

This is a repository copy of *New constraints on data-closeness and needle map consistency for shape-from-shading*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/1987/>

Article:

Worthington, P L and Hancock, E R orcid.org/0000-0003-4496-2028 (1999) New constraints on data-closeness and needle map consistency for shape-from-shading. IEEE Transactions on Pattern Analysis and Machine Intelligence. pp. 1250-1267. ISSN: 0162-8828

<https://doi.org/10.1109/34.817406>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

New Constraints on Data-Closeness and Needle Map Consistency for Shape-from-Shading

Philip L. Worthington and Edwin R. Hancock

Abstract—This paper makes two contributions to the problem of needle-map recovery using shape-from-shading. First, we provide a geometric update procedure which allows the image irradiance equation to be satisfied as a hard constraint. This not only improves the data closeness of the recovered needle-map, but also removes the necessity for extensive parameter tuning. Second, we exploit the improved ease of control of the new shape-from-shading process to investigate various types of needle-map consistency constraint. The first set of constraints are based on needle-map smoothness. The second avenue of investigation is to use curvature information to impose topographic constraints. Third, we explore ways in which the needle-map is recovered so as to be consistent with the image gradient field. In each case we explore a variety of robust error measures and consistency weighting schemes that can be used to impose the desired constraints on the recovered needle-map. We provide an experimental assessment of the new shape-from-shading framework on both real world images and synthetic images with known ground truth surface normals. The main conclusion drawn from our analysis is that the data-closeness constraint improves the efficiency of shape-from-shading and that both the topographic and gradient consistency constraints improve the fidelity of the recovered needle-map.

Index Terms—Shape-from-shading, hard constraints, curvature consistency, gradient consistency, robust statistics.

1 INTRODUCTION

SHAPE-FROM-SHADING (SFS) has been a topic of sustained research activity for over two decades. The process is concerned with recovering local surface orientation from variations in image radiance. In Marr's seminal work detailing a framework for computer vision [21], shape-from-shading is identified as providing a useful route to understanding 3D surface shape from a single image. Moreover, psychophysical motivation for the use of SFS has recently been demonstrated through experiments on the human vision system [17], [18].

From a computational viewpoint, SFS involves solving the image irradiance equation (IIR) to recover a set of surface normals, often described as the needle-map. However, since the IIR is under-constrained, additional constraints such as local smoothness of the needle-map must be invoked for reasons of computational tractability.

One of the most popular approaches is to adopt a regularization framework [6], [10], [12] and iteratively recover the needle-map using an update equation derived by applying variational calculus. The regularization functional must ensure data-closeness by penalizing departures from the IIR, while imposing smoothness on the recovered needle-map. A fundamental obstacle to progress in energy minimization approaches to shape-from-shading is the failure, to-date, to develop a constraint function which is uniquely minimized by a surface which closely matches the true surface. The search for such a functional is hindered by the nature of the minimization framework, which is usually

based upon regularization theory [27]. This requires careful selection of constraint functions and, most significantly, time-consuming experimentation to establish regularization parameters and determine whether or not a given scheme is numerically stable. In many cases, the conditions to ensure numerical stability lead to severe model dominance.

1.1 Related Literature

Many different computational approaches to recovering shape from shading have been proposed, falling into two broad categories; local [25], [6], and global [9], [14], [10], [28]. Local shading analysis involves taking small windows of the image and recovering surface patches which are subsequently quilted together. Local approaches tend to be fast, but often require *a priori* surface height information, such as the elevation of singular points [22], or the linearization of the reflectance map [25]. On real world images, the results of local SFS techniques are often noisy.

In contrast, global approaches attempt to recover the entire surface, either through propagation of height values from singular points of the surface, or by minimization of some energy functional associated with the estimated surface. Global propagation techniques again require *a priori* height data for initialization, whereas global minimization techniques may be developed which require only the image and a light source direction estimate as input. Moreover, global minimization techniques have been shown to be more generally applicable to different types of input images, and more robust to noise than either local techniques or global propagation approaches [36]. With these factors in mind, this paper focuses upon global minimization techniques.

A major criticism of many global approaches has been their tendency to oversmooth the recovered surface; that is, to lose detail, particularly in regions where the true surface is discontinuous. Various approaches have been proposed

• The authors are with the Department of Computer Science, University of York, YO1 5DD, UK. E-mail: {plw, erh}@cs.york.ac.uk.

Manuscript received 22 Dec. 1998; revised 24 Aug. 1999.

Recommended for acceptance by M. Shah.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 108495.

to address this problem. For instance, Ferrie and Legarde [6] have used curvature consistency to augment the more usual constraint on smoothness. Zheng and Chellappa [37] apply an intensity gradient constraint which penalizes discrepancies between the rate of change of reflectance and the rate of change of image intensity. Lee and Kuo [29] describe a scheme based on a triangular-element surface model, in which the smoothness constraint takes the form of a surface stiffness.

Recently, Kimmel and Bruckstein [15] have shown how the apparatus of level-set theory [29] can be used to solve the Euler equations developed by Horn [9] to describe the SFS problem. In this approach, a 3D function is propagated on a grid. A level set of this function tracks the height contours of the recovered shape. This builds upon earlier work by Bruckstein [5], which propagated a single equal-height contour in an image (e.g., a shoreline or horizon in a landscape picture) to find the set of equal-height contours defining the surface. In the level-set approach, the 3D function evolves according to constraints which are analogous to Huygen's wavefront propagation principle, an idea borrowed from recent advances in computational fluid dynamics.

1.2 Paper Outline

Existing work on shape-from-shading can be criticized on two counts. First, the data-closeness constraint invariably plays a relatively weak role in the recovery of the needle-map, while the smoothness error dominates the process. Second, the modeling of the smoothness constraint is extremely simplistic. Most practical schemes opt for a quadratic regularizer [10], [12]. This has the effect not only of over-smoothing fine surface detail, but also of failing to capture the differential or topographic structure of realistic surfaces.

Several pieces of work aim explicitly to address these criticisms. Ferrie and Lagarde [6] impose curvature consistency on the surface as a postprocessing step. Their method attempts to use topographic surface structure to improve the global properties of the surface reconstruction, without resorting to smoothing.

Horn [12] attempts to improve data-closeness by annealing the weight assigned to the smoothness regularizer as convergence is approached. In a more recent paper [32], we have applied the apparatus of robust statistics to modify the strength of the smoothness constraint in a principled manner. Robust error kernels are used in place of a quadratic regularizer. Applying variational calculus to the resulting regularized functional yields needle-map update equations in which data-closeness plays a more significant role.

The novel contribution of this paper is to furnish new constraints on the recovery of the needle-map. The first contribution is to impose data-closeness as a hard constraint. Our starting point is the observation that the IIR constrains the surface normal to fall on a cone of ambiguity. The axis of the cone is the light-source direction. This leads to a simple geometric picture of iterative needle-map recovery which can be visualized as remapping the surface normals onto a cone of possible solutions. The second contribution is to exploit this new framework to explore

more complex constraints on needle-map consistency. Here, we show that both curvature consistency and gradient consistency provide better constraints on the recovery of faithful needle maps than the conventional quadratic smoothness constraint of Horn and Brooks.

The structure of this paper is as follows: As an introduction to the global recovery of shape from shading, we briefly summarize the variational approach of Horn and Brooks [10]. By considering its weaknesses, we identify the main deficiencies of existing global minimization techniques. We then present our framework for SFS and argue for its validity and advantages in terms of addressing these deficiencies. Subsequently, we develop several new constraints for SFS, including image gradient and curvature consistency. Placing these constraints within the new framework, we develop iterative update schemes to recover the needle-map from an image. We compare the performance of our new techniques against the Horn and Brooks variational approach to SFS and other schemes from the literature. Through detailed analysis using synthetic data, we demonstrate the advantages of the new framework. Improved performance on real images is also empirically demonstrated.

2 THE VARIATIONAL APPROACH TO SFS

Central to shape-from-shading is the idea that local regions in an image $E(x, y)$ correspond to illuminated patches of a piecewise continuous surface, $z(x, y)$. The measured brightness $E(x, y)$ will depend on the material properties of the surface, the orientation of the surface at the coordinates (x, y) , and the direction and strength of illumination.

The *reflectance map*, $R(p, q)$ characterizes these properties, and provides an explicit connection between the image and the surface orientation. Surface orientation is described by the components of the surface gradient in the x and y direction, i.e., $p = \frac{\partial z}{\partial x}$ and $q = \frac{\partial z}{\partial y}$. The shape from shading problem is to recover the surface $z(x, y)$ from the intensity image $E(x, y)$. As an intermediate step, we may recover the needle-map, or set of estimated local surface normals, $\mathbf{n}(x, y)$.

Needle-map recovery from a single intensity image is an under-determined problem [22], [12], [2] which requires a number of constraints and assumptions to be made. The common assumptions are that the surface has ideal Lambertian reflectance, constant albedo, and is illuminated by a single point source at infinity. A further assumption is that there are no interreflections, i.e., the light reflected by one portion of the surface does not impinge on any other part.

2.1 The Horn and Brooks Algorithm

The local surface normal may be written as $\mathbf{n} = (-p, -q, 1)^T$, where $p = \frac{\partial z}{\partial x}$ and $q = \frac{\partial z}{\partial y}$. For a light source at infinity, we can similarly write the light source direction as $\mathbf{s} = (-p_l, -q_l, 1)^T$. If the surface is Lambertian the reflectance map is given by

$$R(p, q) = \mathbf{n} \cdot \mathbf{s}. \quad (1)$$

The image irradiance equation [11] states that the measured brightness of the image is proportional to the

radiance at the corresponding point on the surface; that is, just the value of $R(p, q)$ for p, q corresponding to the orientation of the surface. Normalizing both the image intensity, $E(x, y)$, and the reflectance map, the constant of proportionality becomes unity, and the image irradiance equation is simply

$$E(x, y) = R(p, q). \quad (2)$$

Although the image irradiance equation succinctly describes the mapping between the x, y coordinate space of the image and the p, q gradient space of the surface, it provides insufficient constraints for the unique recovery of the needle-map. To overcome this problem, a further constraint must be applied. Usually, the needle-map is assumed to vary smoothly.

The process of smooth surface recovery is posed as a variational problem in which a global error functional is minimized through the iterative adjustment of the needle map. Surface normals are updated with a step size dictated by Euler's equation. Here we consider the formulation of Brooks and Horn [6] which is couched in terms of unit surface normals. Their error functional is defined to be

$$I = \int \int \underbrace{\left(E(x, y) - \mathbf{n} \cdot \mathbf{s} \right)^2}_{\text{BrightnessError}} + \lambda \underbrace{\left(\left\| \frac{\partial \mathbf{n}}{\partial x} \right\|^2 + \left\| \frac{\partial \mathbf{n}}{\partial y} \right\|^2 \right)}_{\text{RegularizingTerm}} dx dy. \quad (3)$$

The terms $\frac{\partial \mathbf{n}}{\partial x}$ and $\frac{\partial \mathbf{n}}{\partial y}$ above are the directional derivatives of the needle-map in the x and y directions respectively. The magnitudes of these quantities are used to measure the smoothness of the surface, with a large value indicating a highly curved region. However, it should be noted that a planar surface has $\frac{\partial \mathbf{n}}{\partial x} = \frac{\partial \mathbf{n}}{\partial y} = \mathbf{0}$ in this case.

The first term of (3) is the brightness error, which encourages data-closeness of the measured image intensity and the reflectance function. The *regularizing term* imposes the smoothness constraint on the recovered surface normals, penalizing large local changes in surface orientation. The constant λ is a Lagrange multiplier. For numerical stability, λ must often be large, resulting in the smoothness term dominating.

Minimization of the functional defined in (3) is accomplished by applying the calculus of variations [7] and solving the resulting Euler equation. The general Euler equation is

$$I_n - \frac{\partial}{\partial x} I_{n_x} - \frac{\partial}{\partial y} I_{n_y} = \mathbf{0}, \quad (4)$$

where I_n is the result of taking the partial derivative of I with respect to the surface normal, \mathbf{n} . Similarly, I_{n_x} and I_{n_y} are partial derivatives of I with respect to the quantities $\frac{\partial \mathbf{n}}{\partial x}$ and $\frac{\partial \mathbf{n}}{\partial y}$, respectively.

Substituting (3) into the general Euler equation yields:

$$\left(E - \mathbf{n} \cdot \mathbf{s} \right) \mathbf{s} + \lambda \nabla^2 \mathbf{n} = \mathbf{0}. \quad (5)$$

Equation (5) is discretized using the following approximation to the Laplacian

$$\left\{ \nabla^2 \mathbf{n} \right\}_{i,j} \approx \frac{4}{\epsilon^2} \left(\bar{\mathbf{n}}_{i,j} - \mathbf{n}_{i,j} \right), \quad (6)$$

where ϵ is the spacing of pixel-sites on the lattice and $\bar{\mathbf{n}}_{i,j}$ is the local 4-neighborhood mean of the surface normals around pixel position i, j ,

$$\bar{\mathbf{n}}_{i,j} = \frac{1}{4} \left(\mathbf{n}_{i,j+1} + \mathbf{n}_{i,j-1} + \mathbf{n}_{i+1,j} + \mathbf{n}_{i-1,j} \right). \quad (7)$$

Discretizing (5) and rearranging, we obtain an update equation to estimate the needle-map at epoch $k+1$ using the estimate at epoch k

$$\mathbf{n}_{i,j}^{(k+1)} = \bar{\mathbf{n}}_{i,j}^{(k)} + \frac{\epsilon^2}{2\lambda} \left(E_{i,j} - \mathbf{n}_{i,j}^{(k)} \cdot \mathbf{s} \right) \mathbf{s}. \quad (8)$$

2.2 Weaknesses of the Horn and Brooks Algorithm

The principal criticism of the Horn and Brooks algorithm, and of similar approaches [14], [30], is its tendency to oversmooth the recovered needle-map. Specifically, in (8), the regularizing term dominates the data term. Since the smoothness constraint is formulated in terms of the directional derivatives of the needle-map, it is trivially minimized by a flat surface. Thus, the conflict between the data and the model leads to a strongly smoothed needle-map and the loss of detail. The problem is exacerbated by the need to select a conservative value for the Lagrange multiplier in order to ensure numerical stability of the scheme [12].

A further criticism of the Brooks and Horn algorithm stems from the use of the Lagrange multiplier, which entails parameter searching. To be confident that the scheme will remain numerically stable for all expected input images, it is often necessary to use a conservative value of the parameter, exacerbating the oversmoothing problem.

In [31], [33], we introduced a novel application of robust M-estimators to SFS in order to address the oversmoothing problem. However, the improvements resulting from this approach are obtained at the expense of introducing an additional parameter to control the kernel width.

Hence, we see clear problems of model domination and time consuming parameter selection in the literature.

3 A NOVEL FRAMEWORK FOR SFS

In this paper, we aim to address the criticisms detailed in the previous section. Specifically, we wish to reduce model dependence without requiring parameters or any information in addition to the image and the light source direction. Note that Pentland [26] developed a method to estimate the light source direction from the image, so in fact only the image is required.

In the variational approach to SFS due to Brooks and Horn [4], the aim is to minimize a dual functional (3). However, the aims of the terms of the functional are different and competing. The smoothness term is minimized by any planar surface. Thus, if the scheme is initialized with the true needle-map, it will tend to "walk-away" as brightness error is traded for smoothness. Horn [12] attempts to address this problem by increasing the relative importance of the brightness error over successive iterations. Here, however, we opt to treat the data-closeness




Fig. 1. Comparison of update process for Horn and Brooks (left) and the new framework with the same smoothness constraint (right). Horn and Brooks allows the updated normal to move away from the cone of ambiguity, sacrificing brightness error for smoothness. The movement in the light source direction is as in the left-hand diagram, where $a = \frac{c^2}{2\lambda} (E - n \cdot s)$ from Equation (8). In contrast, the new framework forces the brightness error to be satisfied by using the rotation matrix Θ to map the smoothness update to the closest point on the cone.

and constraint terms separately [34], [35], in such a way that data-closeness is always guaranteed; that is, the image irradiance equation is treated as a hard constraint on the problem. Hence, we have a valid needle-map which satisfies the image irradiance equation at every iteration. Subject to this constraint, the task becomes iteration toward the true needle-map, using minimization of the smoothness error or some other constraint functional. The need for a weighting parameter is removed entirely.

Geometrically, we consider the image irradiance equation to define a cone of ambiguity about the light source direction for each surface normal. The normals forming the needle-map can only assume directions defined by this cone, as shown in Fig. 1. At each iteration the updated normal is free, as a result of the smoothing process, to lie outside the cone, but is subsequently projected back onto the closest vector lying on the cone.

3.1 Using the Image Irradiance Equation as a Hard Constraint

The new framework requires us to minimize the constraint functional

$$I = \int \int \psi(n(x, y), N(x, y)) dx dy, \quad (9)$$

while satisfying the hard constraint, imposed by the image irradiance equation

$$\int \int (E - n \cdot s) dx dy = 0. \quad (10)$$

Here, $N(x, y)$ is the set of local neighborhood vectors about location (x, y) . For example, in terms of lattice coordinates i, j , the 4-neighborhood of $n_{i,j}$ is defined as:

$$N = \{n_{i+1,j}, n_{i-1,j}, n_{i,j+1}, n_{i,j-1}\}. \quad (11)$$

The function $\psi(n(x, y), N(x, y))$ is a localized function of the current surface normal estimates. The size of the neighborhood may be varied according to the nature of ψ . Clearly, it is possible to incorporate the hard data-closeness constraint directly into ψ , but this needlessly complicates the mathe-

matics. Instead, we choose to impose the constraint after each iteration by mapping the updated normals back to the most similar normal lying on the cone.

If we take the smoothness constraint of Horn and Brooks as an example, then

$$\psi(n, N) = \left\| \frac{\partial n}{\partial x} \right\|^2 + \left\| \frac{\partial n}{\partial y} \right\|^2. \quad (12)$$

Discretizing $\psi(n(x, y), N(x, y))$ directly in terms of forward differences yields the following equation in terms of $n_{i,j}$ and its neighboring normals

$$\psi(n_{i,j}, N_{i,j}) = (n_{i+1,j} - n_{i,j})^2 + (n_{i,j+1} - n_{i,j})^2. \quad (13)$$

In practice, to maintain symmetry, we take the mean of the forward and backward difference results. We cannot use the central difference because this does not include the current normal, $n_{i,j}$. Taking the mean yields

$$\psi(n_{i,j}, N_{i,j}) = \frac{1}{2} \left((n_{i+1,j} - n_{i,j})^2 + (n_{i,j+1} - n_{i,j})^2 + (n_{i,j} - n_{i-1,j})^2 + (n_{i,j} - n_{i,j-1})^2 \right). \quad (14)$$

We can again apply the calculus of variations, as described by (4). Note that, having discretized prior to differentiating, we have only the I_n component. However, the same result can be obtained by performing the discretization after differentiation. The operations are shown in this order to motivate the definition of ψ in terms of the neighborhood normals, $N(x, y)$. Thus, the discrete Euler equation becomes

$$\frac{1}{2} (-2(n_{i+1,j} - n_{i,j}) - 2(n_{i,j+1} - n_{i,j}) + 2(n_{i,j} - n_{i-1,j}) + 2(n_{i,j} - n_{i,j-1})) = 0. \quad (15)$$

Rearranging, and substituting (7), yields

$$n_{i,j}^{k+1} = \bar{n}_{i,j}^k. \quad (16)$$

In other words, we obtain a simple neighborhood averaging, or smoothing action, which if left unchecked will

eventually lead to a flat surface. Note that this update equation could be obtained from the Horn and Brooks update equation simply by setting the brightness error to zero in (8). Equally, we could obtain an identical update equation using the constraint function

$$\psi(\mathbf{n}, \mathbf{N}) = (\mathbf{n} - \bar{\mathbf{n}})^2. \quad (17)$$

Applying the hard constraint, for instance, in terms of a rotation of the update back onto the cone of ambiguity, (16) becomes

$$\mathbf{n}_{i,j}^{k+1} = \Theta \bar{\mathbf{n}}_{i,j}^k, \quad (18)$$

where Θ is a rotation matrix to map the updated normal to the closest normal lying on the cone of ambiguity. Another way to look at this, is that we allow the smoothness constraint to select the direction of the normal estimate in the image plane only, while fixing the angle between the normal estimate and the light source direction.

To achieve the rotation, we define an axis perpendicular to the intermediate update normal, $\bar{\mathbf{n}}_{i,j}^k$, and the light source direction. The axis of rotation is found by taking the cross product of the intermediate update with the light source direction

$$(u, v, w)^T = \bar{\mathbf{n}}_{i,j}^k \times \mathbf{s}. \quad (19)$$

The angle of rotation is the difference between the angle subtended by the intermediate update and the light source, and the apex angle of the cone of ambiguity. Since the image is normalized, the latter angle is simply $\cos^{-1} E$, giving a rotation angle of

$$\theta = -\cos^{-1} \left(\frac{\bar{\mathbf{n}}_{i,j}^k \cdot \mathbf{s}}{\|\bar{\mathbf{n}}_{i,j}^k\| \|\mathbf{s}\|} \right) + \cos^{-1} E. \quad (20)$$

Hence, the rotation matrix is given by

$$\Theta = \begin{pmatrix} c + u^2 c' & -ws + uv c' & vs + uw c' \\ ws + uv c' & c + v^2 c' & -us + vw c' \\ -vs + uw c' & us + vw c' & c + w^2 c' \end{pmatrix},$$

where

$$c = \cos(\theta) \quad c' = 1 - c \quad s = \sin(\theta). \quad (21)$$

Fig. 1 illustrates the update process, and compares it with the Horn and Brooks update of (8).

3.2 Initialization

The new framework forces the surface normals to lie on the cone defined by the image irradiance equation at each iteration. Therefore, we choose an initialization which ensures that the IIR (1) is satisfied at the outset. This differs from the Horn and Brooks algorithm, which is usually initialized by estimating the occluding boundary normals, with all other normals set to point in the light source direction.

We choose to initialize each normal such that its projection onto the image plane lies in the opposite direction to the image gradient direction, as shown in Fig. 2. This results in an initialization with an implicit bias




Fig. 2. The set of surface normals at a point which satisfy the Image Irradiance equation define a cone such that $E \cdot \mathbf{n} \cdot \mathbf{s} = 0$. A normal from this set is chosen such that the direction of its projection to the image plane is opposite to the maximum intensity gradient direction, \mathbf{g} .

toward convex rather than concave surfaces. In other words, bright regions are assumed to correspond to peaks, and the image gradient direction points toward these peaks. Fig. 3 illustrates this bias.

4 ALTERNATIVE CONSTRAINTS

The novel framework developed in the preceding section, using the smoothness constraint of Horn and Brooks, reduces the problem of model dominance, and removes the need to perform exhaustive parameter searches. Our second contribution in this research is to develop new methods of enforcing needle-map consistency. This is motivated by previous work which has shown that the traditional quadratic smoothness regularizer is ill-suited to the SFS problem [33]. Remarkably little attention has been focused upon alternatives to smoothness constraints [6], [37].

Here, we describe three promising approaches. The first is simply to modify the action of needle-map smoothness to more adequately model the properties of real surfaces. Second, we consider the topographic properties of surfaces to develop consistency constraints based upon surface curvature. Finally, we use higher order consistency between the needle-map and image to ensure satisfaction of the image irradiance equation in terms of image gradient as well as intensity.

4.1 Surface Geometry Constraints

This class of constraints includes any which make direct assumptions about the form of the surface. For instance, the quadratic smoothness constraint of Horn and Brooks falls into this category, since it assumes that the surface to be recovered is the smoothest possible surface satisfying the image irradiance equation, (1) subject to the criticisms of Section 2.2. Similarly, the robust regularizer developed in [32], [33] uses a modified version of this assumption. In essence, it considers that the recovered surface should be smooth, except where there is a high probability that a




Fig 3. A 1D bright image patch surrounded by darker regions causes the image gradient at surrounding locations to point toward the bright patch (top). A bright patch corresponds to a region presenting a large area toward the light source direction. Initializing the surface normals to point in the same direction as the negative gradient means that the bright region is always assumed to be a peak, as shown in the first diagram. However, the bottom two diagrams also show valid surfaces which will result in the same intensities, but for which the initialization will not work well. Given the bas-relief ambiguity, there is no way to distinguish between these possibilities, so we define our convention to assume the first case.

discontinuity is present, in which case the smoothing is reduced.

We define the robust regularizer constraint function as

$$\psi(\mathbf{n}, \mathbf{N}) = \rho_{\sigma} \left(\left\| \frac{\partial \mathbf{n}}{\partial x} \right\| \right) + \rho_{\sigma} \left(\left\| \frac{\partial \mathbf{n}}{\partial y} \right\| \right), \quad (22)$$

where $\rho_{\sigma}(\eta)$ is a robust kernel defined on the residual η and with width parameter σ . Applying the calculus of variations, discretizing, and rearranging yields the general update equation

$$\mathbf{n}_{i,j}^{(k+1)} = \Theta \left(\frac{\partial}{\partial x} \rho'_{\sigma} \left(\left\| \frac{\partial \mathbf{n}}{\partial x} \right\| \right) + \frac{\partial}{\partial y} \rho'_{\sigma} \left(\left\| \frac{\partial \mathbf{n}}{\partial y} \right\| \right) \right), \quad (23)$$

where

$$\begin{aligned} \rho'_{\sigma} \left(\left\| \frac{\partial \mathbf{n}}{\partial x} \right\| \right) &= \frac{\partial}{\partial \mathbf{n}_x} \left(\rho_{\sigma} \left(\left\| \frac{\partial \mathbf{n}}{\partial x} \right\| \right) \right) \\ \rho'_{\sigma} \left(\left\| \frac{\partial \mathbf{n}}{\partial y} \right\| \right) &= \frac{\partial}{\partial \mathbf{n}_y} \left(\rho_{\sigma} \left(\left\| \frac{\partial \mathbf{n}}{\partial y} \right\| \right) \right) \end{aligned} \quad (24)$$

While the contents of the parentheses on the right hand side of (23) appear at first glance to be a scalar, in fact the differentiation of the robust kernel by a vector results in another vector. Hence, (23) is dimensionally correct, as demonstrated below using a specific instance of $\rho_{\sigma}(\eta)$.

In [33] we experimented with several robust error kernels, including Li's Adaptive Potential Functions [20], and the Tukey [8] and Huber [13] estimators. However, the sigmoidal derivative M-estimator, a continuous version of Huber's estimator, proved to possess the best properties for handling surface discontinuities, and is defined to be

$$\rho_{\sigma}(\eta) = \frac{\sigma}{\pi} \log \cosh \left(\frac{\pi \eta}{\sigma} \right). \quad (25)$$

Substituting (25) into (23) yields the update equation

$$\begin{aligned} \mathbf{n}_{i,j}^{(k+1)} &= \Theta \left(\left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial x} \right\|^{-1} \tanh \left(\frac{\pi}{\sigma} \left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial x} \right\| \right) \left(\mathbf{n}_{i+1,j}^{(k)} + \mathbf{n}_{i-1,j}^{(k)} \right) \right. \\ &\quad + \left(\frac{\pi}{\sigma} \left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial x} \right\|^{-2} \operatorname{sech}^2 \left(\frac{\pi}{\sigma} \left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial x} \right\| \right) - \right. \\ &\quad \left. \left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial x} \right\|^{-3} \tanh \left(\frac{\pi}{\sigma} \left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial x} \right\| \right) \right) \left(\frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial x} \cdot \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial x} \right) \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial x} \\ &\quad + \left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial y} \right\|^{-1} \tanh \left(\frac{\pi}{\sigma} \left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial y} \right\| \right) \left(\mathbf{n}_{i,j+1}^{(k)} + \mathbf{n}_{i,j-1}^{(k)} \right) \\ &\quad + \left(\frac{\pi}{\sigma} \left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial y} \right\|^{-2} \operatorname{sech}^2 \left(\frac{\pi}{\sigma} \left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial y} \right\| \right) - \right. \\ &\quad \left. \left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial y} \right\|^{-3} \tanh \left(\frac{\pi}{\sigma} \left\| \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial y} \right\| \right) \right) \\ &\quad \times \left(\frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial y} \cdot \frac{\partial^2 \mathbf{n}_{i,j}^{(k)}}{\partial y} \right) \frac{\partial \mathbf{n}_{i,j}^{(k)}}{\partial y} \Bigg). \end{aligned} \quad (26)$$

The robust regularizer approach provides significantly improved results over the simple Horn and Brooks smoothness constraint, but at the expense of introducing the parameter, σ .

4.2 Curvature Consistency

Needle-map smoothness appears to be an over strong and inappropriate constraint for shape from shading. This is primarily because real surfaces are more likely to be *piecewise* smooth; in other words, formed of smooth regions separated by sharp discontinuities in depth or orientation. The oversmoothing problem is exacerbated by the difficulty of formulating the continuous concept of smoothness on a

discrete pixel lattice, as clearly illustrated by the fact that the Horn and Brooks smoothness constraint is trivially minimized by a needle-map corresponding to a planar surface.

In contrast, consider the well defined curvature characteristics of most real world surfaces. Although the curvature classes either side of a depth discontinuity may be completely unrelated, this is not the case for an orientation discontinuity. Orientation discontinuities usually correspond to ruts or ridges. Furthermore, the curvature classes for locations either side of a rut or a ridge should be the most similar classes, either trough or saddle rut for a rut, or dome or saddle ridge for a ridge. This property of smooth variation in class suggests that curvature consistency may be a more appropriate constraint for SFS than smoothness, which strongly penalizes legitimate orientation discontinuities.

The use of a curvature consistency measure was introduced to SFS by Ferrie and Lagarde [6]. They use global consistency of principal curvatures [24] to refine the surface estimate returned by local shading analysis. Curvature consistency is formulated in terms of rotating the local Darboux frame to ensure that the principal curvature directions are locally consistent. Recently, Angelopoulou [1] has recovered estimates of $H - K$ labels using photometric stereo under three different illumination conditions.

One way of representing curvature information is to use $H - K$ labels, but these require us to set four thresholds to define the classes in terms of the mean and Gaussian curvatures. However, we propose to use curvature consistency based upon the shape index of Koenderink and van Doorn [16]. This is a measure which encodes the same curvature class information as $H - K$ labels in a single, continuous representation, and has the further advantage of not requiring any thresholds.

4.2.1 The Shape Index

We reformulate the definition of the shape index in terms of the needle-map. This allows us to use the needle-map directly, rather than needing to reconstruct the surface.

The differential structure of a surface is captured by the Hessian matrix

$$\mathcal{H} = \begin{pmatrix} -\frac{\partial^2 z}{\partial x^2} & -\frac{\partial^2 z}{\partial x \partial y} \\ -\frac{\partial^2 z}{\partial y \partial x} & -\frac{\partial^2 z}{\partial y^2} \end{pmatrix}.$$

The set of surface normals,

$$\mathbf{n} = (-p, -q, 1)^T = \left(-\frac{\partial z}{\partial x}, -\frac{\partial z}{\partial y}, 1 \right)^T$$

can be substituted into (27) to give the local Hessian in terms of the needle-map

$$\mathcal{H} = \begin{pmatrix} \left(\frac{\partial \mathbf{n}}{\partial x} \right)_x & \left(\frac{\partial \mathbf{n}}{\partial x} \right)_y \\ \left(\frac{\partial \mathbf{n}}{\partial y} \right)_x & \left(\frac{\partial \mathbf{n}}{\partial y} \right)_y \end{pmatrix}, \quad (27)$$

where $(\cdots)_x$ and $(\cdots)_y$ denote the x and y components of the parenthesized vector, respectively.

The eigenvalues of the Hessian matrix, found by solving the eigenvector equation $|\mathcal{H} - \kappa \mathbf{I}| = 0$, are the principal




Fig. 4. The shape index scale ranges from -1 to 1 as shown. The shape index values are encoded as a continuous range of grey level values between 1 and 255 , with grey level 0 being reserved for background and flat regions (for which the shape index is undefined).

curvatures of the surface. In terms of surface normals, these are given by

$$\kappa_{1,2} = -\frac{1}{2} \left(\left(\frac{\partial \mathbf{n}}{\partial x} \right)_x + \left(\frac{\partial \mathbf{n}}{\partial y} \right)_y \right) \mp \frac{1}{2} \sqrt{\left(\left(\frac{\partial \mathbf{n}}{\partial x} \right)_x - \left(\frac{\partial \mathbf{n}}{\partial y} \right)_y \right)^2 + 4 \left(\frac{\partial \mathbf{n}}{\partial x} \right)_y \left(\frac{\partial \mathbf{n}}{\partial y} \right)_x}. \quad (28)$$

Koenderink and van Doorn [16] defined the shape index to be

$$\phi = \frac{2}{\pi} \arctan \frac{\kappa_2 + \kappa_1}{\kappa_2 - \kappa_1} \quad \kappa_1 \geq \kappa_2. \quad (29)$$

This may be expressed in terms of surface normals thus,

$$\phi = \frac{2}{\pi} \arctan \frac{\left(\frac{\partial \mathbf{n}}{\partial x} \right)_x + \left(\frac{\partial \mathbf{n}}{\partial y} \right)_y}{\sqrt{\left(\left(\frac{\partial \mathbf{n}}{\partial x} \right)_x - \left(\frac{\partial \mathbf{n}}{\partial y} \right)_y \right)^2 + 4 \left(\frac{\partial \mathbf{n}}{\partial x} \right)_y \left(\frac{\partial \mathbf{n}}{\partial y} \right)_x}}. \quad (30)$$

Fig. 4 shows the range of shape index values, the type of curvature which they represent, and the gray levels used to display different shape index values.

4.2.2 Weighted Mean Using Curvature Consistency

A practical method of using curvature consistency is to use a weighted mean process of the local normals in the update. This follows logically from the update equation obtained using the quadratic smoothness constraint and the robust regularizer above. The idea is to weight in favor of “good” normals in the neighborhood. The selection of a criterion for what constitutes a good normal is an open problem. In this section, we investigate the use of curvature consistency to determine the goodness of each neighborhood normal, while later we investigate the use of image edge strength measurements.

As stated above, since the shape index is a continuous, angular measure, we expect it to vary gradually over a smooth surface. For instance, with reference to Fig. 4, we would not expect the shape index at adjacent pixels to differ by more than one curvature class unless they lie on opposite sides of a surface discontinuity. Since the over smoothing effect of the quadratic smoothness constraint stems directly from the indiscriminate averaging of normals lying across a discontinuity, we anticipate that weighting according to curvature consistency will reduce the problem in a physically principled manner.

We note once again that (18) may be obtained using the constraint function

$$\psi(\mathbf{n}, \mathbf{N}) = (\mathbf{n} - \bar{\mathbf{n}})^2 \quad (31)$$

instead of the constraint on the magnitudes of the directional derivatives of the needle-map used in (12). From this observation, we propose a similar constraint as follows:

$$\psi(\mathbf{n}, \mathbf{N}) = (\mathbf{n} - \tau(\mathbf{N}))^2, \quad (32)$$

where $\tau(\mathbf{N})$ is a weighted mean process defined on the neighborhood normals as

$$\tau(\mathbf{N}) = \frac{\sum_{l \in \mathbf{N}} w_l \mathbf{n}_l}{\sum_{l \in \mathbf{N}} w_l}. \quad (33)$$

To weight on the basis of curvature consistency we choose weights of the form

$$w_l = \exp \left(-\frac{(\phi_l - \mu_\phi)^2}{2\sigma_\phi^2} \right), \quad (34)$$

where μ_ϕ and σ_ϕ are, respectively, the mean and variance of the shape index over the neighborhood.

Hence, a neighborhood normal with a shape index similar to the neighborhood mean gains greater weight in the averaging process than one with an outlying shape index. The latter normal is assumed to lie on the opposite side of a discontinuity from the majority of the neighborhood, and therefore does not contribute significantly to the weighted mean vector.

There is no need to apply the calculus of variations in this case, as the functional based on the constraint of (32) is trivially minimized, within the new framework, by the update equation

$$\mathbf{n}_{i,j}^{(k+1)} = \Theta \tau(\mathbf{N}_{i,j}^k), \quad (35)$$

where, once again, Θ is a rotation matrix to map the intermediate update normal back onto the cone of solutions to the IIR (21).

Normals which are associated with a large difference from the neighborhood mean shape index are given small weights in the update, reflecting the fact that they are expected to lie on the opposite side of a discontinuity from the majority of the neighborhood.

An alternative formulation uses weights calculated using the Median Absolute Deviation (MAD) [8] in place of the mean of the neighborhood shape index values

$$w_l = \exp \left(-\frac{(\phi_l - \text{median}_\phi)^2}{2\text{MAD}_\phi^2} \right). \quad (36)$$

This is intended to deal with the potential problem that, if the neighborhood is evenly split by a discontinuity, the mean shape index may not be representative of either set of normals. Use of the median in this case forces the comparison value to be representative of the set of normals upon one side of the discontinuity or the other, but at the expense of introducing a nonlinear process.

4.2.3 Adaptive Robust Regularizer Using Curvature Consistency

We have also investigated the use of an adaptive robust regularizer in conjunction with curvature consistency [35]. Here, instead of determining the individual weightings of each normal in the neighborhood, we use the shape index statistics to adaptively set the width of the robust kernel, σ (22), to be applied over the neighborhood. The kernel width determines the level of smoothing applied to the neighborhood. If the shape index of the neighborhood varies greatly from the shape index at the center, then strong smoothing is applied, whereas an already smoothly varying shape index pattern receives less attention. Using the robust kernel of (22) reduces the problem of oversmoothing if large differences of shape index are present.

Once again, we have

$$\psi(\mathbf{n}, \mathbf{N}) = \rho_\sigma \left(\left\| \frac{\partial \mathbf{n}}{\partial x} \right\| \right) + \rho_\sigma \left(\left\| \frac{\partial \mathbf{n}}{\partial y} \right\| \right), \quad (37)$$

where, again, \mathbf{N} is the set of local neighborhood normals used to calculate the finite difference approximations to $\frac{\partial \mathbf{n}}{\partial x}$ and $\frac{\partial \mathbf{n}}{\partial y}$.

Instead of a fixed kernel width, σ , used in [32], [33], we use the adaptive value

$$\sigma = \sigma_0 \exp \left[-\left\{ \frac{1}{N} \sum_{l \in \mathbf{N}} \frac{(\phi_l - \phi_c)^2}{\Delta \phi_d^2} \right\}^{\frac{1}{2}} \right], \quad (38)$$

where ϕ_c is the shape index associated with the central normal of the neighborhood, $\mathbf{n}_{i,j}$, and $\Delta \phi_d$ is the difference in shape index between the center values of adjacent curvature classes. The number of neighborhood normals used in calculating the finite difference approximations to $\frac{\partial \mathbf{n}}{\partial x}$ and $\frac{\partial \mathbf{n}}{\partial y}$ is denoted N , and σ_0 is a reference kernel width which we set to unity. Using the scale of Fig. 4, $\Delta \phi_d = \frac{1}{8}$.

If the shape index varies significantly over the neighborhood, a small value of σ results, and the robust regularizer saturates to produce a heavy smoothing effect. In contrast, when the shape index values are already similar, the kernel is widened so that little smoothing occurs.

The update equation is identical to (23), but with the adaptive width parameter, σ , of (38) substituted in.

4.3 Gradient Consistency

Zheng and Chellappa [37] used such an intensity gradient constraint in conjunction with integrability. This constraint ensured that the intensity gradient of the image reconstructed from the needle-map matched the gradient of the input image. However, this constraint is not useful in the new framework since the reconstructed image always precisely matches the input due to the application of the image irradiance equation as a hard constraint. We therefore attempt to design constraints which use the needle-map structure directly.

To formulate a gradient consistency constraint, we desire to ensure that the image gradients in the x and y directions match those of the needle-map. In other words, differentiating the image irradiance (1)

$$\frac{\partial E}{\partial x} = \frac{\partial}{\partial x}(\mathbf{n} \cdot \mathbf{s}) \quad \text{and} \quad \frac{\partial E}{\partial y} = \frac{\partial}{\partial y}(\mathbf{n} \cdot \mathbf{s}). \quad (39)$$

Since the light source direction is constant with respect to x and y , these equations simplify to

$$\frac{\partial E}{\partial x} = \frac{\partial \mathbf{n}}{\partial x} \cdot \mathbf{s} \quad \text{and} \quad \frac{\partial E}{\partial y} = \frac{\partial \mathbf{n}}{\partial y} \cdot \mathbf{s}, \quad (40)$$

respectively. We incorporate these equations into constraint functions of similar forms to those investigated for curvature consistency.

4.3.1 Weighted Mean Using Gradient Consistency

Using (40), we develop a similar constraint function to (32), but with weights defined as

$$w_l = \exp \left[- \left\{ \left(\frac{\partial E}{\partial x} - \left(\frac{\partial \mathbf{n}}{\partial x} \right) \cdot \mathbf{s} \right)^2 + \left(\frac{\partial E}{\partial y} - \left(\frac{\partial \mathbf{n}}{\partial y} \right) \cdot \mathbf{s} \right)^2 \right\} \right]. \quad (41)$$

Hence, a neighborhood normal which has low gradient consistency is weighted less favorably than one which matches the image gradient closely.

4.3.2 Adaptive Robust Regularizer Using Gradient Consistency

Incorporating the gradient constraint into the robust regularizer constraint function described in (22), we define the adaptive kernel width to be the mean of the exponentials of the gradient consistency

$$\sigma = \sigma_0 \frac{1}{N} \sum_{l \in N} \exp \left[- \left\{ \left(\frac{\partial E}{\partial x} - \frac{\partial \mathbf{n}_l}{\partial x} \cdot \mathbf{s} \right)^2 + \left(\frac{\partial E}{\partial y} - \frac{\partial \mathbf{n}_l}{\partial y} \cdot \mathbf{s} \right)^2 \right\} \right]. \quad (42)$$

If the image and needle-map gradients correspond well, σ is large and little smoothing is applied. However, if some of the neighborhood normals exhibit large deviations from image consistency, this causes the kernel to narrow, reducing the influence of the inconsistent normals. If the mean inconsistency is very large, the kernel narrows to the point where all the neighborhood normals are in the saturation region of the kernel, causing the algorithm to revert to the Horn and Brooks smoothing process.

We experiment with two additional methods for adapting the value of σ . In the first case, we use the root mean value of the exponential consistencies as follows:

$$\sigma = \sigma_0 \sqrt{\frac{1}{N} \sum_{l \in N} \exp(-w_l)}, \quad (43)$$

where w_l is defined as in (41). This approach is motivated by the same considerations as above, i.e., to widen the kernel when the image and needle-map gradients correspond well, and narrow it to increase the smoothing when this is not the case. By taking the root mean instead of the simple mean, we reduce the influence of outlying values of w_l on the kernel width.

We also experiment with a scheme which uses the second derivatives of the image irradiance equation. Using this approach, we relate the kernel width to the Laplacian of

the image. Thus, we attempt to obtain consistency on the basis of local edge strength rather than the image gradient. This yields an adaptive kernel width involving the Laplacian of the image

$$\sigma = \sigma_0 \frac{1}{N} \sum_{l \in N} \exp \left[- \left(\nabla^2 E - \left\{ \frac{\partial^2 \mathbf{n}_l}{\partial x^2} + \frac{\partial^2 \mathbf{n}_l}{\partial y^2} \right\} \cdot \mathbf{s} \right) \right]. \quad (44)$$

This edge-based approach has intuitive appeal as an image edge is likely to correspond to a surface discontinuity in many cases. If the mismatch between the local edge strength and the needle-map second derivatives is large, the kernel is narrowed to increase the smoothing, whereas if the image and needle-map are consistent, little smoothing occurs.

4.4 Summary of Investigated Constraints

Table 1 summarizes the constraint functions described in the preceding sections, along with the update equations which result from them. The new schemes are labeled DD1-9 for Data Driven. Clearly, Table 1 is far from exhaustive. Many more constraint functions are possible, both within the categories considered here, and using other properties of surfaces and images. However, the rapid development of several alternatives to the smoothness constraint shows the power of the new framework for experimenting with different constraints.

5 EXPERIMENTS

We have performed extensive experiments upon synthetic and real world images, using the SFS schemes summarized in Table 1. We also compare with the Horn and Brooks algorithm and two recent SFS schemes. These are the local shape-from-shading algorithms of Bichsel and Pentland [3] and Tsai and Shah [28]. The experimental work is divided into three parts. We commence by showing surface reconstructions and needle-maps delivered by the different shape-from-shading algorithms. Next, we investigate the iteration dependence of various performance measures. Finally, we show some results on real-world images.

5.1 Synthetic Data

Fig. 5 illustrates the surfaces which can be recovered from an image of two conjoined spheres. The figure compares the SFS schemes developed in this paper with the alternatives described above. Note that the Bichsel and Pentland algorithm and the Tsai and Shah algorithm recover a surface directly, while the Horn and Brooks algorithm and the new framework recover needle-maps. The latter must be integrated into a surface, which is an open problem in itself. Here, we take the simplest possible approach, summing the needle-map terms in the x and y directions to yield two surfaces, then taking the mean of these. This process is extremely sensitive to noise, so surfaces recovered by this method are often poor. However, we include surfaces in this paper because it is easier to compare surfaces than needle-maps. Conversely, if we calculate the needle-map from the surfaces generated by Tsai and Sha or Bichsel and Pentland and compare on a

TABLE 1
Summary of the Constraint Functions Developed Using the New Framework

Name	Constraint function $\psi(\mathbf{n}, \mathbf{N})$	Update equation
DD1	Needle-Map Smoothness $\left\ \frac{\partial \mathbf{n}}{\partial x} \right\ ^2 + \left\ \frac{\partial \mathbf{n}}{\partial y} \right\ ^2 \left(\equiv (\mathbf{n} - \bar{\mathbf{n}})^2 \right)$	$\mathbf{n}_{i,j}^{(k+1)} = \Theta \bar{\mathbf{n}}_{i,j}^{(k)}$
DD2	Needle-Map Smoothness using Robust Regularizer $\rho_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial x} \right\ \right) + \rho_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial y} \right\ \right)$ where $\rho_\sigma(\eta) = \frac{\sigma}{\pi} \log \cosh \left(\frac{\pi \eta}{\sigma} \right)$	$\mathbf{n}_{i,j}^{(k+1)} = \Theta \left[\frac{\partial}{\partial x} \rho'_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial x} \right\ \right) + \frac{\partial}{\partial y} \rho'_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial y} \right\ \right) \right]$
DD3	Weighted Local Mean using Curvature Consistency $(\mathbf{n} - \tau(\mathbf{N}))^2$ where $\tau(\mathbf{N}) = \frac{\sum_{l \in \mathbf{N}} w_l \mathbf{n}_l}{\sum_{l \in \mathbf{N}} w_l}$ and $w_l = \exp \left(-\frac{(\phi_l - \mu_\phi)^2}{2\sigma_\phi^2} \right)$	$\mathbf{n}_{i,j}^{(k+1)} = \Theta \left(\tau \left(\mathbf{N}_{i,j}^k \right) \right)$
DD4	Weighted Local Mean using Curvature Consistency and Median Statistics $(\mathbf{n} - \tau(\mathbf{N}))^2$ where $\tau(\mathbf{N}) = \frac{\sum_{l \in \mathbf{N}} w_l \mathbf{n}_l}{\sum_{l \in \mathbf{N}} w_l}$ and $w_l = \exp \left(-\frac{(\phi_l - \text{median}_\phi)^2}{2\text{MAD}_\phi^2} \right)$	$\mathbf{n}_{i,j}^{(k+1)} = \Theta \left(\tau \left(\mathbf{N}_{i,j}^k \right) \right)$
DD5	Adaptive Robust Regularizer using Curvature Consistency $\rho_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial x} \right\ \right) + \rho_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial y} \right\ \right)$ where $\rho_\sigma(\eta) = \frac{\sigma}{\pi} \log \cosh \left(\frac{\pi \eta}{\sigma} \right)$ and $\sigma = \sigma_0 \exp \left(\left(\frac{1}{N} \sum_{l \in \mathbf{N}} \frac{(\phi_l - \phi)^2}{\delta \phi_d^2} \right)^{\frac{1}{2}} \right)$	$\mathbf{n}_{i,j}^{(k+1)} = \Theta \left[\frac{\partial}{\partial x} \rho'_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial x} \right\ \right) + \frac{\partial}{\partial y} \rho'_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial y} \right\ \right) \right]$
DD6	Adaptive Robust Regularizer using Gradient Consistency $\rho_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial x} \right\ \right) + \rho_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial y} \right\ \right)$ where $\rho_\sigma(\eta) = \frac{\sigma}{\pi} \log \cosh \left(\frac{\pi \eta}{\sigma} \right)$ and $\sigma = \sigma_0 \frac{1}{N} \sum_{l \in \mathbf{N}} \exp \left(-\left(\left(\frac{\partial E}{\partial x} - \frac{\partial \mathbf{n}}{\partial x} \cdot \mathbf{s} \right)^2 + \left(\frac{\partial E}{\partial y} - \frac{\partial \mathbf{n}}{\partial y} \cdot \mathbf{s} \right)^2 \right) \right)$	$\mathbf{n}_{i,j}^{(k+1)} = \Theta \left[\frac{\partial}{\partial x} \rho'_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial x} \right\ \right) + \frac{\partial}{\partial y} \rho'_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial y} \right\ \right) \right]$
DD7	Weighted Local Mean using Gradient Consistency $(\mathbf{n} - \tau(\mathbf{N}))^2$ where $\tau(\mathbf{N}) = \frac{\sum_{l \in \mathbf{N}} w_l \mathbf{n}_l}{\sum_{l \in \mathbf{N}} w_l}$ and $w_l = \exp \left(-\left(\left(\frac{\partial E}{\partial x} - \frac{\partial \mathbf{n}}{\partial x} \cdot \mathbf{s} \right)^2 + \left(\frac{\partial E}{\partial y} - \frac{\partial \mathbf{n}}{\partial y} \cdot \mathbf{s} \right)^2 \right) \right)$	$\mathbf{n}_{i,j}^{(k+1)} = \Theta \left(\tau \left(\mathbf{N}_{i,j}^k \right) \right)$
DD8	Adaptive Robust Regularizer using Gradient Consistency $\rho_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial x} \right\ \right) + \rho_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial y} \right\ \right)$ where $\sigma = \sigma_0 \frac{1}{N} \sum_{l \in \mathbf{N}} \exp \left(-\sqrt{\sum_{l \in \mathbf{N}} w_l} \right)$ and $w_l = \exp \left(-\left(\left(\frac{\partial E}{\partial x} - \frac{\partial \mathbf{n}}{\partial x} \cdot \mathbf{s} \right)^2 + \left(\frac{\partial E}{\partial y} - \frac{\partial \mathbf{n}}{\partial y} \cdot \mathbf{s} \right)^2 \right) \right)$	$\mathbf{n}_{i,j}^{(k+1)} = \Theta \left[\frac{\partial}{\partial x} \rho'_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial x} \right\ \right) + \frac{\partial}{\partial y} \rho'_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial y} \right\ \right) \right]$
DD9	Adaptive Robust Regularizer using Gradient Consistency $\rho_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial x} \right\ \right) + \rho_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial y} \right\ \right)$ where $\sigma = \sigma_0 \frac{1}{N} \sum_{l \in \mathbf{N}} \exp \left(-\sqrt{\sum_{l \in \mathbf{N}} w_l} \right)$ and $\sigma = \sigma_0 \frac{1}{N} \sum_{l \in \mathbf{N}} \exp \left(-\left(\nabla^2 E - \left(\frac{\partial^2 \mathbf{n}}{\partial x^2} + \frac{\partial^2 \mathbf{n}}{\partial y^2} \right) \cdot \mathbf{s} \right) \right)$	$\mathbf{n}_{i,j}^{(k+1)} = \Theta \left[\frac{\partial}{\partial x} \rho'_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial x} \right\ \right) + \frac{\partial}{\partial y} \rho'_\sigma \left(\left\ \frac{\partial \mathbf{n}}{\partial y} \right\ \right) \right]$

DD1 and DD2 are surface geometric constraints. DD3-5 use curvature consistency, while DD6-9 are based upon gradient consistency.

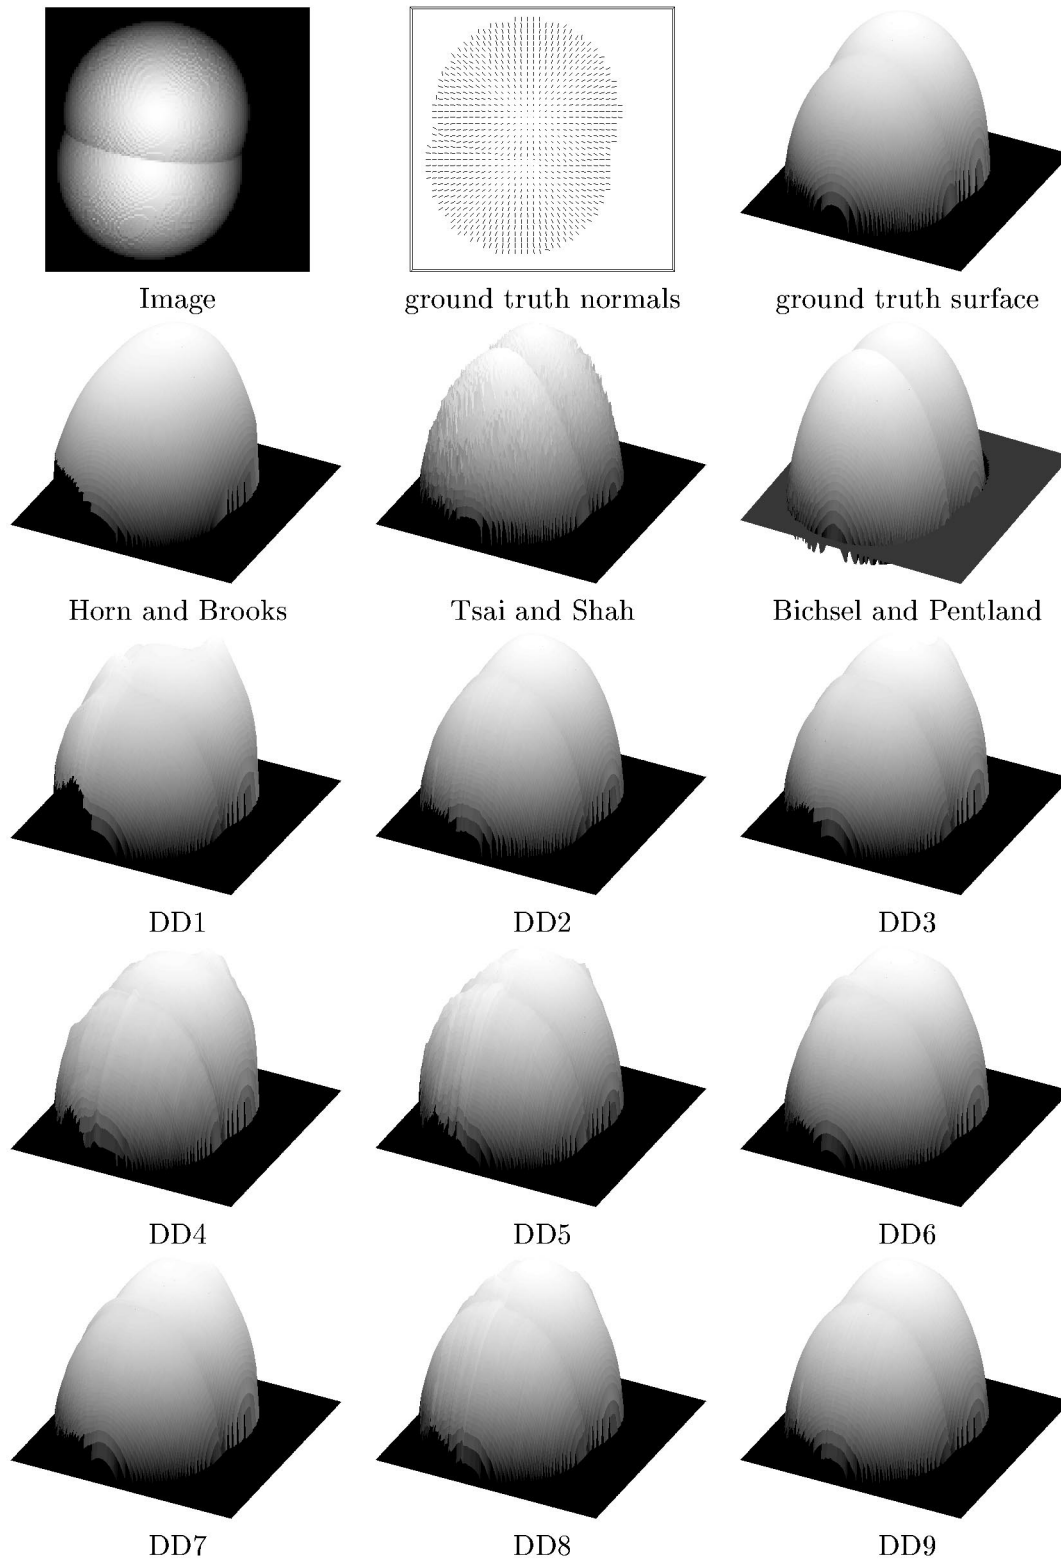


Fig. 5. Surface reconstruction results for a synthetic image of two co-joined spheres.

needle-map basis (as in Figs. 6 and 7), the disadvantage is transferred to these techniques.

Furthermore, to maintain a level playing field for comparisons, we initialize the Horn and Brooks algorithm

using the new initialization of Section 3.2, unless explicitly stated otherwise.

Qualitatively, the most interesting point of comparison between methods is the degree to which junction detail is preserved. The ability to recover discontinuities in the




Fig. 6. Comparison of the various Horn and Brooks-based techniques with Tsai and Shah and Bichsel and Pentland algorithms and of the validity of the new framework and initialization. Top row: Synthetic input images. Second row: Results of applying Bichsel and Pentland algorithm. Third row: Tsai and Shah algorithm. Fourth row: Traditional Horn and Brooks algorithm with occluding boundary initialization. Fifth row: Horn and Brooks algorithm using new, gradient-based initialization. Bottom row: Results of using the smoothness constraint of Horn and Brooks within the new, data-driven framework (DD1).




Fig. 7. Comparison of differences between the recovered needle-maps and ground truth data for the results of Fig. 6. Top row: Synthetic input images. Second row: Difference between ground truth and normals recovered using Bichsel and Pentland algorithm. Third row: Tsai and Shah algorithm. Fourth row: Traditional Horn and Brooks algorithm. Fifth row: Horn and Brooks with new initialization. Bottom row: Horn and Brooks smoothness constraint within the new, data-driven framework (DD1).




Fig. 8. Surface normal error for synthetic images, using DD1, the traditional Horn and Brooks algorithm, and Horn and Brooks with the new initialization.

surface is important given that our primary motivations for developing new constraint functions is to address the problem of oversmoothing. For instance, in Fig. 5, the important feature is the curve defined by the intersection of the two spheres. The Horn and Brooks algorithm loses this detail, while Tsai and Shah is noisy. Bichsel and Pentland preserves the junction well, but distorts the spheres a little. Of the new constraint functions, DD2 (robust smoothness error), DD5 (robust use of curvature consistency with adaptive kernel width), and, particularly, DD9 (robust use of gradient consistency with adaptive kernel width) appear to perform well. Empirically, it appears that the robust kernel approach is more effective than the weighted mean process. Also, the use of gradient consistency seems to yield slightly better results than curvature consistency.

We now turn to the needle-maps delivered by the different algorithms. In Fig. 6, we compare the needle-maps recovered for three objects. These are a pair of co-joined cones and a sphere impaled on an ellipsoid. We compare the new framework with the traditional Horn and Brooks algorithm. For simplicity, we use a quadratic constraint function (i.e., DD1) rather than one of the more complicated curvature consistency or gradient consistency constraints. Also included in the comparison are the Tsai and Shah and Bichsel and Pentland algorithms. For completeness, we investigate the effect of using the Horn and Brooks algorithm with the new initialization described in Section 3.2.

We note that the Tsai and Shah algorithm results in a noisy needle-map despite its relatively good surface recovery, demonstrated in Fig. 5. The traditional Horn and Brooks algorithm, using the occluding boundary initialization, yields needle-maps which lack detail due to oversmoothing, whereas, with the new initialization, it recovers more structure. The new framework captures much more detail, especially around the discontinuity in the cone object. This extra detail is clear, even when using the simple smoothness constraint of Horn and Brooks, although it is also obvious that the quadratic smoothness constraint, i.e., DD1, introduces some errors.

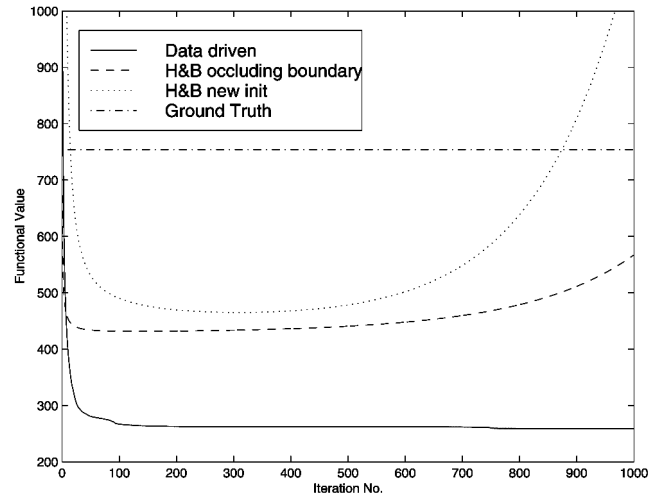


Fig. 9. Constraint functional for synthetic images, using DD1, the traditional Horn and Brooks algorithm, and Horn and Brooks with the new initialization.

5.2 Performance Measures

In this section, we consider the iterative properties of the new shape-from-shading scheme.

5.2.1 Accuracy and Convergence

Figs. 8 and 9 provide an empirical comparison of the normal-error and constraint functional values. We compare for the SFS algorithms which use the Horn and Brooks quadratic smoothness constraint. Fig. 8 is the most important since it shows how the normal error varies with iteration number. This plot demonstrates that the data-closeness constraint gives the smallest normal error. The standard Horn and Brooks algorithm starts from a large initial error and converges slowly. The new initialization process improves the starting point, but convergence is still much slower than the new shape-from-shading method. However, it is worth noting that the minimum normal error is still relatively high, of the order of 0.3 radians per normal on average, or 20° .

Fig. 9 shows the value of the functional used in each case, which is a measure often used in judging convergence in the absence of ground truth data. However, we see that not only do the functional values not possess distinct minima in some cases, but there is little correspondence between the functional value plots and the normal error plot of Fig. 8. Although they begin with the same initialization, the new framework reduces the functional value far below the minimum value achieved by the modified Horn and Brooks algorithm; indeed, even the traditional Horn and Brooks algorithm performs better, in terms of minimizing the constraint functional, than the modified version using gradient initialization. This discrepancy is due to the fact that the occluding boundary initialization yields zero smoothness error everywhere except at the boundary, where the smoothness error is large. The traditional Horn and Brooks algorithm then propagates this boundary error into the object interior, dissipating it along the way. Starting with the new initialization gives nonzero smoothness error in the interior of the object, so the propagation/dissipation




Fig. 10. Comparison of the behavior of various measures for Horn and Brooks, the robust regularizer algorithm of [33], and the new framework using different constraint functions. In each case, the scales of the different measures have been scaled to give a maximum value of unity. Comparison of the form and location of the minima of each measure allows us to qualitatively consider questions of termination (see text).

effect is reduced and the overall smoothness error can only be reduced by reorganization of normals about their respective cones of ambiguity.

These plots lead us to two conclusions. First, the Horn and Brooks smoothness functional is not an appropriate constraint upon the recovery of shape from shading.

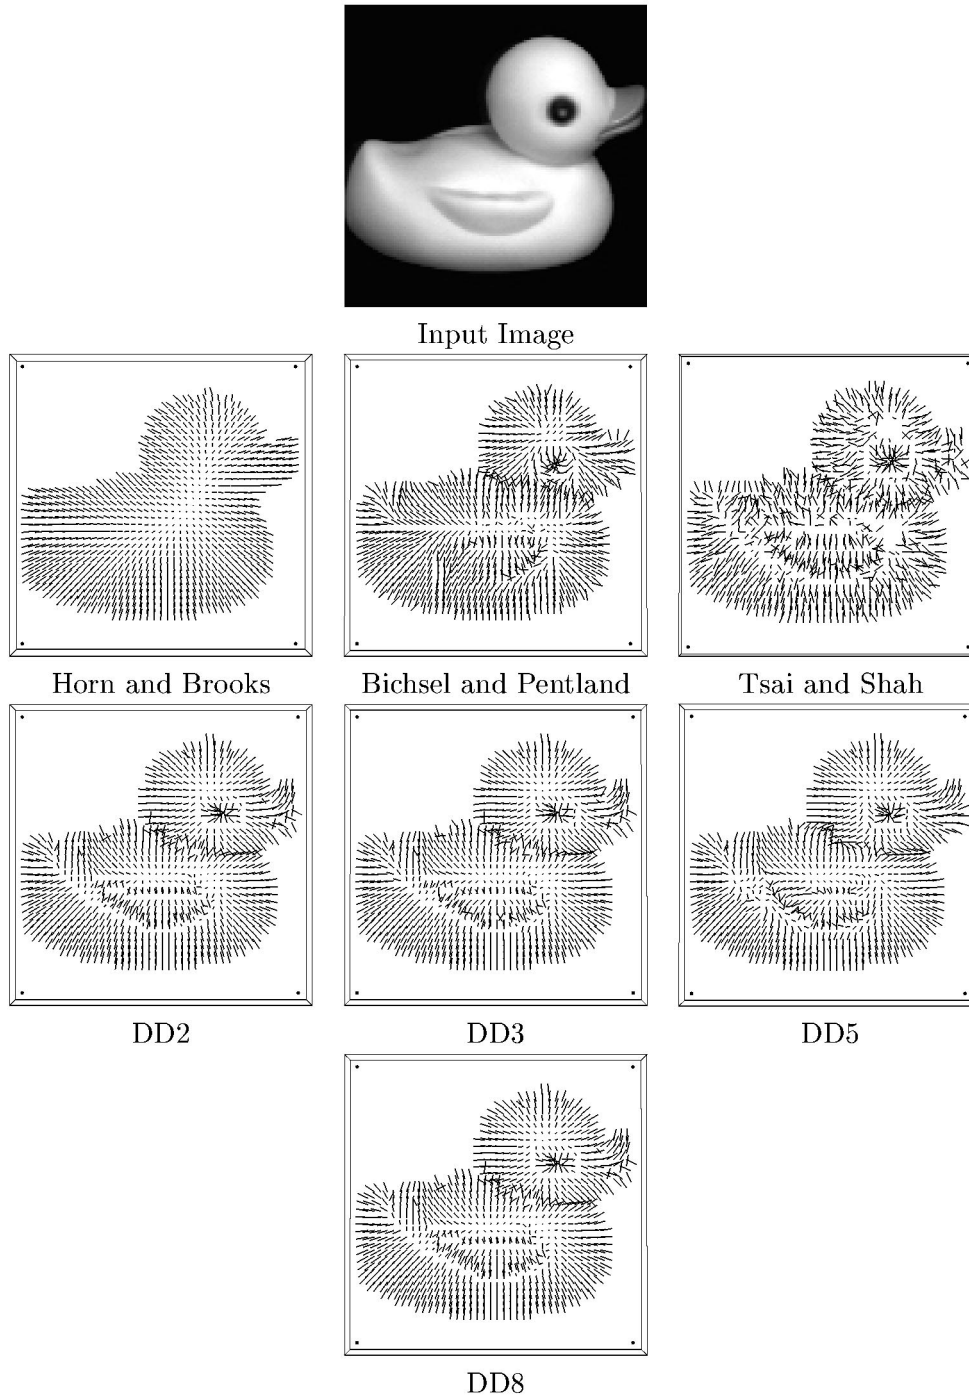


Fig. 11. Needle-maps recovered from toy duck image.

Second, even when using this constraint, the new framework succeeds in minimizing the constraint more effectively than the traditional Horn and Brooks algorithm.

5.2.2 Algorithm Comparison

We now turn our attention to the iterative behavior of the different algorithms listed in Table 1. The results are summarized in Fig. 10. Several quantities are plotted against iteration number for the Horn and Brooks algorithm, the Robust Regularizer approach of [33], and each of

the new framework schemes. In each case, the new initialization of Section 3.2 is used. The measures used are the normal-error, the surface reconstruction error, and the value of the constraint is functional.

Essentially, there are two primary criteria for a suitable SFS constraint function. The first is that it leads to a low minimum value of normal error, while the second, often overlooked, criterion is that there must be some method to stop the algorithm. Ideally, we wish termination to occur as closely as possible to the minimum of the normal error, despite the fact that we will not ordinarily know at which

iteration this will occur. In the absence of ground truth, the obvious termination criterion is the minimization of the functional. However, it is clear from Fig. 10 that the minimization of the constraint functional is not an appropriate stopping criterion. Neither, it appears, is the average change in normals between iterations.

However, it is worth noting that, in the cases of most of the new framework constraint functions, the termination point is not critical since the normal error has only a shallow minimum or continues to fall for many iterations at a diminishing rate. Only curvature consistency measured using the median absolute deviation, i.e., DD4, breaks this pattern, exhibiting strange behavior wherein the normal error increases after initialization. The poor performance of DD4, which walks away from the initial needle-map estimate, is explained by the nonlinear median process incorporated in it, which appears to make it very susceptible to local noise. This effect introduces orientation errors at scales too small to be obvious in the surface reconstruction of Fig. 5.

In the remaining cases, 100-200 iterations usually results in a good approximation to the minimum of the normal error. When the curvature and gradient consistency schemes DD3, DD5, and DD9 are used, the minimum of the functional value and the behavior of the curves suggest that using the minimization of the functional as a termination criterion may be adequate.

Since the normal error is an absolute measure unaffected by which constraint function is used and as all schemes use the same initialization and, so, start with identical values of normal error, we can compare these normalized values directly. Doing so, we see that when the normal smoothness and gradient consistency constraints are modeled using a robust error kernel, i.e., schemes DD2 and DD6, are used, then there is a reduction in the normal error by approximately 57 percent of the initial value. This value is lower than the normal error achieved by the traditional Horn and Brooks algorithm after 1,000 iterations. Hence, use of a robust kernel appears more appropriate than a weighted mean process and gradient consistency performs better, in this case, than curvature consistency.

5.3 Real World Data

In Fig. 11, we illustrate the application of several of the SFS schemes considered in this paper to real data. The image is from the Columbia Object Image Library and contains real world imperfections such as albedo changes, regions of brightness saturation, and deviations from the Lambertian assumption. Only four of the new constraint functions are used, one each from the surface geometric and gradient consistency categories and two curvature consistency techniques. With the exception of DD3, they are all based upon the use of the robust kernel approach.

In Fig. 11, the needle-map recovered by the Horn and Brooks algorithm is seen to be extremely smooth, even losing large-scale details such as the wing of the toy duck. Both Bichsel and Pentland and Tsai and Shah algorithms yield noisy needle-maps when these are calculated from the recovered surfaces. The new framework schemes all perform well, recovering the wing while behaving well over the smooth regions of the surface.

6 CONCLUSIONS

We have presented a novel framework for the shape-from-shading problem which addresses the problem of model dominance encountered using many traditional global algorithms. Specifically, we ensure the fullest use of the data available, namely the input image, by incorporating the image irradiance equation as a hard constraint.

Our second contribution is to develop novel constraints on needle-map consistency. These range from simple modifications to the traditional smoothness constraint, to topographic curvature and gradient consistency. Using extensive experimentation upon synthetic and real shaded images, we have demonstrated that several of these novel constraints perform significantly better than various existing algorithms. Specifically, we find that incorporating an adaptive robust kernel to control local smoothing yields excellent results, especially when used in conjunction with gradient consistency constraints.


Gradient consistency makes only relatively weak assumptions about the surface physics. This is in contrast to the curvature consistency techniques developed here as a result of modeling the topographic properties of surfaces. The slight advantage of the gradient consistency schemes over the more complicated curvature consistency methods suggests that we can obtain good results without recourse to complex arguments about the behavior of real surfaces.

The in depth analysis of the SFS problem presented here explicitly exposes the inadequacy of the traditional quadratic smoothness constraint to the task of SFS recovery, since finding the smoothest possible surface is demonstrably not an appropriate goal. With the new framework, we have demonstrated the development of principled, novel constraint functions modeling the properties of real surfaces. It is hoped that this research will stimulate the search for still better constraints, possessing global minima which more closely correspond to true surfaces.

REFERENCES

- [1] E. Angelopoulou, "Gaussian Curvature from Photometric Scatter Plots," *Proc. IEEE Workshop Photometric Modeling for Computer Vision and Graphics*, 1999.
- [2] P.N. Belhumeur and D.J. Kriegman, "What is the Set of Images of an Object Under All Possible Lighting Conditions?" *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 270-277, 1996.
- [3] M. Bichsel and A.P. Pentland, "A Simple Algorithm for Shape from Shading," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 459-465, 1992.
- [4] M.J. Brooks and B.K.P. Horn, "Shape and Source from Shading," *Proc. Int'l Joint Conf. Artificial Intelligence*, pp. 932-936, 1985.
- [5] A.M. Bruckstein, "On Shape from Shading," *Computer Vision, Graphics, and Image Processing*, vol. 44, pp. 139-154, 1988.
- [6] F.P. Ferrie and J. Lagarde, "Curvature Consistency Improves Local Shading Analysis," *Proc. IEEE Int'l Conf. Pattern Recognition*, vol. I, pp. 70-76, 1990.
- [7] D. Hilbert and R. Courant, *Methods of Math. Physics*. 1953.
- [8] D.C. Hoaglin, F. Mosteller, and J.W. Tukey, eds., *Understanding Robust and Exploratory Data Analysis*. Wiley, New York.: 1983.
- [9] B.K.P. Horn, "Obtaining Shape from Shading Information," *The Psychology of Computer Vision*, P.H. Winston, ed., McGraw Hill, New York.: pp. 115-155, 1975.
- [10] B.K.P. Horn and M.J. Brooks, "The Variational Approach to Shape from Shading," *Computer Vision, Graphics, and Image Processing*, vol. 33, no. 2, pp. 174-208, 1986.
- [11] *Shape from Shading*, B.K.P. Horn and M.J. Brooks, eds. Cambridge, Mass.: MIT Press, 1989.

- [12] B.K.P. Horn, "Height and Gradient from Shading," *Proc. Int'l Joint Conf. Artificial Intelligence*, vol. 5, no. 1, pp. 37-75, 1990.
- [13] P. Huber, *Robust Statistics*. Chichester: Wiley, 1981.
- [14] K. Ikeuchi and B.K.P. Horn, "Numerical Shape from Shading and Occluding Boundaries," *Artificial Intelligence*, vol. 17, no. 3, pp. 141-184, 1981.
- [15] R. Kimmel and A.M. Bruckstein, "Tracking Level Sets by Level Sets: A Method for Solving the Shape from Shading Problem," *Computer Vision and Image Understanding*, vol. 62, no. 1, pp. 47-58, 1995.
- [16] J.J. Koenderink and A.J. van Doorn, "Surface Shape and Curvature Scales," *Image and Vision Computing*, vol. 10, pp. 557-565, 1992.
- [17] J.J. Koenderink and A.J. van Doorn, "Relief Pictorial and Otherwise," *Image and Vision Computing*, vol. 13, no. 5, pp. 321-334, 1995.
- [18] J.J. Koenderink, A.J. van Doorn, C. Christou, and J.S. Lappin, "Perturbation Study of Shading in Pictures," *Perception*, vol. 25, no. 9, pp. 1,009-1,026, 1996.
- [19] K.M. Lee and C.C.J. Kuo, "Shape from Shading with a Linear Triangular Element Surface Model," *IEEE Pattern Analysis and Machine*, vol. 15, no. 8, pp. 815-822, 1993.
- [20] S.Z. Li, "Discontinuous MRF Prior and Robust Statistics: A Comparative Study," *Image and Vision Computing*, vol. 13, no. 3, pp. 227-233, 1995.
- [21] D.C. Marr, *Vision*. San Francisco, Calif.: Freeman, 1982.
- [22] J. Oliensis, "Shape from Shading as a Partially Well-Constrained Problem," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 559-564, 1991.
- [23] J. Oliensis and P. Dupuis, "A Global Algorithm for Shape from Shading," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 692-701, 1993.
- [24] P. Parent and S.W. Zucker, "Curvature Consistency and Curve Detection," *J. Opt. Soc. America A*, vol. 2, no. 13, p. 5, 1985.
- [25] A.P. Pentland, "Local Shading Analysis," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 6, pp. 170-187, 1984.
- [26] A.P. Pentland, "Finding the Illuminant Direction," *J. Optical Society of America*, pp. 448-455, 1982.
- [27] T. Poggio, V. Torre, and C. Koch, "Computational Vision and Regularization Theory," *Nature*, vol. 317, no. 6,035, pp. 314-319, 1985.
- [28] P.S. Tsai and M. Shah, "Shape from Shading Using Linear Approximation," *Image and Vision Computing*, vol. 12, no. 8, pp. 487-498, 1994.
- [29] J.A. Sethian, *Level Set Methods*. Cambridge Univ. Press 1996.
- [30] R. Szeliski, "Fast Shape from Shading," *Computer Vision, Graphics, and Image Processing: Image Understanding*, vol. 53, pp. 129-153, 1991.
- [31] P.L. Worthington and E.R. Hancock, "Shape-from-Shading using Robust Statistics," *Proc. IEEE Int'l Conf. Digital Signal Processing*, 1997.
- [32] P.L. Worthington and E.R. Hancock, "Needle Map Recovery Using Robust Regularizers," *Proc. British Machine Vision Conf., BMVA Press*, vol. I, pp. 31-40, 1997.
- [33] P.L. Worthington and E.R. Hancock, "Needle Map Recovery using Robust Regularizers," *Image and Vision Computing*, vol. 17, no. 8, pp. 545-558, 1999.
- [34] P.L. Worthington and E.R. Hancock, "Data Driven Shape-from-Shading using Curvature Consistency," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. I, pp. 287-292, 1999.
- [35] P.L. Worthington and E.R. Hancock, "Surface Topography from Intensity Images," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 911-917, 1999.
- [36] R. Zhang, P.S. Tsai, J.E. Cryer, and M. Shah, "Analysis of Shape from Shading Techniques," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 377-384, 1994.
- [37] Q. Zheng and R. Chellappa, "Estimation of Illuminant Direction, Albedo, and Shape from Shading," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 680-702, 1991.



Philip L. Worthington received his BA degree in engineering and computing science in 1996 from the University of Oxford, United Kingdom. Since 1996 he has worked toward a DPhil degree in Computer Vision at the University of York, under the supervision of Professor Edwin Hancock. In 1998, he visited NTT Basic Research Laboratories in Japan as a JISTEC/ British Council Young Researcher, working with Dr. Hiroshi Murase on using shape-from-shading in eigenspace-based object recognition. In 1999, he received the University of York Gibbs-Plessey Award to visit United States academic and commercial computer vision laboratories. He is now employed as an engineer at Siemens Roke Manor Research Centre in Hampshire, U.K.



Edwin R. Hancock gained his BSc degree in physics in 1977 and PhD degree in high energy nuclear physics in 1981, both from the University of Durham, United Kingdom. Between 1981 and 1991, he undertook postdoctoral research at the Rutherford Appleton Laboratory. Initially he worked on charm-photo-production experiments performed at the Stanford Linear Accelerator Centre. However, in 1985 he moved into the fields of computer vision and pattern recognition. Between 1982 and 1991, he also held adjunct posts at the Open University and the University of Surrey. He moved to the University of York as a lecturer in computer science in July 1991. In July 1987, he was promoted to senior lecturer. He was further promoted to Reader in July 1998. In December 1998, he was appointed to a chair in computer vision. Professor Hancock leads a group of some 15 faculty and researchers in the areas of computer vision and pattern recognition. He has published about 200 refereed papers in the fields of high energy nuclear physics, computer vision, image processing, and pattern recognition. He was awarded the 1990 Pattern Recognition Society Medal and received an outstanding paper award in 1997. He serves as an associate editor of the journal *Pattern Recognition* and has been a guest editor for the *Image and Vision Computing* journal. He is also guest editor for a forthcoming special edition of the *Pattern Recognition* journal devoted to energy minimization methods in computer vision and pattern recognition. He chaired the 1994 British Machine Vision Conference and has been a program committee member for several national and international conferences.