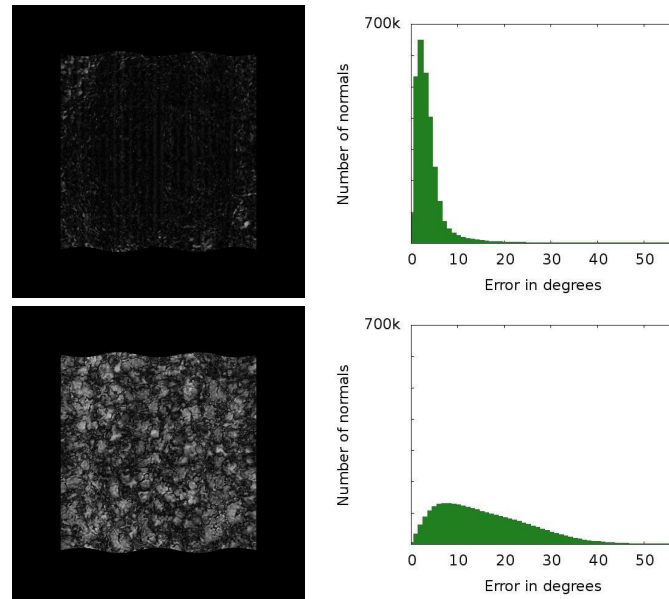


Fig. 8. Error images (left) and their histogram (right). First row: our method. Second row: assuming uniform albedo [13].

13. Hernandez, C., Vogiatzis, G., Brostow, G., Stenger, B., Cipolla, R.: Non-rigid photometric stereo with colored lights. In: Proc. ICCV'07. (2007) 1–8
14. Kontsevich, L., Petrov, A., Vergelskaya, I.: Reconstruction of shape from shading in color images. *J. Opt. Soc. Am. A* **11** (1994) 1047–1052
15. Finlayson, G.D., Hordley, S.D., Lu, C., Drew, M.S.: On the removal of shadows from images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28** (2006) 59–68
16. Bochkhanov, S.: ALGLIB. (URL: <http://www.alglib.net>)
17. Brox, T., Bruhn, A., Papenber, N., Weickert, J.: High accuracy optical flow estimation based on a theory for warping. In: ECCV. (2004) 25–36
18. Sand, P., Teller, S.: Particle video: Long-range motion estimation using point trajectories. In: Proc. CVPR'06. (2006) 2195–2202
19. Chari, V.: High accuracy optical flow. (URL: <http://www.mathworks.com/matlabcentral/fileexchange/17500-high-accuracy-optical-flow>)
20. Ma, W.C., Hawkins, T., Peers, P., Chabert, C.F., Weiss, M., Debevec, P.: Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In: Proc. EGSR'07. (2007)
21. Wenger, A., Gardner, A., Tchou, C., Unger, J., Hawkins, T., Debevec, P.: Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Trans. Graph.* **24** (2005) 756–764
22. Kovési, P.: Shapelets correlated with surface normals produce surface. In: Proc. ICCV'05. (2005) 994–1001
23. Kovési, P.: MATLAB and octave functions for computer vision and image processing. (URL: <http://www.csse.uwa.edu.au/~pk/Research/MatlabFns/>)