

3-D Reconstruction of Shaded Objects from Multiple Images Under Unknown Illumination

Hailin Jin · Daniel Cremers · Dejun Wang ·
Emmanuel Prados · Anthony Yezzi · Stefano Soatto

Received: 6 October 2006 / Accepted: 19 March 2007
© Springer Science+Business Media, LLC 2007

Abstract We propose a variational algorithm to jointly estimate the shape, albedo, and light configuration of a Lambertian scene from a collection of images taken from different vantage points. Our work can be thought of as extending classical multi-view stereo to cases where point correspondence cannot be established, or extending classical shape from shading to the case of multiple views with unknown light sources. We show that a first naive formalization of this problem yields algorithms that are numerically unstable, no matter how close the initialization is to the true geometry. We then propose a computational scheme to overcome this problem, resulting in provably stable algorithms that con-

verge to (local) minima of the cost functional. We develop a new model that explicitly enforces positivity in the light sources with the assumption that the object is Lambertian and its albedo is piecewise constant and show that the new model significantly improves the accuracy and robustness relative to existing approaches.

Keywords Stereoscopic segmentation · Shape from shading · Multi-view stereo · Variational 3D reconstruction · Level set methods · Lighting and appearance reconstruction

H. Jin (✉)
Office of Technology, Adobe Systems Incorporated, 345 Park
Avenue, San Jose, CA 95110, USA
e-mail: hljin@adobe.com

D. Cremers
Dept. of Computer Science, Römerstr. 164, 53117 Bonn,
Germany
e-mail: dcremers@cs.uni-bonn.de

D. Wang · S. Soatto
Computer Science Department, University of California, Los
Angeles, CA 90095, USA

S. Soatto
e-mail: soatto@cs.ucla.edu

E. Prados
INRIA Rhône-Alpes, 655 avenue de l'Europe, 38330
Montbonnot, France
e-mail: emmanuel.prados@inrialpes.fr

A. Yezzi
School of Electrical and Computer Engineering, Georgia Institute
of Technology, Atlanta, GA 30332, USA
e-mail: ayezzi@ece.gatech.edu

1 Introduction

We address the problem of recovering the three-dimensional shape of a scene or object that has diffuse, or “Lambertian”, reflection, seen from multiple images taken from different vantage points. We assume that both intrinsic and extrinsic calibration parameters of the camera that took the images are known. We do not assume that we know the illumination, and explicitly estimate the ambient illumination level and the position and intensity of one or more point light sources.

This problem can be thought of as the multi-view extension of the problem of “shape from shading” (Horn and Brooks 1989; Koenderink and van Doorn 1980; Zheng and Chellappa 1991; Nayar et al. 1991; Oliensis and Dupuis 1993; Zhang et al. 1999; Robles-Kelly and Hancock 2004, 2005; Klette et al. 1998; Durou et al. 2004; Prados 2004) where one makes the assumption that the objects in the scene have homogeneous material, and therefore the shading effects seen on the image are due to their shape, that can be retrieved modulo a bas-relief ambiguity (Belhumeur et al.

1999; Yuille et al. 2003). However, in our model we allow the objects to have piecewise homogeneous materials, a common trait of most man-made objects. Also, we have multiple views available, which allow us to estimate the position and intensity of the light source along with the shape of the observed scene.

The problem also relates to “multi-view stereo”, where one is given multiple views of an object and estimates its shape together with the radiance (“texture map”) of the scene (Kutulakos and Seitz 2000; Faugeras and Keriven 1998; Jin et al. 2003a; Keriven et al. 2006; Seitz et al. 2006). We refer the reader to (Seitz et al. 2006) for evaluations of some of the recent algorithms. However, while in multi-view stereo one has to rely on the gradient being non-zero everywhere, and no knowledge on the reflectance properties of the scene (other than Lambertian) is enforced, we can explicitly make use of the assumption of piecewise constancy of the reflectance, and we also retrieve the lighting configuration, which is not addressed in multi-view stereo. Let us underline here that our scenario is different from other ones which deal with multiple images and varying illumination (Yuille et al. 1999; Chen et al. 2000). More germane to our approach are the work of Samaras (2003) which estimates the lighting direction and the work of Samaras et al. (2000) which has also attempted to combine stereo and shape from shading, although they are different than ours because we do not attempt to establish direct correspondence.

Finally, this problem relates to segmentation, or better to multi-view, or stereoscopic, segmentation (Yezzi and Soatto 2001), in the sense that scenes for which point-to-point correspondence cannot be easily established (for instance because they do not have distinctive enough gradient profiles) are typically easy to segment, and vice-versa. For a recent reformulation of stereoscopic segmentation based on a probabilistic treatment of individual voxels we refer to (Kolev et al. 2006). Our goal is to ultimately be able to integrate all cues, and arrive at a coherent 3-D reconstruction framework that can exploit any available knowledge on the scene, and arrive at a consistent estimate of its shape.

Such an algorithm potentially has important applications in 3-D reconstruction for computer-aided design, manufacturing, mapping, robotic navigation, visual recognition, and computer graphics, specifically image-based modeling and rendering.

To simplify the problem, we assume that the scene of interest is surrounded by a background with uniform radiance and has a piecewise constant albedo. We show how the problem can be formalized as an infinite-dimensional optimization task. Unfortunately the naive algorithm, based on an iterative procedure designed to have the first-order optimality conditions as its fixed point, yields a numerically unstable flow, and therefore cannot be implemented except

for a few special cases. We therefore introduce a different model, based on an auxiliary vector field. This auxiliary vector field, which is directly responsible for shading effects, is then coupled to the field normal to surfaces in the scene via an energy term. This artifice, which can be interpreted as a relaxation, allows us to have a provably stable flow, that has the added benefit of not depending on any derivative of the image, resulting in increased robustness with respect to noise.

Another challenge that we had to overcome, that was a limitation in the prior art, was the enforcing of physical constraints such as the positivity of the light sources. We show that the resulting model yields successful results where prior algorithms failed, as we demonstrate experimentally on synthetic and real image datasets designed to challenge the assumptions underlying our model.

We implement our algorithm in the level set framework, that yields the added benefit of automatically allowing changes of topology both in the shape of the scene, and in the radiance profile, where the boundary between constant materials can split and merge.

This article relates to two of our previous conferences papers, (Jin et al. 2004a) and (Jin et al. 2004b). Jin et al. (2004b) presented the work on recovering shape for piecewise constant radiance without considering the light effects. In Jin et al. (2004a), we recover the shape and the light configuration together and solve the variational problems via level set methods. We then can recover the shape, the lighting conditions and a piecewise constant approximation of the object’s albedo. Moreover, we propose here a physically plausible model of illumination, that explicitly enforces positivity of the sources, without the need to revert to the concept of “negative” light sources introduced in Jin et al. (2004a). We finally improve the algorithm by solving some intermediate steps in closed form.

2 Problem Formulation

In this section we introduce our notation and the basic formalization of the problem, and show how it yields numerically unstable algorithms. This serves as a motivation for our solution, introduced in the next section.

2.1 Notation and Modeling Assumptions

Let S be a smooth two-dimensional surface embedded in \mathbb{R}^3 . We denote with $\mathbf{X} = [X, Y, Z]^T \in \mathbb{R}^3$ the coordinates of a generic point on S with respect to a fixed reference frame. Our goal is to reconstruct the surface S from a set of n images $I_i : \Omega_i \rightarrow \mathbb{R}^+$, $i = 1, \dots, n$, where $\Omega_i \subset \mathbb{R}^2$. The intrinsic and extrinsic calibration parameters for each camera are assumed to be known (see e.g. Chapter 4 of (Ma et al. 2003)), so after some simple pre-processing we can

model each camera as an ideal perspective projection $\pi_i : \mathbb{R}^3 \rightarrow \Omega_i$; $\mathbf{X} \mapsto x_i \doteq \pi_i(\mathbf{X}) = \pi(\mathbf{X}_i) = [X_i/Z_i, Y_i/Z_i]^T$, where $\mathbf{X}_i = [X_i, Y_i, Z_i]^T$ are the coordinates of \mathbf{X} in the i -th camera reference frame. \mathbf{X} and \mathbf{X}_i are related by a rigid body transformation, which can be represented in coordinates by a rotation matrix $R_i \in SO(3)$ and a translation vector $T_i \in \mathbb{R}^3$: $\mathbf{X}_i = R_i\mathbf{X} + T_i$. We assume that there is a background B filling the field of view of each camera, modeled as a sphere with infinite radius. For each camera we define the foreground to be the region $Q_i = \pi_i(S) \subset \Omega_i$ and denote its complement in Ω_i , the background, by Q_i^c . We also define the back-projection $\pi_i^{-1} : \Omega_i \rightarrow \mathbb{R}^3$ of x_i onto S , which is the first intersection of the ray through x_i in the i -th camera with S .

We assume that the scene is made of Lambertian surfaces. For simplicity, we only deal with grayscale images (even if the extension to color images is quite direct), so the reflectance of the scene is described by positive scalar functions (called *albedos*): $\rho : S \rightarrow \mathbb{R}^+$ for the foreground, and $h : B \rightarrow \mathbb{R}^+$ for the background. In what follows we assume that the albedo of the background is constant and that the albedo of the foreground is piecewise constant. Furthermore, we assume that the background is uniformly illuminated which results in a uniform radiance. We denote the radiance value by h with an abuse of notation (for an extension to smooth backgrounds see Jin et al. 2003b). While this may seem like a restrictive assumption, it is a rather good approximation of most man-made objects. Besides, in lack of any assumption about albedo, there is little one can say about the shape of the scene in the presence of changing illumination (Chen et al. 2000).

We assume that illumination can be well approximated by a superposition of two components: One is an *ambient term*, with constant energy E_0 radiated isotropically in all directions. This term approximates inter-reflections, diffuse illumination and short-range effects that would be too complex to model explicitly (although see Langer and Zucker 1994; Stewart and Langer 1997 for a way to exploit shading induced by ambient illumination). Again, it is a reasonable approximation in most well-lit environments including indoors (barring black carpet floors) and outdoors on a cloudy day. The second term consists of a number of distant *point light* sources. For simplicity we only allow one source (e.g. the sun on a clear day) but extension to any number (e.g. spotlights in a theater) presents no theoretical difficulties. Using old theorems by Wiener one can show that any positive distribution on the sphere can be approximated arbitrarily well by such a collection of sources (Vedaldi and Soatto 2006).

Using Lambert’s cosine law (Ma et al. 2003; Horn 1986), we can evaluate the radiance r at each $\mathbf{X} \in S$ via

$$r(\mathbf{X}) = \rho(\mathbf{X})(\langle N(\mathbf{X}), L \rangle \xi(\mathbf{X}) + E_0),$$

where $L \in \mathbb{R}^3$ is a vector pointing in the direction of the light source, with norm equal to its intensity, $N(\mathbf{X}) \in \mathbb{S}_2$ is the outward unit normal to S at \mathbf{X} , $\xi : S \rightarrow \{0, 1\}$ is the visibility of the light and $\rho(\mathbf{X})$ is the albedo of the surface at \mathbf{X} . In the case of convex objects, visibility is given by $\xi = \mathcal{H}(\langle N, L \rangle)$, where \mathcal{H} denotes the Heaviside step function: $\mathcal{H}(x) = 1, \forall x \geq 0$ and $\mathcal{H}(x) = 0$, otherwise. More in general, under the assumptions above, we can write the image formation model as:

$$I(\pi(\mathbf{X})) = \rho(\mathbf{X})(\langle N(\mathbf{X}), L \rangle \xi(\mathbf{X}) + E_0). \tag{1}$$

Let us remark here that we implicitly assume that the brightness of the image I at a pixel x is equal to the irradiance of the associated 3D point located on the retinal plane of the camera (our work can be easily extended to more realistic models, see for example Horn 1986).

Since the albedo ρ is piecewise constant, for simplicity we assume that it partitions S into two regions D_1 and D_2 , which can have several connected components. We denote by C the union of the curves dividing D_1 and D_2 , and assume they are smooth and closed. Therefore, ρ can be described as

$$\rho(\mathbf{X}) = \rho_j \in \mathbb{R} \quad \text{for } \mathbf{X} \in D_j, \quad j = 1, 2.$$

Note that model (1) is not minimal in the sense that there is a multiplicative ambiguity between ρ , L and E_0 as they appear together in products. We shall address this issue later. Extensions of the albedo model to more than two phases are straight forward using respective multiphase level set formulations (Vese and Chan 2002; Brox and Weickert 2006).

2.2 Formalization of the Problem in a Variational Framework

The scene is described by the three-dimensional (3-D) surface S , and by its reflectance, that given the assumptions above is determined by the 2-D curves C on S and the two scalars ρ_1 and ρ_2 . The intensity of the background h and the light sources, L and E_0 , are nuisance parameters, in the sense that they are unknown and not necessarily of direct interest, but they affect the measured images, and we will therefore recover them along with the description of the scene.

In order to recover a 3-D model of the scene, we wish to solve an optimization problem whereby we find, among all possible scenes, the one(s) that generate “virtual” images that most closely resemble the given (measured) ones. In other words, we wish to minimize the discrepancy, for instance the squared \mathcal{L}^2 distance, between the right-hand-side and the left-hand-side of (1). We call such a discrepancy

E_{data} , in the sense that it depends on the measured images

$$\begin{aligned}
 E_{\text{data}} \doteq & \sum_{i=1}^n \int_{\pi_i(D_1)} (I_i(x_i) - \rho_1 \langle N, L \rangle \\
 & \times \xi(\pi_i^{-1}(x_i)) + E_0)^2 d\Omega_i \\
 & + \sum_{i=1}^n \int_{\pi_i(D_2)} (I_i(x_i) - \rho_2 \langle N, L \rangle \\
 & \times \xi(\pi_i^{-1}(x_i)) + E_0)^2 d\Omega_i \\
 & + \sum_{i=1}^n \int_{Q_i^c} (I_i(x_i) - h)^2 d\Omega_i. \tag{2}
 \end{aligned}$$

Among the unknowns, $\rho_1, \rho_2, E_0, L, h, C, S$, the latter two are infinite-dimensional. Therefore, one needs to impose additional assumptions, for instance in the form of regularization, to arrive at a well-posed problem. The most obvious choice is to search for a solution that maximizes the regularity of C and S . That can be easily achieved by adding terms to the energy functionals, for instance

$$E_{\text{smooth}} \doteq \int_S dA \tag{3}$$

that measures the area of S and, when minimized, forces regularity, and similarly for the curves C

$$E_{\text{curv}} \doteq \int_C ds, \tag{4}$$

where dA and ds are the area form of S and the arc length of C respectively. Unfortunately, even considering the composite functional $E_{\text{data}} + \alpha E_{\text{smooth}} + \beta E_{\text{curv}}$, where α and β are positive weighting factors, is not sufficient to yield a well-posed problem, as we illustrate next.

2.3 Stable Formulation by Decoupling the Normal Field

The minimization of the energy functional just described represents an ill-posed problem. To illustrate this phenomenon, consider a simpler version of the energy above, where we neglect the background term and we consider only a constant albedo foreground, $\rho_1 = \rho_2 = \rho$. In other words, we concentrate only on the first line of (2), and for simplicity we neglect the visibility term, so $\xi(\mathbf{X}) = 1, \forall \mathbf{X} \in S$. After some calculations, the reader can verify that the curvature-dependent term of the first-order variation of E_{data} is given by (Jin 2003)

$$\sum_{i=1}^n 2H((I_i - E_0)^2 + 2\rho^2 - 3\langle \rho L, N \rangle^2) - 2\rho^2 \Pi(N \times L), \tag{5}$$

where H denotes the mean curvature and Π the second fundamental form. In regions where the surface faces the light

(and therefore $\langle L, N \rangle \rightarrow 1$) and where the modeled directional intensity ρ^2 exceeds the measured one $(I_i - E_0)^2$, the coefficients of both H and Π are negative. This amounts to a backward heat flow, that models a physically unrealizable process and is well-known to be numerically unstable. Therefore, a local gradient-based minimization algorithm based on the first-order optimality condition is bound to be numerically unstable no matter how close the initial conditions are to the true solution.

The above instability arises due to the strong coupling between surface appearance, as measured by the image, and its normal vector field N . In the presence of noise in the image, with fixed illumination, the surface will bend and ripple to fit the data. To circumvent this instability, we propose a relaxed cost functional in which the normal is decoupled from the surface through an auxiliary unit vector field V , defined as

$$V : S \rightarrow \mathbb{S}_2, \quad \mathbf{X} \mapsto V(\mathbf{X}), \tag{6}$$

where $\mathbb{S}_2 \subset \mathbb{R}^3$ is the unit sphere of dimension 2. V will take the place of the unit normal in modeling the shading effects. We will show that the induced surface flow lacks the potentially unstable curvature-based diffusion terms. The data fitness term now becomes

$$\begin{aligned}
 E_{\text{data}} \doteq & \sum_{i=1}^n \int_{\pi_i(D_1)} (I_i(x_i) \\
 & - \rho_1 \langle V(\pi_i^{-1}(x_i)), L \rangle \xi(\pi_i^{-1}(x_i)) + E_0)^2 d\Omega_i \\
 & + \sum_{i=1}^n \int_{\pi_i(D_2)} (I_i(x_i) \\
 & - \rho_2 \langle V(\pi_i^{-1}(x_i)), L \rangle \xi(\pi_i^{-1}(x_i)) + E_0)^2 d\Omega_i \\
 & + \sum_{i=1}^n \int_{Q_i^c} (I_i(x_i) - h)^2 d\Omega_i. \tag{7}
 \end{aligned}$$

To stay faithful to the physically motivated interaction between the surface normal and the light source direction, we will introduce an indirect coupling between the unit normal and the modeled surface radiance by adding a term to our energy which penalizes the average deviation between the true unit normal of the surface and the unit vector field V which takes its place in the new radiance model. The constraint for V is given by a penalty on the \mathcal{L}^2 distance between V and the unit normal field N on S :

$$\begin{aligned}
 E_{\text{coupling}} &= \frac{1}{2} \int_S |V(\mathbf{X}) - N(\mathbf{X})|^2 dA \\
 &= \int_S (1 - \langle V(\mathbf{X}), N(\mathbf{X}) \rangle) dA, \tag{8}
 \end{aligned}$$

where we have used the fact that the vectors V and N have unit length. The overall cost functional is simply a weighted average of the several costs:

$$E_{\text{total}} = E_{\text{data}} + \alpha E_{\text{smooth}} + \beta E_{\text{curv}} + \gamma E_{\text{coupling}}. \tag{9}$$

When minimizing E_{total} , we need to guarantee that V is always a unit vector field, i.e. $|V(\mathbf{X})|^2 = 1, \forall \mathbf{X} \in S$, and that the radiance coefficients ρ_1, ρ_2 are positive.

Decoupling the effective normal field from the true surface normal allows us to bypass this ill-posed formulation. As the auxiliary field does not depend upon the surface derivatives (as does the true unit normal), the optimality conditions for V (and the ensuing gradient flow) do not involve the second-order derivatives contained in heat flows.

3 Optimal Reconstruction

In order to minimize E_{total} with respect to all the unknowns, we adopt an alternating minimization procedure, evolving one step at a time in the gradient direction of each unknown. For the simplest unknowns, we will be able to compute a closed-form solution for each iteration.

3.1 Updating of the Surface Properties: Geometry and Albedo

We start with minimization with respect to S , that is the most delicate. To facilitate finding the variation of the data fitness term with respect to S , we need to introduce two more terms. Let $\chi_i : S \rightarrow \{0, 1\}$ be the surface visibility function with respect to the i -th camera, i.e. $\chi_i(\mathbf{X}) = 1$ for points on S that are visible from the i -th camera and $\chi_i(\mathbf{X}) = 0$ otherwise. Let σ_i account for the change of coordinates from $d\Omega_i$ to dA , i.e., $\sigma_i = \frac{d\Omega_i}{dA} = \langle \mathbf{X}_i, N_i \rangle / Z_i^3$, where N_i the unit normal N expressed in the i -th camera reference frame, $N_i = R_i N$. For simplicity, we first illustrate the properties of the flow for the case of constant albedo foreground: The data term is

$$\begin{aligned} & \sum_{i=1}^n \int_{Q_i} ((I_i - \langle V, L\xi \rangle - E_0)^2 - (I_i - h)^2) d\Omega_i \\ & + \sum_{i=1}^n \int_{\Omega_i^c} (I_i - h)^2 d\Omega_i \\ & = \sum_{i=1}^n \int_S \chi_i ((I_i - \langle V, L\xi \rangle - E_0)^2 - (I_i - h)^2) \sigma_i dA \\ & + \sum_{i=1}^n \int_{\Omega_i^c} (I_i - h)^2 d\Omega_i, \end{aligned}$$

where the background integral over the complement Ω_i^c of Q_i was expressed as an the integral over the entire domain Ω_i minus an integral over Q_i .

And the gradient descent flow based on the first-order variation is given by (Jin 2003)

$$\begin{aligned} S_t = & \left(\sum_{i=1}^n \frac{1}{Z_i^3} ((I_i - (\langle V, L\xi \rangle + E_0))^2 - (I_i - h)^2) \right. \\ & \cdot \left. \left\langle \nabla \chi_i, R_i^T \mathbf{X}_i \right\rangle - \sum_{i=1}^n 2\chi_i (I_i - \langle V, L\xi \rangle - E_0) \right. \\ & \times \left(\xi L^T \nabla_S V R_i^T \mathbf{X}_i + \sum_{j=1}^l \langle V, L \rangle \langle \nabla \xi, R_j^T \mathbf{X}_j \rangle \right) \\ & \left. + (2H(\alpha + \beta) - \beta \nabla_S \cdot V) \right) N, \tag{10} \end{aligned}$$

where ∇_S is the Laplace–Beltrami operator. Note that the only second-order term (curvature term) in the flow (10) is $2H(\alpha + \beta)N$, therefore the flow is always numerically stable (with a properly chosen time step). Another advantage of flow (10) is that it depends only upon the image values, *not the image gradients*. This property greatly increases the robustness of the resulting algorithm to image noise.

For the more general case, where we have piecewise constant albedo foreground regions separated by the contours C , to facilitate the computing of the gradient descent flow we define the level set function $\phi : S \rightarrow \mathbb{R}$ to represent C via $D_1 = \{X | \phi(X) > 0\}$, $D_2 = \{X | \phi(X) < 0\}$ and $C = \{X | \phi(X) = 0\}$. $\mathcal{H}(\phi)$ is the Heaviside function of the level set function ϕ . The gradient flow of the data term with respect to the surface S is given by (Jin 2003)

$$\begin{aligned} S_t = & \left(\sum_{i=1}^n \frac{\Gamma_i}{Z_i^3} \langle \chi_i X, R_i^T \mathbf{X}_i \right) \\ & - \chi_i \frac{\delta(\phi)}{Z_i^3} ((I_i - \rho_1(\langle V, L \rangle \xi + E_0))^2 \\ & - (I_i - \rho_2(\langle V, L \rangle \xi + E_0))^2) \langle \nabla_S \phi, R_i^T \mathbf{X}_i \rangle \\ & + 2\alpha k + \beta \frac{\delta(\phi)}{|\nabla_S \phi|} \Pi(\nabla_S \phi \times N) \Big) N, \tag{11} \end{aligned}$$

where we have defined

$$\begin{aligned} \Gamma_i \doteq & \mathcal{H}(\phi)(I_i - \rho_1(\langle V, L \rangle \xi + E_0))^2 \\ & + (1 - \mathcal{H}(\phi))(I_i - \rho_2(\langle V, L \rangle \xi + E_0))^2 - (I_i - h)^2. \end{aligned}$$

Similarly one can derive the component of the gradient flow due to the data term for the curves C (Jin et al. 2007):

$$\begin{aligned} C_t = & \left(\sum_{i=1}^n ((I_i - \rho_2(\langle V, L \rangle \xi + E_0))^2 \right. \\ & \left. - (I_i - \rho_1(\langle V, L \rangle \xi + E_0))^2) \sigma_i + \beta k_g \right) \vec{n}, \tag{12} \end{aligned}$$

where \vec{n} is the normal direction of a point in the segmenting curve between regions of different albedo, k is the curvature, k_g is the geodesic curvature. In our level set implementation, we implement the evolution of the level set function ϕ instead of C , as follows (Jin et al. 2007):

$$\begin{aligned} \phi_t = & |\nabla_S \phi| \sum_{i=1}^n ((I_i - \rho_2(\langle V, L \rangle \xi + E_0))^2 \\ & - (I_i - \rho_1(\langle V, L \rangle \xi + E_0))^2) \sigma_i \\ & + \beta \left(\delta_S \phi - \frac{\nabla_S^T \phi \nabla_S^2 \phi \nabla_S \phi}{|\nabla_S \phi|^2} \right). \end{aligned} \tag{13}$$

S and C being now updated, we then update ρ_1 and ρ_2 as follows:

$$\rho_1 = \frac{\sum_{i=1}^n \int_{\pi_i(D_1)} I_i(x_i) d\Omega_i}{\sum_{i=1}^n \int_{\pi_i(D_1)} (\langle V, L \rangle \xi + E_0) d\Omega_i}, \tag{14}$$

$$\rho_2 = \frac{\sum_{i=1}^n \int_{\pi_i(D_2)} I_i(x_i) d\Omega_i}{\sum_{i=1}^n \int_{\pi_i(D_2)} (\langle V, L \rangle \xi + E_0) d\Omega_i}. \tag{15}$$

The next step consists in fixing S, C, ρ_1, ρ_2 and in minimizing the energy with respect to V, L, E_0, h while enforcing normalization and positivity constraints.

3.2 Updating of Radiance Parameters L, E_0 and h

In this step, we fix S, C, V, ρ_1 and ρ_2 and we minimize E_{total} with respect to L, E_0 and h .

Let us start with the updating of h . One can easily verify that the optimal value for h is the mean, on all the images (I_i, Ω_i) , of the intensity of the pixels located in the background parts Q_i^c (let us remind that Q_i^c represents the complement of the region $Q_i = \pi_i(S) \subset \Omega_i$); i.e.

$$h = \frac{\sum_{i=1}^n \int_{Q_i^c} I_i d\Omega_i}{\sum_{i=1}^n \int_{Q_i^c} d\Omega_i}. \tag{16}$$

Now, let us consider the updating of E_0 and L . This step is quite more complicated than the previous one. In effect, a basic optimization with respect to E_0 would clearly involve some changes of its sign. In order to ensure the physical plausibility of this variable at any iteration, we have then to constrain E_0 to be positive

$$E_0 \geq 0$$

during the minimization process. Furthermore in order to increase the consistency and the efficiency of the algorithm, we also minimize our energy with respect to E_0 and L simultaneously.

The minimization with respect to (E_0, L) then consists in a constrained optimization problem. The Kuhn–Tucker conditions (Kuhn and Tucker 1951) provide the appropriate tools for dealing with this kind of problems. These tools provide necessary conditions for a solution to be optimal; we then use these conditions for finding a solution as one usually does in standard optimization. Being given the energy we are minimizing and the enforced constraint, the associated Lagrangian (Kuhn and Tucker 1951) is:

$$\begin{aligned} & \sum_{i=1}^n \int_{\pi_i(D_1)} (I_i(x_i) - \rho_1(\langle V(\pi_i^{-1}(x_i)), L) \\ & \times \xi(\pi_i^{-1}(x_i)) + E_0))^2 d\Omega_i \\ & + \sum_{i=1}^n \int_{\pi_i(D_2)} (I_i(x_i) - \rho_2(\langle V(\pi_i^{-1}(x_i)), L) \\ & \times \xi(\pi_i^{-1}(x_i)) + E_0))^2 d\Omega_i \\ & + \sum_{i=1}^n \int_{Q_i^c} (I_i(x_i) - h)^2 d\Omega_i + \alpha \int_S dA + \beta \int_C ds \\ & + \gamma \int_S (1 - \langle V(\mathbf{X}), N(\mathbf{X}) \rangle) dA + \lambda E_0 = 0 \end{aligned} \tag{17}$$

constrained by $\lambda \geq 0$. The Kuhn–Tucker (Kuhn and Tucker 1951) conditions are: If (E_0, L) is a local minimum, then there exists a non-negative constant λ s.t.

$$\left\{ \begin{aligned} (1) \quad & \sum_{i=1}^n \int_{\pi_i(D_1)} (I_i(x_i) - \rho_1(\langle V, L \rangle \xi + E_0)) \rho_1 V \xi d\Omega_i \\ & + \sum_{i=1}^n \int_{\pi_i(D_2)} (I_i(x_i) - \rho_2(\langle V, L \rangle \xi \\ & + E_0)) \rho_2 V \xi d\Omega_i = 0, \\ (2) \quad & - \sum_{i=1}^n \int_{\pi_i(D_1)} (I_i(x_i) - \rho_1(\langle V, L \rangle \xi + E_0)) \rho_1 d\Omega_i \\ & - \sum_{i=1}^n \int_{\pi_i(D_2)} (I_i(x_i) - \rho_2(\langle V, L \rangle \xi \\ & + E_0)) \rho_2 d\Omega_i + \lambda = 0, \\ (3) \quad & \lambda E_0 = 0. \end{aligned} \right. \tag{18}$$

In the equations above, (1) is derived by differentiating (17) with respect to L and (2) is derived by differentiating (17) with respect to E_0 .

We now focus on the domain of definition of the above equations. If a solution is retrieved within the domain $\{E_0 \mid E_0 > 0\}$, then $\lambda = 0$. Thus according to the first and second equations of (18), we have

$$\left\{ \begin{aligned}
 L &= \left[\rho_1^2 \sum_{i=1}^n \int_{\pi_i(D_1)} V V^T \xi d\Omega_i \right. \\
 &\quad \left. + \rho_2^2 \sum_{i=1}^n \int_{\pi_i(D_2)} V V^T \xi d\Omega_i \right]^{-1} \\
 &\quad \times \left(\rho_1 \sum_{i=1}^n \int_{\pi_i(D_1)} (I_i(x_i) - \rho_1 E_0) V \xi d\Omega_i \right. \\
 &\quad \left. + \rho_2 \sum_{i=1}^n \int_{\pi_i(D_2)} (I_i(x_i) - \rho_2 E_0) V \xi d\Omega_i \right), \quad (19) \\
 E_0 &= \left(\sum_{i=1}^n \int_{\pi_i(D_1)} (I_i(x_i) - \rho_1 \langle V, L \rangle \xi) \rho_1 d\Omega_i \right. \\
 &\quad \left. + \sum_{i=1}^n \int_{\pi_i(D_2)} (I_i(x_i) - \rho_2 \langle V, L \rangle \xi) \rho_2 d\Omega_i \right) \\
 &\quad \times 1 / \left(\sum_{i=1}^n \int_{\pi_i(D_1)} \rho_1^2 d\Omega_i + \sum_{i=1}^n \int_{\pi_i(D_2)} \rho_2^2 d\Omega_i \right),
 \end{aligned} \right.$$

where $[\cdot]^{-1}$ denotes the matrix inverse operator. In order to simplify the above equations, we let

$$\left\{ \begin{aligned}
 A &= \rho_1 \sum_{i=1}^n \int_{\pi_i(D_1)} I_i(x_i) d\Omega_i \\
 &\quad + \rho_2 \sum_{i=1}^n \int_{\pi_i(D_2)} I_i(x_i) d\Omega_i, \\
 V^* &= \rho_1^2 \sum_{i=1}^n \int_{\pi_i(D_1)} V \xi d\Omega_i + \rho_2^2 \sum_{i=1}^n \int_{\pi_i(D_2)} V \xi d\Omega_i, \\
 M &= \rho_1^2 \sum_{i=1}^n \int_{\pi_i(D_1)} V V^T \xi d\Omega_i \\
 &\quad + \rho_2^2 \sum_{i=1}^n \int_{\pi_i(D_2)} V V^T \xi d\Omega_i, \\
 p^* &= \rho_1 \sum_{i=1}^n \int_{\pi_i(D_1)} I_i(x_i) V \xi d\Omega_i \\
 &\quad + \rho_2 \sum_{i=1}^n \int_{\pi_i(D_2)} I_i(x_i) V \xi d\Omega_i, \\
 B &= \rho_1^2 \sum_{i=1}^n \int_{\pi_i(D_1)} d\Omega_i + \rho_2^2 \sum_{i=1}^n \int_{\pi_i(D_2)} d\Omega_i,
 \end{aligned} \right. \quad (20)$$

where $A \in \mathbb{R}$, $V^* \in \mathbb{R}^3$, $M \in M_{3 \times 3}(\mathbb{R})$, $p^* \in \mathbb{R}^3$ and $B \in \mathbb{R}^+$. Let us note here that, even if for all $\mathbf{X} \in S$ the matrix $V(\mathbf{X})V^T(\mathbf{X})$ is not invertible (its rank is 1), it is reasonable to assume that M is. Moreover, as soon as the surface S is visible from any camera, then $B > 0$. By using the above abbreviations, which are all known quantities in this phase of the optimization, we simplify the equations as follows:

$$\begin{cases} L = M^{-1}(p^* - E_0 V^*), \\ E_0 = (A - \langle V^*, L \rangle) / B. \end{cases} \quad (21)$$

By combining the two equations above, we have

$$L = M^{-1}(p^* - (A - \langle V^*, L \rangle) V^* / B),$$

so

$$L = \langle V^*, L \rangle (M^{-1} V^*) / B + M^{-1}(p^* - A V^* / B). \quad (22)$$

In order to simplify this last equation, we introduce two new vector variables, \mathbf{a} and \mathbf{b} :

$$\mathbf{a} = (M^{-1} V^*) / B \in \mathbb{R}^3, \quad (23)$$

and

$$\mathbf{b} = M^{-1}(p^* - A V^* / B) \in \mathbb{R}^3. \quad (24)$$

We have then

$$L = \mathbf{a} V^{*T} L + \mathbf{b}. \quad (25)$$

L is then the solution of the linear system

$$W L = \mathbf{b}, \quad (26)$$

where $W = I d_{3 \times 3} - \mathbf{a} V^{*T}$.

It is therefore a simple matter of linear algebra to retrieve L from the equation above, since all other terms are known.

E_0 can now be derived from L using (21). If it is greater than zero, then it is what we are searching for and so L and E_0 are obtained. Otherwise, we have to look for the minimum such that $E_0 = 0$. This problem is significantly easier. We can directly derive L from (21):

$$L = M^{-1} p^*. \quad (27)$$

3.3 Updating of the Auxiliary Vector Field V

Having fixed S , C , ρ_1 , ρ_2 , L , E_0 and h , we now minimize the energy with respect to V . Recall that we have to guarantee $|V| = 1$, that is to say that V has to stay on the manifold \mathbb{S}_2 (the unit sphere). In our previous conference paper (Jin et al. 2004a), we proposed to update the auxiliary field V by an iterated process. Here we propose a ‘‘closed-form’’ solution.

For any Riemannian manifold $\mathcal{M} \subset \mathbb{R}^m$ and any differentiable energy $E : \mathbb{R}^m \rightarrow \mathbb{R}$, we have the following (necessary) optimality condition: if $x \in \mathcal{M}$ is a local minima of $E|_{\mathcal{M}}$, then $\nabla E(x)$ is orthogonal to the tangent plane $T\mathcal{M}(x)$, i.e.

$$\nabla E(x) \in (T\mathcal{M}(x))^\perp.$$

Here, $\nabla E(x)$ denotes the classical gradient of $E : \mathbb{R}^m \rightarrow \mathbb{R}$ at point x defined in the framework of the differential calculus. Also, the above statement directly results from the chain

rule and the classical optimality condition in the differential calculus framework; see for example (Urakawa 1993; Abraham et al. 1993; Chefid’hotel et al. 2004).

Now, let us remark that for any point V on the unit sphere \mathbb{S}_2 , the vector V is orthogonal to the tangent plane to the sphere at point V , i.e.

$$\forall V \in \mathbb{S}_2, \quad V \in (T\mathbb{S}_2(V))^\perp.$$

In the case where $\mathcal{M} = \mathbb{S}_2$, since the dimension of $(T\mathbb{S}_2(V))^\perp$ is one, then the above optimality condition becomes:

$$\text{if } V \in \mathbb{S}_2 \text{ is a local minima of } E_{|\mathbb{S}_2}, \text{ then } \nabla E(V) \propto V. \quad (28)$$

Beyond this, the terms of E_{total} which depend on V are E_{data} and E_{coupling} . Let us stress that, even if V is indirectly regularized by the coupling term E_{coupling} (which forces V to be close to N which is smooth because of the regularity constraint on the surface S), our energy E_{total} does not directly contain a regularization term for the auxiliary vector field V . The values $V(\mathbf{X})$ of V at different points \mathbf{X} on the surface S are then decorrelated. The optimal vector field $V(\cdot)$ minimizing $E_{\text{data}} + \gamma E_{\text{coupling}}$ (and then E_{total}) is then the vector field V such that for all \mathbf{X} in S , $V(\mathbf{X})$ is given by:

$$\begin{aligned} V(\mathbf{X}) = \arg \min_{\tilde{V} \in \mathbb{S}_2} & \gamma(1 - \langle \tilde{V}, N(\mathbf{X}) \rangle) \\ & + \sum_{i=1}^n \chi_i(\mathbf{X})(I_i(\pi_i(\mathbf{X}) - \rho(\mathbf{X})\langle \tilde{V}, L\xi(\mathbf{X}) \\ & - E_0 \rangle)^2 \sigma_i(\mathbf{X}). \end{aligned} \quad (29)$$

Accordingly to (28) and (29), the vector field $V(\cdot)$ which minimizes E_{total} verifies: for all fixed $\mathbf{X} \in S$, there exists a real scalar ν (depending on \mathbf{X}) such that

$$\begin{aligned} \nu V(\mathbf{X}) = \sum_{i=1}^n & \chi_i(\mathbf{X})(I_i(\pi_i(\mathbf{X})) \\ & - \rho(\mathbf{X})\langle V(\mathbf{X}), L \rangle \xi(\mathbf{X}) + E_0) \\ & \times \rho(\mathbf{X})L\xi(\mathbf{X}) \sigma_i(\mathbf{X}) + \frac{\gamma}{2}N(\mathbf{X}). \end{aligned} \quad (30)$$

If $\xi(\mathbf{X}) = 0$, obviously the solution is $V(\mathbf{X}) = \frac{N(\mathbf{X})}{|N(\mathbf{X})|}$, else we have

$$\begin{aligned} \nu V(\mathbf{X}) = - \left(\sum_{i=1}^n \chi_i(\mathbf{X})\sigma_i(\mathbf{X}) \right) & \rho^2(\mathbf{X}) \langle V(\mathbf{X}), L \rangle L \\ & + \sum_{i=1}^n \chi_i(\mathbf{X})(I_i(\pi_i(\mathbf{X})) - \rho(\mathbf{X})E_0)\rho(\mathbf{X})L\sigma_i(\mathbf{X}) \\ & + \frac{\gamma}{2}N(\mathbf{X}). \end{aligned} \quad (31)$$

If we denote

$$H = \left(\sum_{i=1}^n \chi_i \sigma_i \right) \rho^2 L L^T \in M_{3 \times 3}(\mathbb{R}) \quad (32)$$

and

$$B = \sum_{i=1}^n \chi_i (I_i(x_i) - \rho E_0) \rho L \sigma_i + \frac{\gamma}{2} N \in \mathbb{R}^3 \quad (33)$$

(for simplicity, we have removed the \mathbf{X} in the equations; x_i denotes $\pi_i(\mathbf{X})$) then (31) becomes

$$(\nu Id_{3 \times 3} + H)V = B. \quad (34)$$

Thus V has to verify: there exists ν in \mathbb{R} such that

$$\begin{cases} |[\nu Id_{3 \times 3} + H]^{-1} B| = 1, \\ V = [\nu Id_{3 \times 3} + H]^{-1} B, \end{cases} \quad (35)$$

where $[\cdot]^{-1}$ is the matrix inverse operator (as in Sect. 3.2).

In the above equations, H and B are known. We get ν using Newton methods applied to the first equation of (35). The initial value ν_0 used in this iterative process is determined from the previous value of V (we denote \hat{V}) and other parameters of (30). More precisely, we get ν_0 by solving the following linear minimization problem

$$\begin{aligned} \nu_0 = \arg \min_{\nu} & \left| \sum_{i=1}^n \chi_i (I_i(x_i) - \rho(\langle \hat{V}, L \rangle + E_0)) L \sigma_i \rho \right. \\ & \left. + \frac{\gamma}{2} N - \nu \hat{V} \right|^2. \end{aligned} \quad (36)$$

After calculating ν , we get V from the second equation of (35). Thus the computed V fulfills the requirement of $|V| = 1$ and it minimizes E_{total} on the manifold \mathbb{S}_2 . While this is not strictly speaking a closed-form solution, it is simple and efficient to implement.

3.4 Ambiguities in Illumination

As we previously mentioned in Sect. 2.1, model (1) is not minimal and there is a multiplicative ambiguity. In particular, when minimizing the energy with respect to ρ_1, ρ_2, L, E_0 , we have that if $(\rho_1, \rho_2, E_0, L, C, S)$ is a solution, then so is $(\delta\rho_1, \delta\rho_2, E_0/\delta, L/\delta, C, S)$ with $\delta \in \mathbb{R}^+$.

Although the ambiguity makes it impossible to extract the albedo and illumination exactly, we could still recover $\rho E_0; |L|/E_0; L/|L|$, i.e., reconstruct the parameters up to a scalar transformation, or recover ρ_1/ρ_2 . When we know the maximum albedo on the object, we would still expect to generate the exact albedos and the amplitude of ambient term and point light source as well. In practice, in each iterative step, we fix E_0 to a constant.

4 Experiments

In the first set of experiments, we took 28 calibrated images of a doll figure of approximately uniform albedo standing on a table. The background is dark, and the doll is illuminated both by standard fluorescent overhead lamps and by an additional strong spotlight. Note that the actual environment only approximately satisfies the conditions for which the algorithm is designed, but this is on purpose. In the next experiment we will illustrate the behavior of the algorithm quantitatively on synthetic sequences generated to satisfy the assumptions precisely. Figure 1 shows 4 representative views that show how the light modulates the image intensity from light (front of the head) to dark (upper back). Notice also the non-trivial topology of the object, which is handled automatically by our algorithm as a side benefit of implementing the gradient flow within the level set framework (Osher and Sethian 1988).

The first test consisted in testing the stereoscopic segmentation algorithm (Yezzi and Soatto 2001) on this dataset. This corresponds to iterating the surface evolution in Sect. 3, without directional lighting, i.e. $L = 0$. The underlying assumption is that radiance, not just albedo, is constant, which is patently not the case for the dataset in question. As expected, the algorithm fails to converge to a viable estimate of the shape of the statuette (see Fig. 2). In particular, the darker parts of the legs and back are ascribed to the background.

In the second test, we introduce one light source into the reconstruction problem. Figure 3 shows the results obtained from the algorithm presented in (Jin et al. 2004a) with one punctual light source. In (Jin et al. 2004a), the positivity of



Fig. 1 Example views of the input data set consisting 28 images of a dancer statuette




Fig. 2 Final shape estimated using (Yezzi and Soatto 2001). The algorithm fails to reconstruct the doll, notably the legs and the back, because the assumption of constant radiance of the object is violated

the light source was not enforced, and the underlying image-formation model was

$$I(\pi(\mathbf{X})) = \rho(\mathbf{X}) \left(\sum_i \lambda_i \langle N(\mathbf{X}), L_i \rangle \xi_i(\mathbf{X}) + E_0 \right),$$

where the λ_i and E_0 could be negative. The signs of these variables could change during the optimization process (see Sect. 4.3 of (Jin et al. 2004a)) and the final sign of the lights (returned after the convergence) depends on the initialization. Figure 3 shows two typical examples of reconstruction obtained from the algorithm of (Jin et al. 2004a); according to the initialization of the position of the light source, the light converged towards a positive (left image) or negative sign (right image). In either case the results improve stereoscopic segmentation (Yezzi and Soatto 2001), where light is not explicitly modeled. In order to improve the model (Jin et al. 2004a), one needs two light sources; by using a positive and a negative light, one can obtain reasonable results such as those displayed in Fig. 4. However, since the algorithm in (Jin et al. 2004a) does not impose positivity, the resulting lighting configuration is physically impossible.

Next we tested the full-fledged algorithm described in the previous sections. Contrary to (Jin et al. 2004a), here we




Fig. 3 Shape estimated by using the algorithm (Jin et al. 2004a) with either a positive (left) or a negative (right) directional light. Compared with Fig. 2, we obtained improved reconstructions in both light configurations over that of (Yezzi and Soatto 2001). However, the results are still not satisfactory




Fig. 4 Final shape estimated using the algorithm (Jin et al. 2004a) with a combination of one positive light and one negative light. The algorithm reconstructs the 3D object more accurately (compared with Fig. 3), nevertheless the lighting configuration if not physically plausible

Table 1 The results of the reconstruction results on synthesized data: a ball with radius 10 illuminated by a point light source and ambient light

Reconstruction result	Using positive light	Ground truth
L	(0.000, 0.001, 0.999)	(0.000, 0.000, 1.000)
E_0 Error	1.6%	100.00
Error of Intensity of point light	1.5%	100.00
Shape Error	9%	0 (Note: Radius of ball is 10)

Fig. 5 Visualization of the auxiliary vector field V estimated in the same time as the surface displayed in Fig. 4 via the algorithm (Jin et al. 2004a) (with the combination of one positive light and one negative light)





Fig. 6 Final shape estimated via the proposed algorithm containing the various optimization improvements (in particular the closed-form estimation of the auxiliary vector field) and the addition of the positivity constraint. The improvement on 3D shape details, especially around the hands and lower back, is visible




Fig. 7 Visualization of the auxiliary vector field via the proposed algorithm when we add the positivity constraint for the light and when we estimate the auxiliary vector field in closed-form

enforce the positivity of the light sources. Figure 6 shows the results obtained with a single light source; the improvement is evident from Fig. 3, including the version with two light sources illustrated in Fig. 4. The improvement is particularly clear on details of the surface reconstruction, especially on the hands and lower back (Fig. 7, compare with Fig. 5).

In order to gain a more quantitative understanding of the operation of the algorithm, we tested it on various synthetic sequences, one of which we report in Fig. 8 as an example. It




Fig. 8 Three views of the synthesized constant albedo sphere illuminated by a point light source of intensity $100 \times$ located at $(0, 0, 1)$, and an ambient light of intensity 100




Fig. 9 Final shape estimated by the algorithm. The ground truth shape is shown on the left

is a sphere with constant albedo illuminated by an ambient light and a point light source. The results of the algorithm are shown in Fig. 9, and Table 1 summarizes the numerical results.

In a final experiment, we tested the algorithm on a dataset that contained piecewise constant albedo, shown in Fig. 10. In order to avoid the scale ambiguity, we rescale ρ_1, ρ_2, E_0, L in each evolution to $\delta = E_0/100$; $E_0 = E_0/\delta$; $|L| = |L|/\delta$; $\rho_1 = \rho_1\delta$; $\rho_2 = \rho_2\delta$. Here 100 is just an arbitrary scale. Figure 11 shows the reconstructed shape and partitioning curve for the synthesized painted ball illuminated by directional light. Figure 12 shows the reconstructed shape for the Fish model. We also show the evolution process of the shape and partitioning curve in Fig. 13.




Fig. 10 The upper two images are 2 out of 26 views of a synthetic scene. The scene consists of one sphere painted with a word 'CV' in grey and the rest is white. The sphere is illuminated by ambient light of intensity 100 and a point light source of $100 \times (0, 0.717, 0.717)^T$; The lower two images are 2 out of 30 views of a real scene of a toy fish. Each image size is 390×400 and calibrated with a rig. The fish is green with white stripes. It is illuminated by a point light source from the head of the fish




Fig. 11 The reconstructed shape and partitioning curve between regions with different albedo for the synthesized data in Fig. 10. The two images show the final shape result of the proposed algorithm together with the red segmenting curve. The red straight lines show the coordinate frame for the spheres




Fig. 12 Recovered shape via the algorithm proposed for the real scene shown in Fig. 10. The red curves represents the segmenting curves for different albedo

5 Conclusions

We proposed an image formation model for Lambertian scenes with piecewise constant albedos illuminated by ambient and point light sources. We designed an algorithm to simultaneously infer all model parameters from a collec-




Fig. 13 The evolution process of shape and the segmenting curve for different albedo. The initial surface is an ellipsoid

tion of calibrated images by minimizing a single energy. We solved the resulting infinite-dimensional optimization problem by numerically integrating partial differential equations that converge in steady states to a local minimum of the first variation of the functional. Our work extends the prior art on Stereoscopic Segmentation (Yezzi and Soatto 2001) by explicitly modeling and separating reflectance from illumination. It does so while enforcing physical constraints, such as the positivity of the energy distribution of light sources.

We also addressed a structural problem related to shape estimation under the standard Lambertian shading model, that results in unstable flows when the standard cost functional is minimized in a naive way. We introduced a smooth auxiliary vector field, that is coupled to the normal field via an energy term, and show that this model is not prone to instability like the naive one.

Naturally, the model we proposed is limited by the assumptions of Lambertian reflection and point/diffuse illumination. The assumption of piecewise constant albedo is not very restrictive in theory, as many man-made objects can be well approximated by this model, but in practice when objects become complex the topology and geometry of the albedo boundaries become so difficult that our algorithm may fail to capture subtle variations. One of the potential advantages of our approach is its potential to be integrated with multi-view stereo algorithms, by providing additional constraints in cases where objects with simple albedo do not provide enough constraints to establish point-to-point correspondence and therefore local feature-based approaches fail to provide a dense reconstruction.

Acknowledgements This work is supported by in part by Adobe Systems Incorporated, ONR N00014-03-1-0850:P0001, AFOSR E-16-V91-G2, and German Research Foundation DFG CR 250/1-1.

References

- Abraham, R., Marsden, J. E., & Ratiu, T. (1993). *Applied Mathematical Sciences: Vol. 75. Manifolds, tensor analysis, and applications* (2nd ed.). New York: Springer.
- Belhumeur, P., Kriegman, D., & Yuille, A. L. (1999). The generalized bas relief ambiguity. *International Journal of Computer Vision*, 35, 33–44.
- Brox, T., & Weickert, J. (2006). Level set segmentation with multiple regions. *IEEE Transactions on Image Processing*, 15(10).
- Chefd'hotel, C., Tschumperlé, D., Deriche, R., & Faugeras, O. (2004). Regularizing flows for constrained matrix-valued images. *Journal of Mathematical Imaging and Vision*, 20(1–2), 147–162.

- Chen, H. F., Belhumeur, P. N., & Jacobs, D. W. (2000). In search of illumination invariants. In *Proceedings of the IEEE conference on computer vision and pattern recognition*
- Durou, J.-D., Falcone, M., & Sagona, M. (2004). *A survey of numerical methods for shape from shading*. Research report 2004-2-R, IRIT, January 2004.
- Faugeras, O., & Keriven, R. (1998). Variational principles, surface evolution, pdes, level set methods, and the stereo problem. *IEEE Transactions on Image Processing*, 7(3), 336–344.
- Horn, B. K. P. (1986). *Robot vision*. Cambridge: MIT.
- Horn, B., & Brooks, M. (Eds.). (1989). *Shape from shading*. Cambridge: MIT.
- Jin, H. (2003). *Variational methods for shape reconstruction in computer vision*. PhD thesis, Electrical Engineering Department, Washington University, August 2003.
- Jin, H., Soatto, S., & Yezzi, A. J. (2003a). Multi-view stereo beyond Lambert. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (Vol. I, pp. 171–178), June 2003.
- Jin, H., Yezzi, A. J., Tsai, Y.-H., Cheng, L.-T., & Soatto, S. (2003b). Estimation of 3D surface shape and smooth radiance from 2D images: a level set approach. *Journal of Scientific Computing*, 19(1–3), 267–292.
- Jin, H., Cremers, D., Yezzi, A., & Soatto, S. (2004a). Shedding light on stereoscopic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 36–42).
- Jin, H., Yezzi, A. J., & Soatto, S. (2004b). Region-based segmentation on evolving surfaces with application to 3d shape and radiance estimation. In *Proceedings of the European conference on computer vision* (pp. 114–125), May 2004.
- Jin, H., Yezzi, A., & Soatto, S. (2007, in press). Mumford-shah on the move: region-based segmentation on deforming manifolds with application to 3-D reconstruction of shape and appearance from multi-view images. *Journal of Mathematical Imaging and Vision*.
- Keriven, R., Pons, J.-P., & Faugeras, O. (2006). Multi-view stereo reconstruction and scene flow estimation with a global image-based matching score. To appear in the international journal of computer vision. *International Journal of Computer Vision*
- Klette, R., Kozera, R., & Schläins, K. (1998). *Shape from shading and photometric stereo methods*. Technical Report CITR-TR-20, University of Auckland, New Zealand.
- Koenderink, J., & van Doorn, A. (1980). Photometric invariants related to solid shape. *Optica Acta*, 27(7), 981–996.
- Kolev, K., Brox, T., & Cremers, D. (2006). Robust variational segmentation of 3D objects from multiple views. In K. Franke et al. (Eds.), *Lecture notes in computer science: Vol. 4174. Pattern recognition (proceedings DAGM)* (pp. 688–697). Berlin: Springer.
- Kuhn, H. W., & Tucker, A. W. (1951). Nonlinear programming. In J. Neyman (Ed.), *Proceedings of the second Berkeley symposium on mathematical statistics and probability* (pp. 481–492) Berkeley: University of California Press.
- Kutulakos, K. N., & Seitz, S. M. (2000). A theory of shape by space carving. *International Journal of Computer Vision*, 38(3), 199–218.
- Langer, M. S., & Zucker, S. W. (1994). Shape from shading on a cloudy day. *Journal of Optical Society of America*, 11, 467–478.
- Ma, Y., Soatto, S., Kosecka, J., & Sastry, S. (2003). *An invitation to 3D vision, from images to models*. Berlin: Springer.
- Nayar, S., Ikeuchi, K., & Kanade, T. (1991). Surface reflection: physical and geometrical perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7), 611–634.
- Oliensis, J., & Dupuis, P. (1993). A global algorithm for shape from shading. In *Proceedings of the international conference on computer vision* (pp. 692–710).
- Osher, S., & Sethian, J. (1988). Fronts propagating with curvature-dependent speed: algorithms based on Hamilton–Jacobi equations. *Journal of Computational Physics*, 79, 12–49.
- Prados, E. (2004). Application of the theory of the viscosity solutions to the shape from shading problem. PhD thesis, Univ. of Nice-Sophia Antipolis.
- Robles-Kelly, A., & Hancock, E. R. (2004). Estimating the surface radiance function from single images. In *3DPVT* (pp. 494–501).
- Robles-Kelly, A., & Hancock, E. R. (2005). Surface radiance correction for shape-from-shading. *Pattern Recognition*, 38(10), 1574–1595.
- Samaras, D. (2003). Illumination constraints in deformable models for shape and light direction estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(2), 247–264.
- Samaras, D., Metaxas, D., Fua, P. V., & Leclerc, Y. G. (2000). Variable albedo surface reconstruction from stereo and shape from shading. In *Proceedings of the IEEE international conference on computer vision and pattern recognition* (Vol. I, pp. 480–487).
- Seitz, S., Curless, B., Diebel, J., Scharstein, D., & Szeliski, R. (2006). A comparison and evaluation of multi-view stereo reconstruction algorithms. In *Proceedings of the IEEE international conference on computer vision and pattern recognition* (pp. 519–526).
- Stewart, A. J., & Langer, M. S. (1997). Towards accurate recovery of shape from shading under diffuse lighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(9), 1020–1025.
- Urakawa, H. (1993). *Calculus of variations and harmonic maps*. Providence: American Mathematical Society.
- Vedaldi, A., & Soatto, S. (2006). *Viewpoint invariance for non-planar scenes*. Technical Report TR050012, UCLA CSD.
- Vese, L. A., & Chan, T. F. (2002). A multiphase level set framework for image processing using the Mumford–Shah functional. *International Journal of Computer Vision*, 50(3), 271–293.
- Yezzi, A., & Soatto, S. (2001). Stereoscopic segmentation. In *Proceedings of the international conference on computer vision* (pp. 59–66).
- Yuille, A. L., Snow, D., Epstein, R., & Belhumeur, P. (1999). Determining generative models of objects under varying illumination: shape and albedo from multiple images using svd and integrability. *International Journal of Computer Vision*, 35, 203–222.
- Yuille, A., Coughlan, J. M., & Konishi, S. (2003). Kgbr viewpoint-lighting ambiguity. *Journal of the Optical Society of America A*, 20(1), 24–31.
- Zhang, R., Tsai, P.-S., Cryer, J., & Shah, M. (1999). Shape from shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8), 690–706.
- Zheng, Q., & Chellappa, R. (1991). Estimation of illuminant direction, albedo and shape from shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7), 680–702.