(In Proceedings CVPR'98 pp. 322-329)

Incorporating Illumination Constraints in Deformable Models

Dimitrios Samaras and Dimitris Metaxas
Vision, Analysis and Simulation Technologies Laboratory
Department of Computer & Information Science
University of Pennsylvania
Philadelphia PA 19104-6389
e-mail: samaras@gradient.cis.upenn.edu, dnm@central.cis.upenn.edu

Abstract

We present a method for the integration of illumination constraints within a deformable model framework. These constraints are incorporated as nonlinear holonomic constraints in the Lagrange equations of motion governing the deformation of the model. For improved numerical performance we employ the Baumgarte stabilization method. Our methodology is general and can be used for a broad range of illumination constraints. This approach avoids commonly used approximations in shape from shading, such as linearization, and the use of partial differential equations, which require initial boundary conditions. Furthermore, global and local parameterizations of the deformable models allow an improved estimation of shape from shading. We demonstrate this improvement over previously used approaches through a series of experiments on standardized sets of real and synthetic data [30].

1 Introduction

The integration of visual cues within a physics-based modeling framework has been attempted recently by several researchers [19, 9, 6], due to its potential in improved shape estimation. However, previous attempts have not addressed the use of illumination constraints, such as the ones used in shape from shading, as hard constraints. A major difficulty is the nonlinear nature of the constraints, which leads to a requirement for numerically robust methods. In this paper, we provide a general methodology for the incorporation of illumination constraints within a deformable model framework, and we demonstrate through experiments the improvement in object shape estimation from images.

We first develop a theory for the numerically robust integration of nonlinear holonomic constraints within a deformable model framework, regardless of their origin. We use Lagrange multipliers and a Baumgarte stabilization approach [1]. Subsequently, we demonstrate how any type of illumination constraint can be incorporated into a principled way in our physics-based modeling framework, provided that the illumination law is a differentiable function of the normal or tangents to the surface. This encompasses practically all the illumination models used in computer vision, from the simple Lambertian model to more complex highly nonlinear models [21, 23]. Instead of extracting the shape parameters directly from these illumination constraints (which is not always possible), we use them to provide the necessary generalized forces that will deform our model and estimate the object's 3D shape. Our methodology obviates the need for commonly used approximations (e.g., linearization) to these equations, or the need to solve partial differential equations requiring boundary conditions [4].

We use deformable models or grids with both global and local deformations [19, 20]. During shape estimation, we first fit the model's global parameters given the illumination constraints, and then we refine its shape based on the model's local deformations, using a coarse to fine grid. Use of a deformable model-based approach offers shape flexibility and the additional advantage of the numerically robust computation of the necessary derivatives, producing improved shape estimation results. This is demonstrated in a series of experiments with real and synthetic data. The experiments consist of the standard test images used in the thorough comparative survey by Zhang et al.[30].

2 Previous work

Most of the earlier work on shape from shading (SFS) has been compiled in [13]. The first comprehensive comparative study of a number of SFS algorithms is [30]. They classify SFS algorithms as either global or local, depending on whether they use intensity information across the entire image or only in the local neighborhood of each pixel.

Most SFS algorithms for Lambertian surfaces follow a regularization approach using the image brightness and surface smoothing constraints [4, 14, 27, 17, 15]. Some methods make use of the integrability constraint [8, 12], the intensity gradient constraint [31] and the unit normal constraint [4], while [18] provides a numerically efficient solution to regularization. In the above class of approaches, [15, 10] combine stereo and shading. These methods require good initial depth values, which they obtain from stereo information, and typically result in better estimation results compared to shape from shading alone.

Another class of algorithms [11, 22, 3] propagate height information from known points in the image. These methods require a priori information, and their performance depends on the accuracy of that information. Finally, there are a number of methods that use local information to reconstruct depth. [24] and [16] use a rather restrictive local spherical approximation, while [25] and [29] use linear approximations of the reflectance function.

Most of the work on SFS has utilized the Lambertian reflectance model. [21] provides a hybrid model combining diffuse and specular reflectance. Oren and Nayar [23] studied rough diffuse surfaces that do not follow Lambert's law, and they proposed a much more complex model of diffuse reflection, based on statistical modeling of surface geometry.

As opposed to the previous methods on SFS, we present a model-based approach, where shape from shading is incorporated as a nonlinear hard constraint within a deformable model framework.

3 Deformable Models

In this section we briefly review the general formulation of deformable models [19, 20].

Geometrically, the models used in this paper are parameterized surfaces in space whose intrinsic parameters are $\mathbf{u} = (u, v)$, defined on a domain Ω. The positions $\mathbf{x}(\mathbf{u}, t)$ of points on the model relative to an inertial frame of reference Φ in space are given by

$$\mathbf{x} = \mathbf{c} + \mathbf{R}\mathbf{p}, \qquad (1)$$

where $\mathbf{c}(t)$ is the origin of the model frame, ϕ, and $\mathbf{R}(t)$ is the rotation matrix expressing the orientation of ϕ. $\mathbf{p}(\mathbf{u}, t)$ denotes the positions of points on the model relative to the model frame. We introduce global and local deformations, by further expressing $\mathbf{p}$ as the sum of a reference shape $\mathbf{s}(\mathbf{u}, t)$ and a displacement function $\mathbf{d}(\mathbf{u}, t)$, i.e., $\mathbf{p} = \mathbf{s} + \mathbf{d}$.

For the applications in this paper, we define $\mathbf{s}$ as a geometric primitive parameterized in uv space (in some of the examples the primitive is a superquadric, in others it will only be a grid) that can undergo scaling, tapering and bending deformations (see [19, 20] for details and formulas). We then collect the global deformation parameters that we use to define $\mathbf{s}$ into the vector $\mathbf{q}_s$.

Local displacements $\mathbf{d}$ are computed based on the use of triangular finite elements [19, 20], which provide a tessellation of the deformable model's surface. Associated with every finite element node i is a nodal vector variable $\mathbf{q}_{d,i}$. We collect all the nodal variables into a vector of local degrees of freedom $\mathbf{q}_d = (\ldots, \mathbf{q}_{d,i}^\top, \ldots)^\top$, and we compute the local displacement $\mathbf{d}$ based on the finite element theory as $\mathbf{d} = \mathbf{S}\mathbf{q}_d$. $\mathbf{S}$ is the shape matrix whose entries are the finite element shape functions.

The velocity of points on the model is given by [20],

$$\dot{\mathbf{x}} = \mathbf{L}\dot{\mathbf{q}}, \qquad (2)$$

where $\mathbf{q} = (\mathbf{q}_c^\top, \mathbf{q}_\theta^\top, \mathbf{q}_s^\top, \mathbf{q}_d^\top)^\top$ are the model parameters with $\mathbf{q}_c = \mathbf{c}$ and $\mathbf{q}_\theta$ the model's rotational degrees of freedom expressed as a quaternion, and $\mathbf{L}$ is a Jacobian matrix.

Our goal when fitting the model to visual data is to recover the vector $\mathbf{q}$ which expresses the model's degrees of freedom. The model fitting procedure is done in a physics-based way by enabling the image data to apply traction forces to the surface of the model [19]. Based on Lagrangian dynamics we make our model dynamic in $\mathbf{q}$, and we arrive at the following first order set of motion equations:

$$\dot{\mathbf{q}} + \mathbf{K}\mathbf{q} = \mathbf{f}_q, \qquad (3)$$

where $\mathbf{K}$ is the stiffness matrix, (see [20] for definitions), and where $\mathbf{f}_q(\mathbf{u}, t) = \int \mathbf{L}^\top \mathbf{f}$ are the generalized external forces computed from the 3D forces $\mathbf{f}$ that the data exert on the model. In our model these forces will be computed from edges and the brightness constraint.

In the following sections we will use the above basic deformable model framework to incorporate the nonlinear brightness constraint as a holonomic constraint in (3).

4 Illumination Models

We will now present the lighting models we will use in SFS and we will reformulate them so that they can be introduced as a constraint in a deformable model framework.

4.1 Lambertian Model

This is the simplest lighting model; it models an ideal diffuser. The amount of light reflected from the surface that reaches the viewer depends only on the orientation of the surface relative to the light source. If we assume a point light source at infinity, and orthographic projection along an axis parallel to the optical axis of the observer, the scene radiance is expressed as

$$I_L = B\frac{\rho}{\pi}\cos\theta = B\rho'\ \hat{\mathbf{s}}\cdot\hat{\mathbf{n}}, \tag{4}$$

where θ is the angle between the surface unit normal vector, $\hat{\mathbf{n}} = \mathbf{n}/\|\mathbf{n}\|$, and the unit light source direction vector $\hat{\mathbf{s}}$. B is the strength of the light source and ρ is the constant albedo of the surface, which represents the fraction of the incident light to the surface that is reflected. $\rho' = \rho/\pi$ is the bidirectional reflection distribution function, which is a constant in this model.

The above image irradiance equation (4) can be rewritten as a brightness constraint as follows:

$$C = \hat{\mathbf{s}}\cdot\mathbf{n} - I_L'\|\mathbf{n}\| = 0, \tag{5}$$

where $I_L' = I_L/(B\rho')$, and the values of I_L' range between 0 and 1.. In our application, we want to recover the shape parameters based on the intensity information at n points and therefore we will have an n-component constraint vector $\mathbf{C}$.

$\mathbf{C}$ is a nonlinear constraint with respect to the model parameters, which in the traditional SFS formulations resulted either in nonlinear first-order differential equations (PDE's) or in attempts to linearize the constraint [25, 17, 29]. PDE's require appropriate boundary conditions [4] that are often not available, whereas linearization introduces additional error. We avoid both of these limitations by incorporating the above brightness constraint as a nonlinear holonomic constraint in the deformable model framework, which we will present in a following section. This incorporation depends on the computation of the Jacobian $\mathbf{C_q}$ of the constraint vector $\mathbf{C}$, w.r.t. the model parameters $\mathbf{q}$.

$$\mathbf{C_q} = \hat{\mathbf{S}}\mathbf{N_q} - \mathbf{L_{N_q}}\mathbf{I_L'} - \mathbf{J}_{I_L'}\mathbf{L}\|\mathbf{N}\|. \tag{6}$$

(6) is derived by differentiating each constraint given by (5). $\mathbf{N}$, $\hat{\mathbf{S}}$, $\mathbf{I_L'}$ and $\|\mathbf{N}\|$ are formed by collecting the values of $\mathbf{n}$, $\hat{\mathbf{S}}$, I_L', $\|\mathbf{n}\|$, respectively, for each component of $\mathbf{C}$. $\mathbf{N_q} = \partial\mathbf{N}/\partial\mathbf{q}$, $\mathbf{L_{N_q}} = \partial\|\mathbf{N}\|/\partial\mathbf{q}$. The sampled image values $I_L'(x,y)$ are parameterized by x and y, the coordinates of the sampled model points with positions $\mathbf{x} = (x,y,z)^\top$. We utilize the chain rule to calculate the derivatives of $I_L'(x,y)$, w.r.t $\mathbf{q}$.

$$\frac{\partial\mathbf{I_L'}(x,y)}{\partial\mathbf{q}} = \frac{\partial\mathbf{I_L'}(x,y)}{\partial\mathbf{x}}\frac{\partial\mathbf{x}(x,y,z)}{\partial\mathbf{q}}, \tag{7}$$

so, in (6) $\mathbf{J}_{I_L'} = \partial\mathbf{I_L'}(x,y)/\partial\mathbf{x}$ and $\mathbf{L} = \partial\mathbf{x}(x,y,z)/\partial\mathbf{q}$ is the Jacobian of the model points, described in (2).

(6) requires the computation of the surface normal and its Jacobian. Since we have a parametric representation for our deformable model $\mathbf{s}(u,v) = [x(u,v), y(u,v), z(u,v)]^\top$, the surface tangents are

$$\frac{\partial\mathbf{s}}{\partial u} = \left[\frac{\partial x}{\partial u}, \frac{\partial y}{\partial u}, \frac{\partial z}{\partial u}\right]^\top, \frac{\partial\mathbf{s}}{\partial v} = \left[\frac{\partial x}{\partial v}, \frac{\partial y}{\partial v}, \frac{\partial z}{\partial v}\right]^\top$$

and the surface normal is computed from

$$\mathbf{n} = \frac{\partial\mathbf{s}}{\partial v}\times\frac{\partial\mathbf{s}}{\partial u}. \tag{8}$$

The fact that our approach is model-based offers greater numerical stability with respect to previous non-model-based methods in the computation of the above derivatives, which are very sensitive to noise in the data. This will lead to improved experimental results as will be demonstrated in the experiment section.

4.2 More Accurate Illumination Models

Our approach allows us to use alternative formulations for the brightness constraint, even if they are not invertible w.r.t the normal, as long as the normal, or tangents to the surface, appear in the expression that describes the radiance of the surface.

As an example, a highly complex lighting model that can be accommodated in our formulation is the one proposed by [23], to model more accurately the reflectance of rough diffuse surfaces, such as sand and plaster. In that model, the diffuse radiance of the surface is dependent on both the viewing and illumination directions. The viewing direction at each point is defined with polar and azimuthal angles ϕ_r and θ_r w.r.t to the normal and the illumination direction with ϕ_i and θ_i. The radiance of the surface is then expressed as $L_d = L_1 + L_2$, where the approximations for these components are

$$\begin{aligned} L_1 \;=\; & \frac{\rho}{\pi}E_0\cos\theta_i[C_1(\sigma) \\ & + \cos(\phi_r - \phi_i)C_2(\alpha;\beta;\phi_r - \phi_i;\sigma)\tan\beta \\ & + (1 - |\cos\phi_r - \phi_i|)C_3(\alpha;\beta;\sigma)\tan(\frac{\alpha+\beta}{2})] \end{aligned} \tag{9}$$

$$L_2 = \frac{0.17\rho^2\sigma^2}{\pi(\sigma^2 + 0.13)}E_0\cos\theta_i[1 - \cos(\phi_r - \phi_i)(\frac{2\beta}{\pi})^2], \tag{10}$$

where ρ is the albedo, $E_0\cos\theta_i$ is the irradiance, $\alpha = \max[\theta_r, \theta_i]$ and $\beta = \min[\theta_r, \theta_i]$. The coefficients of the model are $C_1 = 1 - 0.5\frac{\sigma^2}{\sigma^2+0.33}$, $C_2 = 0.45\frac{\sigma^2}{\sigma^2+0.09}\sin\alpha$, if $\cos(\phi_r - \phi_i) \geq 0$, and $C_2 = 0.45\frac{\sigma^2}{\sigma^2+0.09}(\sin\alpha - (\frac{2\beta}{\pi})^3)$ otherwise, and $C_3 = 0.125\frac{\sigma^2}{\sigma^2+0.09}(\frac{4\alpha\beta}{\pi^2})^2$, where σ is the surface roughness parameter.

Despite the complexity of the model, it can still be differentiated w.r.t. the surface normal, since the angles that parameterize it can be converted to equivalent dot product expressions of the normal, that is, if $\hat{\mathbf{s}}$ is the light source direction, $\hat{\mathbf{v}}$ the viewing direction and $\hat{\mathbf{n}}$ the surface normal, $\cos\theta_i = \hat{\mathbf{s}} \cdot \hat{\mathbf{n}}$, $\cos\theta_r = \hat{\mathbf{v}} \cdot \hat{\mathbf{n}}$ and $\cos(\phi_r - \phi_i) = (\hat{\mathbf{s}} - \hat{\mathbf{n}}\cos\theta_i) \cdot (\hat{\mathbf{v}} - \hat{\mathbf{n}}\cos\theta_r)$. Thus, if I_{im} is the measured image value, we can form the constraint

$$C_D = L_d - I_{im} \tag{11}$$

and use it in the same way as in the case of a Lambertian model. Two of the examples in the experiment section were fitted using this model.

4.3 Integration of Lighting Constraints in Deformable Models

Here, we present our method for incorporating the above nonlinear lighting constraints, $\mathbf{C}$, in a deformable model framework. This constraint vector, of the form $\dot{\mathbf{C}}(\mathbf{q}, t) = 0$, is nonlinearly related to the model parameters $\mathbf{q}$; constraints of this type are termed *nonlinear holonomic constraints*. We incorporate these constraints in a deformable model formulation, using the method of Lagrange multipliers. The case of linear non-holonomic constraints was treated in [6], and the case of linear holonomic constraints was treated in [20]. In order to recover the shape parameters based on the constraint information at n points of our model, we will have an $n \times 1$ constraint vector $\mathbf{C}$, which is incorporated in (3) based on the theory of Lagrange multipliers as follows:

$$\dot{\mathbf{q}} = \mathbf{f} - \mathbf{K}_d\mathbf{q} - \mathbf{C}_{\mathbf{q}}^{\top}\boldsymbol{\lambda}, \tag{12}$$

where $\mathbf{f}$ are edge based forces, $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \ldots, \lambda_n]^{\top}$ are the Lagrange multipliers and $\mathbf{C}_{\mathbf{q}}$ is the Jacobian matrix of the constraints $\mathbf{C}$ w.r.t. the shape parameters. We can consider $-\mathbf{C}_{\mathbf{q}}^{\top}\boldsymbol{\lambda}$ to be the vector of generalized forces on the model parameters due to the lighting constraint.

Since (12) has fewer equations than unknowns, to obtain the additional necessary equations we differentiate the constraint equation w.r.t. time $\dot{\mathbf{C}}(\mathbf{q}, t) = 0$ yielding $\mathbf{C}_{\mathbf{q}}\dot{\mathbf{q}} + \mathbf{C}_t = 0$. In our application $\mathbf{C}$ is not directly dependent on time, therefore $\mathbf{C}_t = 0$, which results in

$$\dot{\mathbf{C}}(\mathbf{q}, t) = \mathbf{C}_{\mathbf{q}}\dot{\mathbf{q}} = 0. \tag{13}$$

The above Lagrange multiplier approach works well in practice when the constraints are almost satisfied initially. However, since we will be fitting a deformable model to the data, these constraints will be far from being satisfied initially, and therefore we will use Baumgarte's [1] stabilization method, taking into account that our constraint is nonlinear. This approach is a generalization of the previously developed methodology for linear holonomic constraints [19]. Based on this method we replace $\dot{\mathbf{C}}(\mathbf{q}, t) = 0$ with the following constraint equation

$$\dot{\mathbf{C}} + \alpha\mathbf{C} = 0, \tag{14}$$

where α is a stabilization factor. Based on (13), (14) becomes

$$\mathbf{C}_{\mathbf{q}}\dot{\mathbf{q}} + \alpha\mathbf{C} = 0. \tag{15}$$

Using (12), (15) becomes

$$\mathbf{C}_{\mathbf{q}}(\mathbf{f} - \mathbf{K}_d\mathbf{q} - \mathbf{C}_{\mathbf{q}}^{\top}\boldsymbol{\lambda}) + \alpha\mathbf{C} = 0, \tag{16}$$

which we can solve for $\boldsymbol{\lambda}$ and obtain

$$\boldsymbol{\lambda} = (\mathbf{C}_{\mathbf{q}}\mathbf{C}_{\mathbf{q}}^{\top})^{-1}(\alpha\mathbf{C} + \mathbf{C}_{\mathbf{q}}(\mathbf{f} - \mathbf{K}_d\mathbf{q})). \tag{17}$$

However, we do not explicitly compute the Lagrange multipliers. Instead, we substitute (17) into (12). If we define the vector of all non-constraint (generalized) forces, as $\mathbf{b} = \mathbf{f} - \mathbf{K}_d\mathbf{q}$ then (12) becomes

$$\dot{\mathbf{q}} = \mathbf{b} - \mathbf{C}_{\mathbf{q}}^{\top}(\mathbf{C}_{\mathbf{q}}\mathbf{C}_{\mathbf{q}}^{\top})^{-1}(\alpha\mathbf{C} + \mathbf{C}_{\mathbf{q}}\mathbf{b}). \tag{18}$$

The matrix $\mathbf{C}_{\mathbf{q}}^{+} = \mathbf{C}_{\mathbf{q}}^{\top}(\mathbf{C}_{\mathbf{q}}\mathbf{C}_{\mathbf{q}}^{\top})^{-1}$ is the pseudo-inverse of the matrix $\mathbf{C}_{\mathbf{q}}$, and using this notation in the previous equation we get

$$\dot{\mathbf{q}} = \mathbf{b} - \mathbf{C}_{\mathbf{q}}^{+}(\alpha\mathbf{C} + \mathbf{C}_{\mathbf{q}}\mathbf{b}). \tag{19}$$

In the experiment section we will use (19) with the various lighting constraints presented earlier.

5 Implementation Details

In our implementation, we fit dynamically a deformable surface model of the type described in the previous section. In the model fitting procedure we first estimate the model's global degrees of freedom to obtain a first approximation of the underlying surface. In order to estimate the details of the model's surface we use local deformations. We discretize our model using triangular finite elements with C^0 continuity, as described in [20]. For efficiency, we use data points that coincide with the barycenter of each element. The generation of the finite element grid is done in a coarse to fine fashion. We start with a few elements and we progressively subdivide them in a regular fashion to get finer meshes until the error of fit does not change. This results in significantly faster convergence. In addition, as our model moves closer to the solution, we lower the model's elasticity by decreasing the relative coefficients, in order to refine surface details.

Fitting with global deformations converges faster than fitting with local ones, especially when the global parameters capture a lot of the model's detail, e.g. in the case of models built for specific applications, when there is knowledge of the illuminated object. While achieving interactive rate performance, global deformations provide only gross shape estimation. To avoid loss of detail, when the error rate stops decreasing fitting continues with local deformations, whose rate of convergence is slower (several minutes) and depends on the desired accuracy of approximation and the complexity of the surface. Furthermore the use of global model parameters can help abstract geometric primitives, which could prove useful for the classification of shading patterns. These applications are beyond the scope of this paper.

Finally, the pseudo-inverse of $\mathbf{C_q}$ in (19) is computed through singular value decomposition [26]. Since the number of sampled points is in general far greater than the number of global parameters, the system is over-constrained and our solution is equivalent to a least squares error solution.

6 Results/Evaluation

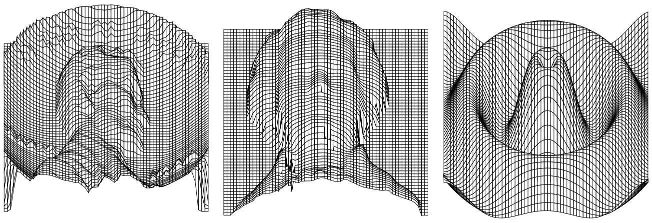

(a) Penny (b) Mozart (c) Sombrero

Figure 1: Original depth maps used to produce synthetic images

The test set consists of five synthetic data sets and five real images. For the synthetic data sets, we have three images for each data set. The images are generated with the following light source vectors: In the first image of each set the direction of the light source is (0,0,1), in the second it is (1,0,1), and in the third it is (5,5,7). We performed experiments on three of the synthetic images[1] that are included in the review of SFS algorithms in [30], so that we could compare the performance of our algorithm with some of the standard algorithms in the field.

In Fig.1 we have the true height maps used to generate the synthetic test images. To the right of each image in Fig.2 (illuminated under the light vectors described above) are the reconstructed height maps. The Penny and Sombrero images were fitted with a hemispherical grid, with global scaling along the z-axis. The Mozart images were fitted with a deformable superquadric that was initially fitted around the contours of the object, using edge forces. This image was suitable for this approach, as the edges between the object and the background are very clear.

In Fig.3 we have the reconstructed height maps of the three real images. The real images that were tested were: Lena, with estimated light source (1.5,0.866,1), Pepper, with estimated light source (-0.707,0.642,1), David, with estimated light source (-0.707,0.707,1). In Lena and Pepper we have albedo discontinuities and self-occlusions, which violate the assumptions of the algorithm. The estimated light sources for those images are the ones given in [30]. We also provide three renderings of each recovered shape; the first under the original light source, the second under a light source with tilt orthogonal to the original and the third under a light source with opposite tilt. The fact that renderings under the light source used for the reconstruction are so close to the original image highlights the advantage of imposing illumination constraints as hard constraints that have to be satisfied exactly.

In all the experiments, a time step of 0.01 is the Euler integration step used in the iterative fitting of the model, and the magnitude of the constraint forces acting on the model is weighted by a factor of 15. We discretize the model to 32×32 nodes for the global parameter fitting and the first iterations of the local parameter fitting. As the fitting proceeds, the mesh size is refined to 64×64 and finally to 128×128. Global parameters and their deformations capture the shape well enough in low resolution grid levels, so we only start using local deformations at higher grid resolutions.

Apart from the height maps, following [30], we also report the average absolute error for each image. For each image, we compare our result to the best result on the same test image as reported in [30]. (The algorithms were those described in [4, 31, 15, 17, 3, 16, 25, 29]). Our results are in almost all cases clearly better than the results computed using the SFS-based only algorithms reviewed in [30], and in no case worse (except from the algorithm in [15] which is not considered, because it is not solely a SFS algorithm, since it relies on good initial values which are provided by stereo).

Since all the algorithms recover relative height, the range of the recovered height values can vary greatly. Therefore before we compare any two surfaces, we need to normalize them to the same range. The average absolute error measure, scales the height values so that the ranges of the two surfaces are equal – we used the same method for this as [30]. This method can be extremely sensitive to noisy outliers – a single point can drastically change the absolute range of height values. For this reason, we evaluate one more error measure,

[1] Space limitations preclude us from reporting results for all five image sets. The three images that we chose were the most structurally complex objects, which also tended to have the highest error rates among the algorithms tested in [30].

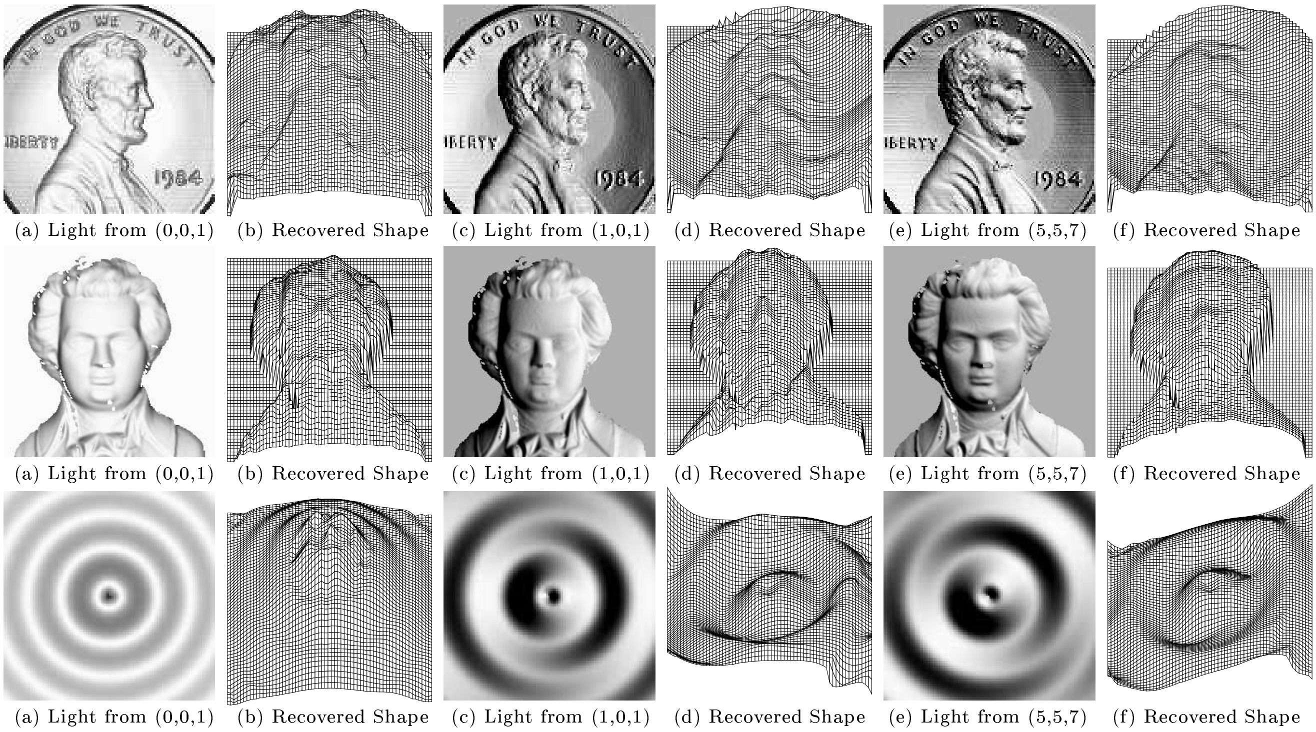

(a) Light from (0,0,1) (b) Recovered Shape (c) Light from (1,0,1) (d) Recovered Shape (e) Light from (5,5,7) (f) Recovered Shape

(a) Light from (0,0,1) (b) Recovered Shape (c) Light from (1,0,1) (d) Recovered Shape (e) Light from (5,5,7) (f) Recovered Shape

(a) Light from (0,0,1) (b) Recovered Shape (c) Light from (1,0,1) (d) Recovered Shape (e) Light from (5,5,7) (f) Recovered Shape

Figure 2: Experiments were run on three data sets of true data: Penny (top), Mozart (middle) and Sombrero (bottom), were rendered each under three light source direction: Left (0,0,1), Center: (1,0,1), Right: (5,5,7)

the best error of fit. In this measure we find the scaling and translation along the Z axis of the two surfaces, so that the distance between them will be minimal, in a least squares sense. These results are reported in Table 1.

Visual inspection of our results shows that our method manages to recover surface information and detail in most cases. Discontinuities are handled well, thanks to the elasticity of the model. A number of the inherent problems of SFS, such as convex/concave ambiguities, problems arising from the violation of the assumptions, rippling parallel to the direction of the light [31, 17], appear in some of the results. The C^0 elements that are used for efficiency reasons, tend to favor flat surfaces. The amount of detail recovered depends on the number of elements in the model. When fitting a fine mesh, local detail will appear in the first few iterations. As the fitting progresses, the deformation will deepen and broaden but without significant loss of detail. This is an advantage of the hard constraints approach. An inspection of the depth error images[2] as in [30], reveals that apart from the the boundaries and areas with detail finer than the discretization of the model, the error tends to concentrate in certain regions of the image. This suggests that in those areas, global or piecewise global approaches could provide further accuracy gains.

Two experiments were performed on a synthetic image of a sphere and Mozart, rendered using the lighting model of [23] described in Sec.4.2. The parameters were $\sigma = 0.543788$ and $\rho = 0.770950$, which are the values that model plaster surfaces [5]. Illumination was at an angle of 22.5 degrees off the viewing axis for the sphere experiment and in the direction (5,5,7) for the Mozart experiment. In both cases the shape was recovered correctly with similar error results as in the case of Lambertian reflectance. Reconstruction results are shown in Fig.4.

7 Conclusion

In this paper we presented a formal method for the integration of Shape from Shading within a deformable model framework. This was done by extending previous approaches of incorporating hard constraints to a Lagrangian framework, to the case of nonlinear holonomic constraints. This methodology was then applied to the Shape from Shading problem. A deformable model with both global and local deformations was employed for the recovery of height, based on the principle of Lambertian illumination. A series of experiments was performed on a standardized image set and

[2] The limited space of a conference paper does not allow us to include them

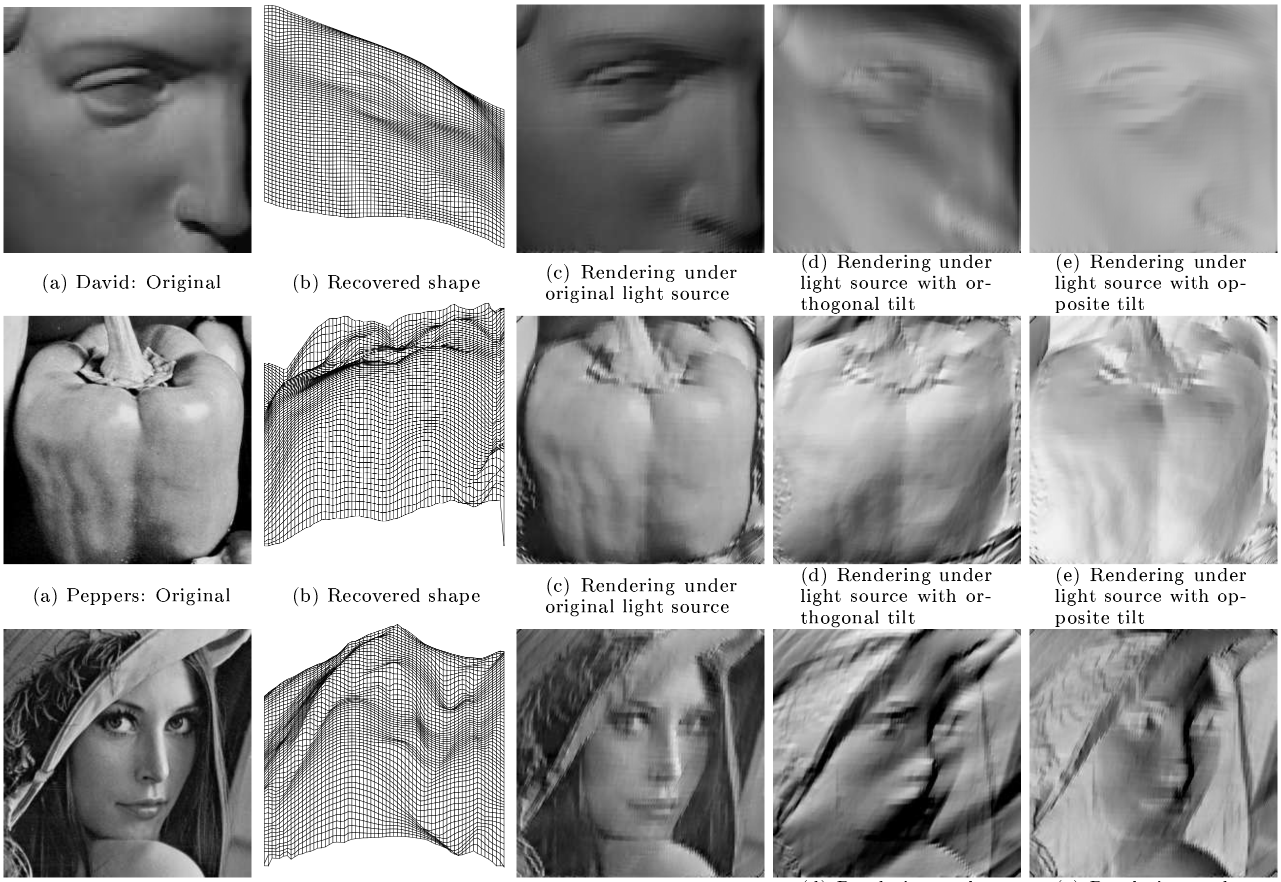

(a) David: Original (b) Recovered shape (c) Rendering under original light source (d) Rendering under light source with orthogonal tilt (e) Rendering under light source with opposite tilt

(a) Peppers: Original (b) Recovered shape (c) Rendering under original light source (d) Rendering under light source with orthogonal tilt (e) Rendering under light source with opposite tilt

(a) Lena: Original (b) Recovered shape (c) Rendering under original light source (d) Rendering under light source with orthogonal tilt (e) Rendering under light source with opposite tilt

Figure 3: Real images: Reconstruction and 3 renderings under the original light source, light source with orthogonal tilt and light source with opposite tilt

the results were compared to some of the better known approaches, thus demonstrating the advantages of the method.

Acknowledgments

The test data sets and images were obtained from the University of Central Florida Computer Vision lab.

This research has been partially supported by an NSF Career Award (NSF-9624604), an ONR-YIP, an NSF-9420397, a Whitaker Foundation Grant, an NSF-9701803, and an ARO-DAAH04-96-1-007.

References

[1] J. Baumgarte. Stabilization of constraints and integrals of motion in dynamical systems. *Computer Methods in Applied Mechanics and Engineering*, 1:1-16, 1972.

[2] A. Barr. Global and Local Deformations of Solid Primitives, *Computer Graphics(18)*, No. 3, 1984,

[3] M. Bichsel and A.P. Pentland. A simple algorithm for shape from shading. In *CVPR 1992* pages 459-465, 1992.

[4] M.J. Brooks and B.K.P. Horn. Shape and source from shading. In *Proceedings of International Joint Conference on Artificial Intelligence*, pages 932-936, 1985.

[5] K.J. Dana, B. van Ginneken, S.K. Nayar, J.J. Koenderink, "Reflectance and Texture of Real-World Surfaces", CVPR 1997

[6] D. DeCarlo and D. Metaxas. The integration of optical flow and deformable models with applications to human face shape and motion estimation. In *CVPR 1996*, pages 231–238, 1996.

[7] D. Forsyth and A. Zisserman. Mutual illumination. In *CVPR 1989*, pages 466-473, 1989.

[8] R.T. Frankot and R. Chellappa. A Method for Enforcing Integrability in Shape from Shading Algorithms. *PAMI*, 10:439-451, 1988.

[9] P. Fua and C. Brechbuhler. Imposing hard constraints on deformable models through optimization in orthogonal subspaces. In *CVIU*, 65(2):148–162, 1997.

Methods	Penny			Sombrero			Mozart		
	S1	S2	S3	S1	S2	S3	S1	S2	S3
Best from [30]	7.2	4.7	4.4	6.9	5.4	5.6	15.1	8.8	7.7
Average Error	6.3	3.9	3.4	4.9	4.5	5.1	8.9	7.1	6.0
Best-fit Error	4.1	3.7	3.4	4.3	3.8	3.9	8.5	7.1	5.8

Table 1: Average Z error for synthetic images. For comparison, "Best from [30]" provides for each image, the best one of the results obtained by the algorithms surveyed in [30]. Thus it includes results from several algorithms — no single algorithm performed best on all the images.

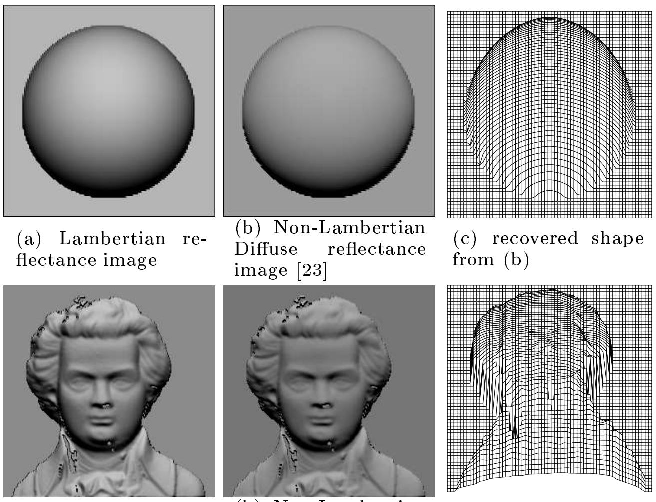

(a) Lambertian reflectance image (b) Non-Lambertian Diffuse reflectance image [23] (c) recovered shape from (b)

(a) Lambertian reflectance image (b) Non-Lambertian Diffuse reflectance image [23] (c) recovered shape from (b)

Figure 4: Recovered shape from the middle image using the Non-Lambertian Diffuse reflectance model described in [23]. To the left is what the same geometry with the same albedo would look like under the lambertian model.

[10] P. Fua and Y.G. Leclerc. Object-Centered Surface Reconstruction: Combining Multi-Image Stereo and Shading, IJCV(16), September 1995.

[11] B.K.P Horn. 1970. Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View. PhD thesis, MIT, 1970.

[12] B.K.P Horn. 1989. Height and gradient from shading. *IJCV*, pages 37-75, 1989.

[13] B.K.P. Horn and M.J. Brooks. eds. Shape from Shading. Cambridge, MA: MIT Press, 1989

[14] K. Ikeuchi and B.K.P Horn. Numerical Shape from shading and occluding boundaries. *Artificial Intelligence*, 17(1-3):141-184, 1981.

[15] Y.G. Leclerc and A.F. Bobick. The Direct Computation of Height from Shading. In *CVPR 1991*, pages 552-558,1991.

[16] C.H. Lee and A. Rosenfeld. Improved methods of estimating shape from shading using the light source coordinate system. *Artificial Intelligence*, 26:125-143, 1985.

[17] K.M. Lee and C.C. Kuo. Shape from Shading with a Linear Triangular Element Surface Model. *PAMI* 15(8) Aug. 1993

[18] S.H. Lai and B.C. Vemuri. Physically-Based Adaptive Preconditioning for Early Vision, PAMI(19), June 1997.

[19] D. Metaxas and D. Terzopoulos. Shape and Nonrigid Motion Estimation Through Physics-based Synthesis. *PAMI–15*(6), pp. 580–591, June 1993.

[20] D. Metaxas, Physics-Based Deformable Models: Applications to Computer Vision, Graphics and Medical Imaging, Kluwer-Academic Publishers, November 1996.

[21] S.K. Nayar, K. Ikeuchi and T. Kanade. Surface Reflection: Physical and Geometrical Perspectives In *PAMI*, 13(7):611–634, July 1991

[22] J. Oliensis and P. Dupois. A global algorithm for shape from shading. In *ICCV 1993*, pages 692-701, 1993.

[23] M. Oren and S.K. Nayar. Diffuse reflectance from rough surfaces. In *ICCV 1993* pages 763-764, 1993.

[24] A.P. Pentland. 1984. Local Shading Analysis. *PAMI* 6:170-187, 1984.

[25] A.P. Pentland. 1988. Shape information from shading: a theory about human perception. In *ICCV 88* pages 404-413, 1988.

[26] W.H. Press, S.A. Teukolsky, W.T. Vettering and B.P. Flannery. *Numerical Recipes in C*, Cambridge University Press, 1992.

[27] R. Szeliski. Fast shape from shading. *CVGIP:IU*, 53:129-153, 1991.

[28] K.E. Torrance and E.M. Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57:1105-1114, 1967.

[29] P.S. Tsai and M. Shah. A simple shape from shading algorithm. In *CVPR 1992* pages 734-736, 1992.

[30] R. Zhang, P. Tsai, J.E. Cryer, M. Shah. Analysis of shape from shading techniques. In *CVPR 1994* pages 377-384, 1994.

[31] Zheng, Chellappa Estimation of illumination direction, albedo, and shape from shading. *PAMI* 13(7):680-702, 1991.