Better Shading for Better Shape Recovery

Moumen T. El-Melegy
Electrical Engineering Department
Assiut University, Assiut 71516, Egypt
moumen@aun.edu.eg

Aly S. Abdelrahim and Aly A. Farag
CVIP Lab, University of Louisville
Louisville, KY 40292, USA
aly.abdelrahim,aly.farag@louisville.edu

Abstract

The basic idea of shape from shading is to infer the shape of a surface from its shading information in a single image. Since this problem is ill-posed, a number of simplifying assumptions have been often used. However they rarely hold in practice. This paper presents a simple shading-correction algorithm that transforms the image to a new image that better satisfies the assumptions typically needed by existing algorithms, thus improving the accuracy of shape recovery. The algorithm takes advantage of some local shading measures that have been driven under these assumptions. The method is successfully evaluated on real data of human teeth with ground-truth 3D shapes.

1. Introduction

Shape from shading (SFS) is a problem that has been studied for about four decades in the vision literature. It aims to recover surface orientation from local variations in measured brightness. To make it solvable, classical SFS methods often assumed an orthographic camera model, distant single point light source, and a Lambertian surface [6, 17, 4]. Unfortunately, those assumptions are not always valid in reality. As such, the reconstruction results of these classical SFS approaches lack accuracy. That is why there have been more recent methods trying to relax some of these assumptions. For example, Prados et al. [9], Tankus et al. [13], and Yuen et al. [16] replaced the orthographic camera model by a pinhole camera model performing a perspective projection, and they assumed that the light source is located at the optical centre. Moreover, a light attenuation term is considered in [9]. These ideas have been further extended by Ahmed et al. [2], Vogel et al. [14], and Abdelrahim et el [1], where the Lambertian reflectance model is replaced by the more realistic model of Oren and Nayar [8] or that of Wolff [10]. On the other hand, various optimization techniques have been employed to solve the SFS problem, see e.g. [4] for a recent survey.

The observation underpinning this paper is that although considerable effort has gone into the development of improved shape-from-shading methods, there is room for further development. Here we propose a novel strategy to improve the surface recovery results of SFS. Instead of trying to develop more sophisticated SFS algorithms that can take into account the more realistic assumptions of the SFS problem, we investigate the possibility of correcting the shading of the given intensity image such that its deviation from the typical SFS assumptions (such as orthographic projection, and Lambertian reflectance model) is reduced. This is quite important as several of the most successful existing methods for SFS are restricted in applicability to Lambertian reflectance and orthographic projection, and cannot be adapted easily to work with other assumptions. Hence, our correction process can be used as a simple tool to preprocess the input image data for such methods, leading to more accurate shape recovery results. To the best of our knowledge, this has not been addressed before in the literature. The only exception [11] that we are aware of presented simplified formulas to perform radiance correction to the Lambertian case. The method is only applicable to the retroreflection case when the object is illuminated in the viewing direction, and requires the specific source reflectance model as well as its parameters be known in advance, which is not always possible in reality. Additionally, the conversion formulas in some cases are not tractable, and a lookup table is used instead.

The shading correction strategy proposed here alleviates many of those mentioned drawbacks. It works blindly on the input image without the knowledge of the original reflectance model, or the illumination direction relative to the viewing direction. It is based on a number of local shading measures derived from the mathematical model of the image shading formation under a typical set of SFS assumptions. These measures deviate from their theoretical average values when an image does not satisfy those put forth assumptions. These derived measures are used to develop a simple algorithm to map an input image to a new, corrected image that better satisfies the assumptions typically needed by the SFS algorithms, thus leading to more accurate shape

recovery.

In a recent critique [5] of the state-of-the-art of SFS, it is stated that the majority of SFS methods have been demonstrated on synthetic data, on which results are not generally predictive of results on real data. In addition, only a few SFS methods have compared a shading based reconstruction of an image of a real object with a depth estimate. Nevertheless, a distinct characteristic of this paper is that the proposed SFS strategy is evaluated on a large set of real images of 3D shapes with known ground-truth depth estimates. This dataset consists of real human teeth images of various tooth types (maxillary deciduous, maxillary molars, mandibular molars, or mandibular third molars), which are compiled for a research project on 3D tooth reconstruction from intra-oral images. The results of the new SFS strategy applied to these images are assessed against ground-truth shapes obtained from CT scans of the real human teeth, and compared with several existing SFS techniques in terms of accuracy.

The rest of this paper is organized as follows. Section 2 develops our new strategy to improve SFS results based on correcting input images. Experimental results are presented in Section 3, and the paper is concluded in Section 4.

2. Shading Correction for Better SFS

The SFS problem requires the recovery of shape of the visible surface from a single intensity image I. Our strategy here to solve this problem poses this question: Is there a mapping $\mathcal{F} : \mathbb{R} \rightarrow \mathbb{R}$ that when applied pixel-wise to the image I, the resulting image $\mathcal{F}(I)$ provides better shape recovery by classical SFS algorithms? An answer to this question calls for addressing two issues. The first one is related to the explicit functional form of this unknown transformation $\mathcal{F}$, and how it can be parameterized. One can address this precisely if the cause(s) of the deviation of the input image from the main assumptions typically put forth in the SFS algorithms is(are) known a priori. For example, if this deviation is attributed to an intensity non-linearity, such as γ-correction, the transformation can be explicitly defined as [7]

$$\mathcal{F}(I) = \varrho \, I^{1/\gamma}, \tag{1}$$

which inverts the process of γ-correction. The constant ϱ is a normalization factor which is determined by the value of γ. Interestingly, it was shown [7] that the same simple parametric form of γ-correction in (1) can also account for a rather complicated camera response function disturbing the common SFS assumptions, regardless of the camera's temporal gain and spatial non-uniformities. Similarly, an explicit form of $\mathcal{F}$ can be devised if this deviation is due to vignetting [7]. However it is more often that the cause of this deviation is unknown, or no explicit form can be envisioned for it (e.g., non-Lambertian surfaces). Therefore, to get around this, we propose to assume a more generic form based on finite-order Taylor-polynomial

$$\mathcal{F}(I(.)) = \sum_{k=0}^{p} c_k \, I^k(.), \tag{2}$$

where $\{c_k \in \mathbb{R}\}$ is a sequence of p-coefficients need to be determined, which brings about the second issue.

The second issue is related to how to find the parameters ($\{c_k\}$) of this explicit form in (2)). A proper objective criterion should be formulated to measure how much an intensity image satisfies the typical SFS assumptions. The smaller this criterion, the better the image meets these assumptions, and accordingly, the better the recovery results of the classical SFS algorithms. Deriving such an objective criterion is rather difficult in general. However, in the next subsection, we formulate a proper objective criterion based a set of local shading measures derived under typical SFS assumptions.

2.1. Local Shading Measures

In the classical SFS problem [6, 17], under a typical set of assumptions, the goal is to find a 3D Lambertian surface $Z(X, Y)$ satisfying the image irradiance equation

$$I(x, y) = \rho \, \hat{\boldsymbol{s}} \cdot \hat{\boldsymbol{n}} + \mu, \tag{3}$$

where $I(x, y)$ is the intensity at the image position (x, y), with ρ being the surface albedo, $\hat{\boldsymbol{n}}$ the surface normal at the 3D point $(X, Y, Z(X, Y))$, $\hat{\boldsymbol{s}}$ the unit vector for the illuminant direction, and μ an ambient lighting term. We assume no presence of interreflections among surface patches. Under the typical assumption of orthographic projection, $x = X$, and $y = Y$. So the surface could be expressed as $Z(x, y)$. It is often convenient to work in spherical coordinates, let's parameterize $\hat{\boldsymbol{s}} = (\cos \tau \sin \gamma, \sin \tau \sin \gamma, \cos \gamma)$, where τ is the tilt angle and γ is the slant angle. Similarly, let's express the surface normal as $\hat{\boldsymbol{n}} = (\cos \varphi \sin \psi, \sin \varphi \sin \psi, \cos \psi)$, where $\varphi = \varphi(X, Y, Z)$ and $\psi = \psi(X, Y, Z)$ are the tilt and slant angles, respectively, of the surface normal. With these representations, (3) becomes

$$I(x, y) = \rho(\cos \tau \sin \gamma \cos \varphi \sin \psi +$$
$$\sin \tau \sin \gamma \sin \varphi \sin \psi + \cos \gamma \cos \psi) + \mu. \tag{4}$$

Following some earlier works (e.g., [17]), assume that at any point (X, Y, Z) the surface could be approximated by a spherical patch:

$$
\begin{aligned}
X &= X_0 + R_0 \sin \psi \cos \varphi, \\
Y &= Y_0 + R_0 \sin \psi \sin \varphi, \\
Z &= Z_0 + R_0 \cos \psi, \tag{5}
\end{aligned}
$$

where the sphere center (X_0, Y_0, Z_0) and radius R_0 depend on the local surface shape, and both may change from one point neighborhood to another. After some computation using (4) and (5), the image derivative I_x at a point (x, y) could be found:

$$
\begin{aligned}
I_x &= \frac{\partial I}{\partial x} = \frac{\partial I}{\partial \varphi} \frac{\partial \varphi}{\partial x} + \frac{\partial I}{\partial \psi} \frac{\partial \psi}{\partial x} \\
&= \frac{\rho}{R_0}(\cos \tau \sin \gamma - \cos \gamma \tan \psi \cos \varphi).
\end{aligned} \quad (6)
$$

Similarly,

$$
\begin{aligned}
I_y &= \frac{\partial I}{\partial y} = \frac{\rho}{R_0}(\sin \tau \sin \gamma - \cos \gamma \tan \psi \sin \varphi), \\
I_{xx} &= \frac{\partial I_x}{\partial x} = -\frac{\rho \cos \gamma}{R_0^2 \cos^3 \psi}(\sin^2 \varphi \cos^2 \psi + \cos^2 \varphi), \\
I_{yy} &= \frac{\partial I_y}{\partial y} = -\frac{\rho \cos \gamma}{R_0^2 \cos^3 \psi}(\cos^2 \varphi \cos^2 \psi + \sin^2 \varphi), \\
I_{xy} &= \frac{\partial I_x}{\partial y} = \frac{\rho \cos \gamma \sin 2\varphi}{2 R_0^2 \cos^3 \psi}(\cos^2 \psi - 1), \\
\nabla^2 I &= I_{xx} + I_{yy} = -\frac{\rho \cos \gamma}{R_0^2 \cos^3 \psi}(\cos^2 \psi + 1).
\end{aligned} \quad (7)
$$

Then these derived quantities are normalized by the image Laplacian $\nabla^2 I$ leading to the following ratios at any image point (x, y):

$$
\begin{aligned}
\mathcal{I}_{xx} &= \frac{I_{xx}}{\nabla^2 I} = \frac{\sin^2 \varphi \cos^2 \psi + \cos^2 \varphi}{\cos^2 \psi + 1}, \\
\mathcal{I}_{yy} &= \frac{I_{yy}}{\nabla^2 I} = \frac{\cos^2 \varphi \cos^2 \psi + \sin^2 \varphi}{\cos^2 \psi + 1}, \\
\mathcal{I}_{xy} &= \frac{I_{xy}}{\nabla^2 I} = \frac{\sin 2\varphi(1 - \cos^2 \psi)}{2(\cos^2 \psi + 1)}.
\end{aligned} \quad (8)
$$

These local measures (8) depend only on the tilt and slant angles of the surface normal at the point; They are not affected by the illuminant direction, ambient lighting, the radius or the center of the spherical approximation in the neighborhood of the point (this minimizes any error in this local spherical approximation assumption). Note also that $\mathcal{I}_{xx} + \mathcal{I}_{yy} = 1$ for any point.

If the statistical distributions of φ and ψ are known, one could obtain the expected estimates of these local measures throughout the whole surface. Such distributions were derived in [17] to solve for the illuminant direction. Here we assume the same distributions in order to estimate the expected values of the ratios in (8). The assumed-*independent* distributions f_φ and f_ψ of φ and ψ, respectively, are [17]

$$
f_\varphi = \frac{1}{2\pi}, 0 \leq \varphi < 2\pi; f_\psi = \cos \psi, 0 \leq \psi \leq \pi/2. \quad (9)
$$

Using (9), one can evaluate the expected value of $\mathcal{I}_{xx}$:

$$
\overline{\mathcal{I}_{xx}} = \mathbb{E}\{\mathcal{I}_{xx}\} = \frac{1}{2\pi} \int_{\varphi=0}^{2\pi} \int_{\psi=0}^{\pi/2} \mathcal{I}_{xx} \cos \psi \, d\varphi \, d\psi = 0.5. \quad (10)
$$

Similarly, $\overline{\mathcal{I}_{yy}} = \mathbb{E}\{\mathcal{I}_{yy}\} = 0.5$, and $\overline{\mathcal{I}_{xy}} = \mathbb{E}\{\mathcal{I}_{xy}\} = 0$. In simple words, the theoretical average values of the local shading measures $\mathcal{I}_{xx}$ and $\mathcal{I}_{yy}$ are 0.5, while that of the local measure $\mathcal{I}_{xy}$ is 0.

The expected values of these local measures can be estimated from the images using the arithmetic means. In practice, we employ the derivatives of a Gaussian function as filters to compute the derivatives of a sampled image function via convolution. For example, $\overline{\mathcal{I}_{xx}}$ is implemented as

$$
\overline{\mathcal{I}_{xx}} = \frac{1}{size(I)} \sum_{x \in I} \sum_{y \in I} \frac{I \circledast G_{xx}}{I \circledast G_{xx} + I \circledast G_{yy}}, \quad (11)
$$

where $\circledast$ denotes convolution, and G_{xx} is the second order partial derivative of a 2D zero mean Gaussian kernel. Image points with zero Laplacian (zero denominator in (11)) are discarded from the summations. One can show, theoretically and experimentally, that these means as computed in (11) have rather small noise sensitivity. Note also that a more robust estimate of the means of these local measures can be obtained by replacing the arithmetic mean in (11)) by the median function.

It is important to stress here that the theoretical averages of the derived local shading measures have been verified for the set of the 3D teeth surfaces obtained from the CT scans of real human teeth that are used in our experimental results in Section 3. Fig. 1 shows the obtained distributions f_φ and f_ψ of the angles φ and ψ of the tooth surface's normals. From those two distributions, the expected values of the shading measures are found as $\overline{\mathcal{I}_{xx}} = 0.507$, $\overline{\mathcal{I}_{yy}} = 0.493$ and $\overline{\mathcal{I}_{xy}} = 0.001$, all of which are quite close to the theoretical values derived in (10).

2.2. Shading Correction Algorithm

The local measures and their theoretical means are derived under typical SFS assumptions. They will start to deviate from these theoretical means when one of more assumptions are unsatisfied. Moreover, it can be shown that common imaging non-linearities, such as intensity and geometric non-linearities (e.g., γ-correction and lens distortion), give rise to deviations in the computed local measures. We here use this to formulate an objective criterion in order to estimate the unknown coefficients in the transformation (2). One way to do this is to define it as

$$
\varepsilon_s = |\overline{\mathcal{I}_{xx}}(\mathcal{F}(I)) - 0.5| + |\overline{\mathcal{I}_{xy}}(\mathcal{F}(I))|, \quad (12)
$$

where $\overline{\mathcal{I}_{xx}}(\mathcal{F}(I))$ denotes the mean of the shading measure computed for the image after transforming its intensities via the mapping $\mathcal{F}$ in (2). Note that we have derived three local shading measures, but only two are independent. The above criterion reflects this fact. In addition, at most two coefficients in (2) can be solved for. As the measures are also invariant to linear transformation of image intensities,

(a)

(b)
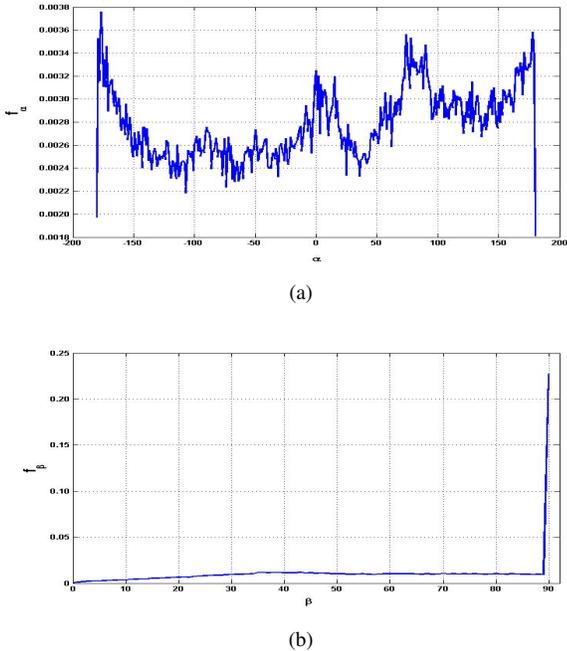
Figure 1. Distributions f_φ (a) and f_ψ (b) of the angles φ and ψ of the tooth surface's normals as computed from the set of 3D teeth surfaces used in building the shape model.

the coefficient c_0 has no affect, and the remaining ones can be determined up to a common scale. Thus for our purpose here, the transformation (2) is written more precisely as $\mathcal{F}(I) = I\,(1 + c'_1\,I + c'_2\,I^2)$, with $c'_1 = c_2/c_1$ and $c'_2 = c_3/c_1$.

The proposed shading correction algorithm thus seeks to estimate the two parameters, namely c'_1 and c'_2, that minimize the criterion (12). In our experience, this criterion often has several local minima, so traditional gradient-based local optimization routines may not work properly to find the global minimum. In order to overcome these drawbacks we use a methodology consisting of two steps: first, determine good initial start values by means of a state-of-the-art global optimization technique and secondly, perform a fine-tuning derivative-free simplex search using the previous result as start value. To determine good initial starting values we use the method of Coupled Simulated Annealing with variance control (CSA) [15]. CSA has already proven to be more effective than multi-start gradient descent optimization [12]. Another advantage of CSA is that it uses the acceptance temperature to control the variance of the acceptance probabilities with a control scheme. This leads to an improved optimization efficiency because it reduces the sensitivity of the algorithm to the initialization parameters while guiding the optimization process to quasi-optimal runs. This initial result is then used as a starting value for a derivative-free simplex search. This extra step is a fine-

tuning procedure resulting in more optimal tuning parameters and hence better performance. In our experiments, a search space of $[-2, 2]$ for either parameter was found satisfactory.

Having estimated the transformation parameters, the intensities of all pixels in the original input image are mapped into new values that better satisfy the typical SFS assumptions. Afterwards, the transformed image can be arbitrarily scaled within the intensity range $[0, 255]$, if required, as this scaling does not affect the local shading measures (or the error criterion).

3. Experimental Results

Several experiments are carried out to evaluate the performance of the proposed shading correction strategy. In order to quantify the gain out of this strategy, it is important to apply it on real data whose ground-truth depth surfaces are known. Qualitative results on the popularly-used test images in the SFS literature (mainly synthetic images) will not be sufficient or very useful [5] for our purpose here. As such, we have constructed a set of images for real 3D surfaces with known depth information based on real human teeth. The accuracy of the SFS outcome is assessed by comparing the resulting surface to the 3D ground-truth surface as obtained from CT scan. This will allow us more rigorous evaluation of the proposed strategy.

We apply the proposed shading correction algorithm on the input test images. The algorithm takes about 25s on a PC with 2.2GHz AMD processor and 4GB RAM. Then the corrected images are fed to a standard SFS algorithm [6] assuming typical assumptions of Lambertian surface, orthographic projection and distant directional light source in order to estimate the surfaces' depth maps. The results are quantitatively compared to the results of the same SFS algorithm [6] without the correction process as well as to the results of a more sophisticated SFS algorithm [2, 9] based on viscosity solutions for Hamilton-Jacobi type formulation of the image irradiance equation. This algorithm was originally developed in [9] for the Lambertian case assuming a perspective camera projection and near-field illumination. Later [2] it has been extended to the Oren-Nayar (O-N) reflectance model which more properly describes the roughness of the real tooth surface. Table 1 summarizes the key differences between the methods under evaluation and comparison.

An ICP-based rigid registration algorithm [3] is used to align the 3D ground-truth surface as obtained from CT scan with the surface recovered from the SFS algorithm. The performance accuracy is assessed using the root mean square (RMS) error between the corresponding 3D points from both surfaces after the 3D alignment.

Table 1. Summary of Methods Under Experimental Comparison.

	Camera	Light source	Reflectance	Corrected	Reference
A	ortho.	infinity	Lamb.	No	[6]
B	pers.	optical center	O-N	No	[2, 9]
C	ortho.	infinity	Lamb.	Yes	*Ours* + [6]

3.1. Dataset

A dataset is constructed from several real invitro teeth of various types (maxillary deciduous, maxillary molars, mandibular molars, or mandibular third molars) from adult subjects with different races, gender and ages. At least 30 teeth of each type have been collected. A Cone-beam CT (KODAK 9000 3D Extraoral) scanner at a resolution of 76 μm is then used to scan the teeth. Figure 2 shows samples of the acquired CT scans. The root part of each tooth is manually taken out from the CT scan, and the visible 3D tooth surface is segmented. Furthermore, the surface roughness of the tooth surfaces which is needed for the O-N diffuse reflection model in Algorithm B is estimated. This is done with the help of a 3D optical surface profiler (NewView 700s from Zygo company), which is based on Scanning White-Light Interferometry technology that offers fast, non-contact, accurate 3D metrology of micro surface features. A near-field imaging setup consisting of a small camera with a built-in bright light source is employed to acquire the 2D images of the occlusal surfaces for all the teeth. The camera has 1/4 Sony CCD, and an image resolution of 2.0 mega pixels.
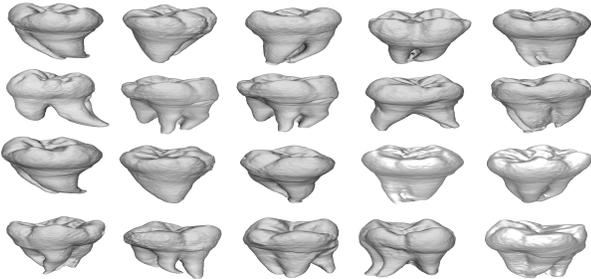
Figure 2. Sample CT scans for various human tooth types. (First row) Maxillary deciduous teeth. (Second row) Maxillary molars. (Third row) Mandibular molars. (Fourth row) Mandibular third molars.

3.2. Results and Comparisons

Sample results from the methods under evaluation are demonstrated in Fig. 3. Fig. 3(a) shows the input test images for four different teeth models (maxillary deciduous, maxillary molar, mandibular molar, and mandibular third molar), while Fig. 3(b) shows the corresponding ground-truth tooth surface as obtained from CT scans. Fig. 3(c)depicts the images after applying the shading correction algorithm to the same images in Fig. 3(a). The outputs from Algorithm A, shown in Fig. 3(d), are rather bumpy with too many peaks as

the algorithm fails to handle the inevitable specularity due to the tooth surface characteristics and the near-illumination setup. As it assumes more realistic settings of perspective projection, near light source and the O-N reflectance model, Algorithm B has provided rather better global shape reconstructions, see Fig. 3(e). However the algorithm is not able to recover the geometrical details of the occlusal surface. Although there are still some bumpy results, our method C shown in Fig. 3(f) has produced better results due to the shading correction pre-process. This can be verified from the root-mean-square (RMS) error between the reconstruction and the ground-truth after performing the 3D rigid alignment, which is given beneath each reconstruction in Fig. 3(d)-(f). Consistently, the proposed strategy (method C) of shading correction followed by a simple classical SFS algorithm provides the smallest RMS errors. The method has an error as low as $50 - 70\%$ of that of the more sophisticated SFS algorithm B but without shading correction. This emphasizes the role of the shading correction pre-process in shape recovery. It is important to stress here that while the improvements are fractions of a millimeter, this is considered significant for dental-related applications such as tooth implant and surface analysis.
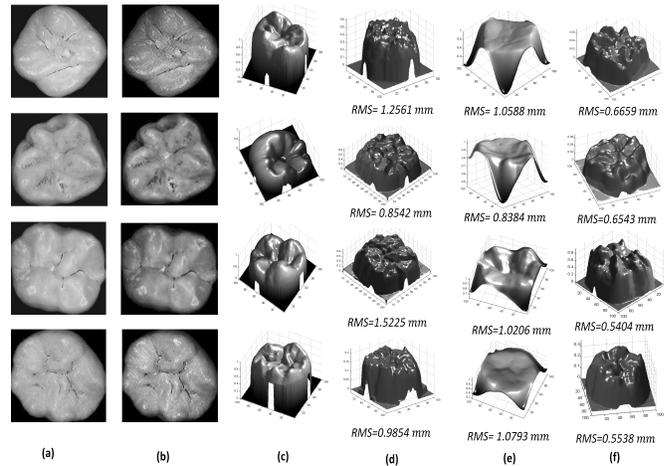
Figure 3. Tooth reconstruction from three different methods. (a) Image acquired by the setup camera.(b) Shading-corrected images for those in (a).(c) Ground-truth occlusal surface generated from a CT scan of the tooth. (d) Reconstruction using Algorithm A. (e) Reconstruction using Algorithm B. (f) Reconstruction using our proposed method (C). Beneath each reconstruction is the root-mean-square (RMS) error when compared to the ground-truth surface from CT.

The overall performance of the various methods on the entire data can be summarized in Table 2, which lists the average RMS error in mm over over all the four types of the dataset teeth (30 per tooth type) for the three methods. It is clear that our method C has considerably better accuracy than the other methods. Although method C relies on the same SFS algorithm [6] used in method A, its overall

error is just 53% of that of method A. It has also an error of 63% of that of the more sophisticated SFS algorithm [2, 9] in method B. This highly emphasizes the gain out of the proposed shading correction algorithm.

Table 2. Overall surface reconstruction accuracy (RMS) in mm

Tooth Type	A	B	C
maxillary deciduous	1.34	1.11	0.74
maxillary molar	0.84	0.85	0.65
mandibular molar	1.60	1.20	0.60
mandibular third molar	1.06	1.01	0.66
Overall	1.24	1.05	0.66

4. Conclusions

This paper has addressed the problem of correcting the shading of a given intensity image to improve the results of shape recovery using classical SFS algorithms. This is a novel alternative strategy to developing more sophisticated SFS algorithms. To the best of our knowledge, this has been seldom addressed before in the literature. We have proposed a new simple shading-correction algorithm that transforms the image to a new image that better satisfies the assumptions typically needed by existing SFS algorithms. It works blindly on the input image without the knowledge of the original reflectance model, or the illumination direction relative to the viewing direction. It is based on a number of local shading measures derived from the mathematical model of the image shading formation under a typical set of SFS assumptions.

The proposed method is evaluated on real teeth images in order to reconstruct the shape of the crowns. Our experimental results have shown that the shading-corrected images have indeed improved the reconstruction accuracy considerably as compared to ground-truth CT scans (even using the same SFS algorithm). Moreover, a simpler SFS algorithm [6] applied to the corrected image has demonstrated a clearly better performance than a more sophisticated PDE-based SFS algorithm [2, 9] applied to the original images. Although in this paper, we focused on the well-known SFS algorithm of Frankot and Chellappa [6] in order to demonstrate the gain out of the proposed shading correction algorithm, the correction algorithm can also be used with any of the classical SFS algorithms to pre-process the input images in order to improve their shape recovery results.

References

[1] A. Abdelrahim, M. El-Melegy, and A. Farag. Realistic 3d reconstruction of the human teeth using shape from shading with shape priors. In *IEEE Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 64–69, 2012.

[2] A. Ahmed and A. Farag. Shape from shading under various imaging conditions. In *Proc. of IEEE CVPR'07, Minneapolis, MN*, pages X1–X8, June 18-23, 2007.

[3] H. Chui and A. Rangarajan. A new point matching algorithm for non-rigid registration. *Comput. Vis. Image Underst.*, 89:114–141, February 2003.

[4] J.-D. Durou, M. Falcone, and M. Sagona. Numerical methods for shape-from-shading: A new survey with benchmarks. *Computer Vision and Image Understanding*, 109(1):22 – 43, 2008.

[5] D. Forsyth. Variable-source shading analysis. *International Journal of Computer Vision*, 91:280–302, 2011.

[6] R. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Trans. on PAMI*, 10(4):439–451, 1988.

[7] A. Litvinov and Y. Schechner. Addressing radiometric non-idealities: a unified framework. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2005.*, volume 2, pages 52–59 vol.2, June 2005.

[8] M. Oren and S. Nayar. Generalization of the lambertian model and implications for machine vision. *International Journal of Computer Vision*, 14(3):227 – 251, 1995.

[9] E. Prados, F. Camilli, and O. Faugeras. A unifying and rigorous shape from shading method adapted to realistic data and applications. *Journal of Mathematical Imaging and Vision*, 25:307–328, 2006.

[10] H. Ragheb and E. Hancock. Surface normals and height from non-lambertian image data. In *Proceedings. 2nd International Symposium on 3D Data Processing, Visualization and Transmission, 3DPVT 2004.*, pages 18 – 25, Sept. 2004.

[11] H. Ragheb and E. R. Hancock. Surface radiance correction for shape from shading. *Pattern Recognition*, 38(10):1574–1595, Oct. 2005.

[12] J. A. Suykens, J. Vandewalle, and B. D. Moor. Intelligence and cooperative search by coupled local minimizers. *International Journal of Bifurcation and Chaos*, 11(8):2133–2144, 2001.

[13] A. Tankus, N. Sochen, and Y. Yeshurun. Shape-from-shading under perspective projection. *International Journal Computer Vision*, 63:21–43, June 2005.

[14] O. Vogel, L. Valgaerts, M. Breu, and J. Weickert. Making shape from shading work for real-world images. In J. Denzler, G. Notni, and H. Se, editors, *Pattern Recognition*, volume 5748 of *Lecture Notes in Computer Science*, pages 191–200. Springer Berlin Heidelberg, 2009.

[15] S. Xavier-de Souza, J. A. K. Suykens, J. Vandewalle, and D. Bolle. Coupled simulated annealing. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 40(2):320–335, 2010.

[16] S. Y. Yuen, Y. Y. Tsui, and C. K. Chow. A fast marching formulation of perspective shape from shading under frontal illumination. *Pattern Recognit. Lett.*, 28:806–824, 2007.

[17] Q. Zheng and R. Chellapa. Estimation of illuminant direction, albedo and shape from shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 13(7):680–702, 1991.