

IMAGE-BASED 3D SURFACE RECONSTRUCTION BY COMBINATION OF SPARSE DEPTH DATA WITH SHAPE FROM SHADING AND POLARISATION

Pablo d'Angelo and Christian Wöhler

DaimlerChrysler Group Research, Machine Perception
P. O. Box 2360, D-89013 Ulm, Germany

KEY WORDS: Industry, Metrology, Application, Polarization, Three-dimensional Reconstruction, Close Range Photogrammetry

ABSTRACT:

In this contribution we describe an image-based framework for 3D surface reconstruction by a combined analysis of reflectance, polarisation, and sparse depth data. An error functional consisting of several error terms related to the measured reflectance and polarisation properties and the depth data is minimised in order to compute a dense surface gradient field and in a subsequent step a dense 3D surface profile. The error terms related to reflectance and polarisation directly depend on the surface gradients, while the depth-related error term describes the deviation between the 3D surface profile implied by the surface gradient field and the measured depth points. Hence, we suggest an optimisation scheme that simultaneously adapts the surface gradients to the measured reflectance and polarisation data and to the surface slopes implied by depth differences between pairs of depth points. To increase the robustness of the optimisation scheme it is implemented as a multi-scale approach, thus providing a result largely independent of the provided initialisation. In our system the sparse depth data are provided by a correlation-based stereo vision algorithm, but in principle arbitrary sources of depth data are possible. We evaluate the algorithm based on synthetic ground truth data, demonstrating that the combined approach increases the accuracy of 3D surface reconstruction, compared to the result obtained by applying either of the techniques alone. Furthermore, we report 3D reconstruction results for a raw forged iron surface and compare them to ground truth depth data obtained by means of a laser focus profilometer. This evaluation yields a depth accuracy (root-mean-square deviation) of our approach of $62 \mu\text{m}$, which is of the same order of magnitude as the intrinsic roughness of the metallic surface.

1 INTRODUCTION

Three-dimensional surface reconstruction is an important topic in various application areas, such as quality inspection and reverse engineering. Many image-based reconstruction methods have been proposed, based on photometric as well as geometric principles. Well known geometric approaches include stereo and structure from motion (Faugeras, 1993), and projection of structured light (Batlle et al., 1998). In practice, even passive methods such as stereo and structure from motion often require structured illumination to artificially produce texture required for a dense reconstruction of the surface (Calow et al., 2002). Reconstruction algorithms based on photometric methods include shape from shading (SfS) and polarisation (Horn and Brooks, 1989; d'Angelo and Wöhler, 2005a; Miyazaki et al., 2003). In contrast to the geometric approaches, they can be used to reconstruct smooth, textureless surfaces without structured illumination.

A combined reconstruction based on geometric and photometric reconstruction methods is desirable, since both approaches complement each other. A number of approaches to combine stereo and shape from shading have been proposed in the literature. Cryer et al. (1995) fuse low-pass filtered stereo depth data and high-pass filtered shape from shading depth data. Samarasinghe et al. (2000) introduce a surface reconstruction algorithm that performs stereo analysis of a scene and uses a minimum description length metric to selectively apply SfS to regions with weak texture. A related approach (Fassold et al., 2004) integrates stereo depth measurements into a variational SfS algorithm and estimates surface shape, light source direction, and diffuse reflectance map.

In this paper we propose a combination of shape from photopolarimetric reflectance (SfPR) with 3D depth measurements from arbitrary sources. Our approach extends a variational SfPR framework (d'Angelo and Wöhler, 2005a) by adding an additional depth error term to the error function. A multi-scale ap-

proach is applied to reconstruct the surface gradient field. In this framework we assume known reflectance functions and light source positions.

2 SHAPE FROM PHOTOPOLARIMETRIC REFLECTANCE

In our scenario, we will assume that the surface $z(x, y)$ to be reconstructed is illuminated by a point light source and viewed by a camera, both situated at infinite distance in the directions \vec{s} and \vec{v} , respectively. The xy plane is parallel to the image plane. Parallel unpolarised incident light and an orthographic projection model are assumed. For each pixel location (u, v) of the image we intend to derive a depth value $z(u, v)$. The surface normal is given in the so-called gradient space by the vector $\vec{n} = (-p, -q, 1)^T$ with $p = \partial z / \partial x$ and $q = \partial z / \partial y$. The incidence angle θ_i is defined as the angle between surface normal \vec{n} and illumination direction \vec{s} , the emission angle θ_e as the angle between surface normal \vec{n} and viewing direction \vec{v} , and the phase angle α as the angle between illumination direction \vec{s} and viewing direction \vec{v} . A measure for the intrinsic reflectivity of the surface is given by the surface albedo $\rho(u, v)$.

In the framework of shape from photopolarimetric reflectance (SfPR), the light reflected from a surface point located at the world coordinates (x, y, z) with corresponding image coordinates (u, v) is described by the observed pixel intensity $I(u, v)$, the polarisation angle $\Phi(u, v)$ (i. e. the direction in which the light is linearly polarised), and the polarisation degree $D(u, v)$. Measurement of polarisation properties is thus limited to linear polarisation while circular or elliptic polarisation is not taken into account. It is assumed that models are available that express these photopolarimetric properties in terms of the surface orientation \vec{n} , illumination direction \vec{s} , and viewing direction \vec{v} . These models may either be physically motivated or empirical (cf. Section 2.2) and are denoted in this paper by R (intensity reflectance), R_Φ

(polarisation angle reflectance), and R_D (polarisation degree reflectance). The aim of surface reconstruction in the presented framework is to determine for each pixel (u, v) the surface gradients $p(u, v)$ and $q(u, v)$, given the illumination direction \vec{s} and the viewing direction \vec{v} , such that the modelled photopolarimetric properties of the pixel correspond to the measured values:

$$I(u, v) = R(p(u, v), q(u, v), \vec{s}, \vec{v}) \quad (1)$$

$$\Phi(u, v) = R_\Phi(p(u, v), q(u, v), \vec{s}, \vec{v}) \quad (2)$$

$$D(u, v) = R_D(p(u, v), q(u, v), \vec{s}, \vec{v}) \quad (3)$$

The reflectance functions (1)–(3) may depend on further, e. g. material-specific, parameters which possibly in turn depend on the pixel coordinates (u, v) , such as the surface albedo $\rho(u, v)$ which influences the intensity reflectance R . A local approach to obtaining the surface gradients $p(u, v)$ and $q(u, v)$ consists of solving the nonlinear system of Eqs. (1)–(3) individually for each pixel location (u, v) either exactly or in the least-mean-squares sense (d'Angelo and Wöhler, 2005b). For integration of large-scale depth information, however, a global optimisation scheme for determining the surface gradient field is more favourable, since it is not straightforward to include global depth constraints into the local approach to estimate $p(u, v)$ and $q(u, v)$.

2.1 Global optimisation scheme

In this section we describe a global approach to adapt the surface gradients $p(u, v)$ and $q(u, v)$ to the observed photopolarimetric properties $I(u, v)$, $\Phi(u, v)$, and $D(u, v)$ by solving the system of equations (1)–(3) (d'Angelo and Wöhler, 2005a). The 3D surface profile $z(u, v)$ is then obtained by integration of the surface gradient field by solving the Poisson equation $\Delta z = p_x + p_y$ (Simchony et al., 1991).

2.1.1 Determination of surface gradients and relative depth

The solving technique is based on the optimisation of a global error function (Horn, 1989; Jiang and Bunke, 1997; d'Angelo and Wöhler, 2005a). One part of this error function is the intensity error term

$$e_I = \sum_{l=1}^L \sum_{u,v} \left[I^{(l)}(u, v) - R(\rho(u, v), p(u, v), q(u, v), \vec{s}^{(l)}, \vec{v}) \right]^2. \quad (4)$$

The number of light sources and thus of acquired images is given by L . We assume orthographic projection, hence $\vec{s}^{(l)}$ and \vec{v} are constants.

As the pixel intensity information alone is not necessarily sufficient to provide an unambiguous solution for the surface gradients $p(u, v)$ and $q(u, v)$, a regularisation constraint e_s is introduced which requires smoothness of the surface, i. e. for example small absolute values of the directional derivatives of the surface gradients. We will therefore make use of the error term

$$e_s = \sum_{u,v} [p_x^2 + p_y^2 + q_x^2 + q_y^2]. \quad (5)$$

(Horn, 1989; Jiang and Bunke, 1997). In the scenarios regarded in this paper, the assumption of a smooth surface is realistic. For wrinkled surfaces, where using Eq. (5) leads to an unsatisfactory result, it can be replaced by the departure from integrability error term described in detail by Horn (1989).

To integrate polarisation angle and degree into the 3D surface reconstruction framework, we define two error terms e_Φ and

e_D which denote the deviations between the measured values and those computed using the corresponding phenomenological model:

$$e_\Phi = \sum_{l=1}^L \sum_{u,v} \left[\Phi^{(l)}(u, v) - R_\Phi(p(u, v), q(u, v), \vec{s}^{(l)}, \vec{v}) \right]^2 \quad (6)$$

$$e_D = \sum_{l=1}^L \sum_{u,v} \left[D^{(l)}(u, v) - R_D(p(u, v), q(u, v), \vec{s}^{(l)}, \vec{v}) \right]^2. \quad (7)$$

Based on the feature-specific error terms e_I , e_Φ , and e_D , a combined error term e is defined which takes into account the reflectance and polarisation properties:

$$e = e_s + \lambda e_I + \mu e_\Phi + \nu e_D. \quad (8)$$

Minimising error term (8) yields the surface gradients $p(u, v)$ and $q(u, v)$ that optimally correspond to the observed reflectance and polarisation properties, where the Lagrange parameters λ , μ , and ν denote the relative weights of the individual reflectance-specific and polarisation-specific error terms. With the discrete approximations $p_x(u, v) = [p(u+1, v) - p(u-1, v)]/2$ and $p_y(u, v) = [p(u, v+1) - p(u, v-1)]/2$ for the second derivatives of the surface and $\bar{p}(u, v)$ as the local average over the four nearest neighbours of pixel (u, v) we obtain an iterative update rule for the surface gradients by setting the derivatives of the error term e with respect to them to zero:

$$\begin{aligned} p_{n+1} = & \bar{p}_n + \lambda \sum_{l=1}^L (I - R(\bar{p}_n, \bar{q}_n)) \frac{\partial R}{\partial p} \Big|_{\bar{p}_n, \bar{q}_n} \\ & + \mu \sum_{l=1}^L (\Phi - R_\Phi(\bar{p}_n, \bar{q}_n)) \frac{\partial R_\Phi}{\partial p} \Big|_{\bar{p}_n, \bar{q}_n} \\ & + \nu \sum_{l=1}^L (D - R_D(\bar{p}_n, \bar{q}_n)) \frac{\partial R_D}{\partial p} \Big|_{\bar{p}_n, \bar{q}_n}. \end{aligned} \quad (9)$$

A corresponding expression for q is obtained in an analogous manner. This derivation is described in more detail in Jiang and Bunke (1997). The initial values $p_0(u, v)$ and $q_0(u, v)$ must be provided based on a-priori knowledge about the surface or on independently obtained depth data (cf. Section 3). The surface profile $z(u, v)$ is then derived from the resulting gradients $p(u, v)$ and $q(u, v)$ by means of numerical integration of the gradient field (Simchony et al., 1991).

The reconstruction is done in a multi-scale approach to speed up convergence and avoid getting stuck in local minima. Reconstruction of the gradient field starts at a low resolution and is repeated on the next pyramid level, using the gradients estimated at the previous level as initial gradient values.

2.2 Determination of empirical photopolarimetric models

For the purpose of determination of empirical reflectance and polarisation models for the surface material the surface normal \vec{n} of a flat sample is adjusted by means of a goniometer, while the illumination direction \vec{s} and the viewing direction \vec{v} are constant over the image. Over a wide range of surface normals \vec{n} , five images are acquired through a linear polarisation filter at orientation angles ω of 0° , 45° , 90° , 135° , and 180° . For each filter orientation ω , an average pixel intensity over an image area containing a flat part of the sample surface is computed. To the measured pixel intensities we fit a sinusoidal function of the form

$$I(\omega) = I_c + I_v \cos(\omega - \Phi) \quad (10)$$

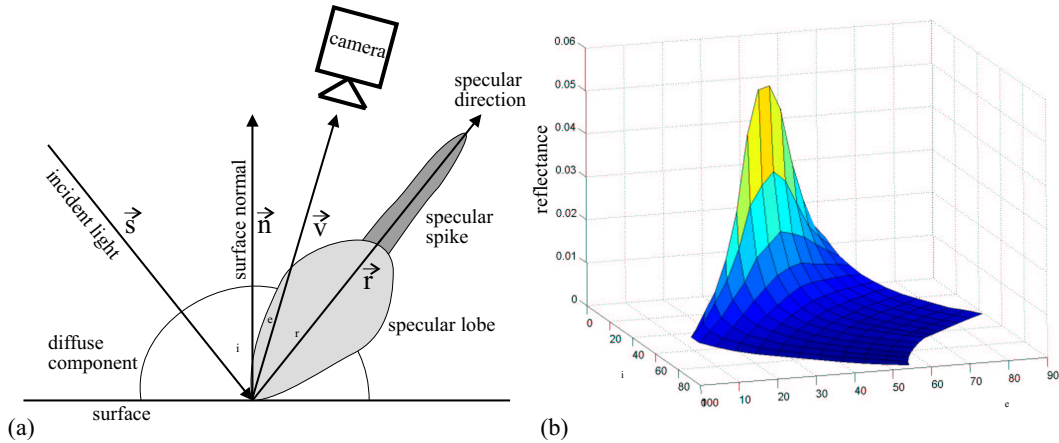


Figure 1: (a) Plot of the three reflectance components. (b) Measured reflectance of a raw forged iron surface for $\alpha = 75^\circ$.

using the linear method described by Rahmann (1999). The filter orientation Φ for which the maximum intensity $I_c + I_v$ is observed corresponds to the polarisation angle. The polarisation degree amounts to $D = I_v/I_c$. In principle, three measurements would be sufficient to determine the three parameters I_c , I_v , and Φ , but the fit becomes less noise-sensitive and thus more accurate when more measurements are used. The parameter I_c corresponds to the intensity reflectance R of the surface.

According to Nayar et al. (1991), the reflectance of a typical rough metallic surface consists of three components: a diffuse (Lambertian) component, the specular lobe, and the specular spike. We model these components by the phenomenological approach

$$R(\theta_i, \theta_e, \alpha) = \rho \left[\cos \theta_i + \sum_{n=1}^N \sigma_n \cdot (\cos \theta_r)^{m_n} \right] \quad (11)$$

with $\cos \theta_r = 2 \cos \theta_i \cos \theta_e - \cos \alpha$ describing the angle between the specular direction \vec{r} and the viewing direction \vec{v} (cf. Fig. 1a). For $\theta_r > 90^\circ$ only the diffuse component proportional to $\cos \theta_i$ is considered. The albedo ρ is assumed to be constant over the image. The shapes of the two specular components are expressed by $N = 2$ terms proportional to powers of $\cos \theta_r$, where the coefficients $\{\sigma_n\}$ denote the strength of the specular components relative to the diffuse component and the parameters $\{m_n\}$ their widths.

The polarisation angle Φ is phenomenologically modelled by an incomplete third-degree polynomial in p and q according to

$$R_\Phi(p, q) = a_\Phi p q + b_\Phi q + c_\Phi p^2 q + d_\Phi q^3. \quad (12)$$

Without loss of generality we assume illumination in the xz plane (zero y component of \vec{s}) and a view along the z axis ($\vec{v} = (0, 0, 1)^T$). Eq. (12) is antisymmetric in q , and $R_\Phi(p, q) = 0$ for $q = 0$, i. e. coplanar vectors \vec{n} , \vec{s} , and \vec{v} . These properties are required for geometrical symmetry reasons as long as an isotropic interaction between the incident light and the surface material can be assumed. The polarisation degree D is modelled by an incomplete second-degree polynomial in p and q according to

$$R_D(p, q) = a_D + b_D p + c_D p^2 + d_D q^2. \quad (13)$$

For rough metallic surfaces, R_D is maximum near the direction of specular reflection. Symmetry in q is imposed to account for isotropic light-surface interaction.

3 INTEGRATION OF DEPTH INFORMATION

Since the obtained solution of SfS and, to a lesser extent, SfPR may be ambiguous as long as single images are regarded, integrating additional information into the surface reconstruction process improves the reconstruction result. For example, a sparse set of 3D points of the object surface can be reconstructed by stereo vision, laser triangulation, or shadow analysis. Previous approaches either merge the results of stereo and SfS (Cryer et al., 1995) or embed the SfS algorithm into stereo (Samaras et al., 2000) or structure from motion algorithms (Lim et al., 2005). For the examples in this paper, a stereo algorithm was used to extract sparse depth information.

3.1 Description of the employed stereo algorithm

A block matching stereo algorithm is used in this paper. We assume that the images are rectified to standard stereo geometry with epipolar lines parallel to the horizontal image axis. The proposed approach is not restricted to this choice since any other source of relative depth information can be used instead.

For each pixel i at position (u, v) in the left image, a corresponding point is searched along the epipolar line in the right image. We use the normalized cross correlation coefficient (**normxcorr**) as similarity measure. A square region of 7 by 7 pixels of the left image (L) is correlated with regions on the corresponding epipolar line in the right image (R) for all candidate disparities d , resulting in an array of correlation coefficients $c_i(d) = \text{normxcorr}(L_{u,v}, R_{u-d,v})$. The disparity with the maximum correlation coefficient $d_i = \text{argmax}_d c_i(d)$ is determined, and a parabola $P(d) = ad^2 + bd + e$ is fitted to the local neighbourhood of the maxima. The disparity d_i is estimated at subpixel accuracy according to $d_i = -b/(2a)$. Only fits with $c_i(d_i) > 0.9$ and $a_i < -0.1$ are used. This ensures that only well localised correspondences are considered for further processing. The coordinates of a point (u_i, v_i) in the left stereo camera coordinate system are then given by $Z_i = bf/d_i$, $X_i = u_i b/d_i$, and $Y_i = v_i b/d_i$. The focal length f and base distance b between the cameras are determined by binocular camera calibration (Krüger et al., 2004).

3.2 Fusion of sparse depth information with SfPR

To incorporate the depth information into the global optimisation scheme presented in Section 2.1, we define a depth error term based on the depth difference between the sparse 3D points and

the integrated gradient field. The depth difference between two 3D points i and j is given by

$$(\Delta z)^{ij} = Z^j - Z^i. \quad (14)$$

The corresponding depth difference of the reconstructed surface gradient field is calculated by integration along a path C^{ij} between the coordinates (u^j, v^j) and (u^i, v^i) :

$$(\Delta z)_{\text{surf}}^{ij} = \int_{C^{ij}} (pdx + qdy). \quad (15)$$

In our implementation the path C^{ij} is approximated by a list of K discrete pixel positions (u_k, v_k) with $k = 1, \dots, K$. While in principle any path C^{ij} between the points i and j is possible, the shortest integration path, a straight line between i and j , is used here. Longer paths tend to produce larger depth difference errors because the gradient field is not guaranteed to be integrable.

Using these depth differences, it is possible to extend the global optimisation scheme introduced in Section 2.1 by adding an error term which minimises the squared distance between all N depth points:

$$e_z = \sum_{i=1}^N \sum_{j=i+1}^N \frac{((\Delta z)^{ij} - (\Delta z)_{\text{surf}}^{ij})^2}{\|(u_i, v_i) - (u_j, v_j)\|_2} \quad (16)$$

The iterative update rule Eq. (9) then becomes

$$p_{n+1}(u, v) = \bar{p}_n(u, v) + \lambda \frac{\partial e_I}{\partial p} + \mu \frac{\partial e_\Phi}{\partial p} + \nu \frac{\partial e_D}{\partial p} + 2\chi \sum_{i=1}^N \sum_{j=i+1}^N \left[\frac{(\Delta z)^{ij} - (\Delta z)_{\text{surf}}^{ij}}{\|(u_i, v_i) - (u_j, v_j)\|_2} \right] \frac{\partial (\Delta z)_{\text{surf}}^{ij}}{\partial p} \Big|_{u, v}. \quad (17)$$

An analogous expression is obtained for q . The derivatives of $(\Delta z)_{\text{surf}}^{ij}$ with respect to p and q may only be nonzero if the pixel (u_k, v_k) belongs to the path C^{ij} and are zero otherwise. They are computed based on the discrete gradient field. The derivative depends on the direction (d_u, d_v) of the integration path at pixel location (u_k, v_k) with $d_u = u_{k+1} - u_k$ and $d_v = v_{k+1} - v_k$:

$$\begin{aligned} \left. \frac{\partial (\Delta z)_{\text{surf}}^{ij}}{\partial p} \right|_{u_k, v_k} &= d_u p(u_k, v_k) \\ \left. \frac{\partial (\Delta z)_{\text{surf}}^{ij}}{\partial q} \right|_{u_k, v_k} &= d_v q(u_k, v_k) \end{aligned} \quad (18)$$

The update of the surface gradient at location (u, v) is then normalised with the number of paths to which the corresponding pixel belongs. Error term (16) will lead to the evaluation of $N(N-1)/2$ lines at each update step and becomes prohibitively expensive for a large number of depth measurements. Therefore only a limited number of randomly chosen lines is used during each update step.

An earlier approach by Wöhler and Hafezi (2005) fuses SfS and shadow analysis using a similar depth difference error term. It is, however, restricted to depth differences along the light source direction. In contrast to the method by Fassold et al. (2004), which directly imposes depth constraints selectively on the sparse set of surface locations with known depth, our approach establishes large-scale surface gradients by computing differences between depth points. Effectively, our method transforms sparse depth data into dense depth difference data as long as a sufficiently large number of paths C^{ij} is taken into account. The influence of the depth error term is thus extended to a large number of pixels.

4 EXPERIMENTAL EVALUATION

To examine the accuracy of 3D reconstruction using the techniques described in Section 3.2, we apply them to synthetically generated surfaces in Section 4.1. In Section 4.2 we regard real-world scenarios of 3D surface reconstruction of metallic surfaces in the domain of industrial quality inspection.

4.1 Synthetic examples

To examine the behaviour of the global optimisation scheme described in Section 3.2, we apply the developed algorithms to the synthetically generated surface shown in Fig. 2a. We assume a perpendicular view on the surface along the z axis, corresponding to $\vec{v} = (0, 0, 1)^T$. The scene is illuminated by a single light source from the positive x direction under an angle of 15° with respect to the horizontal plane. This setting results in a phase angle $\alpha = 75^\circ$. A set of 100 random points has been extracted from the ground truth and is used as depth data Z for the reconstruction.

The reflectance functions of the rough metallic surface measured according to Section 2.2 were used to render the synthetic images shown in Fig. 2c. The reconstruction was performed with synthetic noisy data, where we used Gaussian noise with a standard deviation of 0.001 for I (maximum grey value ~ 0.06), 1° for Φ and 1 pixel for the depth values. Only intensity I , polarisation angle Φ and depth Z have been used during the reconstruction. In the case of rough metallic surfaces, the polarisation degree D contains similar information as the intensity I , has a higher measurement error, and is strongly affected by small-scale variations of the surface roughness (d'Angelo and Wöhler, 2005b), and is therefore not used for reconstruction.

The weights for the error terms according to Eq. (17) were set to $\lambda = 50$ (I in arbitrary units, with a maximum of ~ 0.06), $\mu = 14$ (Φ in radian), $\nu = 0$, and $\chi = 0.5$ (z between 0 and 50 pixels). The surface gradients were initialised with zero values. The 3D reconstruction results obtained with various combinations of error terms are shown in Fig. 2d-f. The reconstruction errors are listed in Table 4.1. It is apparent that the shape from shading reconstruction fails to reconstruct the surface (Fig. 2d), while the surface shape can be reconstructed approximately using intensity and polarisation degree (Fig. 2e). The combined approach (Fig. 2f) shows the smallest error. Table 4.1 also indicates that using intensity, polarisation, and depth leads to better results than either feature alone.

4.2 Real-world example: raw forged iron surface

We have applied our surface reconstruction algorithm to a raw forged iron surface. For the stereo reconstruction of the surface (cf. Section 3.1), we used a vergent stereo setup of two cameras (1032×776 pixels image size, 10° horizontal field of view, 320 mm base distance, average object distance 450 mm). Stereo calibration and image rectification to standard epipolar geometry were performed using the method by Krüger et al. (2004). The disparity values at object distance thus amount to approximately 4000 pixels. Experiments with synthetic data have shown that the standard deviation of the disparity is 0.3 pixels, resulting in an estimated standard error of $30 \mu\text{m}$ of the determined depth values. One of the stereo cameras is equipped with a rotating linear polarisation filter and is used to acquire the images required for SfPR (cf. Section 2.2). Fig. 3a shows the intensity and polarisation angle image, and Fig. 3b shows the triangulated stereo reconstruction result. The stereo reconstruction is very sparse due to the highly non-Lambertian metallic surface and does not extend across the complete surface to be reconstructed.

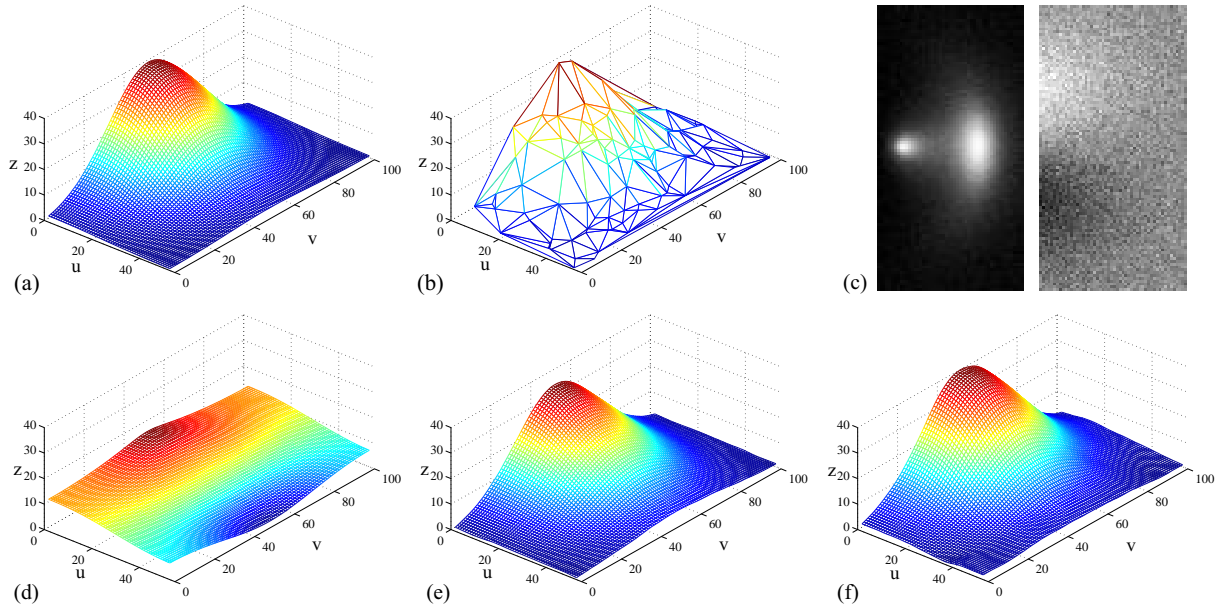


Figure 2: 3D reconstruction of a synthetically generated surface based on a photopolarimetric image and sparse depth values. (a) Ground truth. (b) Noisy 3D data. (c) From the left: noisy intensity and polarisation angle images, based on measured reflectance functions of a raw forged iron surface. The reconstruction result for noisy images of a surface with uniform albedo is shown in (d) using intensity only and in (e) using intensity and polarisation angle. (f) Reconstruction result obtained using the combined SfPR and depth approach.

Table 1: Results on the synthetic ground truth example shown in Fig. 2.

Method	RMS error (without noise)			RMS error (with noise)		
	z [pixels]	p	q	z [pixels]	p	q
I	8.19	0.267	0.508	8.19	0.267	0.508
I, Φ	2.07	0.186	0.039	2.12	0.189	0.058
Z	1.20	0.137	0.102	1.16	0.135	0.136
I, Z	0.80	0.070	0.076	0.79	0.083	0.115
I, Φ, Z	0.46	0.050	0.026	0.50	0.075	0.063

The unknown albedo ρ was computed based on the specular reflections, which appear as regions of maximum intensity I_{spec} and for which we have $\theta_r = 0^\circ$ and $\theta_i = \alpha/2$. Eq. (11) then directly yields the albedo ρ . The reconstructed surface shown in Fig. 3c was computed using $\lambda = 50$, $\mu = 8$, $\nu = 0$, and $\chi = 1$ as error term weights, cf. Eq. (17). A cross-section of the surface was measured with a scanning laser focus profilometer and compared to the corresponding cross-section extracted from the reconstructed 3D profile (Fig. 3d). Although the triangulated depth data RMSE of $80 \mu\text{m}$ along the inspected profile of 14 mm length is already quite low, no small-scale detail of the surface is revealed. When all available photopolarimetric and depth information is used, the RMSE amounts to $62 \mu\text{m}$. Without depth information the SfPR method yields a RMSE of $65 \mu\text{m}$, while intensity information alone results in a much higher RMSE of $300 \mu\text{m}$. If no polarimetric information is available (e.g. when incident light is not polarised by reflection at the surface), the combination between intensity and sparse depth data yields a RMSE of $70 \mu\text{m}$.

5 SUMMARY AND CONCLUSION

In this paper we have presented an image-based method for 3D surface reconstruction relying on the simultaneous evaluation of reflectance, polarisation, and sparse depth data. The reflectance and polarisation properties of the surface material have been ob-

tained by means of a series of images acquired through a linear polarisation filter under different orientations. SfPR and depth difference error terms are minimised using a variational approach, resulting in a surface gradient field. A dense depth map is obtained by numerical integration of the gradient field. A multi-scale approach has been used to improve the convergence behaviour. The proposed method transforms sparse depth data into dense depth difference data. In contrast to previous methods, the influence of the corresponding error term does not remain restricted to a small number of pixels. The presented method has been evaluated based on a synthetically generated surface, and a high accuracy of surface reconstruction has been demonstrated. Furthermore, we have successfully applied our method to the difficult real-world scenario of 3D reconstruction of a surface section of a raw forged iron part, yielding a very reasonable accuracy of $62 \mu\text{m}$ along the inspected profile of 14 mm length. A somewhat lower accuracy of $70 \mu\text{m}$ is obtained when polarisation information is neglected. These measurement errors are of the same order of magnitude as the intrinsic roughness of the metallic surface. We conclude that the suggested approach is a favourable technique for industrial surface inspection systems.

REFERENCES

Battle, J., Mouaddib, E., Salvi, J., 1998. Recent progress in coded structured light as a technique to solve the correspondence problem: a survey. *Pattern Recognition*, 31(7), pp. 963-982.

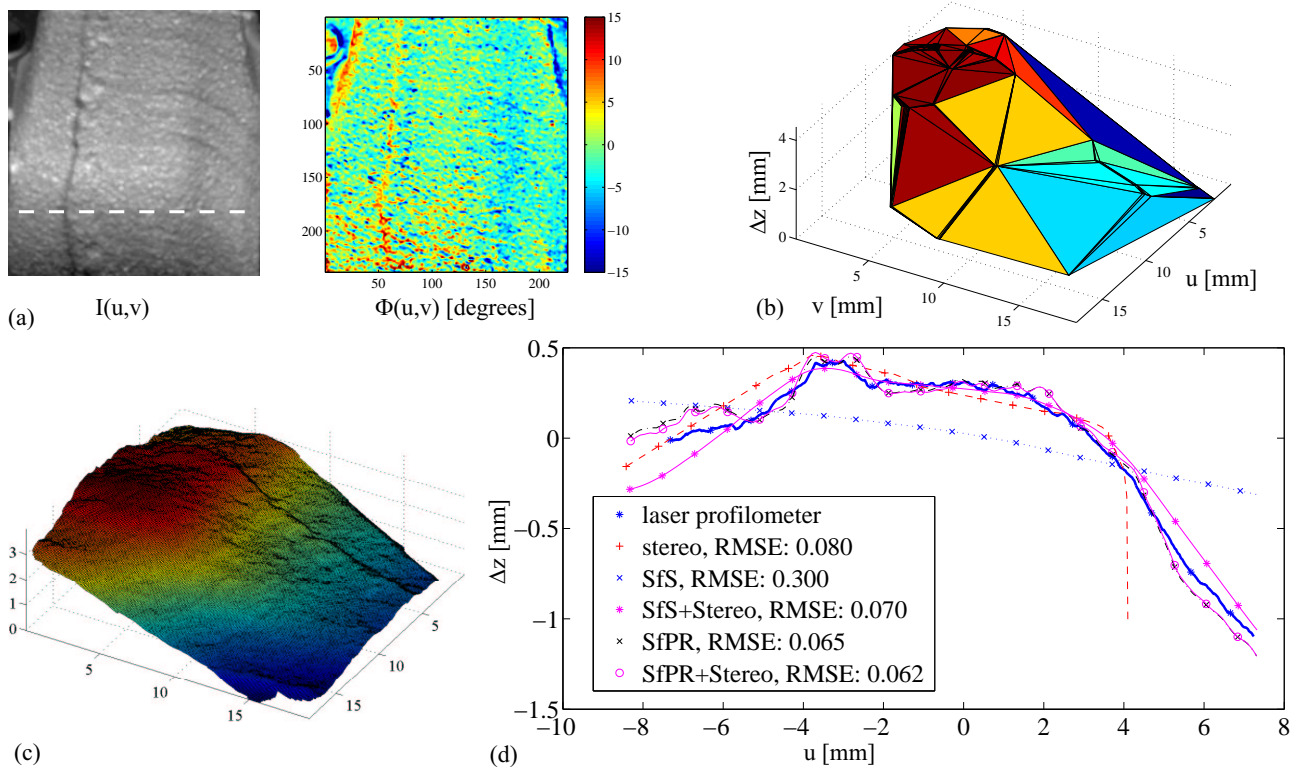


Figure 3: 3D reconstruction of a raw forged iron surface. (a) Reflectance and polarisation angle images. The size of the region of interest is 240×240 pixels. (b) Triangulated stereo reconstruction result. (c) Reconstruction by combined stereo and SfPR. (d) Comparison of the cross-section indicated by the dashed line in (a) to the ground truth measurement obtained with a laser focus profilometer.

Calow R., Gademann G., Krell G., Mecke R., Michaelis B., Riefenstahl N., Walke M., 2002. Photogrammetric measurement of patients in radiotherapy. *ISPRS Journal of Photogrammetry and Remote Sensing*, 56(5-6), pp. 347-359.

Cryer, J.E., Tsai, P.-S., Shah, M., 1995. Integration of shape from shading and stereo. *Pattern Recognition*, 28(7), pp. 1033-1043.

D'Angelo, P., Wöhler, C., 2005a. 3D Reconstruction of Metallic Surfaces by Photopolarimetric Analysis. In: H. Kalviainen et al. (Eds.), *Proc. 14th Scand. Conf. on Image Analysis*, LNCS 3540, Springer-Verlag Berlin Heidelberg, pp. 689-698.

D'Angelo, P., Wöhler, C., 2005b. 3D surface reconstruction based on combined analysis of reflectance and polarisation properties: a local approach. *ISPRS Workshop Towards Benchmarking Automated Calibration, Orientation and Surface Reconstruction from Images*, Beijing, China.

Fassold, H., Danzl, R., Schindler, K., Bischof, H. 2004. Reconstruction of Archaeological Finds using Shape from Stereo and Shape from Shading. *9th Computer Vision Winter Workshop*, Piran, Slovenia, pp. 21-30.

Faugeras, O., 1993. *Three-Dimensional Computer Vision*. MIT Press, Cambridge, Massachusetts.

Horn, B. K. P., Brooks, M. J., 1989. *Shape from Shading*. MIT Press, Cambridge, Massachusetts.

Horn, B. K. P., 1989. Height and Gradient from Shading. MIT technical report 1105A. <http://people.csail.mit.edu/people/bkph/AIM/AIM-1105A-TEX.pdf>

Jiang, X., Bunke, H., 1997. *Dreidimensionales Computervision*. Springer-Verlag, Berlin.

Krüger, L., Wöhler, C., Würz-Wessel, A., Stein, F., 2004. In-factory calibration of multicocular camera systems. *SPIE Photonics Europe (Optical Metrology in Production Engineering)*, Strasbourg, pp. 126-137.

Lim, J., Jeffrey, H., Yang, M., Kriegman, D., 2005. Passive Photometric Stereo from Motion. *IEEE Int. Conf. on Computer Vision*, Beijing, China, vol. II, pp. 1635-1642.

Miyazaki, D., Tan, R. T., Hara, K., Ikeuchi, K., 2003. Polarization-based Inverse Rendering from a Single View. *IEEE Int. Conf. on Computer Vision*, Nice, France, vol. II, pp. 982-987.

Nayar, S. K., Ikeuchi, K., Kanade, T., 1991. Surface Reflection: Physical and Geometrical Perspectives. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(7), pp. 611-634.

Rahmann, S., 1999. Inferring 3D scene structure from a single polarization image. *Conf. on Polarization and Color Techniques in Industrial Inspection*, SPIE Vol. 3826, Munich, Germany, pp. 22-33.

Samaras, D., Metaxas, D., Fua, P., Leclerc, Y.G. Variable Albedo Surface Reconstruction from Stereo and Shape from Shading. *Proc. CVPR 2000*, vol I, pp. 480-487.

Simchony, T., Chellappa, R., Shao, M., 1991. Direct Analytic Methods for Solving Poisson Equations in Computer Vision Problems. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(5), pp. 435-556.

Wöhler, C., Hafezi, K., 2005. A general framework for three-dimensional surface reconstruction by self-consistent fusion of shading and shadow features. *Pattern Recognition*, 38(7), pp. 965-983.