Shape from Specular Reflection and Optical Flow

J. Lellmann · J. Balzer · A. Rieder · J. Beyerer

Abstract Inferring scene geometry from a sequence of camera images is one of the central problems in computer vision. While the overwhelming majority of related research focuses on diffuse surface models, there are cases when this is not a viable assumption: in many industrial applications, one has to deal with metal or coated surfaces exhibiting a strong specular behavior. We propose a novel and generalized constrained gradient descent method to determine the shape of a purely specular object from the reflection of a calibrated scene and additional data required to find a unique solution. This data is exemplarily provided by optical flow measurements obtained by small scale motion of the specular object, with camera and scene remaining stationary. We present a non-approximative general forward model to predict the optical flow of specular surfaces, covering rigid body motion as well as elastic deformation, and allowing for a characterization of problematic points. We demonstrate the applicability of our method by numerical experiments on synthetic and real data.

1 Introduction and Previous Work

We consider the problem of reconstructing a free-form mirror surface by observing the reflected image of a calibrated scene. In a highly controllable industrial environment, one observes via the mirror surface a LC display, whose pixels identify themselves in the camera picture by a time-coded sequence of intensity values (see Fig. 1). The established correspondence serves in finding the surface that has lead to the recorded image series. This method is known as Shape from Specular Reflection or—in the field of optical metrology—as deflectometry, see Kammel (2004), Bonfort (2006) for a comprehensive introduction. It is widely used because it allows detection of small irregularities of the surface under inspection. Real-world applications include high resolution scanning of optical lenses and industrial quality control of coated components. However, as with classical camera-based scene reconstruction, a single measurement generally does not suffice to fully determine the surface structure.

Current literature contains some suggestions on how to overcome the ambiguity. Solem et al. investigate the problem in a variational setting, cf. Solem et al. (2004). Assuming a set of surface points is known, they propose an iterative algorithm to minimize an energy functional consisting of surface normal and point constraints. Kickingereder and Donner (2004) and Bonfort and Sturm (2006) use stereo vision to uniquely recover the surface. In Balzer et al. (2006), the Lambertian behavior of a certain class of objects is employed. The latter approach is closely connected to a re-

J. Lellmann · J. Balzer (✉) · J. Beyerer
Lehrstuhl für Interaktive Echtzeitsysteme, Institut für Technische Informatik (ITEC), Department of Computer Science, University of Karlsruhe, 76128 Karlsruhe, Germany
e-mail: jonathan.balzer@ies.uni-karlsruhe.de

J. Lellmann
e-mail: lellmann@math.uni-heidelberg.de

J. Beyerer
e-mail: juergen.beyerer@iitb.fraunhofer.de

A. Rieder
Institut für Angewandte und Numerische Mathematik, Department of Mathematics, University of Karlsruhe, 76128 Karlsruhe, Germany
e-mail: andreas.rieder@math.uni-karlsruhe.de

proaches are justified by good reconstruction results. However, enforcing uniqueness by means of approximation, accuracy must be compromised to a certain degree. In particular, it has not been checked analytically under which circumstances the approximation error remains within the specified bounds.

In an early predecessor of our work, Oren and Nayar (1995) investigate an abstract setting where only the camera is moving. They do not rely on the classical monocular stereo approach but use a differential multiview method to derive a family of possible surface curves. To our knowledge, the first authors incorporating optical flow measurements were Roth and Black (2006). Their model requires a large distance between the mirror and the reflected scene, leaving room for a more realistic small-scale generalization. This will be one of our main contributions.

This paper is outlined as follows: in the following section, we will state the problem, provide some necessary background in level set theory, and finally show how to formulate the reconstruction problem as a constrained optimization problem. In Sect. 3, we will then derive the governing equations of specular optical flow. We will see that a flow vector field fails to be well-defined under certain circumstances. Our model provides for the characterization of such special points. In Sect. 4, we propose a gradient descent method to solve the general optimization problem of reconstructing the shape of an unknown specular surface in combination with optical flow measurements. Our numerical procedure is sketched and experimental results are presented to underline the practicability of our method.

2 Problem Statement and Mathematical Framework

2.1 Notation and Basic Problem

We frequently use the Euclidean scalar product $\langle \cdot, \cdot \rangle$ which induces the norm $\| \cdot \|$ on $\mathbb{R}^3$. Vectors are distinguished by bold letters $\boldsymbol{x} = (x_1 \ x_2 \ x_3)^\top$, matrices by capitalization $\mathbf{M}$. For example, $\mathbf{I}$ is the identity matrix and $\mathbf{J}_{\boldsymbol{v}}$ the Jacobian of a continuously differentiable vector field $\boldsymbol{v}$. Normalization to unit length is denoted by an additional hat: $\hat{\boldsymbol{x}} = \pi(\boldsymbol{x})$, where $\pi : \mathbb{R}^3 \setminus \{0\} \to S^2$, $\boldsymbol{x} \mapsto \pi(\boldsymbol{x}) := \frac{\boldsymbol{x}}{\|\boldsymbol{x}\|}$, is the projection to the unit sphere S^2. The tangent map of π can be described by the Jacobian of its continuation on $\mathbb{R}^3$: $\mathbf{J}_\pi(\boldsymbol{x}) = \frac{1}{\|\boldsymbol{x}\|} \mathbf{P}_{\boldsymbol{x}}$, where $\mathbf{P}_{\boldsymbol{x}} = \mathbf{I} - \hat{\boldsymbol{x}}\hat{\boldsymbol{x}}^\top$ denotes the orthogonal projection onto $\hat{\boldsymbol{x}}^\perp$ in matrix form.

Figure 2 illustrates the basic setup. We assume a single-viewpoint camera with optical center located in the origin $\boldsymbol{o}$ of the coordinate system. For simplicity, we will generally use a standard pinhole camera model with focal length $f = 1$. The pinhole camera is modelled by the function $\Pi : \Omega \to \Omega_I$, projecting each point $\boldsymbol{x}$ in a subset $\Omega \subseteq \mathbb{R}^3$

Fig. 1 Compact, robot-based sensor head for recording of light map and optical flow data

search area known as Shape from Shading (Horn 1970; Prados et al. 2007). Other methods are based on *regularization by approximation* or *locality*: on the basis of reflected differential geometric entities such as points and curves, Savarese et al. (2005) estimate the Monge form of open neighborhoods around a set of surface points from a single image. Neglecting higher-order surface properties amounts to reducing the problem to a finite-dimensional parameter estimation task, whose outcome will most likely be a unique solution under practical circumstances. Rozenfeld et al. (2007) present a significant enhancement towards industrial application. They notice that for a fixed point tangential directions of feature curves and their images are linked via a 1D homography, containing all desired information on local surface shape (up to a certain order). The approximative nature of the work by Halstead et al. (1996) becomes directly obvious, as the correspondences between scene and image points are only roughly estimated. Other than that, the authors rely on decoupling computation of surface slope and position. This may lead to significant systematic errors, as the problem is severely non-linear, especially in view of the discussion at the end of Sect. 2.1. In particular, a smoothness assumption can *not* resolve the ambiguity. Consequently, the algorithm presented appears to be limited to their application in *videokeratography*.[1] In summary, all of the local ap-

---

[1] Measurement of the human cornea by deflectometric techniques.

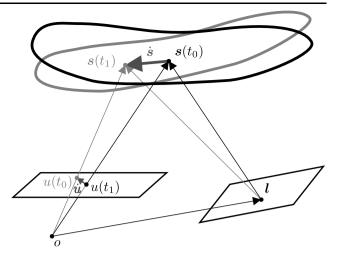**Fig. 2** Basic setup. The scene L is reflected in the specular surface S and observed by the camera in the origin

Fig. 3 Specular optical flow. The reflection point s varies as the specular surface moves relative to camera and diffuse scene

of the field of view[2] onto some u in the image plane $x_3 = 1$, i.e. $u = \Pi(x) = \frac{x}{x_3}$.

Intersecting the ray emanating from the origin through an image point $u \in \Omega_I$ with S, we get the *reflection point* $s(u)$ on S. We generally require that the *surface map* s be well-defined on all of Ω_I, i.e. the surface occupies the whole image. If S were diffuse, the physical image intensity in u would be completely determined by the texture in $s(u)$. But since S is specular, the ray is reflected and advances until it hits the scene surface L in $l(u)$. We denote the reflected ray by r. We also assume that there are no multiple reflections and the *light map* $l : \Omega_I \to L$ is defined on all of Ω_I. We require S to be sufficiently smooth and Ω to be a bounded domain with piecewise smooth boundary $\partial\Omega$. Now, we can formulate the basic problem in an abstract way:

Problem formulation 1 Given the light map $l(u)$, find a corresponding surface map $s(u)$.

The law of reflection connects measurement and surface: supposing the surface ran through an arbitrary point x in the domain Ω, having looked up $l(u)$ for the projection $u = \Pi(x)$ of x onto the image plane, the normal would be

$$\hat{m}(x) = -\frac{\hat{x} + \hat{r}}{\|\hat{x} + \hat{r}\|}, \qquad (2.1)$$

where $r = x - l$. Observe that the normals are forced downward here, i.e. $\langle \hat{m}, x \rangle < 0$. The vector field $\hat{m} : \Omega \to S^2$ restricts the normals of S: if for all $s \in S$

$$\hat{n}(s) = \hat{m}(s) \qquad (2.2)$$

the surface is said to satisfy the *normal condition*. This leads to an equivalent formulation:

Problem formulation 2 Given the normal field $\hat{m}$, find a surface map s satisfying the normal condition (2.2).

Despite its elegance, this way of viewing the problem has not been widely considered. Notable exceptions are the works of Bonfort and Sturm (2006) as well as Hicks and Perline (2004), who also provide a generalized formulation using so-called *m*-distributions.

As the vector field formulation suggests, the problem in its current form does not necessarily admit a unique solution: in Kickingereder and Donner (2004) for example, an adept parametrization transforms (2.2) into a first-order PDE system exhibiting an infinite number of solutions if no initial value is given. The authors state nothing particular about the structure of the solution space, though. The model due to Werling et al. (2007) consists of a system of two quasilinear partial differential equations of first order, whose characteristics are orthogonal to each other in every point of Ω. In fact, it can be shown that—under certain assumptions on the normal field $\hat{m}$—the manifold of solutions to the reconstruction problem is a smooth one-parameter family or "regular curve", parameterizable by the intersection of S with a ray through the origin. It is evidential that in the case where a single surface point is known, a unique solution can be obtained.

2.2 Diffuse Versus Specular Optical Flow

The objective of this paper is to overcome the ambiguity addressed in the last section by measuring how slight displacement of the sought-after surface modulates the light map (or equivalently its inverse). Thus, let us first review the notion of *optical flow*, which lives in the image plane independent of the reflection behavior of the dynamic object. Consider a scene surface containing a fixed point x which is observed at

---

[2]The field of view is the set of all points $x \in \mathbb{R}^3$, $x_3 > 0$, which can be connected to the origin by a straight line intersecting the domain of the image signal Ω_I.

image coordinates $\boldsymbol{u}(t)$ for $t \in T$ with $T = [t_0, t_1]$ denoting a time interval. The optical flow for $\boldsymbol{u} = \boldsymbol{u}(t_0)$ is then defined as the temporal derivative $\dot{\boldsymbol{u}}(t_0)$. Its prominent advantage is that it can be estimated solely from the image sequence if the scene provides distinctive texture and the temporal resolution is sufficiently high (Barron et al. 1994). Being a local property, the optical flow effectively adds information to a *single* image, which allows us to pointwise infer scene information much more easily than with genuine stereo (Horn 1986; Zissermann and Giblin 1989).

Consider a fixed camera observing a dynamic scene, and let $\boldsymbol{w}(\boldsymbol{x})$ denote the velocity vector for each scene point $\boldsymbol{x}$ at time t_0. Note that this includes uniform Euclidean motion, where

$$\boldsymbol{w}(\boldsymbol{x}) = \dot{\boldsymbol{x}}(t_0) = \boldsymbol{\Omega}\boldsymbol{x} + \boldsymbol{v}, \qquad (2.3)$$

and $\boldsymbol{\Omega} = \boldsymbol{\Omega}(t_0)$ is a skew-symmetric matrix s.t. $\boldsymbol{x}(t) = \mathbf{R}(t)\boldsymbol{x} + \boldsymbol{y}(t)$, $\mathbf{R}(t) = e^{\boldsymbol{\Omega}(t)} \in \mathrm{SO}(3)$ and $\boldsymbol{v} = \dot{\boldsymbol{y}}(t_0)$. However, this *velocity field* formulation allows far more complex and possibly non-rigid motion, e.g. waves on a water surface. For a point $\boldsymbol{s}$ on a *diffuse* surface S, the optical flow is given by Horn (1986)

$$\dot{\boldsymbol{u}} = \frac{d}{dt}\Pi(\boldsymbol{s}(t)) = \mathbf{J}_\Pi(\boldsymbol{s})\dot{\boldsymbol{s}}$$

$$= \mathbf{J}_\Pi(\boldsymbol{s})\boldsymbol{w}(\boldsymbol{s}) = \frac{1}{s_3}\begin{pmatrix} 1 & 0 & -u_1 \\ 0 & 1 & -u_2 \\ 0 & 0 & 0 \end{pmatrix}\boldsymbol{w}(\boldsymbol{s}). \qquad (2.4)$$

In the simple case of purely translational motion, i.e. $\boldsymbol{\Omega} = 0$, we have $\boldsymbol{w}(\boldsymbol{s}) \equiv \boldsymbol{v}$ and can easily deduce s_3—and thus $\boldsymbol{s}$—from known $\boldsymbol{u}$, $\boldsymbol{v}$, and $\dot{\boldsymbol{u}}$ (see Hartley and Zisserman 2003 for a survey).

If S is *specular*, the process is far more involved since $\boldsymbol{u}$ is now the projection of the reflection point $\boldsymbol{s}$ of a fixed scene point $\boldsymbol{l}$. Generally, $\boldsymbol{s}$ will move on the surface in time: not the same *surface point* is "seen" although the observed *scene point* does not change (see Fig. 3). This compensating motion superposed by the enforced object trajectory makes up the reflection point trajectory $\boldsymbol{s}(t)$.

Specularities have been widely regarded as a source of error in diffuse surface reconstruction, with many authors proposing methods to detect and discard the respective flow vectors. Roth and Black (2006) recently suggested to take advantage of this extra information while resolving the ambiguities arising in the Shape from Specularities problem. Their model is a direct adoption of what has been described by Chen and Arvo (2000), who perform a thorough differential analysis for the problem of recovering reflection points motion when only the camera is dislocated. Since in their setting, the light map has not been made available by encoding display element locations, they require a large distance

between $\boldsymbol{s}$ and $\boldsymbol{l}$, leaving the reflected ray $\boldsymbol{r}$ in (2.1) approximately constant.

For our metrology setup, this simplification is too restrictive. Firstly, although measurement sensitivity with respect to slope grows quadratically in object distance, we are forced to position the mirror object in proximity of the optical sensor because sampling density on the surface will decrease concurrently. The assumption $\boldsymbol{r} \approx const.$ remains no longer valid. With perfect knowledge of the light map, however, computation of the reflected ray for every surface point is possible, adding further complexity to the modelling task tackled in Sect. 3. A second difference lies in the fact that we consider the case where everything but the mirror surface remains stationary, i.e. camera and diffuse target are rigidly mounted on a compact sensor head which moves relative to the unknown specular object, see Fig. 1. Finally, as the model described in Roth and Black (2006) relies on the truncated Taylor series of the path function, the computation of reflection point trajectories (specifically its tangent vectors) must be approximative in most cases. Since we will directly adopt a differential viewpoint, our model may—under certain circumstances—admit a closed-form analytical solution in search for the tangents to the reflection point trajectories.

2.3 Level Sets and Differential Geometry

The surfaces under concern are assumed to be regular for the rest of the paper, i.e. they can be described by a set $S \subset \mathbb{R}^3$ and an atlas of C^k-diffeomorphisms $h_{\boldsymbol{s}}$, $k \geq 1$, s.t. for every point $\boldsymbol{s} \in S$, there exists a neighborhood $\tilde{U} \subset \mathbb{R}^3$ of $\boldsymbol{s}$ with $h_{\boldsymbol{s}}$ mapping $\tilde{U} \cap S$ to an open set $U \subset \mathbb{R}^2$ (do Carmo 1976). Some inherent drawbacks of this representation, namely the possible need for reparametrization and the strong dependence on the dimension of the underlying space, can be circumvented by using an implicit—or *level set*—representation made popular by Osher and Fedkiw (2002) and Sethian (2005). Here, we express S as the zero set of a function $\varphi \in C^2(V, \mathbb{R})$ s.t. $S \cap V = \varphi^{-1}(\{0\})$. The level set function φ defines a regular surface if $\nabla\varphi \neq 0$ on $\varphi^{-1}(\{0\})$, which we assume throughout.

While a regular surface locally allows for such a representation (for any local graph representation $f : U \subset \mathbb{R}^2 \to \Omega$, set $\varphi(x, y, z) = f(x, y) - z$), in general, no global level set function $\varphi \in C^2(\Omega, \mathbb{R})$ exists. The continuity assumptions on φ enforce orientability, which is why only regular surfaces of the form $S = \partial W$ are allowed with some open $W \subset \mathbb{R}^3$, s.t. $\varphi < 0$ in W and $\varphi > 0$ in $\Omega \backslash W$. We will assume S to satisfy the above requirements.

The following inner geometric quantities are frequently required:

$$\hat{\boldsymbol{n}} = \frac{\nabla\varphi}{\|\nabla\varphi\|}, \qquad \mathbf{K} = -\frac{1}{\|\nabla\varphi\|}\mathbf{P}_{\hat{\boldsymbol{n}}}\nabla^2\varphi\mathbf{P}_{\hat{\boldsymbol{n}}},$$

where $\hat{n}$ is the normal in the direction of $\varphi > 0$ and $\mathbf{K}$ is the symmetric matrix defining the second fundamental form and satisfying $\mathbf{K}\hat{n} = 0$. By $\nabla^2 \varphi = \mathbf{J}_{\nabla\varphi}$, we denote the Hessian of φ. For the mean curvature κ, we get $\kappa = -\operatorname{Tr}\mathbf{K}$ (see Jin 2003).

The *surface gradient* $D_S f(s)$ of a function $f : S \to \mathbb{R}$ is defined as the unique vector $y \in T_s S$ with $\frac{d}{d\tau}\{f(x(\tau))\}|_{\tau=0} = \langle y, \frac{d}{d\tau}x(\tau)|_{\tau=0}\rangle$ for all curves $x \subseteq S$ with $x(0) = s$. Where there is access to the gradient ∇f of a continuation of f on a neighborhood of s, we have $D_S f(s) = \mathbf{P}_{\hat{n}}\nabla f(s)$.

Moving surfaces are naturally written as

$$S(t) = \varphi^{-1}[\{0\}, t],$$

where $\varphi \in C^2(\Omega \times T, \mathbb{R})$, and each $\varphi[\cdot, t]$ itself is a level set function. We will use brackets to distinguish these level set *evolutions* from ordinary level set functions φ. For such an evolution, let $\nabla\varphi := \partial_x \varphi$ and $\dot{\varphi} := \partial_t \varphi$.

Consider a differentiable curve $s : T \to \Omega$ lying on the moving surface, i.e. $\varphi[s(t), t] \equiv 0$. Derivation w.r.t. t results in

$$\dot{\varphi}[s(t), t] = -\nabla\varphi[s(t), t]\dot{s}(t)$$
$$= -\langle \dot{s}(t), \hat{n}[s(t), t]\rangle\|\nabla\varphi[s(t), t]\|.$$

This means that the way the surface evolves depends only on the *normal velocity* field

$$v := \langle \dot{s}, \hat{n}\rangle \tag{2.5}$$

of the parametrization (Solem and Overgaard 2005a). By substituting v, we get the *level set equation*

$$\dot{\varphi} = -v\|\nabla\varphi\|. \tag{2.6}$$

As an important example, consider the special level set function

$$\varphi(x) = e_{l,x_0}(x) := c(x_0) - (\|x\| + \|x - l\|) \tag{2.7}$$

with $s, l \in \mathbb{R}^3$ and $c = c(x_0) := \|x_0\| + \|x_0 - l\|$. The zero set E_{l,x_0} of e_{l,x_0} consists of all points x for which the sum of the distances from o to l via x is the same as via x_0. In the two-dimensional case, it coincides with an ellipse with focal points o and l. Its three-dimensional counterpart is called *prolate spheroid* E_{l,x_0}, being generated by rotating such an ellipse around its major axis through o and l.

For $x \in E_{l,x_0}$, the normal computes as

$$\hat{n}_{l,x_0}(x) = \frac{\nabla e_{l,x_0}(x)}{\|\nabla e_{l,x_0}(x)\|} = -\frac{\hat{x} + \widehat{x - l}}{\|\hat{x} + \widehat{x - l}\|}.$$

Comparing this to (2.1) confirms the fact that such a spheroid will reflect *all* rays emerging from the origin into its second focal point l and vice versa. The law of reflection then reduces to the fact that the surface normal $\hat{n}(s)$ for a point $s \in S$ and the normal $\hat{n}_{l,s}(s)$ of the (unique) spheroid

through s with foci o and l coincide. Whenever S is exactly shaped like $E_{l,s}$ in a neighborhood of s (and assuming there are no occlusions), the scene point l will be visible in a neighborhood of the image point $u = \Pi(s)$. This means we cannot generally assume the injectivity of $l : \Omega_I \to L$, not even locally: features on L might be "infinitely" magnified when observed indirectly through the mirror.

2.4 Stating the Problem

We will now see how the extra information provided by the optical flow can be integrated into the model to help find the original surface. First, observe that solving (2.2) for S can be expressed as minimizing the surface functional

$$J(S) := \frac{1}{2}\int_S \|\hat{n} - \hat{m}\|^2 d\sigma,$$

where $d\sigma$ is a surface element. In level set notation, this can formally be written as

$$J(\varphi) = \frac{1}{2}\int_\Omega \|\hat{n} - \hat{m}\|^2 \|\nabla\varphi\|\delta(\varphi(x))dx,$$

where δ is the Dirac delta distribution. Our extra information shall be contained in a similar functional H. While we will illustrate the procedure only for the optical flow case, H may also include any additional measurements, e.g. from range sensors or analysis of diffuse surface parts. Here, we will use

$$H(S) = \frac{1}{2}\int_S \|\dot{u} - \dot{u}_m\|d\sigma, \tag{2.8}$$

where $\dot{u}$ is the expected optical flow for the surface hypothesis S and $\dot{u}_m(s)$ is the observed flow field at image coordinates $\Pi(s)$, interpolated from measurements. In practice, problems may arise if optical flow data are only partially available, a case we will ignore for the scope of this paper.

Traditionally (cf. Solem et al. 2004; Kickingereder and Donner 2004), one would blend both functionals into a third one, e.g. $E(S) = \alpha J(S) + (1 - \alpha)H(S)$, and then minimize E. We will adopt a different viewpoint inspired by Solem and Overgaard (2005b): let $\mathcal{R}$ be the set of regular surfaces (or even level set surfaces if necessary), and denote by $\mathcal{R}_{\hat{m}} \subset \mathcal{R}$ the subset of all surfaces satisfying the normal condition, i.e. $S \in \mathcal{R}_{\hat{m}} \Leftrightarrow J(S) = 0$. We may now solve the *constrained minimization problem*

$$S = \operatorname*{argmin}_{S \in \mathcal{R}_{\hat{m}}} H(S).$$

This alternative concept provides some benefits: $\mathcal{R}_{\hat{m}}$ has fewer degrees of freedom than $\mathcal{R}$ as mentioned at the end of Sect. 2.1. Furthermore, the surface minimizing the sum functional $E(S)$ above will most likely not satisfy the normal constraint. So it is perfectly reasonable to restrict the

search for a unique solution to $\mathcal{R}_{\hat{m}}$, especially when expecting the deflectometric measurement to be highly reliable in contrast to the regularizing extra information. Assuming $S \in \mathcal{R}_{\hat{m}}$ for the construction of H proves useful in the specular optical flow case, where the law of reflection must hold in order to correctly calculate the flow vectors. Before we determine how to realize the minimization process, we will derive an expression for the expected optical flow $\dot{\boldsymbol{u}}$.

The question remains if constraining the problem in the proposed way generally removes the existing ambiguities. Frankly, it is easy to construct—fair to say rather academic—counterexamples: for instance, it can be shown that for a surface under deflectometric inspection which is rotationally symmetric with respect to the principal ray, all elements of the solution manifold will be rotationally symmetric as well. Meanwhile, rotational motion around the symmetry axis does not affect the light map so that predicted and measured (trivial) flow match for all solution candidates. However, this example also suggests that well-posedness is a matter of the applied object motion. Although we have no theoretical proof for the regularizing effect of specular optical flow data, the method offers an additional process parameter that may be selected such that the constrained functional possesses only a single global minimum, see also the results presented in Sect. 5.

3 Optical Flow on Specular Surfaces

3.1 Calculating Optical Flow

We now solve the forward problem, i.e. computing the optical flow $\dot{\boldsymbol{u}}$ for given S and velocity vector field $\boldsymbol{w}$ describing the motion of each point in space, in particular those lying on the reflecting surface so that, as already mentioned in Sect. 2.2, the surface moves relative to a fixed camera and calibrated scene L.

What we are ultimately looking for is the velocity vector $\dot{\boldsymbol{u}}$ of the image of a scene point $\boldsymbol{l}$ in Ω_I. In the diffuse case, the scene points correspond to texture features pinned to the surface. By tracing the curves they follow under surface motion, we get a (three-dimensional) velocity vector field $\boldsymbol{w}$ on S and by projection the desired flow field as given by (2.4). In contrast, for a specular surface and fixed scene feature point, one has to track the reflection point, via which the feature gets reflected into the optical center of the camera. Its behavior seems intricate: the reflection point may "shift" on the surface or vanish even for non-trivial displacements (see counterexample at the end of the last section).

In the rest of this section, we will explain how to compute the velocity vector of a reflection point, which is constrained physically in two ways: firstly, the reflection point must stay on the surface at all times. Secondly, the law of reflection must always be satisfied with respect to the optical center and the fixed scene feature. Let $\boldsymbol{s}$ denote the trace

of the reflection point curve; each of the constraints dictates the normals at the points $\boldsymbol{s}(t)$. Moving on to infinitesimal time steps, this connection can be exploited to pin down the temporal derivative of $\boldsymbol{s}$ and thus the desired velocity vector.

The following lemma describes the effect of a given surface motion on changes in the normal along $\boldsymbol{s}(t)$. Note that only the normal component v of $\boldsymbol{w}$ is needed. This makes sense: if the surface were diffuse, the tangential velocity component would amount to a *texture shift* within the surface rather than a deformation or displacement. For specular surfaces, these kinds of movements are undetectable, as they do not affect the surface shape; remember the case of rotation around a symmetry axis.

Lemma 1 *Let φ be a sufficiently differentiable level set evolution and $\boldsymbol{s} : T \to \mathbb{R}^3$ a differentiable curve on $S(t) := \varphi^{-1}[\{0\}, t]$, i.e. $\varphi[\boldsymbol{s}(t), t] \equiv 0$. Then,*

$$\frac{d}{dt}\{\hat{\boldsymbol{n}}[\boldsymbol{s}(t), t]\} = -D_S v + \mathbf{J}_{\hat{\boldsymbol{n}}}\mathbf{P}_{\hat{\boldsymbol{n}}}\dot{\boldsymbol{s}},$$

where $D_S v$ is the surface gradient of the normal velocity (2.5).

Proof From the implicit representation, we get a natural extension of $\hat{\boldsymbol{n}}$ and v onto a small neighborhood of a point $\boldsymbol{s}(t_0)$. Thus by the chain rule, we may write for the total derivative

$$\frac{d}{dt}\{\hat{\boldsymbol{n}}[\boldsymbol{s}(t), t]\} = \mathbf{J}_{\hat{\boldsymbol{n}}}[\boldsymbol{s}(t), t]\dot{\boldsymbol{s}}(t) + \partial_t \hat{\boldsymbol{n}}[\boldsymbol{s}(t), t].$$

Decomposing $\dot{\boldsymbol{s}}$ orthogonally into $\dot{\boldsymbol{s}} = \mathbf{P}_{\hat{\boldsymbol{n}}}\dot{\boldsymbol{s}} + \langle \dot{\boldsymbol{s}}, \hat{\boldsymbol{n}} \rangle \hat{\boldsymbol{n}} = \mathbf{P}_{\hat{\boldsymbol{n}}}\dot{\boldsymbol{s}} + v\hat{\boldsymbol{n}}$, we get (in short-hand notation)

$$\frac{d}{dt}\{\hat{\boldsymbol{n}}[\boldsymbol{s}(t), t]\} = \mathbf{J}_{\hat{\boldsymbol{n}}}\mathbf{P}_{\hat{\boldsymbol{n}}}\dot{\boldsymbol{s}} + \mathbf{J}_{\hat{\boldsymbol{n}}}(v\hat{\boldsymbol{n}}) + \partial_t \hat{\boldsymbol{n}}. \qquad (3.1)$$

By $\partial_t \hat{\boldsymbol{n}} = \mathbf{J}_\pi \partial_t \nabla\varphi$, changing the order of differentiation, and substituting $\dot{\varphi} = -v\|\nabla\varphi\|$ from the level set equation (2.6), it follows that

$$\partial_t \hat{\boldsymbol{n}} = \partial_t\{\pi(\nabla\varphi)\} = \mathbf{J}_\pi \nabla\dot{\varphi} = \mathbf{J}_\pi \nabla\{-v\|\nabla\varphi\|\}.$$

By the product rule, $\nabla\{v\|\nabla\varphi\|\} = v\nabla\|\nabla\varphi\| + \|\nabla\varphi\|\nabla v$. With $\nabla\|\boldsymbol{x}\| = \hat{\boldsymbol{x}}$, it follows that

$$\mathbf{J}_\pi \nabla\{v\|\nabla\varphi\|\} = \frac{1}{\|\nabla\varphi\|}\mathbf{P}_{\hat{\boldsymbol{n}}}\left(v\nabla^2\varphi\frac{\nabla\varphi}{\|\nabla\varphi\|} + \|\nabla\varphi\|\nabla v\right).$$

Now, we use $\frac{1}{\|\nabla\varphi\|}\mathbf{P}_{\hat{\boldsymbol{n}}}\nabla^2\varphi = \mathbf{J}_{\hat{\boldsymbol{n}}}$, and thus

$$\partial_t \hat{\boldsymbol{n}} = -\mathbf{J}_{\hat{\boldsymbol{n}}}(v\hat{\boldsymbol{n}}) - \mathbf{P}_{\hat{\boldsymbol{n}}}\nabla v,$$

which yields the assertion in view of (3.1). $\qquad\square$

Now we have all ingredients at hand for the proof of the central result of this section. It states that the reflection point

velocity vector $\dot{s}$ must always solve a certain linear equation system, depending on (local) properties of surface, light map, and the scalar field of normal velocities.

Theorem 2 *Let φ be a level set evolution and s a curve as in Lemma 1 which is additionally a* reflection point curve: *all points on s satisfy the law of reflection* (2.2) *with respect to the origin and a fixed $l \in \mathbb{R}^3$. Then, $\dot{s}$ solves the linear system*

$$(\hat{n}\hat{n}^\top + \mathbf{M} - \mathbf{K})\dot{s} = (\mathbf{I} + \mathbf{J}_{\hat{m}})(v\hat{n}) + D_S v.$$

Here, $\mathbf{M} = -\mathbf{J}_{\hat{m}} = -\frac{1}{\|\nabla e_{l,x}\|}\mathbf{P}_{\hat{n}}\nabla^2 e_{l,s}\mathbf{P}_{\hat{n}}$ *with $e_{l,s}$ from* (2.7).

Proof By definition, s satisfies two constraints:

1. $s(t) \in S(t)$ for all $t \in T$ and, by the law of reflection,
2. $\hat{n}[s(t), t] = \hat{m}(s(t))$.

As shown in Sect. 2.3, the first of the above properties immediately implies

$$\langle \dot{s}, \hat{n} \rangle = v. \tag{3.2}$$

Now by Lemma 1,

$$\frac{d}{dt}\{\hat{n}(s(t), t)\} = -D_S v + \mathbf{J}_{\hat{n}}\mathbf{P}_{\hat{n}}\dot{s}.$$

But the second constraint forces $\hat{n}(s(t), t) = \hat{m}(s(t))$, thus

$$-\mathbf{J}_{\hat{m}}\dot{s} = D_S v - \mathbf{J}_{\hat{n}}\mathbf{P}_{\hat{n}}\dot{s}.$$

As $\mathbf{J}_{\hat{n}}\mathbf{P}_{\hat{n}}\dot{s} = -\mathbf{K}\mathbf{P}_{\hat{n}}\dot{s} = -\mathbf{K}\dot{s}$, we have

$$(-\mathbf{J}_{\hat{m}} - \mathbf{K})\dot{s} = D_S v.$$

Again, decomposing $\dot{s} = \mathbf{P}_{\hat{n}}\dot{s} + \hat{n}\hat{n}^\top\dot{s} = \mathbf{P}_{\hat{n}}\dot{s} + v\hat{n}$, it follows that

$$(-\mathbf{J}_{\hat{m}}\mathbf{P}_{\hat{n}} - \mathbf{K})\dot{s} = \mathbf{J}_{\hat{m}}(v\hat{n}) + D_S v.$$

For $\mathbf{M}$ defined as above, we have $-\mathbf{J}_{\hat{m}}\mathbf{P}_{\hat{n}} = \mathbf{M}$, and thus

$$(\mathbf{M} - \mathbf{K})\dot{s} = \mathbf{J}_{\hat{m}}(v\hat{n}) + D_S v. \tag{3.3}$$

Equations (3.2) and (3.3) form a linear system in $\dot{s}$. Moreover, since $\mathrm{Im}(\mathbf{M} - \mathbf{K}) \subseteq T_s S$ and $\mathbf{J}_{\hat{m}}(v\hat{n})$, $D_S v \in T_s S$, we may multiply (3.2) by $\hat{n}$ and add both equations, thereby transforming the system into the aggregate form

$$(\hat{n}\hat{n}^\top + \mathbf{M} - \mathbf{K})\dot{s} = (\mathbf{I} + \mathbf{J}_{\hat{m}})(v\hat{n}) + D_S v. \quad \square \tag{3.4}$$

Assuming the matrix on the left side of (3.4) is invertible, Theorem 2 allows us to compute, for each point s on the surface and the corresponding *fixed* scene point l, the velocity

vector $\dot{s}$ of the reflection point on the surface. Projecting this into the image plane finally gives the optical flow vector

$$\dot{u} = \frac{1}{s_3}\begin{pmatrix} 1 & 0 & -u_1 \\ 0 & 1 & -u_2 \\ 0 & 0 & 0 \end{pmatrix}\dot{s}(s, v).$$

Note that while this can be derived in the same way as (2.4), here $\dot{s}$ not only depends on s but also on the local surface geometry. A striking consequence of the theorem is that the optical flow is only influenced by surface properties up to second order. This is in close analogy to a result of Savarese et al. (2005) for a related problem.

Observe that the optical flow computation only makes sense if one can presume that the law of reflection holds. This is one of the reasons why our constrained minimization approach is superior: we only need to evaluate flow vectors for surfaces $S \in \mathcal{R}_{\hat{m}}$, which by construction of $\mathcal{R}_{\hat{m}}$ satisfy the law of reflection in every point.

Also note that our model allows arbitrary surface motion, as it depends only on the scalar normal velocity field. In practice, we may want to restrict ourselves to rigid body motion:

Corollary 3 *Under the conditions of Theorem 2, let S move relative to the camera according to an external velocity field w as in Sect. 2.2. Then $\dot{s}$ satisfies*

$$(\hat{n}\hat{n}^\top + \mathbf{M} - \mathbf{K})(\dot{s} - w) = \mathbf{J}_{\hat{m}}w + \mathbf{P}_{\hat{n}}\mathbf{J}_w^\top\hat{n}.$$

Proof Consider an arbitrary point x in the field of view subject to the external motion. Since w is only a function of x, we may attribute to x the unique path $x(t)$ a particle in x follows over time as the solution of the autonomous differential equation $\dot{x} = w(x)$. In particular, even if $x(t) \in S(t)$, $x(t_0) := s$, is the trace of a surface point, this path does not depend on the surface geometry. Without knowledge of S, we can conclude that $\dot{x}(t_0) = w(s)$.

On the other hand, since the whole surface moves according to $w(x)$, such a particle is bound to stay on the surface over time, i.e. $x(t) \in S(t)$ so that the reasoning at the bottom of (2.5) applies, and the normal velocity must be $v = \langle \dot{x}(t_0), \hat{n} \rangle$. Using the result of the last paragraph, we get

$$v = \langle w, \hat{n} \rangle = \hat{n}^\top w.$$

We substitute this into the result of Theorem 2 and apply the product rule to decompose $D_S v = \mathbf{P}_{\hat{n}}\nabla v$:

$(\hat{n}\hat{n}^\top + \mathbf{M} - \mathbf{K})\dot{s}$

$$= (\mathbf{I} + \mathbf{J}_{\hat{m}})(v\hat{n}) + D_S v$$

$$= (\mathbf{I} + \mathbf{J}_{\hat{m}})\hat{n}\hat{n}^\top w + \mathbf{P}_{\hat{n}}\nabla(\hat{n}^\top w)$$

$$= (\hat{n}\hat{n}^\top)w + (\mathbf{J}_{\hat{m}})(\hat{n}\hat{n}^\top)w + \mathbf{P}_{\hat{n}}\mathbf{J}_{\hat{n}}^\top w + \mathbf{P}_{\hat{n}}\mathbf{J}_w^\top \hat{n}.$$

To simplify this further, we use $\mathbf{K} = -\mathbf{J}_{\hat{n}}\mathbf{P}_{\hat{n}}$ and $\mathbf{M} = -\mathbf{J}_{\hat{m}}\mathbf{P}_{\hat{n}}$ as well as the symmetricity of these matrices. We finally get

$$(\hat{n}\hat{n}^\top + \mathbf{M} - \mathbf{K})\dot{s} = (\hat{n}\hat{n}^\top)w + (\mathbf{M} + \mathbf{J}_{\hat{m}})w$$
$$- \mathbf{K}w + \mathbf{P}_{\hat{n}}\mathbf{J}_w^\top \hat{n},$$

which is the assertion in rearranged form.   □

In particular, for an Euclidean rigid body motion we have $w = \mathbf{\Omega}x + v$, so $\mathbf{J}_w = \mathbf{\Omega}$. Also note that by assumption $\hat{n} = \hat{m}$ on S, so any occurrences of $\hat{n}$ may be replaced by the known $\hat{m}$. If the matrix $(\hat{n}\hat{n}^\top + \mathbf{M} - \mathbf{K})$ is regular, $\dot{s}$ is uniquely determined. This is the case if and only if $\mathrm{Rank}(\mathbf{M} - \mathbf{K}) = 2$. To see this, observe that by definition the matrices $\hat{n}\hat{n}^\top$ and $(\mathbf{M} - \mathbf{K})$ operate on the orthogonal subspaces $\mathrm{Span}(\hat{n})$ and $\mathrm{Span}(\hat{n})^\perp$, respectively. This means their ranks add up, and using $\mathrm{Rank}\,\hat{n}\hat{n}^\top = 1$, we get $\mathrm{Rank}(\hat{n}\hat{n}^\top + \mathbf{M} - \mathbf{K}) = 3 \iff \mathrm{Rank}(\mathbf{M} - \mathbf{K}) = 2$. If this regularity condition is satisfied, we will call s a *flow regular* point, otherwise *flow singular*.

For illustration, consider a spheroid $E_{l_0,s}$ with $l_0 := l(\Pi(s)) \in L$ as defined by (2.7). Each point on E reflects l_0 into the origin, meaning that the light map $l : \Omega_I \to L$ cannot be inverted locally as it is not injective. When calculating optical flow, we are essentially tracking the projection $l^{-1}(l_0)$ of the *diffuse scene feature* at l_0. This must fail if l is not at least locally invertible in l_0.

So the natural outcome is that the optical flow vector is not well-defined in points $u = \Pi(s)$ for surfaces that resemble the spheroid $E_{l_0,s}$ in a neighborhood of s. This is a perfectly reasonable constraint since such a surface would locally exhibit "infinite" magnification. But the rank constraint is actually weaker: it also forbids reflection points s where the second order properties of S and $E_{l_0,s}$ match in a single tangential direction $y \in T_s S$. This accounts for the fact that the image of l_0 may split and merge, or even abruptly disappear from the camera image (see for example Fig. 4). We conclude that flow singular points are not only a byproduct of our model but have a real physical counterpart.

Figure 5 shows an exemplary optical flow for a stretched paraboloid rotating around the principal axis of the sensor as well as the "raw" reflection point velocities $\dot{s}$. The result differs drastically from its counterpart for diffuse surfaces, which is rotation symmetric and—in this special case—even independent of the surface shape.

4 Reconstruction of Specular Surfaces

We shall now apply the above results for calculating properties of real-world specular optical flow to the problem introduced in Sect. 2.4. To recapitulate, our setup consists of a camera observing a fixed calibrated scene over a mirror surface subject to known motion. The agreement between measured light map and a candidate surface is cast into a surface functional J. This functional has no unique minimum, which is why we propose a second functional H adding information from specular optical flow measurements. Consequently, we seek a minimum of H within the zero set $\mathcal{R}_{\hat{m}}$ of J.

The analysis of this problem shall be dissected into three parts:

1. We study the unconstrained problem of finding an *arbitrary* surface in $\mathcal{R}_{\hat{m}}$, i.e. one that satisfies the normal condition. A gradient descent method like outlined in Sect. 4.1 is employed. While this does not solve the complete problem, it provides a feasible starting point for the following iterative algorithm.

2. For a large class of functionals such as J, one can generally state the gradient in explicit form. We will briefly review the derivation and supplement it with the necessary natural boundary conditions in Sect. 4.2.

3. Tackling the constrained minimization of H in $\mathcal{R}_{\hat{m}}$, it is instructive to have a closer look at the local structure of $\mathcal{R}_{\hat{m}}$, i.e. on its tangent spaces. In Sect. 4.3, we will see how these can practically be computed. They are applied during the gradient projection method portrayed in Sect. 4.4, during which the gradient descent for H is coupled with a projection step to assert $S \in \mathcal{R}_{\hat{m}}$.

Our focus lies on analyzing the common ingredients of gradient based approaches and showing how the structure of the—at first glance complicated—constraints can be exploited. For efficiency reasons, more elaborate algorithms are definitely desirable and a potential subject of future research. While the gradient projection method is limited, it does provide first numerical results without obscuring the problem structure.

Note that the following analysis is not restricted to H accounting for specular optical flow. In fact, by modifying H one could as well use data from a second camera, from a non-diffuse reflection component or even from tracking diffuse markers on the surface.

4.1 Unconstrained Minimization

We will quickly review how gradient descent works for surface functionals. For an outline of the basic concepts in abstract form, refer to Solem and Overgaard (2005a). A functional analytic survey is given by Burger (2003). The derivation of gradient flows has also been extensively discussed

Fig. 4 Occurrence of flow singular points in concave specular shapes: you may observe a sudden elongation of the designated feature as a result of minor camera displacement
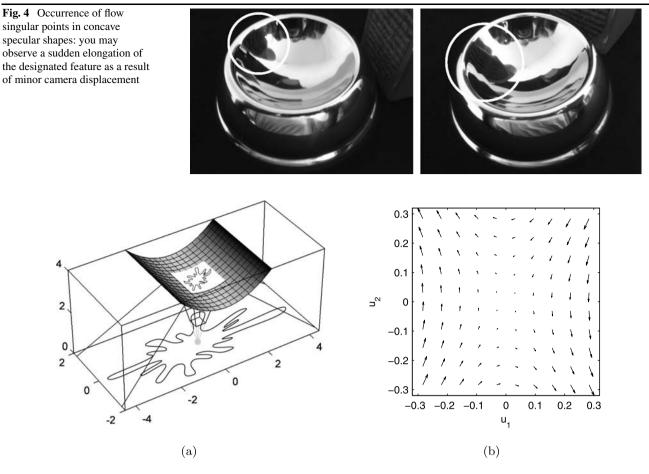

(a)

(b)

Fig. 5 (**a**) Exemplary vector field of reflection point motion $\dot{s}$ for a paraboloid rotating around the z axis and (**b**) corresponding optical flow field $\dot{u}$

by the active contours community (Aubert et al. 2003; Jehan-Besson et al. 2006).

Consider the minimization problem $S = \mathrm{argmin}_{S \in \mathcal{R}} J(S)$ over the set of regular surfaces. In level set notation, J is of the form

$$J(\varphi) = \int_{\Omega} (1 - \langle \hat{\boldsymbol{m}}, \hat{\boldsymbol{n}} \rangle) \|\nabla \varphi\| \delta(\varphi) dx. \qquad (4.1)$$

Note that both $\hat{\boldsymbol{m}}$ and $\hat{\boldsymbol{n}}$ are normalized so that in fact $\frac{1}{2}\|\hat{\boldsymbol{n}} - \hat{\boldsymbol{m}}\|^2 = 1 - \langle \hat{\boldsymbol{m}}, \hat{\boldsymbol{n}} \rangle$. In order to minimize J, we seek a critical point: find φ^* s.t. $dJ(\varphi^*)\eta := \frac{d}{dt}\{J(\varphi^* + t\eta)\}|_{t=0} = 0$ for all test functions $\eta \in C^{\infty}(\Omega)$. Formally differentiating $J(\varphi^* + t\psi)$ w.r.t. t shows this is equivalent to

$$\mathrm{div}(\hat{\boldsymbol{n}}^* - \hat{\boldsymbol{m}})\delta(\varphi^*) = 0, \qquad (4.2)$$

where div is the divergence operator. As will be shown in the next section, this induces a time-dependent process:

$$\dot{\varphi} = \mathrm{div}(\hat{\boldsymbol{n}} - \hat{\boldsymbol{m}})\|\nabla \varphi\|. \qquad (4.3)$$

Any steady state solution of this equation satisfies (4.2) and is thus a stationary point of J. Using the level set equation

(2.6), one can formulate the result in terms of normal velocities:

$$v = -\mathrm{div}(\hat{\boldsymbol{n}} - \hat{\boldsymbol{m}}). \qquad (4.4)$$

This proves to be a very useful tool for deriving properties of evolving surfaces. Under the premise that we can find a corresponding level set evolution φ, we may use the level set representation as a tool to simplify calculations. Afterwards, results are converted back to a representation-independent form—which must be possible if the surface evolution is well defined, i.e. depends only on the *zero set* of the level set function. In the following sections we will take advantage of this approach.

4.2 Deriving Natural Boundary Conditions

In current vision-related level set literature, it is common practice (Goldlücke and Magnor 2004; Solem and Overgaard 2005a) to postulate that S is closed and fully contained in the computational domain Ω, i.e. $S \cap \partial \Omega = \emptyset$. The surface is then represented by a discretization of the actual

level set function, and standard Neumann boundary conditions $\frac{d\varphi}{d\hat{o}} = 0$ are applied ($\hat{o}$ refers to the outside normal of $\partial\Omega$). Using these homogeneous boundary conditions is justified, provided that S is at a sufficiently large distance from $\partial\Omega$.

However, this case requires normal measurements for every point in Ω—even outside the actual surface—which in practice are rarely available. In our setup, the camera generally sees only a small part of the surface, thus forcing S and $\partial\Omega$ to intersect. This renders the above approach useless: the Neumann condition would then force $\langle\hat{n}, \hat{o}\rangle = 0$, which will generally be violated by the sought-after surface. We shall therefore extend the proof of the normal speed equation (4.4), as shown in Solem and Overgaard (2005a), by formally deriving the natural boundary condition.

Let E be the level set version of a surface error functional of the form

$$E(\varphi) = \int_\Omega g(\boldsymbol{x}, \hat{n}) \|\nabla\varphi\|\delta(\varphi)d\boldsymbol{x}$$

with $g : \mathbb{R}^3 \times S^2 \to \mathbb{R}_0^+$ integrable over S. Note that this includes definition (4.1) of the normal functional J if we set $g(\boldsymbol{x}, \hat{n}) = \frac{1}{2}\|\hat{n} - \hat{m}(\boldsymbol{x})\|$. We denote by $\boldsymbol{g}_{\hat{n}} := D_S\hat{n}$ the surface Jacobian of $\hat{n}$. The following theorem gives natural boundary conditions which must hold at every stationary point of E and thus also at the minimizing surface. This effectively allows us to limit the search space to surfaces fulfilling the boundary condition.

Theorem 4 *Let the level set surface φ be a critical point of the functional E. Then for sufficiently smooth $\partial\Omega$ the natural boundary condition*

$$\langle\boldsymbol{g}_{\hat{n}} + g\hat{n}, \hat{o}\rangle = 0 \tag{4.5}$$

on $\varphi^{-1}[\{0\}, t] \cap \partial\Omega$ follows. Generally, if φ fulfills this condition, we have the following expression for the Gâteaux derivative $dE(\varphi)\eta$ of E w.r.t. η:

$$dE(\varphi)\eta = -\int_\Omega \eta\,\mathrm{div}(\boldsymbol{g}_{\hat{n}} + g\hat{n})\delta(\varphi)dx \tag{4.6}$$

for all $\eta \in C^\infty(\Omega)$.

Proof For a normal variation $\varphi^\varepsilon := \varphi + \varepsilon\eta$, we have

$$dE(\varphi)\eta = \frac{d}{d\varepsilon}E(\varphi^\varepsilon)\Big|_{\varepsilon=0}$$

$$= \int_\Omega \frac{d}{d\varepsilon}\{g(\boldsymbol{x}, \hat{n})\|\nabla\varphi^\varepsilon\|\delta(\varphi^\varepsilon)\}\Big|_{\varepsilon=0} d\sigma.$$

As shown in Solem and Overgaard (2005a), this leads to

$$dE(\varphi)\eta = \int_\Omega \langle\nabla\eta, \boldsymbol{g}_{\hat{n}} + g\hat{n}\rangle\delta(\varphi)dx + \int_\Omega \eta g\|\nabla\varphi\|\delta'(\varphi)dx.$$

Applying Gauss' theorem to the first integral yields

$$dE(\varphi)\eta = \int_{\partial\Omega} \eta\langle\boldsymbol{g}_{\hat{n}} + g\hat{n}, \hat{o}\rangle\delta(\varphi)d\sigma$$

$$- \int_\Omega \eta\,\mathrm{div}\{(\boldsymbol{g}_{\hat{n}} + g\hat{n})\delta(\varphi)\}dx$$

$$+ \int_\Omega \eta g\|\nabla\varphi\|\delta'(\varphi)dx$$

$$= \int_{\partial\Omega} \eta\langle\boldsymbol{g}_{\hat{n}} + g\hat{n}, \hat{o}\rangle\delta(\varphi)d\sigma$$

$$- \int_\Omega \eta\,\mathrm{div}(\boldsymbol{g}_{\hat{n}} + g\hat{n})\delta(\varphi)dx$$

$$- \int_\Omega \eta\langle\boldsymbol{g}_{\hat{n}} + g\hat{n}, \nabla\varphi\rangle\delta'(\varphi)dx$$

$$+ \int_\Omega \eta g\|\nabla\varphi\|\delta'(\varphi)dx.$$

Observe that $\langle\boldsymbol{g}_{\hat{n}}, \nabla\varphi\rangle = \langle\boldsymbol{g}_{\hat{n}}, \hat{n}\rangle\|\nabla\varphi\| = 0$, since the surface gradient $\boldsymbol{g}_{\hat{n}}$ by definition lies in the tangent plane of S. Also $\langle\hat{n}, \nabla\varphi\rangle = \|\nabla\varphi\|$, so the remaining terms in the last line cancel each other leaving

$$dE(\varphi)\eta = \int_{\partial\Omega} \eta\langle\boldsymbol{g}_{\hat{n}} + g\hat{n}, \hat{o}\rangle\delta(\varphi)d\sigma$$

$$- \int_\Omega \eta\,\mathrm{div}(\boldsymbol{g}_{\hat{n}} + g\hat{n})\delta(\varphi)dx.$$

Now, let φ be a critical point of E. Constraining the choice of η to functions vanishing outside a compact domain, we see that the right side vanishes for all such η. Since this applies to all η, it forces $\mathrm{div}(\boldsymbol{g}_{\hat{n}} + g\hat{n})\delta(\varphi) \equiv 0$. But this means the left integral must vanish itself for all η. Again, this applies to all $\eta \in C^\infty(\Omega)$, resulting in the natural boundary condition $\langle\boldsymbol{g}_{\hat{n}} + g\hat{n}, \hat{o}\rangle\delta(\varphi) \equiv 0$ on $\partial\Omega$.

On the other hand, if φ satisfies this boundary restriction, we have

$$dE(\varphi)\eta = -\int_\Omega \eta\,\mathrm{div}(\boldsymbol{g}_{\hat{n}} + g\hat{n})\delta(\varphi)dx. \qquad \square$$

The term $\mathrm{div}(\boldsymbol{g}_{\hat{n}} + g\hat{n})$ is the so-called *shape gradient* of E, so in view of (4.4), our minimization scheme is of steepest descent type. For a thorough introduction from this perspective, see Delfour and Zolesio (2001), Burger (2003), Charpiat et al. (2005).

4.3 The Space of Admissible Gradients

At the constrained minimization of the specular optical flow error functional H from (2.8), we face two difficulties:

1. Following Theorem 2, H contains derivatives of φ up to second order, while the above derivation only permits first order derivatives.

2. We must respect the normal condition, i.e. we may not leave $\mathcal{R}_{\hat{m}}$ during the time advancement step.

Let us consider the first point: while it is possible to derive a gradient descent process similar to (4.3), this requires computing surface derivatives of up to fourth order. Following the common Lagrangian approach, the gradient must then be projected onto a—yet to be determined—space consisting of all gradients v of surface evolutions through the current surface that obey the normal condition. As the manifold of solutions to the reconstruction problem is assumed to be a "regular curve", its tangent space—which we denote by $\mathcal{A}_{\hat{m}}$—will clearly be one-dimensional. Thus, the gradient projection reduces to the computation of a directional derivative of H, which can be approximated numerically. Sparing the cumbersome calculation of the complete gradient, both difficulties addressed above are resolved this way.

Under moderate assumptions, the tangent space $\mathcal{A}_{\hat{m}}$ matches the solution space of some partial differential equation:

Proposition 5 *Let $\varphi : \Omega \times T \to \mathbb{R}$ be a sufficiently differentiable level set evolution respecting the normal condition at each point in time, i.e. $\hat{n}[s, t] = \hat{m}(s), t \in T, s \in S(t)$. Then for all $t \in T$ and $s \in S(t)$, the following relationship holds:*

$$D_S v + \mathbf{J}_{\hat{m}}(v\hat{m}) = 0. \tag{4.7}$$

Here $D_S v$ denotes the surface gradient of v on S.

For a fixed point in time, let $\tilde{\mathcal{A}}_{\hat{m}}$ consist of all v satisfying (4.7). Now note that $\mathcal{A}_{\hat{m}}$ is closed under scalar multiplication. The linear structure of (4.7) confirms that $\tilde{\mathcal{A}}_{\hat{m}}$ is a vector space. Clearly, $\tilde{\mathcal{A}}_{\hat{m}}$ must contain $\mathcal{A}_{\hat{m}}$. From the Lipschitz-continuity of $\hat{m}$, it can be deduced by means of general PDE theory that the solution of (4.7) will actually be unique up to a scalar, i.e. $\dim \tilde{\mathcal{A}}_{\hat{m}} = 1$. The equality in dimension implies that both spaces coincide whenever $\mathcal{A}_{\hat{m}} \neq \{0\}$, a case that has been excluded by our regularity assumption. The projection is exact!

Proof of Proposition 5 For a sufficiently differentiable φ and $t_0 \in T$, $x_0 \in S(t_0)$, we can locally find a curve $s : T \to \mathbb{R}^3$ on S with $s(t_0) = x_0$. This can be seen making the ansatz $s(t) = x + \hat{n}(x_0, t_0)\lambda(t), \lambda(t_0) = 0$, and using an implicit function argument to find $\lambda \in C^1(T_\varepsilon, \mathbb{R})$ in a neighborhood $T_\varepsilon := (t_0 - \varepsilon, t_0 + \varepsilon)$ of t_0.

From Lemma 1, we get $D_S v + \frac{d}{dt}\{\hat{n}[s(t), t]\} - \mathbf{J}_{\hat{n}}\mathbf{P}_{\hat{n}}\dot{s} = 0$. Using the normal condition $\hat{n}[s(t), t] = \hat{m}(s(t))$, it follows that

$$D_S v + \mathbf{J}_{\hat{m}}\dot{s} - \mathbf{J}_{\hat{m}}\mathbf{P}_{\hat{n}}\dot{s} = 0.$$

Decompose s orthogonally, then $\mathbf{J}_{\hat{m}}\dot{s} = \mathbf{J}_{\hat{m}}\mathbf{P}_{\hat{n}}\dot{s} + \mathbf{J}_{\hat{m}}(\hat{n}\hat{n}^\top)\dot{s}$ and

$$D_S v + \mathbf{J}_{\hat{m}}(\hat{n}\hat{n}^\top)\dot{s} = 0.$$

Substituting $\mathbf{J}_{\hat{m}}(\hat{n}\hat{n}^\top)\dot{s} = (\mathbf{J}_{\hat{m}}\hat{m})(\hat{m}^\top\hat{m})v = \mathbf{J}_{\hat{m}}(v\hat{m})$ yields the assertion. $\qquad\square$

4.4 The Minimization Algorithm

Let us now outline the complete gradient descent procedure for solving the constrained minimization problem for general H.

1. Choose a start value $S(t_0 = 0)$ for the surface iteration.
2. Advance S in time according to (4.4) until $J(S(t)) < \varepsilon_n$.
3. While $H(S(t)) \geqslant \varepsilon_H$, repeat:
   (a) For $S(t)$, find a basis vector b of the discrete version of $\tilde{\mathcal{A}}_{\hat{m}}$.
   (b) Numerically, approximate the directional derivative $dH(t)b$. The gradient of H projected onto $\tilde{\mathcal{A}}_{\hat{m}}$ is then $\tilde{v} = b\{dH(S(t))b\}$.
   (c) Advance $t \to t + \delta_t$ and $S(t) \to S(t + \delta_t)$ according to the normal speed given by $\tilde{v}$.

As with all local methods, one has to rely on a sufficiently accurate initial guess. This guess could be based on an estimate of the camera-surface distance, on the expected shape in a quality control environment, or on the output of a coarse but robust method.

Depending on the specific method used, the time discretization step may cause an increase in the normal error J during phase 3. In our prototype, we use a simple forward Euler step with some additional repetitions of step 2 if the normal error exceeds a certain upper bound. For more implementation details, see Lellmann (2006).

Note that the formulation works for any H, not only for the optical flow case. Also note that although the derivation of the proposed method is based on level set calculus, we are free to use any other surface representation for the implementation. As an example, we shall briefly discuss the case of an explicit parametrization in two dimensions. Extending the method to three dimensions is straightforward. Here we have $f : X \times T \to \mathbb{R}^+$ for an interval $X = [x_l, x_r]$ and $\Omega = X \times \mathbb{R}^+$. The surface is then simply the graph of f, $S(t) = \{(x, f[x, t])^\top \mid x \in X\}$. Some important normal related properties are

$$\hat{n} = \frac{1}{\sqrt{(f')^2 + 1}}\begin{pmatrix} f' \\ -1 \end{pmatrix},$$

$$\mathbf{J}_{\hat{n}} = \frac{1}{[(f')^2 + 1]^{\frac{3}{2}}}\begin{pmatrix} f'' & 0 \\ f'f'' & 0 \end{pmatrix}, \qquad \operatorname{div}\hat{n} = \frac{f''}{[(f')^2 + 1]^{\frac{3}{2}}}.$$

Consider a curve s on S, i.e. $s(t) = (x(t), f[x(t), t])^\top$. By $v = \langle \dot{s}, \hat{n} \rangle$, we get

$$v = \frac{1}{\sqrt{(f')^2 + 1}}\left\langle \begin{pmatrix} \dot{x} \\ f'\dot{x} + \dot{f} \end{pmatrix}, \begin{pmatrix} f' \\ -1 \end{pmatrix} \right\rangle = \frac{-\dot{f}}{\sqrt{(f')^2 + 1}}.$$

The normal velocity field v translates into the temporal derivative of f by $\dot{f} = -v\sqrt{f'^2 + 1}$, in direct analogy to the level set equation $\dot{\varphi} = -v\|\nabla\varphi\|$. The normal boundary condition reads

$$\left\langle \frac{1}{\sqrt{(f')^2 + 1}}\begin{pmatrix} f' \\ -1 \end{pmatrix} - \hat{\boldsymbol{m}}, \begin{pmatrix} -1 \\ 0 \end{pmatrix} \right\rangle = 0,$$

where $\hat{\boldsymbol{m}} = (\hat{m}_1, \hat{m}_2)^\top$. This scalar equation is equivalent to

$$f' = \hat{m}_1\sqrt{f'^2 + 1}. \tag{4.8}$$

Assuming that the normal field $\hat{\boldsymbol{m}}$ is compatible with the surface model in the sense that through every point $\boldsymbol{x}$, one can actually find a surface in this parametrization s.t. $\hat{\boldsymbol{m}}(\boldsymbol{x}) = \hat{\boldsymbol{n}}(\boldsymbol{x})$, the case $\hat{m}_2 = 0$ can be excluded. We finally obtain

$$f'(x_l) = \frac{\hat{m}_1(x_l)}{|\hat{m}_2(x_l)|}$$

as the natural boundary condition. The same result applies to the right boundary. Note that, in this special case, this is equivalent to forcing $\hat{\boldsymbol{n}} = \hat{\boldsymbol{m}}$ on $S \cap \partial\Omega$, i.e. on the border the minimizing surface normals coincide exactly with those of the measured normal field.

5 Experimental Results

To verify the feasibility of the presented minimization process, we created a MATLAB implementation of the algorithm. For the evaluation of our method, we used the above explicit parametrization for two-dimensional Ω combined with a simple finite difference scheme and forward Euler step method for solving the PDE. At the m-th iteration we have $f^m = (f_1^m, \ldots, f_N^m)$ and $\dot{f}^m = (\dot{f}_1^m, \ldots, \dot{f}_N^m)$. Using central differences for $\dot{f}^m$, the discrete equivalent of the space $\tilde{\mathcal{A}}_{\hat{\boldsymbol{m}}}$ is the kernel of a matrix of the form

$$\mathbf{A} = \begin{pmatrix} \boldsymbol{a}_{1,1} & \boldsymbol{a}_{1,2} & & \\ & \boldsymbol{a}_{2,1} & \boldsymbol{a}_{2,2} & \\ & & \ddots & \ddots \\ & & & \boldsymbol{a}_{N-1,1} & \boldsymbol{a}_{N-1,2} \end{pmatrix} \in \mathbb{R}^{(2N-2)\times N},$$

where each $\boldsymbol{a}_{i,j}$ represents an $\mathbb{R}^2$ vector. This kernel is found by decomposing $\mathbf{A}$ using SVD and inspecting the singular values. Looking at the defining (4.7), we expect rows $2k - 1$ and $2k$ to constitute one scalar equation thus limiting the rank of $\mathbf{A}$ to $N - 1$. In this special two-dimensional case, one could actually parameterize v by arc length of the surface curve $S(t)$. This transforms (4.7) into a homogenous linear ODE for v, so that all solutions are scalar multiples of an arbitrary non-zero solution v_0. Thus, the kernel of $\mathbf{A}$ should be one-dimensional, which was confirmed in all test cases. This supports the assumption about the simple structure of $\mathcal{R}_{\hat{\boldsymbol{m}}}$. Figure 6(a) shows a visualization of $\mathcal{R}_{\hat{\boldsymbol{m}}}$ for a sample curved surface. This surface is common to all following illustrations.

In order to obtain good data for the normal field, the discrete simulation values of the light map $\boldsymbol{l}$ are interpolated using cubic B-splines. For the optical flow data, we use a piecewise linear interpolation, which simplifies the exclusion of singular points. To estimate the error, we recorded the values of the error functionals J and H. Figures 6 and 7 show the steps taken as well as the error evolution for the reconstruction of a sample surface. As can be seen, the surface converges within expectations for the simple discretization.
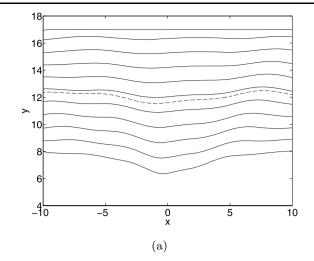
Notwithstanding these encouraging results, one might ask if the optical flow is actually a good candidate for selecting a unique solution from $\mathcal{R}_{\hat{\boldsymbol{m}}}$. To this matter, we evaluated the optical flow functional H for a number of solution candidates in $\mathcal{R}_{\hat{\boldsymbol{m}}}$, see Fig. 8. While there are some—possibly discretization-caused—local minima, the specular optical flow error exhibits a distinct global minimum at the reference surface.

Extending the method to three dimensions is straightforward. Figure 9 shows the result of a simulation for a curved three-dimensional surface. The corresponding error evolution is shown in Fig. 10.

Real data were recorded by the compact sensor head depicted in Fig. 1. It consists of an industrial PC unit, a camera, and LC display, the latter two components being thoroughly calibrated with respect to each other. Since the flow-generating object motion is commanded to the robot controller in camera coordinates, a hand-eye calibration was also required. To demonstrate the regularizing effect of specular optical flow data, we investigated our method for the simple example of a planar mirror.

The LCD was imaged over the test surface *before* and *after* a translation of 50 μm in the direction of the principle ray. Once again, variability of process parameters proved to be very convenient: for flow estimation, we utilized an implementation of the standard Lucas-Kanade algorithm. This is a local method so that density of the estimated flow field may become poor in large homogenous regions. Although this is most likely uncritical in our application (theoretically, a single flow vector might suffice for restoring uniqueness), the effect can be compensated by diligent selection of the displayed pattern.

The reconstruction of the plane mirror is shown in Fig. 11(a). Attaching a diffuse checkerboard pattern to it, a standard extrinsic calibration technique provided for some ground truth. Figure 11(b) shows the mean curvature of the result surface, which should ideally be zero in this case. The moderate discontinuities in the corners of the xy-domain, where the outer normal to Ω is not well-defined, result from discretization of the natural boundary conditions and
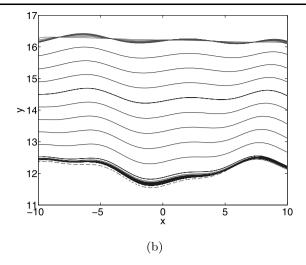
(a)

(b)

Fig. 6 (**a**) The space $\mathcal{R}_{\hat{m}}$ visualized by a set of surfaces satisfying the normal condition for a curved surface (*dashed*). Here, Ω_I is the $z = 1$ plane, $\Omega = (-10, 10) \times \mathbb{R}^+$, and the scene L is the $z = 0$ plane. (**b**) Reconstruction process for a curved surface starting with $f^0 \equiv 16$.

Curves represent snapshots of f^m every 5 steps for a total of 500 iterations. The normal adjustment takes place in the first few iterations at the *top* of the figure. The constrained minimization (*black*) converges to the correct solution (*dashed*)
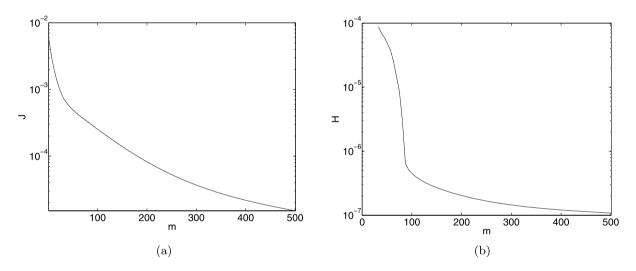

(a)

(b)

Fig. 7 Error progression for (**a**) normal and (**b**) specular optical flow error functionals J and H, respectively

the curvature operator. The normal error decreases exponentially towards the threshold ε_n, which is maintained by intermediate normal adjustments during the projection phase, see Fig. 12. Mean relative normal and distance error turned out to be 1.81 % and 6.23%, confirming the higher sensitivity of deflectometric methods with respect to slope.

6 Conclusion

Based on the theory of level set evolutions, we contributed a very general model for describing optical flow fields which originate from the imaging of a scene over a time-variant specular object. We could additionally derive necessary and sufficient conditions for these modelling equations to be invertible. Particularly, it followed that the model works for all cases where an actual measurement is possible.

A novel view of the reconstruction problem was put forward thereupon: light map measurements were treated as a constraint for the minimization of an additional functional which—in our case—represented the optical flow error. This could then be solved by finding the steady-state solution of a partial differential equation under the normal constraint. To cope with partial measurements, natural boundary conditions were derived. To solve the PDE, a gradient projection type algorithm was proposed, which made it necessary
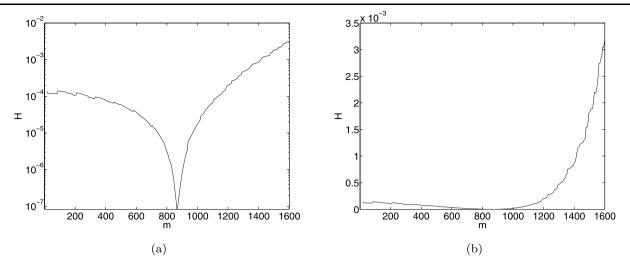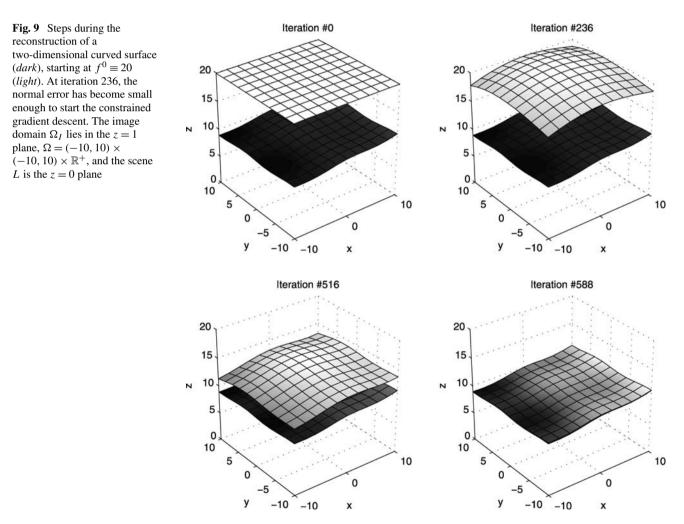
(a)   (b)

Fig. 8 The (**a**) logarithmic and (**b**) linear specular optical flow error functional evaluated for several surfaces, all obeying the normal constraint, see also Fig. 6. The minimum uniquely identifies the desired surface

Fig. 9 Steps during the reconstruction of a two-dimensional curved surface (*dark*), starting at $f^0 \equiv 20$ (*light*). At iteration 236, the normal error has become small enough to start the constrained gradient descent. The image domain Ω_I lies in the $z = 1$ plane, $\Omega = (-10, 10) \times (-10, 10) \times \mathbb{R}^+$, and the scene L is the $z = 0$ plane

to take a closer look at the space of admissible gradients. Finally, the results of a reference implementation were presented, suggesting good convergence. An additional experiment helped to justify the choice of an optical flow-based method for pinpointing the unique solution to the reconstruction problem.
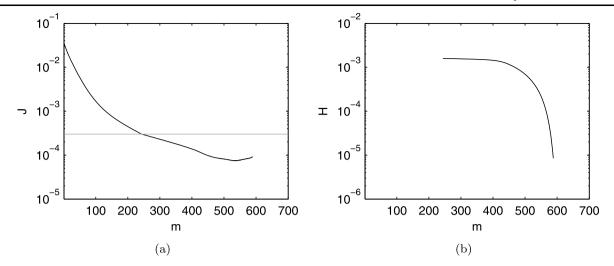
Fig. 10 Error evolution for (**a**) normal and (**b**) optical flow error functionals J and H. The normal error threshold ε_n is marked with a gray line
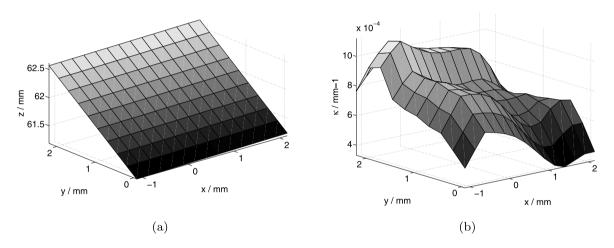

Fig. 11 (**a**) Reconstruction result in the case of the planar mirror. (**b**) Validation by virtue of mean curvature
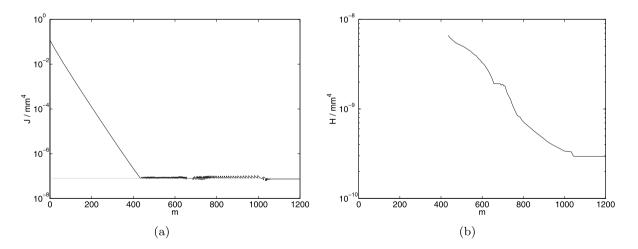

Fig. 12 Development of (**a**) normal and (**b**) flow error J, respectively H. The threshold ε_n for the beginning of the projection phase is marked in (**a**) by a *horizontal line*

In the future, adopting a more sophisticated integration process will certainly improve accuracy as well as convergence speed. Finally, an analytical answer to the question of well-posedness in the case of the flow-regularized gradient descent has yet to be found.

References

Aubert, G., Barlaud, M., Faugeras, O., & Jehan-Besson, S. (2003). Image segmentation using active contours: Calculus of variations or shape gradients? *SIAM Journal on Applied Mathematics*, *63*(6), 2128–2154.

Balzer, J., Werling, S., & Beyerer, J. (2006). Regularization of the deflectometry problem using shading data. In *Proceedings of the SPIE optics east*.

Barron, J. L., Fleet, D. J., & Beauchemin, S. S. (1994). Performance of optical flow techniques. *International Journal of Computer Vision*, *12*(1), 43–77.

Bonfort, T. (2006). *Reconstruction de surfaces réfléchissantes à partir d'images*. Ph.D. thesis, Institut National Polytechnique de Grenoble.

Bonfort, T., & Sturm, P. (2003). Voxel carving for specular surfaces. In *ICCV* (Vol. 1, pp. 591–596).

Burger, M. (2003). A framework for the construction of level set methods for shape optimization and reconstruction. *Interfaces and Free Boundaries*, *5*, 301–329.

Charpiat, G., Keriven, R., Pons, J.-P., & Faugeras, O. (2005). Designing spatially coherent minimizing flows for variational problems based on active contours. In *ICCV* (Vol. 2, pp. 1403–1408).

Chen, M., & Arvo, J. (2000). Theory and application of specular path perturbation. *ACM Transactions on Graphics*, *19*(1), 246–278.

Delfour, M., & Zolesio, J.-P. (2001). *Shapes and geometries*. Philadelphia: SIAM.

do Carmo, M. (1976). *Differential geometry of curves and surfaces*. Prentice-Hall: Englewood Cliffs.

Goldlücke, M., & Magnor, M. (2004). Weighted minimal hypersurfaces and their applications in computer vision. In *ECCV* (Vol. 2, pp. 366–378).

Halstead, M., Barsky, B., Klein, S., & Mandell, R. (1996). Reconstructing curved surfaces from specular reflection patterns using spline surface fitting of normals. In *SIGGRAPH '96* (pp. 335–342). New York: ACM Press.

Hartley, R., & Zisserman, A. (2003). *Multiple view geometry in computer vision* (2nd ed.). Cambridge: Cambridge University Press.

Hicks, R. A., & Perline, R. K. (2004) The method of vector fields for catadioptric sensor design with applications to panoramic imaging. In *CVPR* (Vol. 2, pp. 143–150), Jun, 2004.

Horn, B. (1970). *Shape from Shading: A method for obtaining the shape of a smooth opaque object from one view*. Ph.D. thesis, Department of Electrical Engineering, MIT.

Horn, B. (1986). *Robot vision*. Cambridge: MIT Press.

Jehan-Besson, S., Herbulot, A., Barlaud, M., & Aubert, G. (2006). *Handbook of mathematical models in computer vision*. New York: Springer (Chap. 19, pp. 309–323).

Jin, H. (2003). Estimation of 3d surface shape and smooth radiance from 2d images: A level set approach. *Journal of Scientific Computing*, *19*(1–3), 267–292.

Kammel, S. (2004). *Deflektometrische untersuchung spiegelnd reflektierender freiformflächen*. Ph.D. thesis, University of Karlsruhe.

Kickingereder, R., & Donner, K. (2004). Stereo vision on specular surfaces. In *Proceedings of IASTED conference on visualization, imaging, and image processing* (pp. 335–339).

Lellmann, J. (2006). *Mathematical modelling and analysis of the deflectometry problem: level set based reconstruction of specular free-form surfaces from optical flow*. German title: Mathematische Modellierung und Analyse des Deflektometrieproblems: Rekonstruktion reflektierender Freiformflächen anhand des optischen Flusses auf der Grundlage von Level Sets. Diploma thesis, University of Karlsruhe.

Oren, M., & Nayar, S. K. (1995). A theory of specular surface geometry. In *ICCV* (pp. 740–747).

Osher, S., & Fedkiw, R. (2002). *Level set methods and dynamic implicit surfaces*. New York: Springer.

Prados, E., Camilli, F., & Faugeras, O. (2007). A viscosity solution method for shape-from-shading without image boundary data. *Mathematical Modelling and Numerical Analysis*, *40*(2), 393–412.

Roth, S., & Black, M. (2006). Specular flow and the recovery of surface structure. In *Proceedings of IEEE conference on computer vision and pattern recognition (CVPR)* (Vol. 2, pp. 1869–1876).

Rozenfeld, S., Shimshoni, I., & Lindenbaum, M. (2007). Dense mirroring surface recovery from 1d homographies and sparse correspondences. In *Proceedings of the international conference on CVPR*, Minneapolis.

Savarese, S., Chen, M., & Perona, P. (2005). Local shape from mirror reflections. *International Journal of Computer Vision*, *64*(1), 31–67.

Sethian, J. (2005). *Level set methods and fast marching methods* (2nd ed.). Cambridge: Cambridge University Press.

Solem, J., & Overgaard, N. (2005a). A gradient descent procedure for variational dynamic surface problems with constraints. In *Proceedings of the workshop on variational, geometric, and level set methods in computer vision* (pp. 332–343).

Solem, J., & Overgaard, N. (2005b). A geometric formulation of gradient descent for variational problems with moving surfaces. In *Scale-space* (pp. 419–430).

Solem, J., Aanaes, H., & Heyden, A. (2004). A variational analysis of shape from specularities using sparse data. In *3DPVT* (pp. 26–33).

Werling, S., Balzer, J., & Beyerer, J. (2007). Initial value estimation for robust deflectometric reconstruction. In *Proceedings of 8th international conference on optical 3-d measurement techniques*, Zurich.

Zissermann, A., Giblin, P., & Blake, A. (1989). The information available to a moving observer from specularities. *Image and Vision Computing*, *7*, 38–42.