Optical Flow on Evolving Surfaces with an Application to the Analysis of 4D Microscopy Data

Clemens Kirisits¹, Lukas F. Lang¹, and Otmar Scherzer^{1,2}

¹Computational Science Center, University of Vienna, Nordbergstr. 15, 1090 Vienna, Austria ²Radon Institute of Computational and Applied Mathematics, Austrian Academy of Sciences, Altenberger Str. 69, 4040 Linz, Austria

May 21, 2013

Abstract

We extend the concept of optical flow to a dynamic non-Euclidean setting. Optical flow is traditionally computed from a sequence of flat images. It is the purpose of this paper to introduce variational motion estimation for images that are defined on an evolving surface. Volumetric microscopy images depicting a live zebrafish embryo serve as both biological motivation and test data.

Keywords: Computer Vision, biomedical imaging, optical flow, variational methods, evolving surfaces, zebrafish, laser-scanning microscopy.

1 Introduction

Advances in laser-scanning microscopy and fluorescent protein technology have increased resolution of microscopy imaging up to a single cell level [11]. They allow for four-dimensional (volumetric time-lapse) imaging of living organisms and shed light on cellular processes during early embryonic development. Understanding cellular development often requires estimation and analysis of cell motion. However, the amount of data captured is tremendous and therefore manual analysis is not an option.

The specific biological motivation for this work is to understand the motion and division behaviour of fluorescently labelled endodermal cells of a zebrafish embryo. The marked cells develop on the surface of the embryo's yolk, where they form a non-contiguous monolayer [17]. Loosely speaking, they only sit next to each other but not on top of each other. Moreover, the yolk deforms over time; see Fig. 1.



Figure 1: Sequence of embryonic zebrafish images. The curved mesh represents a section of the yolk's surface. Depicted are frames no. 30, 45, 55, and 60 of the entire sequence. All dimensions are in micrometer (μ m). See Sec. 4.1 for more details on the microscopy data.

We take these biological facts into account and restrict our attention to the analysis of cell motion on the yolk's surface. With this approach it is possible to reduce the amount of data by one space dimension. The resulting problem consists in the estimation of motion of brightness patterns that are restricted to an itself moving surface. We approach this problem by adapting the classical concept of optical flow to the present setting, where the image domain is both non-Euclidean and dynamic. Note that due to the monolayer structure cell occlusions cannot occur. This makes the optical flow field a more reliable approximation to the true motion field.

Our contributions in the field of optical flow are as follows. First, we formulate the optical flow problem on an evolving two-dimensional manifold and give two equivalent ways of linearising the brightness constancy assumption (Secs. 2.1 and 2.2). One uses a parametrisation of the evolving surface, the other one is parameter-independent. Second, we use a generalisation of the Horn-Schunck model to regularise the optical flow field (Sec. 2.3). For a given global parametrisation of the evolving surface, we solve the associated Euler-Lagrange equations in the parameter domain with a finite difference scheme (Sec. 3). Finally, we apply this technique to obtain qualitative results from the afore-mentioned zebrafish data (Sec. 4). Our experiments show that the optical flow is an appropriate tool for analysing these data. It is capable of estimating global trends as well as individual cell movements and, in particular, it is able to indicate cell division events.

1.0.1 Related work.

Optical flow is the apparent motion in a sequence of images. Its estimation is a key problem in Computer Vision. Horn and Schunck [5] were the first to propose a variational approach assuming constant brightness of moving points and spatial smoothness of the velocity field. Since then, a vast number of modifications has been developed. See [1] for a recent survey.

Miura [13] observed that until 2005 optical flow has been mostly disregarded as a method for motion extraction in cell biological data. Since then, a few articles have explored this direction: Melani et al. [12] and Hubený et al. [6] extended variational optical flow methods to volumetric images to obtain 3D displacement fields. In the former article, the resulting algorithm is also applied to zebrafish microscopy data. Quelhas et al. [15] use optical flow to detect cell divisions in a live plant root. However, they work with 2D (plus time) data only. Therefore, their approach suffers from errors caused by 3D off-plane motion.

Clearly, certain natural scenarios are more accurately described by a velocity field on a non-flat surface rather than on a flat domain. With applications to robot vision, Imiya et al. [7, 16] considered optical flow for spherical images. In a more general setting, Lefèvre and Baillet [10] extended the Horn-Schunck method to 2-Riemannian manifolds and showed well-posedness. They solve the numerical problem with finite elements on a surface triangulation. In all of the above works the underlying imaging surface is fixed over time, while in this paper it is not.

2 Optical Flow on Evolving Surfaces

2.1 Brightness Constancy

Let $\mathcal{M}_t \subset \mathbb{R}^3$, $t \in I = [0, T)$, be a compact smooth two-dimensional manifold evolving smoothly over time. We assume the velocity to be unknown. Moreover, denote by \tilde{f} a scalar time-dependent quantity defined on the surface

$$\tilde{f}: \bigcup_{t\in I} \left(\mathcal{M}_t \times \{t\}\right) \to \mathbb{R}.$$

We begin with a Lagrangian specification of the optical flow field. That is, for every starting point $x_0 \in \mathcal{M}_0$ we seek a trajectory where the data \tilde{f} are conserved. More precisely, we want to find a function

$$\gamma\colon \mathcal{M}_0\times I\to \bigcup_{t\in I}\mathcal{M}_t,$$

such that

- 1. $\gamma(\boldsymbol{x}_0, t) \in \mathcal{M}_t$ for all $t \in I$, for all $\boldsymbol{x}_0 \in \mathcal{M}_0$,
- 2. $\gamma(\cdot, t)$ is a diffeomorphism between \mathcal{M}_0 and \mathcal{M}_t for all $t \in I$,
- 3. $\gamma(\cdot, 0) = \mathrm{Id}_{\mathcal{M}_0},$

is fulfilled and which satisfies a "brightness" constancy assumption (BCA)

$$f(\boldsymbol{x}_0, 0) = f(\gamma(\boldsymbol{x}_0, t), t), \text{ for all } (\boldsymbol{x}_0, t) \in \mathcal{M}_0 \times I.$$
(1)

In classical optical flow computations it is common practice to linearise the BCA by taking its time derivative and to solve the resulting equation for the Eulerian unknown $\dot{\gamma}^{,1}$ We also take this route, but differentiation of \tilde{f} is more involved. Observe, for example, that for an arbitrary $t_0 \in I$ and $\boldsymbol{x} \in \mathcal{M}_{t_0}$ the usual partial derivative

$$\partial_t \tilde{f}(\boldsymbol{x}, t_0) = \lim_{h \to 0} \frac{1}{h} \left(\tilde{f}(\boldsymbol{x}, t_0 + h) - \tilde{f}(\boldsymbol{x}, t_0) \right)$$

is not well-defined, simply because, in general, \boldsymbol{x} is not an element of \mathcal{M}_{t_0+h} for all $h \neq 0$.

In the next section we linearise (1) in two different ways. First, we use a global parametrisation to pull the data back to a fixed reference domain and linearise afterwards. In our second approach we borrow some notions from continuum mechanics [2] to directly linearise (1).

2.2 Linearisation

Linearisation after pull-back. Let $\Omega \in \mathbb{R}^2$ be a compact domain and

$$\boldsymbol{x} \colon \Omega \times I \to \mathbb{R}^3, \ (x_1, x_2, t) = (x, t) \mapsto \boldsymbol{x}(x, t) \in \mathcal{M}_t$$

be a parametrisation of the evolving surface. Denote by f the coordinate representation of \tilde{f} , that is,

$$f(x,t) = f(\boldsymbol{x}(x,t),t)$$
(2)

and let

$$\beta \colon \Omega \times I \to \Omega$$

¹To simplify expressions we use Newton's notation for those time derivatives that correspond to actual velocities, for example $\dot{\gamma} = \partial_t \gamma$.



Figure 2: Commutative diagram describing the relation between unknowns β and γ .

be the coordinate counterpart of γ . This means, if we let $\boldsymbol{x}_0 = \boldsymbol{x}(x_0, 0)$, then $\beta(x_0, t)$ gives the coordinates of $\gamma(\boldsymbol{x}_0, t) \in \mathcal{M}_t$ in Ω (see Fig. 2). In other words, we have the identity

$$\gamma(\boldsymbol{x}(x_0, 0), t) = \boldsymbol{x}(\beta(x_0, t), t), \quad \text{for all } (x_0, t) \in \Omega \times I.$$
(3)

Now, from (1), (2) and (3) we get

$$\begin{split} f(x_0,0) &= \tilde{f}(\boldsymbol{x}_0,0) \\ &= \tilde{f}(\gamma(\boldsymbol{x}_0,t),t) \\ &= \tilde{f}(\boldsymbol{x}(\beta(x_0,t),t),t) \\ &= f(\beta(x_0,t),t), \end{split}$$

which is a coordinate version of the BCA. After differentiation with respect to t it becomes $\nabla^2 t + \hat{c} + 0 + c = 0$

$$\nabla^2 f \cdot \beta + \partial_t f = 0, \tag{4}$$

where $\nabla^2 = (\partial_1, \partial_2)^{\top}$ is the two-dimensional spatial gradient. Note that the last equation is nothing but the classical optical flow constraint (OFC) for Euclidean data f and a displacement field $\dot{\beta}$.

Direct linearisation. We turn to our second derivation. While, as pointed out above, the partial derivative $\partial_t \tilde{f}$ is undefined in general, it does make sense to differentiate \tilde{f} following the surface movement. Let \boldsymbol{y} be a point on \mathcal{M}_{t_0} and $\boldsymbol{\xi}: t \mapsto \boldsymbol{\xi}(t) \in \mathcal{M}_t$ an arbitrary smooth trajectory through the evolving surface satisfying $\boldsymbol{\xi}(t_0) = \boldsymbol{y}$. Now we can compute

$$\left. \frac{\mathrm{d}}{\mathrm{d}t} \tilde{f}(\xi(t), t) \right|_{t=t_0} = \lim_{h \to 0} \frac{1}{h} \left(\tilde{f}(\xi(t_0 + h), t_0 + h) - \tilde{f}(\boldsymbol{y}, t_0) \right)$$

to obtain a valid derivative of \tilde{f} . Since this time derivative only depends on the vector $\boldsymbol{v} = \dot{\xi}(t_0)$, we denote it by $d_t^{\boldsymbol{v}} \tilde{f}$. A natural candidate for a trajectory along which to differentiate is given by the parametrisation $\xi(t) = \boldsymbol{x}(x,t)$. Another possible choice would be a trajectory that is normal to \mathcal{M}_{t_0} . The resulting normal time derivative is accordingly denoted by $d_t^{\boldsymbol{v}} \tilde{f}$.

Finally, we also need the surface gradient $\nabla_{\mathcal{M}} \tilde{f}$. If F is a smooth extension of \tilde{f} to an open neighbourhood of $\boldsymbol{y} \in \mathcal{M}_{t_0}$ in \mathbb{R}^3 , then the surface gradient of Fat \boldsymbol{y} is defined as the projection of the three-dimensional spatial gradient $\nabla^3 F$ onto the tangent plane to \mathcal{M}_{t_0}

$$\nabla_{\mathcal{M}}F = \nabla^3 F - (\nabla^3 F \cdot \hat{\boldsymbol{n}})\hat{\boldsymbol{n}},$$

where $\hat{\boldsymbol{n}}$ is the unit normal to \mathcal{M}_{t_0} . The surface gradient only depends on the values of F on the surface; see e.g. [4, p. 389]. Thus, $\nabla_{\mathcal{M}}\tilde{f} = \nabla_{\mathcal{M}}F$ is well-defined.

The spatial and temporal derivatives of f introduced above are related in a simple way. As shown in [2], they satisfy the equality

$$d_t^{\hat{\boldsymbol{x}}} \tilde{f} = \nabla_{\mathcal{M}} \tilde{f} \cdot \dot{\boldsymbol{x}} + d_t^{\boldsymbol{n}} \tilde{f} = \nabla_{\mathcal{M}} \tilde{f} \cdot \dot{\boldsymbol{x}}_{tan} + d_t^{\boldsymbol{n}} \tilde{f},$$
(5)

where $\dot{\boldsymbol{x}}_{tan}$ is the tangential surface velocity, that is, the projection of $\dot{\boldsymbol{x}}$ onto the tangent plane to \mathcal{M}_{t_0} . This decomposition of $d_t^{\dot{\boldsymbol{x}}} \tilde{f}$ into normal and tangential components is clearly valid for any trajectory in place of \boldsymbol{x} , and therefore in particular for the unknown γ . This means we can use (5) in order to differentiate the BCA (1) with respect to t. The resulting OFC reads

$$\nabla_{\mathcal{M}}\tilde{f}\cdot\dot{\gamma}_{\mathrm{tan}} + \mathrm{d}_{t}^{\boldsymbol{n}}\tilde{f} = 0.$$
(6)

Discussion. We conclude this section with a brief comparison of the two OFCs derived above. We start by showing how to obtain (4) from (6) and vice versa. To this end we again assume the existence of a global parametrisation and rewrite all quantities in (6) in terms of \boldsymbol{x} . First observe that, by (3), the velocity of γ equals the surface velocity $\dot{\boldsymbol{x}}$ plus a purely tangential component

$$\dot{\gamma} = \dot{x} + J\dot{\beta}$$

where $J = (\partial_1 \boldsymbol{x} \ \partial_2 \boldsymbol{x})$ is the Jacobian matrix of \boldsymbol{x} with respect to \boldsymbol{x} . On the other hand, by (5), the normal time derivative is equal to the time derivative of \tilde{f} following \boldsymbol{x} minus its tangential component

$$\mathbf{d}_t^{\boldsymbol{n}}\tilde{f} = \mathbf{d}_t^{\boldsymbol{\dot{x}}}\tilde{f} - \nabla_{\mathcal{M}}\tilde{f}\cdot\boldsymbol{\dot{x}}.$$

Using the last two equations to rewrite the left-hand side of (6) yields

$$\nabla_{\mathcal{M}}\tilde{f}\cdot\dot{\gamma} + \mathbf{d}_{t}^{\boldsymbol{n}}\tilde{f} = \nabla_{\mathcal{M}}\tilde{f}\cdot\left(\dot{\boldsymbol{x}}+J\dot{\beta}\right) + \mathbf{d}_{t}^{\dot{\boldsymbol{x}}}\tilde{f} - \nabla_{\mathcal{M}}\tilde{f}\cdot\dot{\boldsymbol{x}}$$
$$= \nabla_{\mathcal{M}}\tilde{f}\cdot J\dot{\beta} + \mathbf{d}_{t}^{\dot{\boldsymbol{x}}}\tilde{f},$$

which is already the left-hand side of (4) in terms of \tilde{f} . It only remains to observe that $d_t^{\dot{x}}\tilde{f} = \partial_t f$ and to replace the surface gradient $\nabla_{\mathcal{M}}\tilde{f}$ by its coordinate expression $Jg^{-1}\nabla^2 f$, where $g = J^{\top}J$ is the coefficient matrix of the Riemannian metric; see e.g. [9].

We highlight the qualitative difference between the constraints (4) and (6). Note that in the former the unknown is $\dot{\beta}$, while in the latter it is $\dot{\gamma}_{tan} = \dot{x}_{tan} + J\dot{\beta}$. This means that (4) constrains the motion relative to the tangential surface velocity \dot{x}_{tan} , while (6) constrains the absolute tangential motion.

The nature of our microscopy data suggests a simple global parametrisation (see Sec. 3). We therefore pull the data back to the Euclidean plane and solve (4). However, equation (6) is independent of any parametrisation. It can thus serve as a starting point for alternative numerical approaches.

2.3 Regularisation

From now on we fix an arbitrary $t_0 \in I$ and turn to the actual solution of the parametrised OFC for $(u^1(x), u^2(x))^{\top} = u(x) = \dot{\beta}(x, t_0)$. Recall that with this notation u contains the coefficients of the tangential vector field $\boldsymbol{u} = J\dot{\beta}$ with respect to the tangential basis $(\partial_1 \boldsymbol{x}, \partial_2 \boldsymbol{x})$ of \mathcal{M}_{t_0} . Note also that, by fixing t_0 , there is no more time-dependence in our problem which makes it effectively an optical flow problem on a static surface. Hence we omit any reference to t_0 from now on and write \mathcal{M} instead of \mathcal{M}_{t_0} .

The sought vector field is underdetermined by the OFC alone. We overcome this by minimising a functional that penalises violation of the OFC while imposing an additional smoothness restriction on \boldsymbol{u} . More precisely, we adopt a recent extension of the original quadratic Horn-Schunck regularisation to a Riemannian setting [10]. Basically, they propose to minimise

$$\mathcal{E}(u) = \frac{\alpha}{2} \left\| \nabla^2 f \cdot u + \partial_t f \right\|_{L^2(\mathcal{M})}^2 + \frac{1}{2} \left\| Du \right\|_{L^2(\mathcal{M})}^2.$$
(7)

Here, $\alpha > 0$ is the regularisation parameter and $Du = (D_j u^i)$ is the 2×2 matrix containing the coefficient functions of the covariant derivatives

$$abla_j oldsymbol{u} = \sum_{i=1}^2 D_j u^i \partial_i oldsymbol{x}, \ j = 1, 2,$$

of \boldsymbol{u} . Using the Christoffel symbols Γ_{jk}^i (see Sec. 3) associated to the parametrisation \boldsymbol{x} the coefficients are given by

$$D_j u^i = \partial_j u^i + \sum_{k=1}^2 \Gamma^i_{jk} u^k, \ i, j = 1, 2.$$

Rewriting (7) as an integral over the coordinate domain, we arrive at the functional

$$\mathcal{E}(u) = \frac{1}{2} \int_{\Omega} \left[\alpha \left(\nabla^2 f \cdot u + \partial_t f \right)^2 + \|Du\|_F^2 \right] \sqrt{\det g} \, \mathrm{d}x, \tag{8}$$

where $\|\cdot\|_F$ is the Frobenius norm.

3 Numerical Solution

We solve the problem of minimising functional \mathcal{E} via its associated Euler-Lagrange equations. Regarding the integrand of \mathcal{E} as a function $G(x, u, \nabla^2 u^1, \nabla^2 u^2)$, they read

$$\begin{split} G_{u^1} &= \partial_1 G_{\partial_1 u^1} + \partial_2 G_{\partial_2 u^1} \\ G_{u^2} &= \partial_1 G_{\partial_1 u^2} + \partial_2 G_{\partial_2 u^2}, \end{split}$$

where subscripts of G denote partial derivatives. The resulting pair of linear PDEs is of the form

$$\Delta u^1 = \nabla^2 u^1 \cdot c + \nabla^2 u^2 \cdot d + u \cdot b_1 + a_1$$

$$\Delta u^2 = \nabla^2 u^2 \cdot c + \nabla^2 u^1 \cdot d + u \cdot b_2 + a_2.$$
(9)

The coefficient vectors a, b_1, b_2, c, d are rather lengthy functions of the data f and metric tensor g, which is why we do not write them out in full here. Letting $\Omega = (0, 1)^2$ for simplicity, the natural boundary conditions of the variational problem are

$$\partial_j u^i + \sum_k \Gamma^i_{jk} u^k = 0, \text{ for } x_j \in \{0, 1\},$$
 (10)

where $i, j \in \{1, 2\}$. In case of a flat manifold, e.g. $\mathcal{M} = \Omega$, the Euler-Lagrange equations (9) reduce to those of the original Horn-Schunck functional and the boundary conditions become the usual homogeneous Neumann ones. For more details on the calculus of variations we refer to [3].

Due to the nature of the microscopy data (see Sec. 4.1 and Fig. 1), the manifold \mathcal{M}_t modelling the deforming yolk is a surface with boundary that is most easily parametrised as the graph of a function $z : \Omega \times I \to \mathbb{R}$. Hence, we set $\mathbf{x}(x_1, x_2, t) = (x_1, x_2, z(x_1, x_2, t))^{\top}$. Accordingly, for the metric we get

$$g = I_2 + \nabla^2 z \nabla^2 z^\top, \qquad \det g = 1 + |\nabla^2 z|^2,$$

where $I_2 \in \mathbb{R}^{2 \times 2}$ is the identity matrix. The Christoffel symbols turn out to be

$$\Gamma_{jk}^{i} = \frac{1}{2} \sum_{m=1}^{2} g^{mi} \left(\partial_{j} g_{km} + \partial_{k} g_{mj} - \partial_{m} g_{jk} \right) = \frac{\partial_{i} z \, \partial_{jk} z}{\det g}.$$

Partial derivatives of z and of the projected data f were approximated by central differences. The system (9) with boundary conditions (10) was then solved with a standard finite difference scheme. In the following section numerical results are presented.

4 Experiments

4.1 Data

As mentioned before, the biological motivation for this work are cellular image data of a zebrafish embryo. Endoderm cells expressing green fluorescent protein were recorded via confocal laser-scanning microscopy resulting in timelapse volumetric (4D) images; see [11] for the imaging techniques. This type of image shows a high contrast at cell boundaries and a low signal-to-noise ratio in general. Our videos were obtained during the gastrula period, which is an early stage in the animal's developmental process and takes place approximately five to ten hours post fertilisation. In short, the fish forms on the surface of a spherical-shaped yolk; see e.g. [8] for many illustrations and detailed explanations. For the biological methods such as the fluorescence marker and the embryos used in this work we refer to [14]. The important aspect about endodermal cells is that they are known to form a monolayer during gastrulation [17], meaning that the radial extent is only a single cell. This crucial fact allows for the straightforward extraction of a surface together with a twodimensional image of the stained cells. Since only a cuboid region of approximately $860 \times 860 \times 340 \,\mu\text{m}^3$ of the pole region is captured by the microscope, this surface can easily be parametrised; cf. Sec. 3. The spatial resolution of the Gaussian filtered images is 512×512 pixels and all intensities are given in the interval [0, 1]. Our sequence contains 77 frames recorded in intervals of 240 s with clearly visible cellular movements and cell divisions.

4.2 Numerical results

In the following we present qualitative results and demonstrate the feasibility of our approach. For every subsequent pair of frames we minimised the functional (8) as outlined in Sec. 3. We chose grid size as well as temporal displacement as h = 1 and the regularisation parameter was set to $\alpha = 10$. For demonstration purpose we make use of the standard flow colour-coding [1], which maps (normalised) flow vectors to a colour space defined inside the unit circle. It is easy to see that the same colours are valid all over the manifold due to the parametrisation.

As representative candidates for this discussion we chose the displacement field between frames 57 and 58 for the following reasons. First, the surface is distinctly developed. Second, a considerable number of cells is present in the image, and third, the interval contains cell divisions. Figure 3, left, shows the colour-coded tangential vector field and the colour space whereas Fig. 3, right, displays the same motion field as computed in the parameter space.² A visual inspection of the dataset shows that cells tend to move towards the embryo's body axis, which roughly runs along the main diagonal in Fig. 3, right. Clearly, the velocity field is sufficiently smooth and suggests this behaviour in an adequate manner on a large scale. The expected change in orientation along the body axis is well represented by the colour shift from orange-yellow below the main diagonal to purplish blue in the region above. On the contrary, the choice of the regularisation parameter ensures that individual movements are well preserved as can be observed from the image.

Figure 4 gives a detailed view of the section outlined by a (red) rectangle

²Some figures may appear in colour only in the online version.



Figure 3: Optical flow field between frames 57 and 58 of the sequence. Colours indicate direction whereas darkness of a colour indicates the length of the vector. Note that the colour circle has been enlarged for better visibility.

in Fig. 3, right. This section was chosen because it depicts a cell division. Figure 4, left, and Fig. 4, right, display the frames before and after the event, respectively. Moreover, in Fig. 4, left, the velocity field is shown. From the raw data we observed that when a cell actually splits, the two daughter cells drift apart in a 180° angle with respect to the mother cell. The displacement field clearly shows the anticipated pattern caused by the diverging daughter cells. In Fig. 3, right, the event is point up by two areas which are coloured mutually opposite with respect to the colour space. Our results suggest that cell division can be indicated reasonably well by our model. Both implementation and data are available on our website.³

5 Conclusion

Aiming at efficient motion analysis of 4D cellular microscopy data, we generalised the Horn-Schunck method to videos defined on evolving surfaces. The biological fact that the observed cells move along an itself deforming surface allows for motion estimation in 2D (plus time). In the course of this work, we presented two ways to linearise the brightness constancy assumption and showed that one could be obtained from the other and vice versa. The resulting optical flow constraint was solved by means of quadratic regularisation and verified on the basis of the afore-mentioned data. Our qualitative results suggest that both global trends as well as individual movements including cell division are well shown in the surface velocity field. However, so far we only laid the basic groundwork in terms of a mathematical model.

³http://www.csc.univie.ac.at



Figure 4: Detailed view of a cell division occurring between frames 57 (left) and 58 (right). All vectors are scaled and only every fourth vector is shown. Intensities are interpolated for smooth illustration

Acknowledgements. We thank Pia Aanstad from the University of Innsbruck for sharing her biological insight and for kindly providing the microscopy data. This work has been supported by the Vienna Graduate School in Computational Science (IK I059-N) funded by the University of Vienna. In addition, we acknowledge the support by the Austrian Science Fund (FWF) within the national research networks "Photoacoustic Imaging in Biology and Medicine" (project S10505-N20, Reconstruction Algorithms for PAI) and "Geometry + Simulation" (project S11704, Variational Methods for Imaging on Manifolds).

References

- S. Baker, D. Scharstein, J. P. Lewis, S. Roth, M. J. Black, and R. Szeliski. A Database and Evaluation Methodology for Optical Flow. *Int. J. Comput. Vision*, 92(1):1–31, November 2011.
- [2] P. Cermelli, E. Fried, and M. E. Gurtin. Transport relations for surface integrals arising in the formulation of balance laws for evolving fluid interfaces. J. Fluid Mech., 544:339–351, 2005.
- [3] R. Courant and D. Hilbert. Methods of mathematical physics. Vol. I. Interscience Publishers, Inc., New York, N.Y., 1953.
- [4] D. Gilbarg and N. Trudinger. Elliptic Partial Differential Equations of Second Order. Classics in Mathematics. Springer Verlag, Berlin, 2001. Reprint of the 1998 edition.

- [5] B. K. P. Horn and B. G. Schunck. Determining optical flow. Artificial Intelligence, 17:185–203, 1981.
- [6] J. Hubený, V. Ulman, and P. Matula. Estimating large local motion in livecell imaging using variational optical flow. In VISAPP: Proc. of the Second International Conference on Computer Vision Theory and Applications, pages 542–548. INSTICC, 2007.
- [7] A. Imiya, H. Sugaya, A. Torii, and Y. Mochizuki. Variational analysis of spherical images. In A. Gagalowicz and W. Philips, editors, *Computer Analysis of Images and Patterns*, volume 3691 of *Lecture Notes in Computer Science*, pages 104–111. Springer Berlin, Heidelberg, 2005.
- [8] C. B. Kimmel, W. W. Ballard, S. R. Kimmel, B. Ullmann, and T. F. Schilling. Stages of embryonic development of the zebrafish. *Devel. Dyn.*, 203(3):253–310, 1995.
- [9] J. M. Lee. *Riemannian manifolds*, volume 176 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997. An introduction to curvature.
- [10] J. Lefèvre and S. Baillet. Optical flow and advection on 2-Riemannian manifolds: A common framework. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(6):1081–1092, June 2008.
- [11] S. G. Megason and S. E. Fraser. Digitizing life at the level of the cell: high-performance laser-scanning microscopy and image analysis for in toto imaging of development. *Mech. Dev.*, 120(11):1407–1420, 2003.
- [12] C. Melani, M. Campana, B. Lombardot, B. Rizzi, F. Veronesi, C. Zanella, P. Bourgine, K. Mikula, N. Peyriéras, and A. Sarti. Cells tracking in a live zebrafish embryo. In *Proceedings of the 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS* 2007), pages 1631—1634, 2007.
- [13] K. Miura. Tracking Movement in Cell Biology. In J. Rietdorf, editor, Microscopy Techniques, volume 95 of Advances in Biochemical Engineering/Biotechnology, pages 267–295. Springer, 2005.
- [14] T. Mizoguchi, H. Verkade, J. K. Heath, A. Kuroiwa, and Y. Kikuchi. Sdf1/Cxcr4 signaling controls the dorsal migration of endodermal cells during zebrafish gastrulation. *Development*, 135(15):2521–2529, 2008.
- [15] P. Quelhas, A. M. Mendonça, and A. Campilho. Optical flow based arabidopsis thaliana root meristem cell division detection. In A. Campilho and M. Kamel, editors, *Image Analysis and Recognition*, volume 6112 of *Lecture Notes in Computer Science*, pages 217–226. Springer Berlin Heidelberg, 2010.

- [16] A. Torii, A. Imiya, H. Sugaya, and Y. Mochizuki. Optical Flow Computation for Compound Eyes: Variational Analysis of Omni-Directional Views. In M. De Gregorio, V. Di Maio, M. Frucci, and C. Musio, editors, Brain, Vision, and Artificial Intelligence, volume 3704 of Lecture Notes in Computer Science, pages 527–536. Springer Berlin, Heidelberg, 2005.
- [17] R. M. Warga and C. Nüsslein-Volhard. Origin and development of the zebrafish endoderm. *Development*, 126(4):827–838, February 1999.