Specular flow and Weingarten map

Alexey Pak

Vision and Fusion Laboratory
Institute for Anthropomatics
Karlsruhe Institute of Technology (KIT), Germany
apak@ies.uni-karlsruhe.de

Technical Report IES-2015-13

Abstract: If a moving camera observes a specular surface in which some environment is reflected, the pixel values per se do not directly characterise the surface. However, the associated optical flow, or specular flow (SF), as it is called in this situation, is an environment-agnostic observable that depends on the surface position, orientation, and curvature. The derivation of the SF in the limit of an infinitely remote environment has been published earlier by the authors, but in a relatively opaque coordinate-dependent form. In this report, we present a simpler and a more general derivation of the SF as a function of the surface structure, where the crucial part is played by the so-called Weingarten map. This result allows us to formulate the conditions when the SF diverges, and to derive a simple formula to relate the Gaussian curvature of the surface to the SF.

1 Introduction

The shape-from-specular flow (SFSF) approach has been suggested [RB06] as a convenient tool to estimate the 3D shape of moving reflecting objects while abstracting away from the details of the environment. One first estimates the optical flow (OF) field – the apparent displacements of prominent image elements between the subsequent video frames – and uses it to reconstruct the shape. If the camera observes a curved mirror, the motion field may exhibit smooth distortions, infinities, discontinuities etc., which may warrant the adaptation of the common OF algorithms. Such OF is commonly referred to as specular flow (SF) [AZBS11, CnVA+09].

The first SFSF methods were able to estimate the parameters of simple known shapes [RB06]. Later, the problem was theoretically re-formulated in terms of global variational reconstruction [LBRB08]. More recently, Adato et al [AVBSZ07, AVZBS10, CnVA+09, VZGBS11] provided a practical general solution for a special setup where a telecentric camera is fixed with respect to the object and the infinitely remote textured environment undergoes a global rigid rotation. The flow field is obtained from the camera images with the common OF algorithms, and the resulting system of coupled linear partial differential equations is discretized and solved with standard tools. The improvements introduced in that series of papers lead to significant relaxation of the original requirements to the system calibration and the needed data, and reduced the problem to a system of linear partial differential equations (PDEs) [CnVA+09].

While very elegant, such an approach is not yet suitable for the common industrial settings. A more realistic setup would use e.g. a pinhole camera, moving along some trajectory with respect to the object that reflects a distant static environment. In this formulation, the problem has an additional dimensional parameter that is missing in the orthographic model – the distance between the camera and the object. The motion of isolated specularities relative to the fixed surface features in this case has been studied in [BB91], but no attempt has been made to recover the arbitrary surface shapes from the observed data. More recently, [Pak14a] introduced the paraxial formalism for the SF in perspective projection but did not present the final results.

In what follows we reproduce the derivation of the main results of [Pak14a] with the help of an alternative approach developed in [BB91]. We further analyze the structure of the resulting SF equations and in particular outline the relation between the Gaussian curvature of the surface and the observed specular flow.

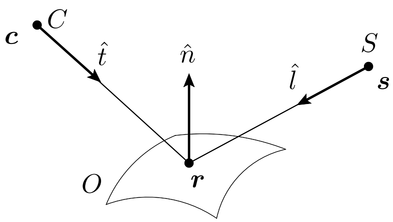

Figure 2.1: Geometry of a camera's view ray after a single specular reflection.

2 Notation

Let us assume an observation setup shown in Fig. 2.1. A camera C with the projection center located at point $\boldsymbol{c}$ observes a reflection of a light source S located at $\boldsymbol{s}$. A light ray originates at S and follows the (unit) direction $\hat{l}$ until it hits the specular object O at some surface point $\boldsymbol{r}$, and then reflects towards the camera. The unit direction from the camera to the reflection point is $\hat{t}$, and the normalized surface normal vector at the reflection point is $\hat{n}$. The distances from the camera and the source to the reflection point are λ and μ, respectively, i.e.

$$\boldsymbol{r} = \boldsymbol{c} + \lambda\,\hat{t} = \boldsymbol{s} + \mu\,\hat{l}. \tag{2.1}$$

In what follows, we will heavily use a linear operator $\pi[\hat{d}]$ that projects an arbitrary vector $\boldsymbol{x}$ on the subspace orthogonal to $\hat{d}$ (some unit vector such as $\hat{n}$, $\hat{t}$, or $\hat{l}$):

$$\pi[\hat{d}]\cdot\boldsymbol{x} \equiv \boldsymbol{x}_{\perp\hat{d}} = \boldsymbol{x} - \hat{d}\,(\boldsymbol{x}\cdot\hat{d}), \;\text{ or }\; \Big(\pi[\hat{d}]\Big)_{ij} = \delta_{ij} - d_i d_j, \tag{2.2}$$

where δ_{ij} is the Kronecker symbol and d_i, d_j are the components of $\hat{d}$. For the product of two projectors, we reserve the special notation $\Pi = \pi[\hat{n}]\cdot\pi[\hat{t}]/\big(\hat{n}\cdot\hat{t}\big)$. We will also need a projector $\mathcal{P}$ whose action on a vector $\boldsymbol{x}$ is as follows:

$$\mathcal{P}\cdot\boldsymbol{x} = \frac{\hat{n}\times\big(\boldsymbol{x}\times\hat{t}\big)}{\hat{n}\cdot\hat{t}}, \;\text{ or }\; (\mathcal{P})_{ij} = \delta_{ij} - \frac{t_i\,n_j}{\hat{n}\cdot\hat{t}}.$$

The latter expression may be obtained by contracting antisymmetric tensors[1] in the more explicit form $(\mathcal{P})_{ij} = \epsilon_{ikl}\;n_k\;\epsilon_{ljm}\;t_m/\big(\hat{n}\cdot\hat{t}\big)$. From the above definitions and the matrix forms of the projectors, one easily finds their trivial properties:

$$\begin{gathered}\pi[\hat{d}]\cdot\boldsymbol{x} = \boldsymbol{0} \;\text{ for }\; \boldsymbol{x}\parallel\hat{d},\;\; \pi[\hat{d}]\cdot\boldsymbol{x} = \boldsymbol{x} \;\text{ for }\; \boldsymbol{x}\perp\hat{d}, \\ \mathcal{P}\cdot\boldsymbol{x} = \boldsymbol{0} \;\text{ for }\; \boldsymbol{x}\parallel\hat{t},\;\; \mathcal{P}\cdot\boldsymbol{x} = \boldsymbol{x} \;\text{ for }\; \boldsymbol{x}\perp\hat{n},\;\; \mathcal{P}\cdot\pi[\hat{t}] = \mathcal{P}, \text{ and } \Pi\cdot\mathcal{P} = \Pi.\end{gathered} \tag{2.3}$$

Since the SF is related to motion, we use the "dotted"notation for time derivatives. For example, the camera's velocity is $\dot{\boldsymbol{c}} \equiv \partial\boldsymbol{c}/\partial t$. In a scene of Fig. 2.1, in general, if the camera moves, only the source location remains constant, i.e. $\dot{\boldsymbol{s}} \equiv \boldsymbol{0}$.

[1] A fully antisymmetric tensor, or a Levi-Civita symbol, is defined here as follows: $\epsilon_{123} = 1$, and $\epsilon_{ijk} = -\epsilon_{jik} = -\epsilon_{ikj}$. It is commonly used to define vector products: $(\boldsymbol{a}\times\boldsymbol{b})_i = \epsilon_{ijk}a_j b_k$.

Finally, we introduce the Weingarten map W that connects the variation of a point $\boldsymbol{r}$ on a surface and the respective change in the normal vector $\hat{n}$. In particular,

$$\dot{\hat{n}} = W \cdot \dot{\boldsymbol{r}}. \tag{2.4}$$

The coordinate representation of W is frame-dependent, but its eigenvalues and eigenvectors are the intrinsic invariant properties of the surface at the point $\boldsymbol{r}$.

3 Specular flow for a moving pinhole camera

With the introduced notation, we are now in the position to derive the specular flow perceived by the moving camera (the object O and the source S remain static) in terms of the camera motion $\dot{\boldsymbol{c}}$ and the local surface parameters at point $\boldsymbol{r}$. Our guideline will be to prevent the proliferation of the source parameters to the result.

According to Fig. 2.1, the directions $\hat{t}$, $\hat{s}$, and $\hat{n}$ satisfy the reflection law, i.e.

$$\hat{t} + \hat{l} = \nu\, \hat{n} \tag{3.1}$$

for some ν. Let us differentiate Eq. (3.1) with respect to time and apply $\pi[\hat{n}]$:

$$\pi[\hat{n}] \cdot \left(\dot{\hat{t}} + \dot{\hat{l}}\right) = \pi[\hat{n}] \cdot \left(\dot{\nu}\, \hat{n} + \nu\, \dot{\hat{n}}\right) = \nu\, \dot{\hat{n}}. \tag{3.2}$$

The latter equality holds since $\|\hat{n}\| = 1$, thus $\dot{\hat{n}} \cdot \hat{n} = 0$ and $\pi[\hat{n}] \cdot \dot{\hat{n}} = \dot{\hat{n}}$. Next, let us multiply Eq. (3.1) with $\hat{n}$ and use $\hat{n} \cdot \hat{t} = \hat{n} \cdot \hat{l}$:

$$\nu = \hat{n} \cdot \left(\hat{t} + \hat{l}\right) = 2\hat{n} \cdot \hat{t}. \tag{3.3}$$

Substituting ν from Eq. (3.3) into Eq. (3.2), we arrive at

$$\dot{\hat{n}} = \pi[\hat{n}] \cdot \left(\dot{\hat{t}} + \dot{\hat{l}}\right) / \left(2\hat{n} \cdot \hat{t}\right). \tag{3.4}$$

Next, let us differentiate the first equality in Eq. (2.1) and project it using $\pi[\hat{t}]$:

$$\pi[\hat{t}] \cdot \dot{\boldsymbol{r}} = \pi[\hat{t}] \cdot \left(\dot{\boldsymbol{c}} + \dot{\lambda}\, \hat{t} + \lambda\, \dot{\hat{t}}\right) = \pi[\hat{t}] \cdot \dot{\boldsymbol{c}} + \lambda\, \dot{\hat{t}}, \text{ or}$$

$$\dot{\hat{t}} = \lambda^{-1} \pi[\hat{t}] \cdot \left(\dot{\boldsymbol{r}} - \dot{\boldsymbol{c}}\right). \tag{3.5}$$

Again, we exploit that $\dot{\hat{t}} \cdot \hat{t} = 0$ and $\pi[\hat{t}] \cdot \dot{\hat{t}} = \dot{\hat{t}}$. By analogy, we may differentiate the second equality of Eq. (2.1), project it with $\pi[\hat{l}]$, and use $\dot{\boldsymbol{s}} \equiv \boldsymbol{0}$ to find

$$\dot{\hat{l}} = \mu^{-1} \pi[\hat{l}] \cdot \dot{\boldsymbol{r}}. \tag{3.6}$$

Now we may substitute Eqs. (3.5) and (3.6) into Eq. (3.4):

$$\dot{\hat{n}} = \pi[\hat{n}] \cdot \left[\lambda^{-1}\pi[\hat{t}] \cdot (\dot{\boldsymbol{r}} - \dot{\boldsymbol{c}}) + \mu^{-1}\pi[\hat{l}] \cdot \dot{\boldsymbol{r}}\right] / \left(2\hat{n} \cdot \hat{t}\right). \tag{3.7}$$

This expression involves the operator products $\pi[\hat{n}] \cdot \pi[\hat{t}]$ and $\pi[\hat{n}] \cdot \pi[\hat{l}]$. In order to eliminate the source-dependent operator $\pi[\hat{l}]$, we must prove that

$$\pi[\hat{n}] \cdot \pi[\hat{l}] \cdot \dot{\boldsymbol{r}} = \pi[\hat{n}] \cdot \pi[\hat{t}] \cdot \dot{\boldsymbol{r}}. \tag{3.8}$$

Indeed, since $\dot{\boldsymbol{r}}$ is tangential, and $\hat{n}$ is orthogonal to the mirror surface, then $\dot{\boldsymbol{r}} \cdot \hat{n} = 0$. Further, the reflection law Eq. (3.1) implies that $\hat{t} = \boldsymbol{t}_{\parallel\hat{n}} + \boldsymbol{t}_{\perp\hat{n}}$ and $\hat{l} = \boldsymbol{t}_{\parallel\hat{n}} - \boldsymbol{t}_{\perp\hat{n}}$, where $\hat{n} \cdot \boldsymbol{t}_{\perp\hat{n}} = 0$ and $\boldsymbol{t}_{\parallel\hat{n}} = (\nu/2)\,\hat{n}$. Combined, this leads to the equality

$$\hat{t} \cdot \dot{\boldsymbol{r}} = \boldsymbol{t}_{\perp\hat{n}} \cdot \dot{\boldsymbol{r}} = -\hat{l} \cdot \dot{\boldsymbol{r}}. \tag{3.9}$$

Using Eq. (2.2) and (3.9), we than establish that

$$\begin{aligned}
&\pi[\hat{n}] \cdot \pi[\hat{t}] \cdot \dot{\boldsymbol{r}} = \pi[\hat{n}] \cdot \left[\dot{\boldsymbol{r}} - \hat{t}\left(\hat{t} \cdot \dot{\boldsymbol{r}}\right)\right] = \dot{\boldsymbol{r}} - \boldsymbol{t}_{\perp\hat{n}}\left(\hat{t} \cdot \dot{\boldsymbol{r}}\right), \text{ and}\\
&\pi[\hat{n}] \cdot \pi[\hat{l}] \cdot \dot{\boldsymbol{r}} = \pi[\hat{n}] \cdot \left[\dot{\boldsymbol{r}} - \hat{l}\left(\hat{l} \cdot \dot{\boldsymbol{r}}\right)\right] = \dot{\boldsymbol{r}} + \boldsymbol{t}_{\perp\hat{n}}\left(\hat{l} \cdot \dot{\boldsymbol{r}}\right) = \dot{\boldsymbol{r}} - \boldsymbol{t}_{\perp\hat{n}}\left(\hat{t} \cdot \dot{\boldsymbol{r}}\right).
\end{aligned}$$

This proves Eq. (3.8) and allows us to transform Eq. (3.7) into

$$\dot{\hat{n}} = \pi[\hat{n}] \cdot \pi[\hat{t}] \cdot \left[\lambda^{-1}\left(\dot{\boldsymbol{r}} - \dot{\boldsymbol{c}}\right) + \mu^{-1}\dot{\boldsymbol{r}}\right] / \left(2\hat{n} \cdot \hat{t}\right). \tag{3.10}$$

Our primary quantity of interest is the SF, defined as $\boldsymbol{f} \equiv \dot{\hat{t}}$, i.e. it is the variation of the sight ray direction caused by the motion of the camera. From Eq. (3.5),

$$\lambda\,\boldsymbol{f} = \pi[\hat{t}] \cdot (\dot{\boldsymbol{r}} - \dot{\boldsymbol{c}}). \tag{3.11}$$

Applying $\mathcal{P}$ on Eq. (3.11), using Eq. (2.3) and the fact that $\dot{\boldsymbol{r}} \perp \hat{n}$, we find:

$$\lambda\mathcal{P} \cdot \boldsymbol{f} = \mathcal{P} \cdot \pi[\hat{t}] \cdot (\dot{\boldsymbol{r}} - \dot{\boldsymbol{c}}) = \dot{\boldsymbol{r}} - \mathcal{P} \cdot \dot{\boldsymbol{c}}, \quad \text{or} \quad \dot{\boldsymbol{r}} = \mathcal{P} \cdot (\lambda\boldsymbol{f} + \dot{\boldsymbol{c}}). \tag{3.12}$$

Now we may plug $\dot{\hat{n}}$ of Eq. (2.4) and $\dot{\boldsymbol{r}}$ of Eq. (3.12) into Eq. (3.10) and simplify:

$$\begin{aligned}
&\dot{\hat{n}} = W \cdot \dot{\boldsymbol{r}} = W \cdot \mathcal{P} \cdot (\lambda\boldsymbol{f} + \dot{\boldsymbol{c}}) = \frac{1}{2}\,\Pi \cdot \left[\kappa\,\mathcal{P} \cdot (\lambda\boldsymbol{f} + \dot{\boldsymbol{c}}) - \frac{\dot{\boldsymbol{c}}}{\lambda}\right], \text{ and}\\
&\lambda^2\left[2W - \kappa\Pi\right] \cdot \mathcal{P} \cdot \boldsymbol{f} = \left[-2\lambda W \cdot \mathcal{P} + \lambda\kappa\Pi \cdot \mathcal{P} - \Pi\right] \cdot \dot{\boldsymbol{c}}.
\end{aligned}$$

(In the above formula, we introduce $\kappa = \lambda^{-1} + \mu^{-1}$). Using again the properties of the projectors Eq. (2.3), we arrive at the final invariant SF equation:

$$\lambda^2 \left[2W - \kappa\Pi\right] \cdot \mathcal{P} \cdot \boldsymbol{f} = \left[(\lambda\kappa - 1)\,\Pi - 2\lambda W \cdot \mathcal{P}\right] \cdot \dot{\boldsymbol{c}}. \tag{3.13}$$

In case the source S is located very far away from the mirror and the camera ($\mu \to \infty$), then $\kappa = \lambda^{-1}$, and we obtain an even simpler expression

$$\left[2\lambda W - \Pi\right] \cdot \mathcal{P} \cdot \boldsymbol{f} = -2W \cdot \mathcal{P} \cdot \dot{\boldsymbol{c}}.$$

4 Two-dimensional SF equation

First of all, we notice that Eq. (3.13) is linear in the SF and the camera velocity, as should be expected from the lowest-order approximation. Further, since both $\dot{\boldsymbol{c}}$ and $\boldsymbol{f}$ are multiplied with $\mathcal{P}$ or Π, their components parallel to $\hat{t}$ do not play any role, and Eq. (3.13) is equivalent to a system of two equations. We therefore cannot e.g. solve Eq. (3.13) for $\boldsymbol{f}$, since the matrix $[2W - \kappa\Pi] \cdot \mathcal{P}$ has no inverse.

In order to move further, let us assume some coordinate frame such that $\boldsymbol{c} = (0, 0, 0)^T$ and $\hat{t} = (0, 0, 1)^T$. The SF and the camera motion in this frame are $\boldsymbol{f} = (f_1, f_2, f_3)^T$ and $\dot{\boldsymbol{c}} = (\dot{c}_1, \dot{c}_2, \dot{c}_3)^T$. Since only the components of $\boldsymbol{f}$, $\dot{\boldsymbol{c}}$ orthogonal to $\hat{t}$ (i.e. the first and the second) matter, we may re-write Eq. (3.13) as follows:

$$\lambda^2 \left[2W_{(2)} - \kappa\Pi_{(2)}\right] \cdot \boldsymbol{f}_{(2)} = \left[(\lambda\kappa - 1)\,\Pi_{(2)} - 2\lambda W_{(2)}\right] \cdot \dot{\boldsymbol{c}}_{(2)}. \tag{4.1}$$

Here $\boldsymbol{f}_{(2)} = (f_1, f_2)^T$, $\dot{\boldsymbol{c}}_{(2)} = (\dot{c}_1, \dot{c}_2)^T$, and all operators with subscript (2) are 2×2 matrices obtained by dropping the last row and the last column from the 3×3 representation. In particular, $\mathcal{P}_{(2)}$ becomes a unit matrix and drops out, and $W_{(2)}$ is a matrix known as the second fundamental form of the surface.

Since the SF equation is two-dimensional, we may at any camera position record at most two independent SF vectors $\boldsymbol{f}_{(2)}^{(1)}$ and $\boldsymbol{f}_{(2)}^{(2)}$, corresponding to the two independent motion vectors $\dot{\boldsymbol{c}}_{(2)}^{(1)}$ and $\dot{\boldsymbol{c}}_{(2)}^{(2)}$. Having made two such observations, we combine them into the matrices $F = \left(\boldsymbol{f}_{(2)}^{(1)}, \boldsymbol{f}_{(2)}^{(2)}\right)$ and $\dot{C} = \left(\dot{\boldsymbol{c}}_{(2)}^{(1)}, \dot{\boldsymbol{c}}_{(2)}^{(2)}\right)$ so that

$$\lambda^2 \left[2W_{(2)} - \kappa\Pi_{(2)}\right] \cdot \tilde{F} = (\lambda\kappa - 1)\,\Pi_{(2)} - 2\lambda W_{(2)}, \tag{4.2}$$

where $\tilde{F} = F \cdot \left(\dot{C}\right)^{-1}$ is a normalized SF corresponding to the two orthogonal unit camera motions. Now the criterion for the existence of a finite SF is straightforward: if the matrix $\left[2W_{(2)} - \kappa\Pi_{(2)}\right]$ is singular, the SF diverges. This happens, when κ is an eigenvalue of the matrix $2W_{(2)} \cdot \Pi_{(2)}^{-1}$. Such condition is more complex than the identification of "parabolic points", commonly believed to cause infinite SF, and involves (via κ) the positions of the camera and the source.

In case $\mu \to \infty$ Eq. (4.2) simplifies to:

$$\left[2\lambda W_{(2)} - \Pi_{(2)}\right] \cdot \tilde{F} = -2W_{(2)}. \tag{4.3}$$

5 Gaussian curvature and the SF

Let us re-write Eq. (4.2) as

$$2\lambda W_{(2)} \cdot \left[\lambda\tilde{F} + I_{(2)}\right] = \Pi_{(2)} \cdot \left[(\lambda\kappa - 1)I_{(2)} + \lambda^2\kappa\tilde{F}\right], \tag{5.1}$$

where $I_{(2)}$ is a 2×2 unit matrix. If we introduce explicitly the components of the unit normal vector $\hat{n} = (n_1, n_2, n_3)^T$, the projector becomes

$$\Pi_{(2)} = \frac{1}{n_3} \begin{pmatrix} 1 - n_1^2 & -n_1 n_2 \\ -n_1 n_2 & 1 - n_2^2, \end{pmatrix},$$

which implies $\det \Pi_{(2)} = 1$. Taking the determinant of Eq. (5.1), we end up with

$$\det W_{(2)} = \det\left[(\lambda\kappa - 1)I_{(2)} + \lambda^2\kappa\tilde{F}\right] \det\left[2\lambda\left(\lambda\tilde{F} + I_{(2)}\right)\right]^{-1},$$

or, in the limit $\mu \to \infty$,

$$\det W_{(2)} = \det\left[\tilde{F}\right] \det\left[2\left(\lambda\tilde{F} + I_{(2)}\right)\right]^{-1}. \tag{5.2}$$

Eq. (5.2) is quite remarkable: $\det W_{(2)}$ is the Gaussian curvature, which is an intrinsic property of the surface at point $\boldsymbol{r}$, independent of how the surface is embedded in the 3D space. The right-hand side depends only on the measured SF matrix $\tilde{F}$ and the distance λ. Let us now consider three special cases:

- If the surface is nearly flat, the observed SF should be close to zero, and $\|\lambda\tilde{F}\| \ll 1$. The distance-dependent term then drops out, and we get

  $$\det W_{(2)} = \frac{1}{4}\det\tilde{F}.$$

  This equation suggests a direct method to measure small residual curvatures of nearly planar surfaces.

- If instead $\|\lambda\tilde{F}\| \gg 1$, which happens e.g. for a camera positioned near the focus point of a spherical mirror of radius 2λ (i.e. all camera sight rays are reflected in the same direction), the SF-dependence disappears, as expected:

$$\det W_{(2)} = \frac{1}{4\lambda^2}.$$

- If the surface is developable, i.e. it may be locally "unfolded" into a flat sheet without stretching (this happens when it is locally a cylinder), its Gaussian curvature is zero. According to Eq. 5.2, this implies that $\det \tilde{F} = 0$. The two recorded SF vectors then must be collinear, or, alternatively, there must exist a direction of the camera motion such that the respective SF vanishes. Intuitively, if the camera moves along the cylinder axis, then indeed the reflection of the infinitely remote environment must remain static on the sensor, and the SF is zero.

On a more general level, the derived relations between the SF and the curvature allow us to differentiate the SF-based reconstruction from the other optical shape measurement techniques as follows. The triangulation-based methods (as well as interferometry) are directly sensitive to the point position in space (or the zeroth-order derivative). In deflectometry, the primary measured quantity is the normal vector (the first-order surface derivatives). The SF, as shown above, is directly sensitive to the second order derivatives which enter the Weingarten map and the curvatures. Therefore, the error profiles and the dominant error sources of any SF-based methods will be different from the alternatives and may justify their use in some applications.

6 Explicit coordinate form

In order to obtain the explicit expressions for the Weingarten map and the normal vector, we need to assume some parametrization of the reflecting surface near $\boldsymbol{r}$ in terms of two intrinsic surface coordinates u and v. Given some function $\boldsymbol{r}(u, v)$,

$$W_{(2)} = \frac{1}{EG - F^2} \begin{pmatrix} eG - fF & fG - gF \\ fE - eF & gE - fF \end{pmatrix},$$

where $E = \boldsymbol{r}_u \cdot \boldsymbol{r}_u$, $F = \boldsymbol{r}_u \cdot \boldsymbol{r}_v$, $G = \boldsymbol{r}_v \cdot \boldsymbol{r}_v$, $e = \boldsymbol{r}_{uu} \cdot \hat{n}$, $f = \boldsymbol{r}_{uv} \cdot \hat{n}$, $g = \boldsymbol{r}_{vv} \cdot \hat{n}$, and the subscripts stand for the partial derivatives such as $\boldsymbol{r}_{uv} \equiv \partial_u \partial_v \boldsymbol{r}$.

One convenient parametrization identifies u and v with the sensor coordinates of a pinhole camera located at $\boldsymbol{c}$ such that its central direction coincides with $\hat{t}$, and the u- and v-axes are collinear with the global x- and y-axes, respectively. The surface geometry near $\boldsymbol{r}(0,0)$ is described by a two-dimensional depth map $s(u,v)$:

$$\begin{aligned} \boldsymbol{r}(u,v) &= s(u,v)\,(u,v,1)^T, \text{where} \\ s(u,v) &= s_0 + us_u + vs_v + \frac{u^2}{2}s_{uu} + \frac{v^2}{2}s_{vv} + uv\,s_{uv}. \end{aligned}$$

This is precisely the parametrization suggested in [Pak14a]. If one expresses $W_{(2)}$ and the components of the normal vector $\hat{n} = (\boldsymbol{r}_v \times \boldsymbol{r}_u)\,/\,|\boldsymbol{r}_v \times \boldsymbol{r}_u|$ in terms of the shape parameters s_0, ..., s_{uv}, and substitutes them into Eq. (4.1) in the limit $\mu \to \infty$, one arrives at the explicit form of variables and the operators in Eq. (4.3):

$$\begin{aligned} \lambda &= s_0,\ \ W_{(2)} = s^{-1}\left(s_0^2 + s_u^2 + s_v^2\right)^{-3/2} M, \text{ where} \\ M &= s_0^2 \begin{pmatrix} s_{uu} & s_{uv} \\ s_{uv} & s_{vv} \end{pmatrix} - 2s_0 \begin{pmatrix} s_u^2 & s_u s_v \\ s_u s_v & s_v^2 \end{pmatrix} \\ &\quad + \begin{pmatrix} s_v\left(s_v s_{uu} - s_u s_{uv}\right) & s_v\left(s_v s_{uv} - s_u s_{vv}\right) \\ -s_u\left(s_v s_{uu} - s_u s_{uv}\right) & -s_u\left(s_v s_{uv} - s_u s_{vv}\right) \end{pmatrix}, \text{ and} \\ \Pi_{(2)} &= s^{-1}\left(s_0^2 + s_u^2 + s_v^2\right)^{-1/2} \begin{pmatrix} -s_0^2 - s_v^2 & s_u s_v \\ s_u s_v & -s_0^2 - s_u^2 \end{pmatrix}. \end{aligned}$$

Solving Eq. (4.1) with these operators for the components of $\boldsymbol{f}_{(2)}$ as a function of $\dot{\boldsymbol{c}}_{(2)}$, one exactly reproduces $\boldsymbol{f}(0,0)$ in Eq. (3.1) of that reference. Our present calculation is therefore equally powerful, but provides more insight into the structure of the SF equations.

7 Conclusion

In this report, we provided a detailed derivation of how the SF in a point depends on the observation parameters and the properties of the reflecting surface. Unlike previous results, the presented equations are explicitly invariant and compact, allowing a deeper insight into their structure. In particular, we formulated the simple condition for the SF to diverge, and highlighted the relation between the SF and the Gaussian curvature of the surface. The future work will be related to finding ways to recover the 3D shape of the surface based on the observed SF.

Bibliography

[AVBSZ07] Yair Adato, Yuriy Vasilyev, Ohad Ben-Shahar, and Todd Zickler. Toward a theory of shape from specular flow. Proc. ICCV, 2007.

[AVZBS10] Yair Adato, Yuriy Vasilyev, Todd Zickler, and Ohad Ben-Shahar. Shape from specular flow. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32:2054 – 2070, 2010.

[AZBS11] Yair Adato, Todd Zickler, and Ohad Ben-Shahar. A polar representation of motion and implications for optical flow. Proc. CVPR, 2011.

[BB91] Andrew Blake and Heinrich Bülthoff. Shape from specularities: Computation and psychophysics. *Phil. Trans. R. Soc. Lond. B*, 331:237–252, 1991.

[CnVA+09] Guillermo D. Cañas, Yuriy Vasilyev, Yair Adato, Todd Zickler, Steven J. Gortler, and Ohad Ben-Shahar. A linear formulation of shape from specular flow. Proc. ICCV, pages 191–198, 2009.

[LBRB08] J. Lellmann, J. Balzer, A. Rieder, and J. Beyerer. Shape from specular reflection and optical flow. *Int J. Comput. Vis.*, 80:226–241, 2008.

[Pak14a] Alexey Pak. Local shape from specular flow. Proc. 2013 Joint Workshop of Fraunhofer IOSB and Institute for Anthropomatics, Vision and Fusion Laboratory, pages 53–59, 2014.

[RB06] Stefan Roth and Michael J. Black. Specular flow and the recovery of surface structure. Proc. CVPR, pages 1869–1876, 2006.

[VZGBS11] Yuriy Vasilyev, Todd Zickler, Steven Gortler, and Ohad Ben-Shahar. Shape from specular flow: Is one flow enough? Proc. CVPR, pages 2561–2568, 2011.