



universität
wien

DISSERTATION

Titel der Dissertation

Optical Flow on Evolving Manifolds with an Application
to the Analysis of Fluorescence Microscopy Data

Verfasser

Dipl.-Ing. Lukas Lang, BSc.

angestrebter akademischer Grad

Doktor der Technischen Wissenschaften (Dr. techn.)

Wien, 2015

Studienkennzahl lt. Studienblatt:	A 786 880
Dissertationsgebiet lt. Studienblatt:	Informatik, IK: Computational Science
Betreuer:	Univ.-Prof. Dipl.-Ing. Dr. Otmar Scherzer
Zweitbetreuerin:	Univ.-Prof. Dr. Monika Henzinger

Abstract

Motion estimation is an omnipresent goal in image analysis and computer vision. An important task within is optical flow computation in a sequence of images. It addresses the issue of inferring a vector field from intensity variations, thereby describing the displacements of moving objects. Typically, optical flow is computed in the plane but is readily generalised to non-Euclidean settings allowing, for instance, for cell motion analysis in time-lapse microscopy data.

Today, fluorescence microscopy enables high-resolution observations of biological model organisms, such as the zebrafish, on the scale of single cells. Despite its importance for tissue and organ formation, only little is known about cell migration and proliferation patterns during the zebrafish's early embryonic development. Many of the questions raised involve estimating cell motion.

In view of increasing spatial as well as temporal resolutions resulting in tremendous amounts of data, manual analysis through visual inspection by humans is impracticable. Therefore, automated cell motion estimation is key to the large-scale analysis of above-mentioned data. Optical flow delivers quantitative methods and leads to insights into underlying cellular mechanisms and the dynamic behaviour of cells.

The primary biological motivation for this thesis is the desire to analyse cell motion in a living zebrafish embryo during early embryogenesis. The data at hand depict endodermal cells expressing a green fluorescence protein. Laser-scanning microscopy allows recording (volumetric time-lapse) 4D images of these labelled cells without capturing the background.

During early development these cells float on a so-called *monolayer*, meaning that they form a round surface in a single layer. We exploit this situation and model this layer as a two-dimensional surface deforming over time. The main idea of this thesis is to conceive cell motion only on this evolving surface. As a direct consequence, one is able to reduce the spatial dimension of the data, resulting in more efficient motion estimation of afore-mentioned microscopy data. We formulate the problem of cell motion analysis as a variational optical flow problem on evolving two-dimensional manifolds. Naturally, this surface is subject to geometrical approximations.

In the first part, we focus on the embryo's changing geometry and assume that the cells' layer deforms over time. To this end, we translate the original (Tikhonov-regularised) Horn-Schunck functional and the spatio-temporal extension by Weickert and Schnörr to this non-Euclidean and dynamic setting.

In the second part of this thesis, we pay close attention the topology of the embryo's surface. First, we assume that it is a static round sphere and investigate several vector field decomposition functionals. In particular, we follow recent trends in image decomposition and study $u + v$ and hierarchical

decomposition models for optical flow. The chosen numerical method solves the problem in a finite-dimensional space spanned by tangential vector spherical harmonics and is advantageous in two ways. It provides great flexibility with respect to the regularisation functionals and, as a by-product, yields a Helmholtz decomposition of the flow field.

Second, we consider a more appropriate geometrical model for the zebrafish's embryo, namely evolving sphere-like surfaces. These surfaces can be parametrised from the 2-sphere and constitute a more natural approximation of the embryo's shape. We extend the optical flow functional of Lefèvre and Baillet for surfaces embedded in \mathbb{R}^3 to this new setting. The variational problem is solved by means of a Galerkin method based on tangential vector spherical harmonics. In order to find the sphere-like surfaces from cell microscopy data we devise a method for surface interpolation by means of scalar spherical harmonics expansion.

Finally, we present numerical results on the basis of the afore-mentioned cell microscopy data of a live zebrafish and picture the results in a visually adequate manner.

Preface

This manuscript is a cumulative dissertation. It is a collection of four articles, all of which have been published or submitted to journals. They were written in the course of three years and are the result of a fruitful collaborative effort. The purpose of this preface is to serve as a guide to this dissertation.

This thesis is concerned with the analysis of cell motion in fluorescence microscopy data of a zebrafish by means of variational optical flow. It is structured in three parts.

The first part, the preamble, is intended as an introductory read to this thesis and contains three chapters. In Chap. 1, we introduce challenges in the analysis of fluorescence microscopy data and discuss the biological motivation for this work. Moreover, we present the particular datasets which triggered our interest in computationally feasible methods for their analysis and outline our solution strategy, which exploits a biological fact about the recorded cells. In Chap. 2, we briefly introduce variational optical flow as a reliable and well-established method for motion estimation. Most importantly, in Chap. 3, we discuss the problem in an abstract mathematical setting and present it in a unified notation, explain our efforts to solve it, and summarise the contributions of this thesis.

The second part contains the four publications. They are arranged in chapters and are ordered chronologically. Three articles [50, 51, 52] resulted from joint work with Clemens Kirisits and Otmar Scherzer, whereas [52] is an extended journal version of [50]. They constitute Chaps. 4–6. It needs to be mentioned that these articles were submitted previously as part of a thesis project in [49]. In Chaps. 4 and 5, we transfer optical flow to a new dynamic non-Euclidean setting of evolving surfaces, whereas in Chap. 6 we consider various regularisation functionals for the decomposition of optical flow on the static sphere. Another article [55] was written together with Otmar Scherzer and is presented in Chap. 7. There, we consider the scenario of evolving sphere-like surfaces.

The third part is the appendix. It contains a single bibliography for the whole manuscript, a German translation of the abstract, and a CV with a focus on academia.

At this point let me take the chance to thank a great number of people who contributed, knowingly or unknowingly, to this thesis and who have helped me completing it.

First and foremost I would like to express my sincere gratitude to my supervisor Otmar Scherzer. He has guided me over the years, shared his expertise, and provided valuable advice. I appreciate your patience, your encouragement, and your great sense of humour. For the close and fruitful collaboration I am deeply grateful to my colleague and co-author Clemens Kirisits. I have benefited a great deal from your knowledge and enjoyed the numerous enlightening discussions we had during this time.

My sincere gratefulness goes to all my colleagues at the Scherzer group. I am highly indebted to you for all your contributions. Be it with stimulating discussions, by proofreading my manuscripts, or in administrative manners. I am thankful for the enjoyable time we spent including numerous activities asides from work, all the fun moments, and the great laughs we shared. Furthermore, I am truly grateful to our la Pavoni coffee machine for its loyalty during this period.

I thank Pia Aanstad from the University of Innsbruck for sharing her biological insight and for kindly providing the microscopy data which motivated this work. Moreover, I would like to thank Jean-Francois Aujol, Willi Freeden, and Radu Ioan Boț for agreeing to review this thesis.

I acknowledge the financial support of Monika Henzinger as head of the Vienna Graduate School in Computational Science IK I059-N, funded by the University of Vienna. Further, I was supported by the Austrian Science Fund (FWF) within the FSP S117 – “Geometry and Simulation”.

Lastly, let me express my gratefulness to my family for their enduring encouragement in my academic pursuit and their never-ending support. I also would like to thank all my close friends for their care and all the entertaining moments. Many more to come.

Most of all I am deeply grateful to Marlene for being such a loving, supportive, and encouraging person. Thank you.

Lukas Lang
University of Vienna
June, 2015

Contents

Abstract	v
Preface	vii
I Preamble	1
1 Challenges in the Analysis of Fluorescence Microscopy Data	3
2 Variational Optical Flow	7
3 Contributions of the Thesis	11
3.1 Optical Flow on Evolving Surfaces	13
3.2 Decomposition of Optical Flow on the Sphere	15
3.3 Optical Flow on Evolving Sphere-Like Surfaces	17
3.4 Discussion and Further Research	18
II Publications	21
4 Optical Flow on Evolving Surfaces with an Application to the Analysis of 4D Microscopy Data	23
5 Optical Flow on Evolving Surfaces with Space and Time Regularisation	35
6 Decomposition of Optical Flow on the Sphere	59
7 Optical Flow on Evolving Sphere-Like Surfaces	83
III Appendix	117
Bibliography	119
Zusammenfassung	127
Curriculum Vitae	129

Part I

Preamble

Chapter 1

Challenges in the Analysis of Fluorescence Microscopy Data

In the past decade, a tremendous progress in imaging of biological model organisms, such as the zebrafish, has been made. High-resolution observations of entire developing animals became possible. Today, they can be studied in more detail than ever before. See for instance [43, 44, 45, 46]. Fluorescence microscopy allows to record time-lapse images on the scale of single cells, see e.g. [43, 60, 74].

Nevertheless, increasing spatial as well as temporal resolutions result in vast amounts of data. Therefore, analysis by visual examination, which is usually carried out by trained biologists, is nearly impracticable, particularly because biological studies typically require a considerable number of samples. Thus, automated analysis of such microscopy data poses new challenges to the imaging community and is key to their large-scale interpretation. See for instance [64].

The afore-mentioned zebrafish is a popular model organism, see e.g. [89, Chap. 2]. We point the reader to [48] for many illustrations and an exhaustive discussion of its developmental process. Despite its importance for tissue and organ formation, only little is known about cell migration and proliferation patterns during early embryonic development [3, 74]. An urging question concerning this stage is therefore:

How do cells move and proliferate during early development?

Understanding the complex movement patterns of cells and their life cycle — including for instance cell division — is of immense interest, see e.g. [66].

The answer to the above question involves computing the cells' velocities, respectively their complete trajectories. In Fig. 1.1 we visualise the approximate movements of endodermal cells in some of our datasets. They have been computed by a simple heuristic that iteratively matches pairs of closest cell centres in consecutive frames and simultaneously allows for cell divisions. The cells' centres are typically indicated by local maxima in image intensity and can reliably be found by Gaussian filtering followed by thresholding.

In this thesis we aim at developing novel methods for the quantitative analysis of cell motion during early embryogenesis and try to assist biologists with answering the above question. By taking into consideration a crucial fact

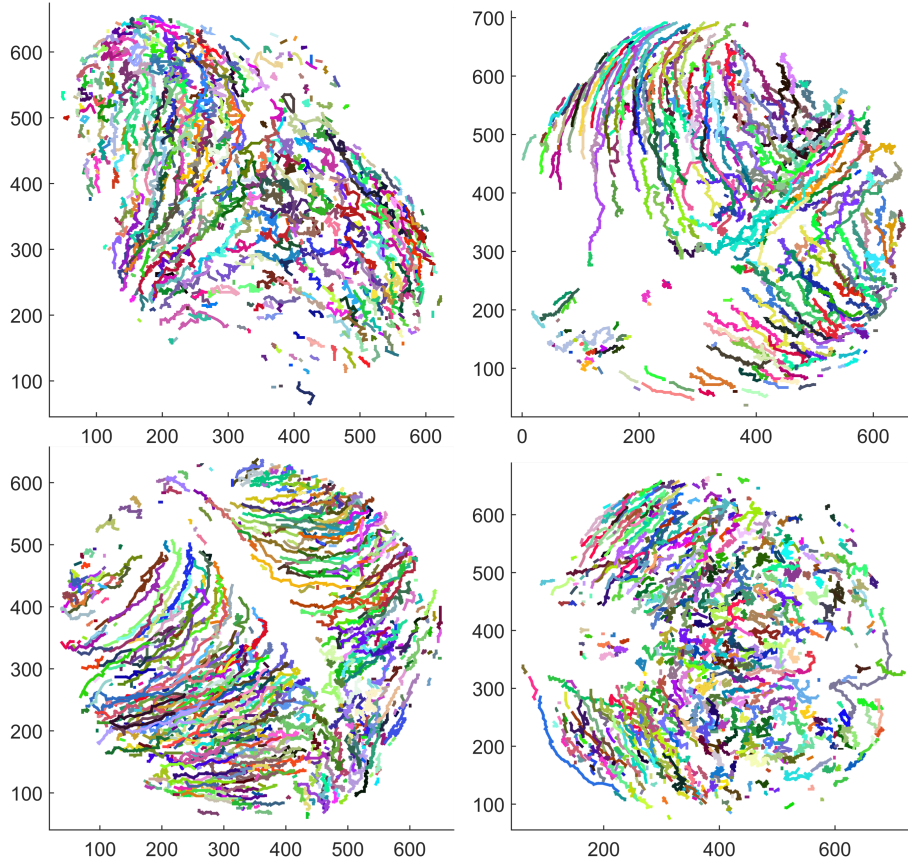


Figure 1.1: Approximate trajectories of endodermal cells in a top view. The colour of each track is assigned randomly and serves for better distinction. Over time, the cells float from the outside towards the embryo’s body axis. It is (roughly) aligned with the diagonal from bottom left to top right in the top row, and from top left to bottom right and almost horizontally in the bottom row. The trajectories visualise the complex motion patterns, show distinguished global features, and highlight individual fluctuations. All dimensions are in micrometer (μm).

about the biological nature of endodermal cells, computationally reasonable approaches can be achieved.

The data at hand consist of volumetric time-lapse (four-dimensional) images of a living zebrafish embryo during the gastrula period. These videos were recorded approximately five to ten hours after fertilisation by means of confocal laser-scanning microscopy, a common imaging modality that allows *in vivo* imaging of organisms, see e.g. [68, 69].¹ They depict endodermal cells expressing a green fluorescence protein [16]. As a consequence, these labelled

¹The basic principle of confocal laser-scanning microscopy is to scan the specimen point by point and to shift the focal plane, thereby capturing a 3D image of the specimen. However, as light contributions from out of focus points are filtered, a light source of high intensity (such as a laser) has to be used.




Figure 1.2: Two consecutive frames of the volumetric zebrafish microscopy images recorded during early embryogenesis. Endodermal cells floating on the upper hemisphere of the animal are shown. Fluorescence response is indicated by blue colour and is proportional to the recorded intensity. All dimensions are in micrometer (μm).

cells can be recorded without capturing the background and therefore allow for a separate treatment. Figure 1.2 depicts two frames of the captured sequence, containing only the upper hemisphere of the animal embryo.

The crucial feature of endodermal cells is the fact that they form a so-called *monolayer* during early morphogenesis, see [83]. It means that these cells do not sit on top of each other but float side by side. In Fig. 1.2 one can see the nuclei of cells forming a round surface in a single layer. Observe also the noise present and the sparsity of the data.

This monolayer can be regarded as a surface and allows for the straightforward extraction of a two-dimensional image sequence. We exploit this situation and model this layer as an evolving surface. The main intention of this thesis is to consider cell motion only on this, potentially moving, two-dimensional manifold. As a result we are able to reduce the spatial dimension of the data allowing for more efficient motion estimation, simultaneously preserving as much information as possible.

Optical flow is a popular and well-established framework for dense motion estimation. It delivers necessary quantitative methods for cell motion analysis and sheds light on the underlying cellular mechanisms and the dynamic behaviour of cells. See for example [4, 40, 61, 72, 74] and the references therein.

For the task of quantitative motion estimation in time-lapse image sequences optical flow has been considered before, see e.g. [2, 14, 21]. Moreover, three-dimensional cell motion has been computed, for instance, in [40, 41, 61]. In [4, 72, 74], the optical flow was employed for the analysis of cell motion in microscopy data. In particular, Schmid et al. [74] followed a similar approach to ours by exploiting the monolayer structure and modelling the embryo as a round sphere. There, the optical flow was computed in map projections.

The central theme of this thesis is to use variational optical flow for the estimation of endodermal cell motion. We develop mathematical models for optical flow on evolving manifolds such as the zebrafish endoderm. As opposed to computing the optical flow in three-dimensional Euclidean space, this approach has the advantage of reducing the spatial dimension. Under reasonable assumptions on the geometry it preserves most information relevant to cell motion. Furthermore, it outputs dense vector fields and, together with

appropriate visualisation techniques, yields insights into global trends and the collective motion of cells. Since the optical flow computes the cells' velocities, which under certain assumptions equal the true velocities, one can easily obtain (approximate) trajectories.

Chapter 2

Variational Optical Flow

Motion estimation is an ubiquitous goal in the analysis of image sequences. Adding a temporal dimension to a recorded scene allows perceivable motion. Naturally, objects move through the physical world at different velocities and their appearance in a recording is typically linked to the sampling frequency. Furthermore, we can distinguish different types of motions. For instance, some of the cells in Fig. 1.1 follow smooth trajectories whereas others rather seem to follow Brownian motion.

By the nature of the human eye, motion is naturally perceived by variations in intensity. In addition, it sometimes adds a sense of depth to a scene, e.g. when objects are occluded. However, when observing motion through a camera — or more precisely as the projection onto an image plane — only the so-called *apparent motion* is observable through intensity variations. This is mainly known as the *optical flow* [38, 39]. Under certain conditions it equals the two-dimensional projection or even the true three-dimensional motion in the recorded scene, see e.g. [81]. Typical counterexamples are changes in illumination or the well-known rotating *barber's pole*. We refer to Figs. 5.12 and 5.13 in [8, Sec. 5.3] for illustrations of the discussed.

Typically, optical flow models are based on the assumption of *constant brightness*. Given a sequence of images, it presumes that a point moving through the recorded scene does not change its appearance over time. In other words, its intensity remains constant. This so-called *brightness constancy assumption* is reasonable as long as there are no changes in illumination and no occlusions occur.

In order to formulate the idea more precisely, let us consider a sequence of images

$$f : I \times \Omega \rightarrow \mathbb{R},$$

where $I := [0, T] \subset \mathbb{R}$ is a real time interval and $\Omega \subset \mathbb{R}^2$ a planar domain. The value of $f(t, \xi)$ is the recorded image intensity at a time $t \in I$ and at a point $\xi \in \Omega$. Following a smooth trajectory

$$\gamma(\cdot, \xi_0) : I \rightarrow \Omega$$

which starts at $\xi_0 \in \Omega$, the assumption of constant brightness requires that the value of $f(t, \gamma(t, \xi_0))$ remains constant over time. In other words,

$$f(t, \gamma(t, \xi_0)) = f(0, \xi_0), \tag{2.1}$$

for all times $t \in I$. Assuming that $f \in C^1(I \times \Omega)$, it follows that

$$\frac{d}{dt}f(t, \gamma(t, \xi_0)) = \partial_t f(t, \gamma(t, \xi_0)) + \nabla_{\mathbb{R}^2} f(t, \gamma(t, \xi_0)) \cdot \partial_t \gamma(t, \xi_0) = 0$$

must hold true for all $\xi_0 \in \Omega$ and for all times $t \in I$. Here we have denoted by d/dt the total and by ∂_t the partial derivative with respect to time t , and by $\nabla_{\mathbb{R}^2}$ the (spatial) gradient of \mathbb{R}^2 .

For simplicity, let us denote by

$$\mathbf{v}(t, \gamma(t, \xi_0)) := \partial_t \gamma(t, \xi_0)$$

the velocity of a point moving along γ . Then, the problem of optical flow estimation is to find a time-dependent vector field $\mathbf{v} : I \times \Omega \rightarrow \mathbb{R}^2$ satisfying

$$\partial_t f + \nabla_{\mathbb{R}^2} f \cdot \mathbf{v} = 0 \quad (2.2)$$

for all $(t, \xi) \in I \times \Omega$. For the sake of brevity we have omitted the arguments.

Under suitable assumptions and with appropriate initial and boundary conditions, equation (2.2) is an equivalent Eulerian specification of assumption (2.1) and is termed *optical flow equation*. From that, the original Lagrangian specification in terms of γ can be recovered by solving the initial value problem

$$\begin{aligned} \partial_t \gamma(t, \xi_0) &= \mathbf{v}(t, \gamma(t, \xi_0)), \quad \text{for all } t \in I, \\ \gamma(0, \xi_0) &= \xi_0, \end{aligned} \quad (2.3)$$

or, in other words, by computing the integral curves. Let us mention that (2.2) is linear in the unknown \mathbf{v} and, after discretisation, is reasonable for sufficiently small displacements.

The optical flow equation (2.2) is underdetermined and uniqueness of a solution is not guaranteed. One can easily see that if a vector field $\mathbf{v}(t, \xi)$ solves (2.2), adding a flow $c \nabla_{\mathbb{R}^2} f(t, \xi)^\perp$ which is orthogonal to the image gradient solves the equation equally well for every $c \in \mathbb{R}$, and even for every $c : I \times \Omega \rightarrow \mathbb{R}$. In other words, from (2.2) one can only infer the movement along the direction of the image gradient, that is, perpendicular to the level lines of f . This issue is called *aperture problem*. See also the illustrations in [8, Sec. 5.3]. A different problem arises whenever $\partial_t f \neq 0$ but $\nabla_{\mathbb{R}^2} f = 0$. Then, equation (2.2) can not be fulfilled.

For the above reasons, estimating the optical flow can also be interpreted as an *ill-posed* inverse problem. We refer to [49] for a further discussion with a focus on this matter. For a general treatment of inverse problems and regularisation theory we refer to [25] and to [73], which has a particular focus on imaging.

We have seen that in order to solve the optical flow equation (2.2), one has to overcome the above-mentioned problems. For introductory material on this subject we refer to [8] and for a comparison of various techniques to [10, 27]. Furthermore, the benchmark framework created by Baker et al. [9] is noteworthy. The corresponding website of the Middlebury College lists currently leading methods.¹

¹<http://vision.middlebury.edu/flow/>

In this thesis, however, we employ a *variational approach* for dense optical flow estimation. A reasonable first idea is to consider solutions to the linear *least squares problem*

$$\min_{\mathbf{v}} \|\partial_t f + \nabla_{\mathbb{R}^2} f \cdot \mathbf{v}\|_{L^2(I \times \Omega)}^2. \quad (2.4)$$

As discussed before, a solution in general might not exist or, if it exists, it might not be unique.

However, the so-called *normal flow*, given by

$$\mathbf{v}^\dagger = \begin{cases} -\frac{\partial_t f}{|\nabla_{\mathbb{R}^2} f|^2} \nabla_{\mathbb{R}^2} f, & \text{if } \nabla_{\mathbb{R}^2} f \neq 0, \\ 0, & \text{else,} \end{cases}$$

is a solution to (2.4). Nevertheless, for obvious reasons, the normal flow \mathbf{v}^\dagger is of little practical use.² As a consequence, one needs to restrict the space of solutions to desirable ones.

A common approach to obtain *well-posedness* of the optical flow problem, and with it uniqueness of a solution, is *Tikhonov regularisation* [78]. It consists of finding the unique minimiser of

$$\|\partial_t f + \nabla_{\mathbb{R}^2} f \cdot \mathbf{v}\|_{L^2(I \times \Omega)}^2 + \alpha \mathcal{R}(\mathbf{v}).$$

Here, $\mathcal{R}(\mathbf{v})$ is a regularisation functional and $\alpha > 0$ is a regularisation parameter balancing the two terms. Typically, the first term is called *data term*, whereas the second is called *smoothness term*. Loosely speaking, the latter incorporates prior information about favoured solutions.

A natural choice for $\mathcal{R}(\mathbf{v})$ is the squared $H^1(I \times \Omega)$ Sobolev seminorm. It penalises first derivatives with respect to space and time and favours spatial as well as temporal regularity of the solution. Simply put, the sought velocity field should vary only smoothly in a small neighbourhood as nearby points are likely to originate from the same object.

In their seminal work, however, Horn and Schunck [39] did not consider temporal derivatives. They proposed to compute the minimiser of

$$\|\partial_t f + \nabla_{\mathbb{R}^2} f \cdot \mathbf{v}\|_{L^2(\Omega)}^2 + \alpha |\mathbf{v}|_{H^1(\Omega)}^2.$$

and considered one pair of frames only. In the sequel we refer to it as the *Horn-Schunck functional*. We highlight that, although relatively easy to solve numerically, the functional entails isotropic regularity and does not allow discontinuities in the flow field. Well-posedness of the above functional was first shown by Schnörr [75], making additional assumptions on ∇f .

Solutions to the above functional can be computed by solving the corresponding set of Euler-Lagrange equations following the calculus of variations [19, Chap. IV]. This system of second-order elliptic partial differential equations is given by

$$\begin{aligned} (\partial_t f + \nabla_{\mathbb{R}^2} f \cdot \mathbf{v}) \nabla_{\mathbb{R}^2} f - \alpha \Delta_{\mathbb{R}^2} \mathbf{v} &= 0 \quad \text{for all } \xi \in \Omega, \\ \partial_{\mathbf{n}} \mathbf{v} &= 0 \quad \text{for all } \xi \in \partial\Omega. \end{aligned}$$

²The reader is encouraged to picture a moving object attaining constant intensity, i.e. $\nabla_{\mathbb{R}^2} f = 0$, in the interior.

Here, $\partial_{\mathbf{n}}$ is the normal derivative along the outward unit normal \mathbf{n} to Ω .

The above-mentioned spatio-temporal model was first pursued by Weickert and Schnörr [86] in a slight generalisation. The form relevant to this thesis reads

$$\|\partial_t f + \nabla_{\mathbb{R}^2} f \cdot \mathbf{v}\|_{L^2(I \times \Omega)}^2 + \alpha |\mathbf{v}|_{H^1(I \times \Omega)}^2$$

and is particularly convenient when trying to recover the trajectories from the estimated flow, in other words, when solving problem (2.3). In [86] it is mentioned that well-posedness of the Horn-Schunck functional readily translates to this spatio-temporal setting. There, minimisation was done by applying an explicit (Euler forward) scheme to the associated steepest descent equations. Compared to estimating the optical flow between each pair of frames it is computationally more demanding, though.

Let us conclude this chapter with a few remarks. While the above-discussed functionals are relatively straightforward to minimise, they come with several drawbacks. For instance, they do not allow for discontinuities in the flow field. As a remedy, several improvements have been proposed over the last years. For discontinuity-preserving functionals we refer e.g. to [6, 7] and [36]. A unifying framework including spatial as well as temporal regularisation functionals was proposed in [85]. Moreover, we refer to [84] for a survey of numerous optical flow data and regularisation terms. With regard to efficient numerical methods we mention primal dual methods developed in [17].

Finally, let us highlight that for certain data it is beneficial to use different data terms as, for instance, the assumption of brightness is often violated in practice. With applications to fluid motion estimation we point e.g. to [18, 20, 88] where conservation of mass is assumed and to the framework in [32], which discusses physical processes.

Chapter 3

Contributions of the Thesis

In this thesis we model the volumetric microscopy data as a time-dependent non-negative function

$$f^\delta : I \times \Omega \subset \mathbb{R}^3 \rightarrow [0, \infty),$$

where $I = [0, T]$, as before, is the time interval of a recorded sequence. The value of f^δ directly corresponds to the fluorescence response of the observed cells. Here, Ω is the cuboid region captured by the confocal laser-scanning microscope, see Fig. 1.2 in Chap. 1.

Adapting the concept of optical flow to this new setting is challenging in several ways:

1. We desire computational feasibility for real-life data.
2. It should be based on plausible (geometric) assumptions.
3. We seek reasonable approximations of the cells' true velocities.

A straightforward idea, if permitted by the nature of the data, is to project the data f^δ to \mathbb{R}^2 and then solve the optical flow equation

$$\partial_t P f^\delta + \nabla_{\mathbb{R}^2} P f^\delta \cdot \mathbf{v} = 0$$

in the plane, as discussed before. Here, P denotes an appropriate projection operator to \mathbb{R}^2 , such as taking the maximum intensity along the z -coordinate. This approach is feasible whenever only a section of the sphere-shaped embryo is considered, see e.g. the data in [61, Fig. 2], and the velocities of the cells can be approximated reasonably well by their projections to the plane. It is most closely related to the discussion in Chap. 2.

Another idea is to generalise optical flow to the three-dimensional Euclidean space. It requires solving

$$\partial_t f^\delta + \nabla_{\mathbb{R}^3} f^\delta \cdot \mathbf{V} = 0$$

for the vector field $\mathbf{V} : I \times \Omega \rightarrow \mathbb{R}^3$. Although this approach is computationally costly given the sparsity of the data, we stress that it is capable of estimating the true velocities in \mathbb{R}^3 . This approach was pursued, for instance, in [2, 40, 41, 61] for various kinds of data.

For the sake of completeness, let us also mention a conceptually different approach. In [3], the goal of dense motion estimation is abandoned. Instead, cells are segmented and the optical flow is computed based on super-voxels in order to achieve a reduction in the amount of data. For an exhaustive discussion of various other methods on we refer to the survey [64].

The central idea of this thesis, however, is to consider cell motion with regard to a restriction of f^δ to the developing monolayer, or an approximation thereof. The temporal evolution of the data can then be tracked by solving an optical flow problem on this — possibly moving — two-dimensional manifold $\mathcal{M} \subset \mathbb{R}^3$. For the time being, let us define this restriction (literally) by

$$\hat{f} := f^\delta|_{\mathcal{M}}.$$

In the following, we will take more sophisticated attempts such as projecting the maximum image intensity along surface normals within a narrow band.

In this thesis we are mainly concerned with the following problem: Given an image sequence $\hat{f} : \mathcal{M} \rightarrow \mathbb{R}$ living on a manifold, we intend to solve

$$\partial_t \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{v}} = 0$$

on a suitable approximation \mathcal{M} of the zebrafish embryo. Here, $\hat{\mathbf{v}} : \mathcal{M} \rightarrow T\mathcal{M}$ is the sought tangent vector field, $\nabla_{\mathcal{M}}$ denotes the surface gradient, and dot denotes the standard inner product of \mathbb{R}^3 . We highlight that, while for static surfaces embedded in \mathbb{R}^3 this is a sound generalisation of the optical flow equation (2.2), care must be taken for moving surfaces, see Sec. 3.1, and Chaps. 4 and 5.

In their seminal work, Imiya et al. [42, 79] discussed the optical flow problem in a spherical setting. Later, Lefèvre and Baillet [58] adapted the Horn-Schunck functional (cf. Chap. 2) to surfaces embedded in \mathbb{R}^3 and established well-posedness.

The corresponding variational problem is to minimise

$$\|\partial_t \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{v}}\|_{L^2(\mathcal{M})}^2 + \alpha \|\hat{\mathbf{v}}\|_{H^1(\mathcal{M}, T\mathcal{M})}^2. \quad (3.1)$$

The conceptual difference compared to the Euclidean case is not just the domain but the regularisation functional. It involves covariant derivatives, which are necessary in order to correctly penalise tangent vector fields, see e.g. Example 1 in Sec. 5.3.

For a tangent vector field $\hat{\mathbf{v}}$ which is smoothly extended to \mathbb{R}^3 and denoted by $\hat{\hat{\mathbf{v}}}$, and a tangent vector $\hat{\mathbf{u}}$, we define the covariant derivative $\nabla_{\hat{\mathbf{u}}} \hat{\mathbf{v}}$ as the tangential part of the directional derivative of $\hat{\hat{\mathbf{v}}}$ along $\hat{\mathbf{u}}$ in the embedding space. That is,

$$\nabla_{\hat{\mathbf{u}}} \hat{\mathbf{v}} = P_{\mathcal{M}} \nabla_{\mathbb{R}^3} \hat{\hat{\mathbf{v}}}(\hat{\mathbf{u}}),$$

where $P_{\mathcal{M}}$ is the orthogonal projector onto the tangent space.

At this point, let us highlight that, by virtue of the Hairy-Ball Theorem, for the 2-sphere and for closed surfaces homeomorphic to \mathcal{S} the H^1 Sobolev seminorm is indeed a norm. The reason is that the only covariantly constant vector field is $\hat{\mathbf{v}} = 0$. In the above-mentioned settings well-posedness is based on this fact. For further discussion we refer to Lefèvre and Baillet [58] and to Bauer et al. [12].

However, in this thesis we are concerned with appropriate models of the embryo's shape. Clearly, the chosen surface is subject to geometric approximations. In addition, estimating this surface from microscopy data is an interesting problem on its own, see Sec. 3.3 and Chap. 7.

In this thesis, we consider three major scenarios. First, we devote ourselves to evolving surfaces $\mathcal{M} \subset \mathbb{R}^4$. As mentioned before, the optical flow equation requires a considerable adaptation to this new setting. We consider only a rectangular section of the embryo, represent the surface by a time-varying height field, and solve the variational optical flow problem in the coordinate domain.

In the second scenario, we follow the ideas of Schmidt et al. [74] and approximate the embryo with the static 2-sphere \mathcal{S} . However, our approach is substantially different as we directly solve the variational problem on the sphere.

Finally, in the third scenario, we consider evolving sphere-like surfaces. These can be parametrised from the 2-sphere \mathcal{S} and, most conveniently, the associated variational problem can be dealt with on the sphere. The advantage of this representation is that the sphere-like geometry is a very natural candidate to consider for the approximation of the zebrafish embryo.

In the following, we sketch the above-mentioned scenarios, briefly discuss the contributions to the analysis of said microscopy data, and summarise our results.

Concerning the notation, in Chap. 2 we have seen a scalar (intensity) function $f : I \times \Omega \rightarrow \mathbb{R}$ and a vector field $\mathbf{v} : I \times \Omega \rightarrow \mathbb{R}^2$, both defined in the plane. In this chapter, we adjust this notation to non-Euclidean settings. We keep denoting vector fields by boldface letters, however, now distinguish between lower and upper case boldface letters. The former denote tangent vector fields whereas the latter denote general vector fields in \mathbb{R}^3 . As a notational convention, we indicate scalar functions living on the 2-sphere by \tilde{f} and those defined on a general manifold \mathcal{M} by \hat{f} , respectively. Accordingly, tangent vector fields are denoted by $\tilde{\mathbf{v}}$ and $\hat{\mathbf{v}}$, respectively.

3.1 Optical Flow on Evolving Surfaces

In Chaps. 4 and 5 we suppose that the zebrafish embryo can be represented as a surface which is evolving over time as the animal develops. The goal of these chapters is to generalise the idea of constant brightness portrayed in Chap. 2 to this setting.

We define an evolving surface as

$$\mathcal{M} := \bigcup_{t \in I} (\{t\} \times \mathcal{M}_t) \subset \mathbb{R}^4$$

and assume it is given in terms of a Lagrangian specification $\phi : I \times \mathcal{M}_0 \rightarrow \mathbb{R}^3$ such that $\phi(0, \cdot)$ is the identity and $\phi(t, \cdot)$ is a diffeomorphism between $\mathcal{M}_t \subset \mathbb{R}^3$ and \mathcal{M}_0 for every $t \in I$.

In addition, we are given a function \hat{f} , its domain being \mathcal{M} , arising from the restriction of the microscopy data to \mathcal{M} . For a time $t \in I$,

$$\hat{f}(t, \cdot) : \mathcal{M}_t \rightarrow \mathbb{R}$$

is then an image on the surface. Adapting the idea of constant brightness — as discussed in Chap. 2 — to this new setting requires that along a smooth trajectory $\gamma(\cdot, x) : t \mapsto \gamma(t, x) \in \mathcal{M}_t$ that starts at $x \in \mathcal{M}_0$ and always stays on the surface, we must have

$$\hat{f}(t, \gamma(t, x)) = \hat{f}(0, x). \quad (3.2)$$

However, in order to proceed as before and obtain a linearisation of the stated constraint one needs to define a meaningful derivative with respect to time.

In Chaps. 4 and 5 we introduce derivatives along trajectories following the moving surface. To this end, we consider the surface's velocity $\partial_t \phi = \hat{\mathbf{V}}$, its domain being $\bigcup_{t \in I} (\{t\} \times \mathcal{M}_t) \subset \mathbb{R}^4$. Let us mention that $\hat{\mathbf{V}}$ is in general not tangent to \mathcal{M}_t , $t \in I$, and hence in our notation denoted by a boldface capital letter.

Then,

$$d_t^{\hat{\mathbf{V}}} \hat{f}(t_0, x) := \left. \frac{d}{dt} \hat{f}(t, \phi(t, x_0)) \right|_{t=t_0} \quad (3.3)$$

is the time derivative of \hat{f} at $x = \phi(t_0, x_0)$ along the trajectory $\phi(\cdot, x_0)$. As a consequence, one can deduce that

$$d_t^{\hat{\mathbf{V}}} \hat{f} = d_t^{\hat{\mathbf{N}}} \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{V}}$$

holds, see [15]. Here, $d_t^{\hat{\mathbf{N}}} \hat{f}(t_0, x)$ denotes the time derivative of \hat{f} in normal direction. It is defined analogously to (3.3) following a trajectory $\psi_{\hat{\mathbf{N}}}$ through $x \in \mathcal{M}_{t_0}$ for which $\partial_t \psi_{\hat{\mathbf{N}}}(t_0, x)$ is orthogonal to the tangent plane $T_x \mathcal{M}_{t_0}$ at x .

The above relation immediately allows us to formulate the above idea of constant brightness (3.2) along γ . For simplicity, let us define the velocity of a point following the trajectory γ by $\hat{\mathbf{M}} := \partial_t \gamma$. Then, along a moving surface we require that

$$d_t^{\hat{\mathbf{M}}} \hat{f} = d_t^{\hat{\mathbf{N}}} \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{M}} = 0 \quad (3.4)$$

must hold. Equation (3.4) is a *generalised optical flow equation*. In Fig. 3.1 we sketch the various trajectories through the evolving surface and their corresponding velocities.

As a consequence, one can derive that the total motion $\hat{\mathbf{M}} = \partial_t \gamma$ of a cell moving along a trajectory γ can be decomposed into

$$\hat{\mathbf{M}} = \hat{\mathbf{V}} + \hat{\mathbf{v}}, \quad (3.5)$$

where $\hat{\mathbf{v}}$ is a tangent vector field and $\hat{\mathbf{V}} = \partial_t \phi$ is the surface velocity. The above relation states that the total velocity $\hat{\mathbf{M}}$ along a level line of constant intensity is the sum of a tangential velocity $\hat{\mathbf{v}}$ relative to the prescribed surface velocity $\hat{\mathbf{V}}$. See Fig. 3.1.

Solving the generalised optical flow equation (3.4), however, is inconvenient as $\psi_{\hat{\mathbf{N}}}$ and, in further consequence, $d_t^{\hat{\mathbf{N}}}$ is unknown or hard to estimate from real microscopy data.

One remedy, which is pursued in Chaps. 4 and 5, is to relate (3.4) and (3.5). In further consequence, we arrive at the *parametrised optical flow equation*

$$d_t^{\hat{\mathbf{V}}} \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{v}} = 0.$$




Figure 3.1: Illustration of trajectories through the evolving surface. Their corresponding velocities and the tangential velocity \hat{v} are shown in grey.

Solving for the optical flow then means finding a (time-varying) vector field \hat{v} that is tangent to the surface at all times and satisfies the above equation at every point $x \in \mathcal{M}$ on the moving surface. From a numerical point of view this is particularly convenient as the equation can easily be restated in terms of the coordinate domain, see Chap. 5, or more general, in terms of a reference manifold such as \mathcal{S} , see Sec. 3.3 and Chap. 7.

Having derived a suitable optical flow equation for moving surfaces, we generalise the spatio-temporal model by Weickert and Schnörr [86] to this non-Euclidean and dynamic setting. We minimise the functional

$$\int_I \left(\|d_t^{\hat{V}} \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{v}\|_{L^2(\mathcal{M}_t)}^2 + \alpha \|\hat{v}\|_{H^1(\mathcal{M}_t, T\mathcal{M}_t)}^2 + \beta \|P_{\mathcal{M}} d_t^{\hat{V}} \hat{v}\|_{L^2(\mathcal{M}_t, T\mathcal{M}_t)}^2 \right) dt,$$

where $P_{\mathcal{M}} d_t^{\hat{V}} \hat{v}$, akin to the covariant derivative, is the projection of the temporal derivative (3.3) of \hat{v} along the trajectory ϕ to the tangent space of \mathcal{M} .

We mention that, in general, there exist infinitely many Lagrangian specifications ϕ for a given evolving surface \mathcal{M} . The actual surface velocity, however, might be unknown or cannot be estimated from the data, as it is the case in this thesis. As a remedy we impose a surface velocity by choosing a specification ϕ which is both convenient from an implementation point of view and relatively natural given the characteristics of the data.

Moreover, we again stress that the sought tangent vector field \hat{v} depends on the chosen ϕ , or equivalently on \hat{V} , and should be interpreted with care. The actual trajectories γ though can be reconstructed by finding the integral curves of (3.5). We refer to Chap. 5 for this precise approach.

3.2 Decomposition of Optical Flow on the Sphere

While in Chaps. 4 and 5 the zebrafish embryo is assumed to change its shape over time, in Chap. 6 we consider a spherical and static approximation of the embryo. We choose a sphere $\mathcal{S} \subset \mathbb{R}^3$ of appropriate radius and assume to be given suitable projections of the microscopy data onto \mathcal{S} as

$$\tilde{f} : I \times \mathcal{S} \rightarrow \mathbb{R}.$$

In further consequence, the optical flow problem asks to solve

$$\partial_t \tilde{f} + \nabla_S \tilde{f} \cdot \tilde{\mathbf{w}} = 0$$

on a spherical geometry and the variational problem then reads

$$\|\partial_t \tilde{f} + \nabla_S \tilde{f} \cdot \tilde{\mathbf{w}}\|_{L^2(S)}^2 + \alpha |\tilde{\mathbf{w}}|_{H^s(S, TS)}^2,$$

for an appropriately chosen parameter $s \in \mathbb{R}$.

Numerically, we propose to solve the problem by projection to a finite-dimensional space spanned by tangential vector spherical harmonics. It provides great flexibility with respect to the chosen norm $H^s(S, TS)$ and allows for the straightforward generalisation to multiple regularisation terms.

In this part of the thesis we consider several decomposition models for optical flow. The goal is to compute not just one but several vector fields capturing different structural parts of the total motion $\tilde{\mathbf{w}}$.

The first setting, referred to as $\tilde{\mathbf{u}} + \tilde{\mathbf{v}}$ decomposition, is inspired by the idea of Meyer [62] and aims at separating cartoon and texture parts. The former should contain large-scale structural parts and therefore should be piecewise smooth, whereas the latter should contain the high-frequency oscillations. Consequentially, given two real numbers $r, s \in \mathbb{R}$ such that $r > 0$ and $r > s$, the idea is to minimise

$$\|\partial_t \tilde{f} + \nabla_S \tilde{f} \cdot (\tilde{\mathbf{u}} + \tilde{\mathbf{v}})\|_{L^2(S)}^2 + \alpha |\tilde{\mathbf{u}}|_{H^r(S, TS)}^2 + \beta |\tilde{\mathbf{v}}|_{H^s(S, TS)}^2,$$

where $\alpha, \beta > 0$ are regularisation parameters balancing the three terms.

The main idea is that $\tilde{\mathbf{u}} \in H^r(S, TS)$ should contain the collective motion that dictates the movements of a large portion of cells, whereas $\tilde{\mathbf{v}} \in H^s(S, TS)$ should indicate individual deviations, such as cell divisions.

In the second part of this chapter, we consider hierarchical decomposition models. Pioneered by Tadmor et al. [77], they provide a multiscale description of the input data. We transfer this idea to the optical flow setting. The goal is to output an arbitrarily fine multiscale description of the total motion $\tilde{\mathbf{w}}$.

Given two non-increasing sequences (α_k) and (s_k) , we devise an iterative scheme and compute

$$\tilde{\mathbf{v}}_k = \begin{cases} \arg \min_{\tilde{\mathbf{v}}} \|\partial_t \tilde{f} + \nabla_S \tilde{f} \cdot \tilde{\mathbf{v}}\|_{L^2(S)}^2 + \alpha_1 |\tilde{\mathbf{v}}|_{H^{s_1}(S, TS)}^2, & \text{if } k = 1, \\ \arg \min_{\tilde{\mathbf{v}}} \|\partial_t \tilde{f} + \nabla_S \tilde{f} \cdot \left(\tilde{\mathbf{v}} + \sum_{i=1}^{k-1} \tilde{\mathbf{v}}_i \right)\|_{L^2(S)}^2 + \alpha_k |\tilde{\mathbf{v}}|_{H^{s_k}(S, TS)}^2, & \text{if } k > 1. \end{cases}$$

The intention is that the sequence

$$\left\{ \tilde{\mathbf{v}}^{(k)} := \sum_{i=1}^k \tilde{\mathbf{v}}_i : k \in \mathbb{N} \right\}$$

of solutions captures cell motion on arbitrarily fine scales.

The third model considered in this chapter is a Helmholtz decomposition of the total motion. It is based on the fact that every tangential vector field on




Figure 3.2: Schematic illustration of a cut through the surfaces \mathcal{S} and \mathcal{M}_t intersecting the origin. The sphere-like surface \mathcal{M}_t is a deformation of the sphere along the depicted radial line.

\mathcal{S} has a unique decomposition into a curl-free and divergence-free vector field. Accordingly, the total motion $\tilde{\mathbf{w}}$ is decomposed into $\tilde{\mathbf{u}} + \tilde{\mathbf{v}}$ such that

$$\begin{aligned}\nabla_{\mathcal{S}} \cdot (\tilde{\mathbf{u}} \times \tilde{\mathbf{N}}) &= 0, \\ \nabla_{\mathcal{S}} \cdot \tilde{\mathbf{v}} &= 0,\end{aligned}$$

where $\tilde{\mathbf{N}}$ is the outward surface unit normal to \mathcal{S} . Conveniently, the chosen numerical method by projection to tangential vector spherical harmonics automatically yields such a Helmholtz decomposition.

3.3 Optical Flow on Evolving Sphere-Like Surfaces

The third part is mainly motivated by the characteristic shape of the zebrafish endoderm during gastrulation, see Chap. 1 and the references therein. We take into consideration the topology of the embryo and model the monolayer as a closed 2-manifold which is deforming over time. Based on what has been said in Secs. 3.1 and 3.2, we take a step towards better approximation of the true shape.

In Chap. 7, we consider a closed surface $\mathcal{M}_t \subset \mathbb{R}^3$ for which the mapping

$$(t, x) \mapsto \tilde{\rho}(t, x)x, \quad x \in \mathcal{S} \quad (3.6)$$

is a diffeomorphism between the 2-sphere and \mathcal{M}_t for every time $t \in I$ and refer to it as *evolving sphere-like surface*. Figure 3.2 depicts the scenario for a better understanding.

We adapt the Thikonov-regularised quadratic optical flow functional (3.1) to this new setting and compute the minimiser of

$$\|d_t^{\tilde{\mathbf{V}}} \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{v}}\|_{L^2(\mathcal{M}_t)}^2 + \alpha |\hat{\mathbf{v}}|_{H^s(\mathcal{M}_t, T\mathcal{M}_t)}^2.$$

Numerically, the problem is solved by means of a Galerkin method based on vector spherical harmonics. However, beforehand these vector fields need to be “transferred” to the sphere-like surface. In other words, the sought tangent vector field is expanded in terms of the pushforward of vector spherical

harmonics. Conveniently, the simplicity of the chosen parametrisation allows to rewrite the problem and numerically treat it on the 2-sphere.

Until now, the question of how to actually find this geometrical approximation of the zebrafish embryo has been neglected. As part of this chapter, we treat it as a variational problem. We propose surface interpolation with squared H^s Sobolev seminorm regularisation in order to compute the parametrisation (3.6) of the evolving sphere-like surface. We minimise

$$\|\tilde{\rho} - \tilde{\rho}^\delta\|_{L^2(\mathcal{S})}^2 + \beta \|\tilde{\rho}\|_{H^s(\mathcal{S})}^2,$$

where $\beta \in \mathbb{R}$ again is a regularisation parameter. Approximate cell centres serve as sample points $\tilde{\rho}^\delta$ of the surface. Scalar spherical harmonics provide the necessary flexibility with respect to the chosen space $H^s(\mathcal{S})$ and allow to easily meet the smoothness requirements of the manifold.

3.4 Discussion and Further Research

Aiming at computationally feasible cell motion analysis in 4D fluorescence microscopy data of a living zebrafish, we employed variational optical flow in order to estimate dense velocity fields. We exploited the monolayer property of the developing zebrafish (cf. Chap. 1) and addressed the problem in two conceptually different ways.

First, we paid close attention to the dynamic characteristics of the embryo. To this end, we considered evolving surfaces and developed a generalised optical flow equation suitable for such changing geometries. Moreover, we showed that the approximate velocities of cells and from that their approximate trajectories can be estimated. The key is to compute the tangential velocities relative to the presumed surface evolution.

Second, we focused on a simple geometry and approximated the embryo with a round sphere. In contrast to the first model it represents a global model of the animal. While this neglects the true shape, it allows to easily experiment with various regularisation functionals. To this end, we suggested several decomposition models, each dividing the total motion of cells into different aspects and leading to insights about the collective behaviour and individual fluctuations of cells.

In a further generalisation we translated the idea of changing geometries to sphere-like surfaces. Such surfaces can be parametrised easily from the 2-sphere. They constitute a more realistic model of the embryo's shape and also are advantageous from a numerical perspective.

Assuming that the geometry can be represented suitably well, our approaches give reasonable approximations of the cells' velocities. However, several interesting questions remain open.

First, it remains to discuss and evaluate the obtained results together with trained biologists. A particularly challenging goal is to extract cell divisions from the computed motion fields. From a mathematical point of view this problem is interesting as such divisions constitute discontinuities in the flow field. A first step would be to employ discontinuity-preserving regularisation functionals paired with efficient numerical methods (see also the references in Chap. 2). Moreover, data terms could be tailored to appropriately model

the data recorded with confocal laser-scanning microscopy and to reflect the physical properties of cell tissue.

Finally, let us point out that the visualisation of evolving surfaces together with an image is challenging in general and even more so if motion fields are added. We feel that, in order to provide biologists with satisfying charts and figures, further research should be done.

All in all, we think that optical flow together with an appropriate model of the embryo's shape leads to insights into the movement patterns of cells, delivers quantitative methods for cell motion estimation, and can hopefully help answering the question raised in Chap. 1:

How do cells move and proliferate during early development?

Part II

Publications

Chapter 4

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

Authors & Contributions The authors are Clemens Kirisits, Lukas Lang, and Otmar Scherzer. The development of this article was a gradual, cooperative process, and each of the authors made significant contributions to every aspect of the paper.

Publication Status Published [50]: C. Kirisits, L. F. Lang, and O. Scherzer. Optical flow on evolving surfaces with an application to the analysis of 4D microscopy data. In A. Kuijper, K. Bredies, T. Pock, and H. Bischof, editors, *SSVM'13: Proceedings of the fourth International Conference on Scale Space and Variational Methods in Computer Vision*, volume 7893 of *Lecture Notes in Computer Science*, pages 246–257, Berlin, Heidelberg, 2013. Springer-Verlag.

The final publication is available at <http://link.springer.com/>.

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,
Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria

²Radon Institute of Computational and Applied Mathematics,
Austrian Academy of Sciences, Altenberger Str. 69, 4040 Linz, Austria

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.

4.1 Introduction

Advances in laser-scanning microscopy and fluorescent protein technology have increased resolution of microscopy imaging up to a single cell level [60]. 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 [83]. Loosely speaking, they only sit next to each other but not on top of each other. Moreover, the yolk deforms over time; see Fig. 4.1.

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




Figure 4.1: 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 (μm). See Sec. 4.4 for more details on the microscopy data.

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. 4.2 and 4.2). 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. 4.2). 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. 4.3). Finally, we apply this technique to obtain qualitative results from the afore-mentioned zebrafish data (Sec. 4.4). 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.

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 [39] 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 [9] for a recent survey.

Miura [65] 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. [61] and Hubený et al. [41] 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. [72] 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. [42, 79] considered optical flow for spherical images. In a more general setting, Lefèvre and Baillet [58] 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.

4.2 Optical Flow on Evolving Surfaces

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

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

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

$$\gamma: \mathcal{M}_0 \times I \rightarrow \bigcup_{t \in I} \mathcal{M}_t,$$

such that

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

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

$$\tilde{f}(\mathbf{x}_0, 0) = \tilde{f}(\gamma(\mathbf{x}_0, t), t), \text{ for all } (\mathbf{x}_0, t) \in \mathcal{M}_0 \times I. \quad (4.1)$$

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 $\mathbf{x} \in \mathcal{M}_{t_0}$ the usual partial derivative

$$\partial_t \tilde{f}(\mathbf{x}, t_0) = \lim_{h \rightarrow 0} \frac{1}{h} (\tilde{f}(\mathbf{x}, t_0 + h) - \tilde{f}(\mathbf{x}, t_0))$$

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

In the next section we linearise (4.1) 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 [15] to directly linearise (4.1).

Linearisation

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

$$\mathbf{x}: \Omega \times I \rightarrow \mathbb{R}^3, \quad (x_1, x_2, t) = (x, t) \mapsto \mathbf{x}(x, t) \in \mathcal{M}_t$$

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

$$f(x, t) = \tilde{f}(\mathbf{x}(x, t), t) \quad (4.2)$$

and let

$$\beta: \Omega \times I \rightarrow \Omega$$

be the coordinate counterpart of γ . This means, if we let $\mathbf{x}_0 = \mathbf{x}(x_0, 0)$, then $\beta(x_0, t)$ gives the coordinates of $\gamma(\mathbf{x}_0, t) \in \mathcal{M}_t$ in Ω (see Fig. 4.2). In other words, we have the identity

$$\gamma(\mathbf{x}(x_0, 0), t) = \mathbf{x}(\beta(x_0, t), t), \text{ for all } (x_0, t) \in \Omega \times I. \quad (4.3)$$

Now, from (4.1), (4.2) and (4.3) we get

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

which is a coordinate version of the BCA. After differentiation with respect to t it becomes

$$\nabla^2 f \cdot \dot{\beta} + \partial_t f = 0, \quad (4.4)$$

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 $\dot{\beta}$.

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

$$\begin{array}{ccc}
\Omega & \xrightarrow{\beta(\cdot, t)} & \Omega \\
\mathbf{x}(\cdot, 0) \downarrow & & \downarrow \mathbf{x}(\cdot, t) \\
\mathcal{M}_0 & \xrightarrow{\gamma(\cdot, t)} & \mathcal{M}_t
\end{array}$$

Figure 4.2: Commutative diagram describing the relation between unknowns β and γ .

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 \mathbf{y} be a point on \mathcal{M}_{t_0} and $\xi: t \mapsto \xi(t) \in \mathcal{M}_t$ an arbitrary smooth trajectory through the evolving surface satisfying $\xi(t_0) = \mathbf{y}$. Now we can compute

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

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

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 $\mathbf{y} \in \mathcal{M}_{t_0}$ in \mathbb{R}^3 , then the surface gradient of F at \mathbf{y} is defined as the projection of the three-dimensional spatial gradient $\nabla^3 F$ onto the tangent plane to \mathcal{M}_{t_0}

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

where $\hat{\mathbf{n}}$ 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. [31, p. 389]. Thus, $\nabla_{\mathcal{M}} \tilde{f} = \nabla_{\mathcal{M}} F$ is well-defined.

The spatial and temporal derivatives of \tilde{f} introduced above are related in a simple way. As shown in [15], they satisfy the equality

$$\begin{aligned}
d_t^{\mathbf{x}} \tilde{f} &= \nabla_{\mathcal{M}} \tilde{f} \cdot \dot{\mathbf{x}} + d_t^{\mathbf{n}} \tilde{f} \\
&= \nabla_{\mathcal{M}} \tilde{f} \cdot \dot{\mathbf{x}}_{\text{tan}} + d_t^{\mathbf{n}} \tilde{f},
\end{aligned} \tag{4.5}$$

where $\dot{\mathbf{x}}_{\text{tan}}$ is the tangential surface velocity, that is, the projection of $\dot{\mathbf{x}}$ onto the tangent plane to \mathcal{M}_{t_0} . This decomposition of $d_t^{\mathbf{x}} \tilde{f}$ into normal and tangential components is clearly valid for any trajectory in place of \mathbf{x} , and therefore in particular for the unknown γ . This means we can use (4.5) in order to differentiate the BCA (4.1) with respect to t . The resulting OFC reads

$$\nabla_{\mathcal{M}} \tilde{f} \cdot \dot{\gamma}_{\text{tan}} + d_t^{\mathbf{n}} \tilde{f} = 0. \tag{4.6}$$

Discussion. We conclude this section with a brief comparison of the two OFCs derived above. We start by showing how to obtain (4.4) from (4.6) and

vice versa. To this end we again assume the existence of a global parametrisation and rewrite all quantities in (4.6) in terms of \mathbf{x} . First observe that, by (4.3), the velocity of γ equals the surface velocity $\dot{\mathbf{x}}$ plus a purely tangential component

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

where $J = (\partial_1 \mathbf{x} \ \partial_2 \mathbf{x})$ is the Jacobian matrix of \mathbf{x} with respect to x . On the other hand, by (4.5), the normal time derivative is equal to the time derivative of \tilde{f} following \mathbf{x} minus its tangential component

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

Using the last two equations to rewrite the left-hand side of (4.6) yields

$$\begin{aligned} \nabla_{\mathcal{M}} \tilde{f} \cdot \dot{\gamma} + d_t^n \tilde{f} &= \nabla_{\mathcal{M}} \tilde{f} \cdot (\dot{\mathbf{x}} + J\dot{\beta}) + d_t^{\mathbf{x}} \tilde{f} - \nabla_{\mathcal{M}} \tilde{f} \cdot \dot{\mathbf{x}} \\ &= \nabla_{\mathcal{M}} \tilde{f} \cdot J\dot{\beta} + d_t^{\mathbf{x}} \tilde{f}, \end{aligned}$$

which is already the left-hand side of (4.4) in terms of \tilde{f} . It only remains to observe that $d_t^{\mathbf{x}} \tilde{f} = \partial_t f$ 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. [56].

We highlight the qualitative difference between the constraints (4.4) and (4.6). Note that in the former the unknown is $\dot{\beta}$, while in the latter it is $\dot{\gamma}_{\text{tan}} = \dot{\mathbf{x}}_{\text{tan}} + J\dot{\beta}$. This means that (4.4) constrains the motion relative to the tangential surface velocity $\dot{\mathbf{x}}_{\text{tan}}$, while (4.6) constrains the absolute tangential motion.

The nature of our microscopy data suggests a simple global parametrisation (see Sec. 4.3). We therefore pull the data back to the Euclidean plane and solve (4.4). However, equation (4.6) is independent of any parametrisation. It can thus serve as a starting point for alternative numerical approaches.

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 $\mathbf{u} = J\dot{\beta}$ with respect to the tangential basis $(\partial_1 \mathbf{x}, \partial_2 \mathbf{x})$ 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 \mathbf{u} . More precisely, we adopt a recent extension of the original quadratic Horn-Schunck regularisation to a Riemannian setting [58]. Basically, they propose to minimise

$$\mathcal{E}(u) = \frac{\alpha}{2} \|\nabla^2 f \cdot u + \partial_t f\|_{L^2(\mathcal{M})}^2 + \frac{1}{2} \|Du\|_{L^2(\mathcal{M})}^2. \quad (4.7)$$

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

$$\nabla_j \mathbf{u} = \sum_{i=1}^2 D_j u^i \partial_i \mathbf{x}, \quad j = 1, 2,$$

of \mathbf{u} . Using the Christoffel symbols Γ_{jk}^i (see Sec. 4.3) associated to the parametrisation \mathbf{x} the coefficients are given by

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

Rewriting (4.7) as an integral over the coordinate domain, we arrive at the functional

$$\mathcal{E}(u) = \frac{1}{2} \int_{\Omega} \left[\alpha (\nabla^2 f \cdot u + \partial_t f)^2 + \|Du\|_F^2 \right] \sqrt{\det g} \, dx, \quad (4.8)$$

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

4.3 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

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

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

$$\begin{aligned} \Delta u^1 &= \nabla^2 u^1 \cdot c + \nabla^2 u^2 \cdot d + u \cdot b_1 + a_1 \\ \Delta u^2 &= \nabla^2 u^2 \cdot c + \nabla^2 u^1 \cdot d + u \cdot b_2 + a_2. \end{aligned} \quad (4.9)$$

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

$$\partial_j u^i + \sum_k \Gamma_{jk}^i u^k = 0, \quad \text{for } x_j \in \{0, 1\}, \quad (4.10)$$

where $i, j \in \{1, 2\}$. In case of a flat manifold, e.g. $\mathcal{M} = \Omega$, the Euler-Lagrange equations (4.9) 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 [19].

Due to the nature of the microscopy data (see Sec. 4.4 and Fig. 4.1), 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 $z : \Omega \times I \rightarrow \mathbb{R}$. 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

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

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

$$\Gamma_{jk}^i = \frac{1}{2} \sum_{m=1}^2 g^{mi} (\partial_j g_{km} + \partial_k g_{mj} - \partial_m g_{jk}) = \frac{\partial_i z \partial_{jk} z}{\det g}.$$

Partial derivatives of z and of the projected data f were approximated by central differences. The system (4.9) with boundary conditions (4.10) was then solved with a standard finite difference scheme. In the following section numerical results are presented.

4.4 Experiments

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 time-lapse volumetric (4D) images; see [60] 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. [48] 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 [66]. The important aspect about endodermal cells is that they are known to form a monolayer during gastrulation [83], meaning that the radial extent is only a single cell. This crucial fact allows for the straightforward extraction of a surface together with a two-dimensional 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. 4.3. 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.

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 (4.8) as outlined in Sec. 4.3. 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 [9], 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 4.3, left, shows the colour-coded tangential vector field and the colour space whereas Fig. 4.3, 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. 4.3, 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.

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




Figure 4.3: 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.

Figure 4.4 gives a detailed view of the section outlined by a (red) rectangle in Fig. 4.3, right. This section was chosen because it depicts a cell division. Figure 4.4, left, and Fig. 4.4, right, display the frames before and after the event, respectively. Moreover, in Fig. 4.4, 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. 4.3, right, the event is pointed 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

4.5 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.

Acknowledgements. We thank Pia Aanstad from the University of Innsbruck for sharing her biological insight and for kindly providing the microscopy

³<http://www.csc.univie.ac.at>




Figure 4.4: 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

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).

Chapter 5

Optical Flow on Evolving Surfaces with Space and Time Regularisation

Authors & Contributions The authors are Clemens Kirisits, Lukas Lang, and Otmar Scherzer. The development of this article was a gradual, cooperative process, and each of the authors made significant contributions to every aspect of the paper.

Publication Status Published [52]: C. Kirisits, L. F. Lang, and O. Scherzer. Optical flow on evolving surfaces with space and time regularisation. *J. Math. Imaging Vision*, 52(1):55–70, 2015.

The final publication is available at <http://link.springer.com/>.

Optical Flow on Evolving Surfaces with Space and Time Regularisation

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

¹Computational Science Center, University of Vienna,
Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria

²Radon Institute of Computational and Applied Mathematics,
Austrian Academy of Sciences, Altenberger Str. 69, 4040 Linz, Austria

Abstract

We extend the concept of optical flow with spatiotemporal regularisation to a dynamic non-Euclidean setting. Optical flow is traditionally computed from a sequence of flat images. The purpose of this paper is 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: biomedical imaging, Computer Vision, evolving surfaces, optical flow, spatiotemporal regularisation, variational methods.

5.1 Introduction

Motivation

Advances in laser-scanning microscopy and fluorescent protein technology have increased resolution of microscopy imaging up to a single cell level [60]. They allow for four-dimensional (volumetric time-lapse) imaging of living organisms and shed light on cellular processes during early embryonic development. Understanding cellular processes often requires estimation and analysis of cell motion. However, the amount of data that is recorded is tremendous and therefore in many cases automated image analysis is necessary.

The specific biological motivation for this work is to understand the motion and division behaviour of fluorescently labelled endodermal cells of a zebrafish embryo. Although of considerable importance for developmental biology, knowledge about the migration patterns of these cells is scarce [74]. The dataset under consideration consists of volumetric time-lapse images taken by a laser-scanning microscope. The recorded sequence depicts a cuboid section $S \subset \mathbb{R}^3$ of said zebrafish embryo, whose endodermal cells express a fluorescent protein. We model this sequence by a scalar function

$$\bar{F} : [0, T] \times S \rightarrow \mathbb{R}$$

that assigns to every pair $(t, x) \in [0, T] \times S$ a nonnegative value $\bar{F}(t, x)$ proportional to the fluorescence response of point x at time t .

Optical flow methods are used regularly to estimate cellular motion, see Sec. 5.1. Applying them directly to our data \bar{F} to obtain a dense 3D velocity field

$$\mathbf{m} : [0, T] \times S \rightarrow \mathbb{R}^3$$

is possible but problematic from a computational point of view [4], even more so if temporal regularisation is to be included. We propose a solution to this by adapting our model according to biological facts about the nature of the marked cells.

Endodermal cells develop on the surface of the embryo’s yolk, where they form a non-contiguous monolayer [83]. Loosely speaking, they only sit next to each other but not on top of each other. Moreover, the yolk’s shape is roughly spherical and deforms over time. This means that the yolk’s surface can be modelled by an embedded two-dimensional manifold $\mathcal{M}_t \subset \mathbb{R}^3$, the subscript indicating dependence on time. In practice, \mathcal{M}_t can be approximated by fitting piecewise polynomials, for instance, to the cell centres.¹ Consequently it is possible to reduce the data dimension by only considering the restriction F of \bar{F} to this moving surface; see Fig. 5.2. More details on the acquisition and preprocessing of the microscopy data are given in Sec. 5.5. This dimension reduction, in turn, necessitates the development of an optical flow model for data defined on an evolving surface, which is the main contribution of this article.

Let t_0 be a fixed instant of time and $x_0 \in \mathcal{M}_{t_0}$. Assume a cell located at x_0 , indicated by a relatively high value of $F(t_0, x_0)$, moves with velocity $\mathbf{m}(t_0, x_0)$. On the other hand, suppose the yolk’s surface has velocity $\mathbf{V}(t_0, x_0)$. The purely tangential vector

$$\mathbf{u}(t_0, x_0) = \mathbf{m}(t_0, x_0) - \mathbf{V}(t_0, x_0) \quad (5.1)$$

describes the cell’s velocity relative to \mathbf{V} . Put differently, the total observed velocity \mathbf{m} of a cell is the sum of the surface velocity \mathbf{V} and the cell’s tangential velocity \mathbf{u} . Compare Fig. 5.1. While the former is a quantity extrinsic to the surface the latter is intrinsic. A motion estimation method dealing with the full 4D dataset \bar{F} would directly try to calculate \mathbf{m} for all $(t, x) \in [0, T] \times S$. The method proposed in this article, however, only computes the tangential field \mathbf{u} for a given surface velocity \mathbf{V} . The total velocity can then be recovered by adding the two vector fields.

In practice the true velocity of a moving surface might not be known and might even be impossible to determine from available data. This is also the case for the microscopy data considered in this paper. Our solution consists in picking one surface velocity \mathbf{V} that is consistent with \mathcal{M}_t , of which there are infinitely many in general, and to estimate the tangent field \mathbf{u} relative to this chosen surface velocity. While the resulting \mathbf{u} must be interpreted with care, it is reasonable to assume that the sum $\mathbf{u} + \mathbf{V}$ is close to the true total velocity \mathbf{m} . The selected surface velocity ideally strikes a balance between being easy to implement while being not too unnatural. While modelling the optical flow

¹Sometimes it is possible to already capture the yolk’s surface with the microscope in a second sequence of images. We do not, however, use such additional data in this article.




Figure 5.1: Sketch of a cell (indicated by a black ellipse) moving along a trajectory γ on a moving surface. The cell's velocity is given by $\partial_t \gamma = \mathbf{m}$, which can be decomposed into surface velocity \mathbf{V} and relative tangential motion \mathbf{u} .

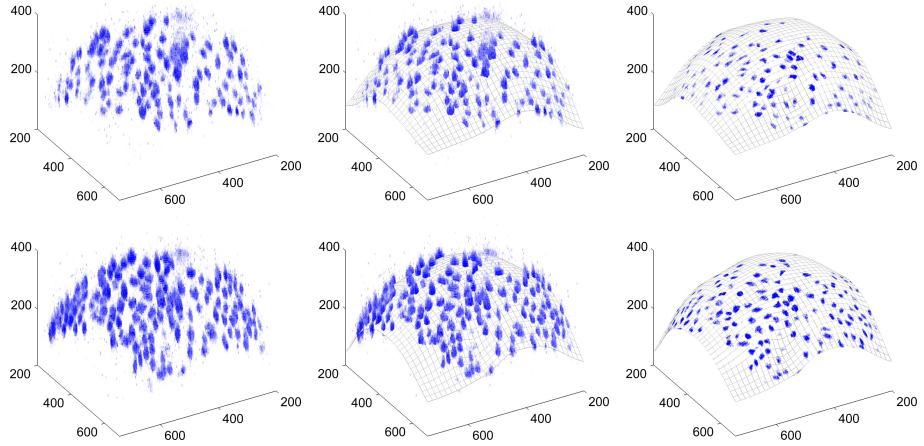


Figure 5.2: Frame no. 50 (top row) and 61 (bottom row) of the embryonic zebrafish image sequence. The left images illustrate the raw volumetric data \bar{F} . Intensity corresponds to fluorescence response. In the middle images, the curved mesh represents surfaces fitted to the cell's centres. The right images depict only the surface and the extracted two-dimensional image F . All dimensions are in micrometer (μm). For more details on the microscopy data and the preprocessing steps see Sec. 5.5.

on an evolving surface is the main novelty of this article, from the viewpoint of our particular application, it can be regarded as a subproblem making the computation of 3D velocities feasible, namely by reducing the data dimension while keeping as much accuracy as possible.

Contribution

The contributions of this article are as follows. First, we formulate the optical flow problem on an evolving two-dimensional manifold and derive a generalised optical flow constraint. Second, we translate the classical functional by Horn and Schunck [39] and its spatiotemporal extension by Weickert and Schnörr [86]

to the setting of moving manifolds. The associated Euler-Lagrange equations are solved with a finite difference scheme requiring a global parametrisation of the moving manifold. Finally, we apply this technique to obtain qualitative results from the aforementioned zebrafish data. Our experiments show that the optical flow is an appropriate tool for analysing these data. It is capable of visualising global trends as well as individual cell movements. In particular, the computed flow field can indicate cell divisions, while its integral curves approximate cell trajectories.

Finally, we address a point raised in the recent publication by Schmid et al. [74], who also analysed endodermal cell dynamics in a zebrafish embryo. They approximated the surface by a sphere, used different map projections to reduce the amount of data by one dimension, and subsequently computed cell motion in the plane. They acknowledge, however, the need for more exact, and supposedly slower, imaging techniques that do not discard any 3D information. While our approach still requires the volume data to be projected onto a surface and thus is faster than comparable 3D approaches, it does not require the surface to be very simple — e.g. spherical or planar — or static.

This article is structured as follows. In the next subsection we review related literature. Section 5.2 is devoted to providing the necessary mathematical background, notations, and definitions. Sections 5.3 and 5.4 introduce our variational model of optical flow on evolving surfaces and contain the continuous and discretised optimality conditions, respectively. In Sec. 5.5 we explain our microscopy data and the necessary preprocessing steps, summarise our approach, and finally present numerical results.

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 [39] 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 have been developed. See [9, 84] for recent surveys.

Using optical flow to extract motion information from cell biological data has gained popularity over the last decade. See, for example [2, 4, 14, 21, 41, 61, 65, 72, 74]. In these works displacement fields are computed either from 3D images or from 2D projections of the 3D data. While projections can suffer from inaccuracies [72, 74], the extraction of dense velocities from volumetric time-lapse data poses computational challenges [4]. In the present article we avoid both of these problems.

Many 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. [42, 79] considered optical flow for spherical images. Lefèvre and Baillet [58] extended the Horn-Schunck method to general 2-Riemannian manifolds, showed well-posedness, and applied it to brain imaging data. They solved 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.

A preliminary version of this paper appeared in [50]. The main differences to the present article are as follows. First, our current implementation allows us to regularise spatiotemporally as well as only spatially. In [50] we only treated

spatial regularisation. Second, the spatial regularisation functional has been improved in the sense that it is now parametrisation invariant. We have also conducted new experiments with the cell microscopy data and, in contrast to [50], computed approximate cell trajectories. Finally, we added some recent references.

5.2 Notation and Background

Whenever convenient we make use of the Einstein summation convention. Every index that appears exactly twice in an expression, once as a sub- and once as a superscript, is summed over.

Evolving Surfaces

Let $\mathcal{M} = (\mathcal{M}_t)_{t \in I}$ be a family of compact smooth 2-manifolds $\mathcal{M}_t \subset \mathbb{R}^3$ indexed by a time interval $I = [0, T]$. Each \mathcal{M}_t is assumed to be oriented by the unit normal field $\mathbf{N}(t, \cdot)$. For every $t \in I$ and $x \in \mathcal{M}_t$ the orthogonal projector onto the tangent plane $T_x \mathcal{M}_t$ is given by

$$\mathbf{P}(t, x) := \text{Id} - \mathbf{N}(t, x)\mathbf{N}(t, x)^\top. \quad (5.2)$$

We call \mathcal{M} an *evolving surface*, if there is a smooth function

$$\phi : I \times \mathcal{M}_0 \rightarrow \mathbb{R}^3$$

such that $\phi(t, \cdot)$ is a diffeomorphism between \mathcal{M}_0 and \mathcal{M}_t for every t , and $\phi(0, \cdot)$ is the identity on \mathcal{M}_0 . Note that ϕ cannot be unique in general. With every ϕ there is associated a surface velocity. Denote the inverse of $\phi(t, \cdot)$ by $\phi_t^{-1}(\cdot)$. The surface velocity at a point $x \in \mathcal{M}_t$ is then defined by

$$\mathbf{V}(t, x) := \partial_t \phi(t, \phi_t^{-1}(x)). \quad (5.3)$$

In contrast to ϕ the domain of \mathbf{V} is not $I \times \mathcal{M}_0$, but rather the 3-manifold

$$\bar{\mathcal{M}} := \bigcup_{t \in I} (\{t\} \times \mathcal{M}_t) \subset \mathbb{R}^4.$$

In other words, \mathbf{V} is a Eulerian specification of \mathcal{M} , while ϕ is a Lagrangian one. Even though different functions ϕ, ϕ' give rise to different velocities \mathbf{V}, \mathbf{V}' , the normal velocity of \mathcal{M} is independent of the choice of ϕ . That is, $\mathbf{V} \cdot \mathbf{N} = \mathbf{V}' \cdot \mathbf{N}$. We provide a short proof of this statement in Proposition 1 in the Appendix. Given a Eulerian specification \mathbf{V} of \mathcal{M} , we can obtain, at least locally, a Lagrangian one by solving the ordinary differential equation (5.3) for ϕ with initial condition $\phi(0, x_0) = x_0$. From now on we consider ϕ and \mathbf{V} fixed. See Sec. 5.5 for the specific ϕ and \mathbf{V} we use in the numerical computations.

Let $\mathbf{x}_0 : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}^3$ be a parametrisation of \mathcal{M}_0 mapping local coordinates $\xi = (\xi^1, \xi^2)$ to points $x = (x^1, x^2, x^3)$ of Euclidean space. By composing ϕ and \mathbf{x}_0 we obtain a parametrisation of the evolving surface \mathcal{M}

$$\mathbf{x} : I \times \Omega \rightarrow \mathbb{R}^3, \quad \mathbf{x}(t, \xi) = \phi(t, \mathbf{x}_0(\xi)). \quad (5.4)$$

With this convention we always have $\partial_t \mathbf{x} = \mathbf{V}$. Differentiation with respect to ξ^i will be denoted by ∂_i . The set $\{\partial_1 \mathbf{x}(t, \xi), \partial_2 \mathbf{x}(t, \xi)\}$ forms a basis of

$T_{\mathbf{x}(t,\xi)}\mathcal{M}_t$. Note that this basis is not orthonormal in general. Using dot notation for the standard inner product of \mathbb{R}^3 , the components of the first fundamental form $g = (g_{ij})$ are given by

$$g_{ij} = \partial_i \mathbf{x} \cdot \partial_j \mathbf{x}. \quad (5.5)$$

The elements of its inverse are denoted by upper indices $g^{-1} = (g^{ij})$.

Let $F : \bar{\mathcal{M}} \rightarrow \mathbb{R}$ be a scalar function and $f : I \times \Omega \rightarrow \mathbb{R}$ its coordinate representation,² that is

$$F(t, \mathbf{x}(t, \xi)) = f(t, \xi).$$

The integral of F over the evolving surface is then given by

$$\int_I \int_{\mathcal{M}_t} F \, dA \, dt := \int_I \int_{\Omega} f \sqrt{\det g} \, d\xi \, dt,$$

where dA denotes the surface measure.

We refer to [15], [24] and the references therein for more information on evolving surfaces. Eulerian and Lagrangian coordinates can be read up in Sec. 2.1 of [11], for example.

Derivatives on Evolving Surfaces

Spatial Derivatives. The spatial differential operators introduced below are not different from those on static manifolds. Therefore $t \in I$ can be considered fixed in this paragraph.

The surface gradient $\nabla_{\mathcal{M}} F$ of F is the tangent vector field which points in the direction of greatest increase of F . In local coordinates it is given by

$$\nabla_{\mathcal{M}} F = g^{ij} \partial_i f \partial_j \mathbf{x}, \quad (5.6)$$

where we omitted the arguments $(t, \mathbf{x}(t, \xi))$ on the left and (t, ξ) on the right hand side, respectively. The surface gradient is just the tangential part of the \mathbb{R}^3 gradient. More precisely, if \hat{F} is a smooth extension of F to an open neighbourhood of \mathcal{M}_t in \mathbb{R}^3 , then

$$\nabla_{\mathcal{M}} F = P \nabla_{\mathbb{R}^3} \hat{F}.$$

Note that the last expression does not depend on the choice of \hat{F} .

Similarly, for two tangent vector fields \mathbf{u}, \mathbf{v} on \mathcal{M}_t the covariant derivative $\nabla_{\mathbf{v}} \mathbf{u}$ of \mathbf{u} along \mathbf{v} is the tangential part of the conventional directional derivative of \mathbf{u} along \mathbf{v} . That is

$$\nabla_{\mathbf{v}} \mathbf{u} = P \nabla_{\mathbb{R}^3} \hat{\mathbf{u}}(\mathbf{v}),$$

where $\hat{\mathbf{u}}$ is an extension of \mathbf{u} as above and $\nabla_{\mathbb{R}^3} \hat{\mathbf{u}}(\mathbf{v})$ is the Jacobian of $\hat{\mathbf{u}}$ applied to \mathbf{v} . Let $\mathbf{u} := u^i \partial_i \mathbf{x}$ and $\mathbf{v} := v^i \partial_i \mathbf{x}$ be their representations in the coordinate basis. The covariant derivative then reads

$$\nabla_{\mathbf{v}} \mathbf{u} = \left(v^i \partial_i u^j + v^i u^k \Gamma_{ik}^j \right) \partial_j \mathbf{x}. \quad (5.7)$$

²Distinguishing between a surface quantity and its coordinate representation is often avoided. We decided, however, to make this distinction for the data F , and only for F , as we found it helpful especially in Sec. 5.3.

The Christoffel symbols Γ_{ik}^j are defined by the action of ∇ on the coordinate basis

$$\nabla_{\partial_i \mathbf{x}} \partial_k \mathbf{x} = \Gamma_{ik}^j \partial_j \mathbf{x}. \quad (5.8)$$

An explicit expression for the Christoffel symbols in terms of the first fundamental form is given by

$$\Gamma_{ik}^j = \frac{1}{2} g^{mj} (\partial_i g_{km} + \partial_k g_{mi} - \partial_m g_{ik}).$$

Recall that the coordinate basis is in general not orthonormal. In Sec. 5.4, however, we want to rewrite the regularisation functional in terms of an orthonormal basis in order to simplify subsequent calculations. Therefore we now make the little extra effort of expressing the covariant derivative $\nabla_{\mathbf{v}} \mathbf{u}$ in terms of an arbitrary, possibly non-coordinate, frame $\{\mathbf{e}_1, \mathbf{e}_2\}$. Writing $\mathbf{u} = w^i \mathbf{e}_i$ and $\mathbf{v} = z^i \mathbf{e}_i$ in this basis, the corresponding formula reads

$$\nabla_{\mathbf{v}} \mathbf{u} = \left(\nabla_{\mathbf{v}} w^j + z^i w^k \tilde{\Gamma}_{ik}^j \right) \mathbf{e}_j. \quad (5.9)$$

For scalar functions like w^j the covariant derivative $\nabla_{\mathbf{v}} w^j$ is just the directional derivative along \mathbf{v} . It can be computed by using linearity of the covariant derivative with respect to its lower argument

$$\nabla_{\mathbf{v}} w^j = \nabla_{v^i \partial_i \mathbf{x}} w^j = v^i \nabla_{\partial_i \mathbf{x}} w^j = v^i \partial_i w^j.$$

The $\tilde{\Gamma}_{ik}^j$ are the symbols associated to the new frame $\{\mathbf{e}_1, \mathbf{e}_2\}$. In analogy to (5.8), they are defined by

$$\nabla_{\mathbf{e}_i} \mathbf{e}_k = \tilde{\Gamma}_{ik}^j \mathbf{e}_j. \quad (5.10)$$

For an orthonormal frame $\{\mathbf{e}_1, \mathbf{e}_2\}$ the following transformation law describes the relation between the two types of symbols

$$\tilde{\Gamma}_{ik}^j = \delta^{jp} \alpha_p^h g_{hm} (\alpha_i^\ell \partial_\ell \alpha_k^m + \alpha_i^\ell \alpha_k^n \Gamma_{\ell n}^m), \quad (5.11)$$

where α_i^j is the $\partial_j \mathbf{x}$ -coordinate of \mathbf{e}_i , that is, $\mathbf{e}_i = \alpha_i^j \partial_j \mathbf{x}$ and δ^{jp} is the Kronecker delta. We give a short derivation of the equation above in Lemma 3 in the Appendix.

The covariant derivative of \mathbf{u} at a point (t, ξ) is a linear operator on $T_{\mathbf{x}(t, \xi)} \mathcal{M}_t$, mapping tangent vectors \mathbf{v} to tangent vectors $\nabla_{\mathbf{v}} \mathbf{u}$. Its 2-norm (Frobenius norm) can be computed via

$$\|\nabla \mathbf{u}(t, \xi)\|_2^2 = |\nabla_{\mathbf{e}_1} \mathbf{u}(t, \xi)|^2 + |\nabla_{\mathbf{e}_2} \mathbf{u}(t, \xi)|^2, \quad (5.12)$$

where $\{\mathbf{e}_1, \mathbf{e}_2\}$ now is an arbitrary orthonormal basis of the tangent space $T_{\mathbf{x}(t, \xi)} \mathcal{M}_t$, that is, $\mathbf{e}_i \cdot \mathbf{e}_j = \delta_{ij}$. Note that, if \mathbf{x} is a global parametrisation, then we can obtain a frame $\{\mathbf{e}_1, \mathbf{e}_2\}$ which is orthonormal everywhere by Gram-Schmidt orthonormalisation of the coordinate basis $\{\partial_1 \mathbf{x}, \partial_2 \mathbf{x}\}$.

For a thorough treatment of the concepts introduced in this section we refer to [23, 56]. More basic differential geometry texts are [22, 54], for example.




Figure 5.3: Sketch of different trajectories through the evolving surface giving rise to different temporal derivatives. Corresponding velocities are depicted in grey.

Temporal Derivatives. Let $x \in \mathcal{M}_{t_0}$. Denote by $\psi : t \mapsto \psi(t) \in \mathcal{M}_t$ a trajectory through \mathcal{M} with $\psi(t_0) = x$. We define the time derivative of F following ψ at x as³

$$d_t^\psi F(t_0, x) := \left. \frac{d}{dt} F(t, \psi(t)) \right|_{t=t_0}. \quad (5.13)$$

There are a few special cases of this derivative that are worth mentioning. Let ψ_N be a trajectory for which the vector $\partial_t \psi(t_0)$ is orthogonal to $T_x \mathcal{M}_{t_0}$. The corresponding derivative is called normal time derivative and denoted by

$$d_t^N F(t_0, x) := \left. \frac{d}{dt} F(t, \psi_N(t)) \right|_{t=t_0}. \quad (5.14)$$

Every Lagrangian coordinate system ϕ of \mathcal{M} engenders a time derivative like (5.13) in a natural way. For $x = \phi(t, y) \in \mathcal{M}_t$ the time derivative of F following ϕ is defined by

$$d_t^V F(t_0, x) := \left. \frac{d}{dt} F(t, \phi(t, y)) \right|_{t=t_0}. \quad (5.15)$$

We choose the notation $d_t^N F$ and $d_t^V F$, because the derivative (5.13) in fact only depends on the velocity of ψ at x , see Lemma 1. Finally, if \mathcal{M} is parametrised according to (5.4), which we assume from now on, then $d_t^V F = \partial_t f$. For illustration see Fig. 5.3.

We stress that if V is the physical surface velocity, then d_t^V is the natural time derivative for functions defined on $\bar{\mathcal{M}}$, since it measures the temporal change along trajectories $\phi(\cdot, y)$ of surface points. These trajectories are *not* cell trajectories in general. They coincide only if the cells do not move by themselves and all the motion is surface motion.

Lemma 1. *With the definitions from above, we have*

$$d_t^V F = d_t^N F + \nabla_{\mathcal{M}} F \cdot V.$$

³Note that this composition of F with ψ is necessary, because the conventional partial derivative $\partial_t F(t_0, x)$ is meaningless in general.

Proof. The main idea in this derivation from [15] is to consider the normally constant extension \hat{F} of F : Let $\mathcal{N} \subset \mathbb{R}^4$ be an open neighbourhood of $\bar{\mathcal{M}}$. If \mathcal{N} is chosen sufficiently small, it is possible to define a function $\hat{F} : \mathcal{N} \rightarrow \mathbb{R}$ that is smooth, constant on normal lines through \mathcal{M}_t for every t , and agrees with F on $\bar{\mathcal{M}}$. Therefore

$$\begin{aligned} \frac{d}{dt} F(t, \phi(t, y)) &= \frac{d}{dt} \hat{F}(t, \phi(t, y)) \\ &= \partial_t \hat{F} + \nabla_{\mathbb{R}^3} \hat{F} \cdot \partial_t \phi \\ &= d_t^{\mathbf{N}} F + \nabla_{\mathcal{M}} F \cdot \mathbf{V} \end{aligned}$$

The last equality holds because, by construction, $\nabla_{\mathbb{R}^3} \hat{F}$ equals $\nabla_{\mathcal{M}} F$ and

$$d_t^{\mathbf{N}} F = \frac{d}{dt} \hat{F}(t, \psi_{\mathbf{N}}(t)) = \partial_t \hat{F} + \nabla_{\mathbb{R}^3} \hat{F} \cdot \partial_t \psi_{\mathbf{N}} = \partial_t \hat{F}.$$

Finally, note that by definition (5.3) we have $\partial_t \phi = \mathbf{V}$. \square

Note that, since $\nabla_{\mathcal{M}} F$ is tangential, $d_t^{\mathbf{V}} F$ actually only depends on the tangential part $\mathbf{P}\mathbf{V}$ of \mathbf{V} . Here \mathbf{P} is the orthogonal projector defined in (5.2).

Let \mathbf{u} be a tangent vector field on the evolving surface \mathcal{M} , that is, a function $\mathbf{u} : \bar{\mathcal{M}} \rightarrow \mathbb{R}^3$ such that

$$\mathbf{u}(t, \cdot) : \mathcal{M}_t \rightarrow T\mathcal{M}_t$$

for all t . In analogy to the covariant derivative (5.7) and to (5.15), we define the following time derivative

$$\nabla_t \mathbf{u} = \mathbf{P} d_t^{\mathbf{V}} \mathbf{u}, \quad (5.16)$$

where application of $d_t^{\mathbf{V}}$ to \mathbf{u} is understood componentwise. Again we have $d_t^{\mathbf{V}} \mathbf{u} = \partial_t \mathbf{u}$. A normal time derivative for \mathbf{u} could be defined as well but will not be needed in the sequel. As in the scalar case, $\nabla_t \mathbf{u}$ can be considered the natural time derivative for a tangent vector field \mathbf{u} , if \mathbf{V} is the physical surface velocity. By setting

$$\nabla_t \partial_i \mathbf{x} = \Gamma_{0i}^j \partial_j \mathbf{x} \quad (5.17)$$

we arrive at the following expression for $\nabla_t \mathbf{u}$ in local coordinates

$$\nabla_t \mathbf{u} = \left(\partial_t u^j + u^i \Gamma_{0i}^j \right) \partial_j \mathbf{x}.$$

The new symbols have the explicit representation

$$\Gamma_{0i}^j = g^{jk} \partial_{ti} \mathbf{x} \cdot \partial_k \mathbf{x}, \quad (5.18)$$

which can be verified by taking inner products of both sides of (5.17) with the coordinate basis vectors.

Again, in order to simplify calculations later on, we want to express this derivative in terms of an orthonormal frame $\{\mathbf{e}_1, \mathbf{e}_2\}$. We have

$$\nabla_t \mathbf{u} = \left(\partial_t w^j + w^i \tilde{\Gamma}_{0i}^j \right) \mathbf{e}_j, \quad (5.19)$$

where the symbols $\tilde{\Gamma}_{0i}^j$ are defined as before and satisfy an analogous transformation law

$$\tilde{\Gamma}_{0i}^j = \delta^{jp} \alpha_p^h g_{hm} \left(\partial_t \alpha_i^m + \alpha_i^k \Gamma_{0k}^m \right). \quad (5.20)$$

The derivation is analogous to (5.11) and can be found in Lemma 3 in the Appendix.

5.3 Model Statement

Generalised Optical Flow Equation

We assume to be given an evolving surface \mathcal{M} together with a known Lagrangian specification ϕ or, equivalently, a Eulerian one \mathbf{V} . In addition we are given scalar data F on \mathcal{M} which we want to track over time.

Our optical flow model is based on the so-called brightness constancy assumption. For every $x \in \mathcal{M}_0$ we seek a trajectory $\gamma(\cdot, x) : t \mapsto \gamma(t, x) \in \mathcal{M}_t$ along which F is constant. In other words, we assume existence of a Lagrangian specification γ of \mathcal{M} such that

$$F(t, \gamma(t, x)) = F(0, x). \quad (5.21)$$

This implies that the time derivative of F following γ has to vanish identically. We deduce from Lemma 1 that the following generalised optical flow equation has to hold

$$d_t^{\mathbf{N}} F + \nabla_{\mathcal{M}} F \cdot \partial_t \gamma = 0, \quad (5.22)$$

where $d_t^{\mathbf{N}} F$ is the normal time derivative as defined in (5.14) and $\nabla_{\mathcal{M}} F$ is the surface gradient of F , cf. (5.6).

Let us continue the discussion of Sec. 5.1. According to our definition of γ , a cell located at $x_0 \in \mathcal{M}_{t_0}$ moves with velocity

$$\partial_t \gamma(t_0, \gamma_{t_0}^{-1}(x_0)) = \mathbf{m}(t_0, x_0) = \mathbf{u}(t_0, x_0) + \mathbf{V}(t_0, x_0), \quad (5.23)$$

where $\gamma_{t_0}^{-1}$ is the inverse of $\gamma(t_0, \cdot)$, \mathbf{m} is the total observed velocity of a cell as introduced in Sec. 5.1 and \mathbf{u} is its velocity relative to \mathbf{V} . The second equality above is due to decomposition (5.1). According to our assumptions at the beginning of this section, we consider \mathbf{V} as given so that the actual unknown is \mathbf{u} .

The remaining part of this subsection is devoted to rewriting (5.22) in terms of local coordinates. First, we give an interpretation of the coordinates u^i of \mathbf{u} with respect to the basis $\{\partial_1 \mathbf{x}, \partial_2 \mathbf{x}\}$. Let $\beta = (\beta^1, \beta^2) : I \times \Omega \rightarrow \Omega$ be the coordinate counterpart of γ , defined by the equation

$$\gamma(t, \mathbf{x}_0(\xi)) = \mathbf{x}(t, \beta(t, \xi)).$$

See also Fig. 5.4. Taking time derivatives on both sides and dropping arguments yields

$$\mathbf{m} = \mathbf{V} + \partial_t \beta^i \partial_i \mathbf{x},$$

since $\partial_t \mathbf{x} = \mathbf{V}$. We can conclude that $u^i = \partial_t \beta^i$, which means that (u^1, u^2) is just the 2D velocity of the parametrised trajectory β . It remains to rewrite (5.22) in terms of u^1 and u^2 .

Lemma 2. *The optical flow equation (5.22) is equivalent to*

$$d_t^{\mathbf{V}} F + \nabla_{\mathcal{M}} F \cdot \mathbf{u} = 0.$$

In local coordinates it reads

$$\partial_t f + u^i \partial_i f = 0.$$

$$\begin{array}{ccc}
\Omega & \xrightarrow{\beta(t, \cdot)} & \Omega \\
\mathbf{x}_0(\cdot) \downarrow & & \downarrow \mathbf{x}(t, \cdot) \\
\mathcal{M}_0 & \xrightarrow{\gamma(t, \cdot)} & \mathcal{M}_t
\end{array}$$

Figure 5.4: Commutative diagram describing the relation between β and γ .

Proof. We prove the assertion in two steps. First we show that

$$\mathrm{d}_t^{\mathbf{N}} F + \nabla_{\mathcal{M}} F \cdot \partial_t \gamma = \mathrm{d}_t^{\mathbf{V}} F + \nabla_{\mathcal{M}} F \cdot \mathbf{u},$$

and afterwards rewrite the right hand side in local coordinates.

By Lemma 1 the normal time derivative can be written as

$$\mathrm{d}_t^{\mathbf{N}} F = \mathrm{d}_t^{\mathbf{V}} F - \nabla_{\mathcal{M}} F \cdot \mathbf{V}$$

The other summand of (5.22) rewrites as

$$\nabla_{\mathcal{M}} F \cdot \partial_t \gamma = \nabla_{\mathcal{M}} F \cdot (\mathbf{V} + \mathbf{u}).$$

Note that \mathbf{V} is not assumed to be normal to \mathcal{M}_t , so that the term $\nabla_{\mathcal{M}} F \cdot \mathbf{V}$ does not vanish in general. However, it does appear twice with opposite signs. Finally recall that $\mathrm{d}_t^{\mathbf{V}} F = \partial_t f$ and by the definition of the first fundamental form

$$\begin{aligned}
\nabla_{\mathcal{M}} F \cdot \mathbf{u} &= g^{ij} \partial_i f \partial_j \mathbf{x} \cdot u^k \partial_k \mathbf{x} \\
&= g^{ij} g_{jk} \partial_i f u^k \\
&= \partial_i f u^i.
\end{aligned}$$

□

It is worth noting that the parametrised optical flow equation has precisely the same form as the classical 2D equation.

Regularisation

Directly solving the optical flow equation in the new setting is just as ill-posed as it is in the classical setting. We use variational regularisation to overcome this. In particular, we propose to minimise the following quadratic spatiotemporal functional to recover a vector field \mathbf{u} describing the tangential motion of data F .

$$\int_I \int_{\mathcal{M}_t} \left((\mathrm{d}_t^{\mathbf{V}} F + \nabla_{\mathcal{M}} F \cdot \mathbf{u})^2 + \lambda_0 |\nabla_t \mathbf{u}|^2 + \lambda_1 \|\nabla \mathbf{u}\|_2^2 \right) \mathrm{d}A \mathrm{d}t \quad (5.24)$$

Here $\lambda_0 \geq 0$ and $\lambda_1 > 0$ are regularisation parameters. Recall from Sec. 5.2 that \mathbf{u} is temporally regularised according to the assumed surface motion \mathbf{V} . Functional (5.24) is a generalisation of the one presented in [86] for the Euclidean setting.

Moreover, if $\lambda_0 = 0$, minimisation of (5.24) is equivalent to minimising

$$\int_{\mathcal{M}_t} \left((d_t^Y F + \nabla_{\mathcal{M}} F \cdot \mathbf{u})^2 + \lambda_1 \|\nabla \mathbf{u}\|_2^2 \right) dA \quad (5.25)$$

for every instant $t \in I$ separately. If $\mathcal{M}_t = \mathcal{M}_0$ for all t , the functional reduces to that of [58]. The spatial regularisation term as defined in (5.12) is independent of the chosen parametrisation. This is an improvement over the functional chosen in [50].

Example 1. *We end this section with a brief explanation, from an applied point of view, of why we regularise with covariant derivatives. Consider as a toy manifold the non-moving unit circle $\mathcal{M}_t = \mathcal{S}^1 \subset \mathbb{R}^2$ with parametrisation $\mathbf{x}(\theta) = (\cos \theta, \sin \theta)^\top$, $\theta \in [0, 2\pi)$ and tangent basis $\{\partial_\theta \mathbf{x}\}$. Consider the tangent vector field $\mathbf{u} = c \partial_\theta \mathbf{x}$, where $c \neq 0$ is a fixed number. This field would describe a uniform translation of data F along the circle, and thus should not be penalised by a regularisation term that enforces spatial smoothness. But while conventional differentiation does not yield a vanishing vector field*

$$\partial_\theta \mathbf{u} = c \partial_{\theta\theta} \mathbf{x} = -c \mathbf{x},$$

covariant differentiation does

$$\nabla_\theta \mathbf{u} = P \partial_\theta \mathbf{u} = -c P \mathbf{x} = -c(\mathbf{x} - \mathbf{x} \mathbf{x}^\top \mathbf{x}) = 0.$$

Here we used the fact that $\mathbf{N} = \mathbf{x}$ and $\mathbf{x}^\top \mathbf{x} = 1$.

An analogous argument explains our penalisation of $\nabla_t \mathbf{u}$ of γ instead of the unprojected derivative $\partial_t \mathbf{u}$.

5.4 Euler-Lagrange Equations

To simplify matters from now on we will assume having a global parametrisation \mathbf{x}_0 of \mathcal{M}_0 and thus a global parametrisation \mathbf{x} of the whole evolving surface, cf. (5.4). In addition, we express the functional (5.24) in an orthonormal non-coordinate basis $\{\mathbf{e}_1, \mathbf{e}_2\}$ with

$$\mathbf{e}_i = \alpha_i^j \partial_j \mathbf{x}. \quad (5.26)$$

This leads to wearisome calculations at first, which however pay off eventually when we compute the optimality conditions for the coordinates of \mathbf{u} with respect to this frame. Note that an orthonormal coordinate basis does not exist in general [56].

In this section we use the following notational convention. First, we identify t with ξ^0 . In addition, Latin indices are always understood to run over the set $\{1, 2\}$, while Greek indices are reserved for $\{0, 1, 2\}$.

Rewriting Functional (5.24)

Let

$$\mathbf{u} = w^i \mathbf{e}_i \quad (5.27)$$

be the representation of the unknown \mathbf{u} in the orthonormal frame (5.26). It follows that $w^j = w^i \alpha_i^j$. Recall from (5.9), (5.19) that the derivatives of \mathbf{u} read

$$\begin{aligned}\nabla_{\mathbf{e}_i} \mathbf{u} &= \left(\alpha_i^k \partial_k w^j + w^k \tilde{\Gamma}_{ik}^j \right) \mathbf{e}_j, \\ \nabla_t \mathbf{u} &= \left(\partial_t w^j + w^i \tilde{\Gamma}_{0i}^j \right) \mathbf{e}_j.\end{aligned}$$

If we set $\alpha_\mu^0 = \delta_\mu^0$ and $\alpha_0^\mu = \delta_0^\mu$, the coefficients of \mathbf{e}_j above can be rewritten using the unified notation

$$D_\mu w^j = \alpha_\mu^\nu \partial_\nu w^j + w^i \tilde{\Gamma}_{\mu i}^j,$$

where $\mu = 0, 1, 2$ and $j = 1, 2$. Consequently, defining the operator $D = (D_0, D_1, D_2)^\top$, the integrand of the regularisation term becomes a weighted 2-norm of the matrix $Dw = (D_\mu w^j)_{\mu j}$. The parametrised version of energy functional (5.24) now takes the following compact form

$$\int_0^T \int_\Omega \left((\partial_t f + w^j \alpha_j^i \partial_i f)^2 + \sum_{\mu, j} \lambda_\mu (D_\mu w^j)^2 \right) \sqrt{\det g} \, d\xi \, dt, \quad (5.28)$$

where $\lambda_1 = \lambda_2$ and g is the first fundamental form as introduced in (5.5). Observe that the simple form of the regulariser originates from representing $\nabla_{\mathbf{e}_i} \mathbf{u}$ and $\nabla_t \mathbf{u}$ in an orthonormal basis. This also simplifies the computation of the optimality conditions.

Optimality System

Denote the interior of $I \times \Omega \subset \mathbb{R}^3$ by D . Functional (5.28) takes the general form

$$\mathcal{E}(w) = \int_D L(w, \nabla w) \, d\xi,$$

where the Lagrangian L is a smooth function of all w^i and $\partial_\mu w^i$. Denote partial derivatives of L by subscripts. A minimiser (w^1, w^2) of \mathcal{E} has to satisfy the following second-order elliptic system

$$\begin{aligned}L_{w^m} &= \sum_\mu \partial_\mu L_{\partial_\mu w^m}, \quad \text{in } D, \\ 0 &= \sum_\mu n_\mu L_{\partial_\mu w^m}, \quad \text{on } \partial D,\end{aligned} \quad (5.29)$$

for $m = 1, 2$ and where $n = (n_0, n_1, n_2)^\top$ is the outward normal to D . The derivatives of the Lagrangian read

$$\begin{aligned}L_{w^m} &= \sqrt{\det g} \left(\alpha_m^i \partial_i f (w^j \alpha_j^k \partial_k f + \partial_t f) + \sum_{\mu, j} \lambda_\mu \tilde{\Gamma}_{\mu m}^j D_\mu w^j \right), \\ L_{\partial_\nu w^m} &= \sqrt{\det g} \sum_\mu \lambda_\mu \alpha_\mu^\nu D_\mu w^m.\end{aligned}$$

System (5.29) in terms of derivatives of w together with explicit formulas for all coefficients can be found in the Appendix. For more details on variational calculus we refer to [19, 30].




Figure 5.5: Eleven point stencil arising from the discretisation.

Remark 1. If \mathcal{M} is a fixed plane, then $\alpha_\mu^\nu = \delta_\mu^\nu$ and all connection symbols vanish. Consequently, the boundary conditions become standard Neumann ones and system (5.29) reduces to the one derived in [39] or [86], respectively.

Discretisation and Numerical Aspects

We solve the Euler-Lagrange equations (5.29) with a standard finite difference scheme. The spatiotemporal domain D is assumed to be the unit cube $(0, 1)^3$ and is approximated by a Cartesian grid with spacing of h_σ in the direction of ξ^σ , where $h_1 = h_2$. Grid points are denoted by p . Thus, $w_p^m := w^m(p)$ refers to the numerical approximation of w^m at $p \in D$. Partial derivatives of the unknowns are approximated using central differences. They read

$$\partial_\sigma w^m(p) \approx \frac{1}{2h_\sigma} \left(w_{N_\sigma^+}^m(p) - w_{N_\sigma^-}^m(p) \right),$$

$$\partial_{\sigma\sigma} w^m(p) \approx \frac{1}{h_\sigma^2} \left(w_{N_\sigma^+}^m(p) - 2w_p^m + w_{N_\sigma^-}^m(p) \right),$$

and

$$\partial_{\nu\sigma} w^m(p) \approx \frac{1}{4h_\nu h_\sigma} \left(w_{N_{\nu\sigma}^{++}}^m(p) - w_{N_{\nu\sigma}^{+-}}^m(p) - w_{N_{\nu\sigma}^{-+}}^m(p) + w_{N_{\nu\sigma}^{--}}^m(p) \right),$$

where the symbols $N_\sigma^\pm(p)$ and $N_{\nu\sigma}^{\pm\pm}(p)$ denote the neighbours of w_p^m in the grid along coordinates σ and ν, σ , respectively. From the choice of the discrete derivatives an eleven-point stencil is obtained; see Fig. 5.5. Derivatives of the data f and the surface parametrisation \mathbf{x} are handled likewise, using central differences in the interior and inward differences at the boundaries.

However, the resulting (sparse) linear system is underdetermined from equations (5.29) alone, because the approximations used for the mixed derivatives of w^m refer to points not occurring in any boundary condition. Thus, at every grid point $p \in C \subset \partial D$ with

$$C := (\{\xi^1 = 0\} \cup \{\xi^1 = 1\}) \cap (\{\xi^2 = 0\} \cup \{\xi^2 = 1\}) \quad (5.30)$$

additional boundary conditions are needed. At these points we set $n = (0, \pm 1, \pm 1)^\top$ in the boundary condition (5.29), which is a vector pointing

in the direction of the undetermined grid neighbour. This leads to expressions of the form $\pm\partial_1 w^m \pm \partial_2 w^m$, which, interpreted as a directional derivative, can be approximated by

$$\frac{1}{2\sqrt{2}h_\sigma} \left(w_{\mathcal{N}_{ij}^{\pm\pm}(p)}^m - w_{\mathcal{N}_{ij}^{\mp\mp}(p)}^m \right).$$

5.5 Experiments

Zebrafish Microscopy Data

As mentioned before, the biological motivation for this work are cellular image sequences of a zebrafish embryo. Endoderm cells expressing green fluorescent protein were recorded via confocal laser-scanning microscopy resulting in time-lapse volumetric (4D) images. See e.g. [60] for the imaging techniques.

The microscopy images 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, which itself deforms over time. Detailed explanations and numerous illustrations can be found in [48]. For the biological methods such as the fluorescence marker and the embryos used in this work we refer to [66].

The captured area is approximately $540 \times 490 \times 340 \mu\text{m}^3$ and shows the pole region of the yolk. Figure 5.2, left column, depict two frames of the raw data. The sequence contains 77 frames recorded in intervals of 240 s with clearly visible cellular movements and cell divisions. The spatial resolution of the data is $512 \times 512 \times 44$ voxels. Intensities are in the range $[0, 1]$. In the following we denote by

$$\bar{F}^\delta \in [0, 1]^{77 \times 512 \times 512 \times 44}$$

the unprocessed microscopy data approximating \bar{F} from Sec. 5.1.

The important aspect about endodermal cells is that they are known to form a monolayer during gastrulation [83]. In other words, the radial extent is only a single cell. This crucial fact allows for the straightforward extraction of a surface together with an image of the stained cells. Figure 5.2 illustrates the idea for two particular frames.

Preprocessing and Acquisition of Surface Data

In this section, we relate the mathematical concepts introduced in Sec. 5.2 to the 4D microscopic images. We give a concrete global parametrisation suitable for this type of data and discuss the necessary preprocessing steps leading to an approximation of the evolving surface $\bar{\mathcal{M}}$ together with an approximation of the scalar quantity F .

The first step is to extract approximate cell centres from the raw microscopy data. As the positions of cells are characterised by local maxima in intensity they can be reliably obtained as follows. For every frame, a Gaussian filter is applied to the volumetric data \bar{F}^δ . Then, local maxima with respect to intensity are computed and treated as cell centres whenever they exceed a certain threshold.




Figure 5.6: Sequence of embryonic zebrafish images. Depicted are frames no. 46 to 60 of the entire sequence (aligned left to right, top to bottom).

Next we fit a surface to the cell positions. For every frame, this is done by least squares fitting of a piecewise linear function combined with first-order regularisation. From that we get a height field $z^\delta \in \mathbb{R}^{77 \times 512 \times 512}$ which completely describes the discrete evolving surface. Finally, the numerical approximation f^δ of f is calculated by linear interpolation of \bar{F}^δ and evaluation at surface points determined by z^δ .

The combination of all processing steps described above turns the original 4D array \bar{F}^δ into two three-dimensional arrays

$$\begin{aligned} f^\delta &\in [0, 1]^{77 \times 512 \times 512}, \\ z^\delta &\in \mathbb{R}^{77 \times 512 \times 512}. \end{aligned}$$

Figure 5.2, right column, illustrates both surfaces and the obtained images for two particular frames. In Fig. 5.6, a segment of the sequence is shown.

Let us quickly relate z^δ to the quantities introduced in Sec. 5.2. The mapping

$$(t, \xi^1, \xi^2) \mapsto (\xi^1, \xi^2, z^\delta(t, \xi^1, \xi^2)),$$

where (t, ξ^1, ξ^2) ranges over a $77 \times 512 \times 512$ grid, is the discrete parametrisation. The corresponding ϕ is the function that identifies surface points with identical (ξ^1, ξ^2) coordinates. Thus, the surface motion \mathbf{V} occurs only in direction of x^3 . However, we stress that this particular parametrisation was chosen due to the lack of knowledge about the true motion of material points on the surface.

Solving for the Velocity Fields

After the preprocessing of the microscopy data as explained above, the following steps lead to the desired solution:

1. From the parametrisation compute approximations of $\partial_i \mathbf{x}$, g , Γ_{ij}^k , α_i^j , $\tilde{\Gamma}_{\mu j}^k$ as explained in Sec. 5.2. Like all other quantities the α_i^j are functions of space and time. They can be computed, for example, by Gram-Schmidt orthonormalisation of the coordinate basis $\{\partial_1 \mathbf{x}, \partial_2 \mathbf{x}\}$.

2. Discretise optimality system (5.33) as described in Sec. 5.4.
3. Compute coefficients (5.34) of discretised optimality system from the quantities calculated in step 1.
4. Solve resulting linear system for unknowns w , see Sec. 5.5.
5. Compute relative tangential velocity \mathbf{u} via (5.27) and recover total velocity $\mathbf{m} = \mathbf{u} + \mathbf{V}$.
6. Finally, cell trajectories can be approximated by computing the integral curves of \mathbf{m} , see (5.32) in Sec. 5.5.

Visualisation

In order to illustrate the computed tangential velocity fields we apply the standard flow colour-coding from [9].⁴ This coding turns \mathbb{R}^2 vector fields into colour images according to a particular 2D colour space.

However, we are interested in visualising tangential vector fields on an embedded manifold, which are functions with values in \mathbb{R}^3 . To be able to apply the colour-coding mentioned above we turn the computed optical flow fields \mathbf{u} into \mathbb{R}^2 vector fields in the following way

$$\mathbf{u} \mapsto \frac{|\mathbf{u}|}{|\mathbf{P}_{x^3}\mathbf{u}|} \mathbf{P}_{x^3}\mathbf{u}, \quad (5.31)$$

where $\mathbf{P}_{x^3} : (x^1, x^2, x^3) \mapsto (x^1, x^2, 0)$ is the orthogonal projection onto the x^1 - x^2 plane. If the scaling factor $\frac{|\mathbf{u}|}{|\mathbf{P}_{x^3}\mathbf{u}|}$ were omitted, the new vector field would simply be the original one as viewed from above. The reason for including this scaling are vectors having a large x^3 -component, which would otherwise seem unnaturally short. Finally, the image resulting from the colour-coding of vector field (5.31) is painted back on the surface. Figure 5.7 illustrates the colour-coded tangential vector field \mathbf{u} and the colour space. In all figures the surface is slightly smoothed for better visual effect.

Numerical results

We conducted four experiments with different parameter settings and minimised functional (5.24) as outlined in Sec. 5.5. Due to a low cell density near the boundaries we only worked with a part of the whole dataset. The grid dimensions were $(N_0, N_1, N_2) = (30, 370, 370)$. Accordingly, grid spacing was set to $h_\sigma = 1/N_\sigma$. Our implementation was done in Matlab and all experiments were performed on an Intel Xeon E5-1620 3.6GHz workstation with 128GB RAM. We used the Generalized Minimal Residual Method (GMRES) to solve the resulting linear system. As a termination criterion we chose a relative residual of 0.02 and a maximum number of 2000 iterations with a restart every 30 iterations. The resulting runtime was approximately two hours. In Table 5.1, the parameters for all experiments are listed, and the resulting running times and relative residuals are given. Implementation and data are available on our website.⁵

⁴Some figures may appear in colour only in the online version.

⁵<http://www.csc.univie.ac.at>




Figure 5.7: Optical flow field between frames 60 and 61 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. Parameters are $\lambda_0 = c/100$ and $\lambda_1 = \lambda_2 = c$, where $c := 0.5$.

No.	λ_0	$\lambda_1 = \lambda_2$	Runtime	Rel. residual
1	c	c	2.05 h	0.075
2	$c/10$	c	2.07 h	0.086
3	$c/100$	c	2.09 h	0.103
4	$c/100$	$c/10$	2.14 h	0.016

Table 5.1: Runtimes and relative residuals of the experiments. For convenience, we define $c := 0.5$.

Regularisation. In a first experiment, we compared different regularisation parameters. They were chosen such that individual movements of cells are well preserved and the velocity field is sufficiently homogeneous both in time and space. Figure 5.8 depicts these results. A visual inspection of the dataset shows that cells tend to move towards the embryo's body axis which roughly runs from the bottom left to the top right corner in Fig. 5.7, right. This behaviour is clearly visible from the obtained velocity fields. In Fig. 5.9, we show the optical flow field for the sequence depicted in Fig. 5.6.

Cell Trajectories. In order to reconstruct the paths travelled by individual cells, we computed the integral curves of \mathbf{m} . By (5.23), for every starting point $x_0 \in \mathcal{M}_0$ the trajectory $\gamma(\cdot, x_0)$ is the solution of the following ordinary differential equation

$$\begin{aligned} \partial_t \gamma(t, x_0) &= \mathbf{m}(t, \gamma(t, x_0)), \\ \gamma(0, x_0) &= x_0, \end{aligned} \tag{5.32}$$

where \mathbf{m} is the total velocity of a cell; cf. Sec. 5.2. As discussed in Sec. 5.5, a local maximum of F at $x_0 \in \mathcal{M}_0$ indicates the approximate position of a cell. Hence, we chose local maxima as initial values and approximated (5.32) by solving the projection of

$$\begin{aligned} \hat{\gamma}(t+1, x_0) &= \hat{\gamma}(t, x_0) + s\mathbf{m}(t, \hat{\gamma}(t, x_0)) \\ \hat{\gamma}(0, x_0) &= x_0, \end{aligned}$$

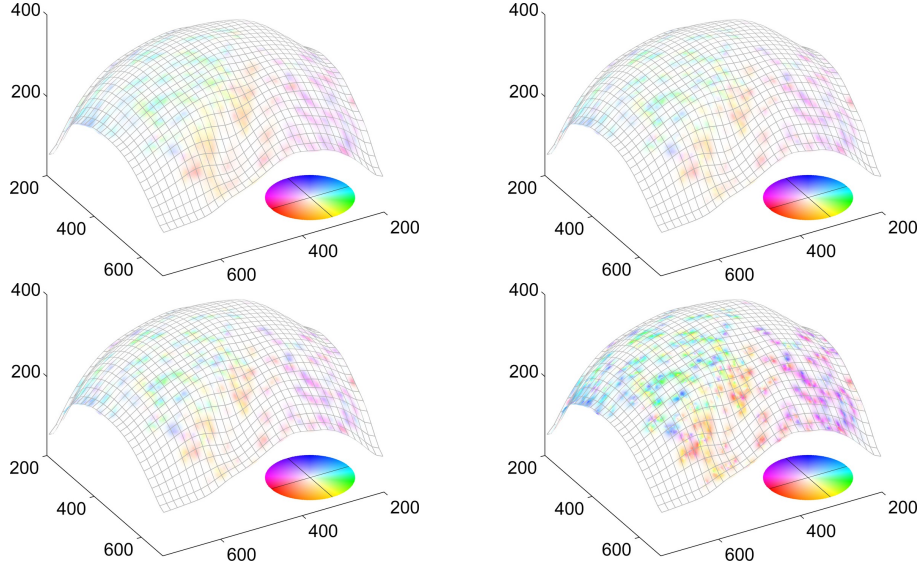


Figure 5.8: Resulting velocity field \mathbf{u} between frames 60 and 61 obtained with different regularisation parameters. Denote $c := 0.5$. Top left: $\lambda_0 = \lambda_1 = \lambda_2 = c$. Top right: $\lambda_0 = c/10$ and $\lambda_1 = \lambda_2 = c$. Bottom left: $\lambda_0 = c/100$ and $\lambda_1 = \lambda_2 = c$. Bottom right: $\lambda_0 = c/100$ and $\lambda_1 = \lambda_2 = c/10$.




Figure 5.9: Sequence of colour-coded tangential velocity fields. Depicted are the same frames as in Fig. 5.6. Parameters are $\lambda_0 = c/100$ and $\lambda_1 = \lambda_2 = c/10$.




Figure 5.10: Integral curves for frames $\{40, \dots, 60\}$ for the identical regularisation parameters as in Fig. 5.8. The colour gradient of a trajectory from yellow to green (bright to dark) indicates temporal progress. Local intensity maxima at the first frame serve as initial values. The embryo's body axis runs from bottom left to top right.

to the x^1 - x^2 -plane, because it allows for a better illustration. The parameter s is a step size and was chosen as $s := 10$. Figure 5.10 shows the projection $P_{x^3} \hat{\gamma}$ of the computed curves for several values of the regularisation parameters. The effect on the smoothness of the trajectories is clearly visible.

Cell Divisions. Figure 5.11 shows two cell divisions in more detail. The displacement field clearly resembles the splitting of the mother cell and the diverging daughter cells. Our results suggest that cell divisions can be indicated reasonably well by the proposed model.




Figure 5.11: Detailed view of two cell divisions occurring between frames 41 and 44 (top, left to right) and frames 55 and 58 (bottom, left to right). Parameters are $\lambda_0 = c/100$ and $\lambda_1 = \lambda_2 = c/10$. Vectors are scaled and only every second vector is shown. Data intensities are interpolated for smooth illustration.

5.6 Conclusion

Aiming at an accurate and efficient motion analysis of 4D cellular microscopy data, we generalised both the Horn-Schunck and Weickert-Schnörr functionals to images defined on evolving surfaces. The resulting optical flow constraint was solved by means of quadratic regularisation and verified on the basis of real data. Our experimental results suggest that cell movements including divisions are well captured by our model.

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).

Appendix

We first sketch a proof about the statement from Sec. 5.2 that the normal velocity of an evolving surface is independent of ϕ .

Proposition 1. *Let ϕ be a Langrangian specification of \mathcal{M} and \mathbf{V} the corresponding velocity as defined in (5.3). Then $\mathbf{V} \cdot \mathbf{N}$ is independent of the chosen specification.*

Proof. We can represent $\bar{\mathcal{M}}$ locally as the level set of a real-valued function $G(t, x)$, whose gradient does not vanish, see e.g. [57, Prop. 5.16]. We now express $\mathbf{V} \cdot \mathbf{N}$ solely in terms of G and thereby prove the assertion. Observing that the composition of G with ϕ is constant, we calculate

$$0 = \frac{d}{dt} G(t, \phi(t, x_0)) = \partial_t G + \nabla_{\mathbb{R}^3} G \cdot \mathbf{V} = \partial_t G + |\nabla_{\mathbb{R}^3} G| \mathbf{V} \cdot \mathbf{N}.$$

The second equality holds, because $\nabla_{\mathbb{R}^3} G$ is normal to the surface. We conclude that

$$\mathbf{V} \cdot \mathbf{N} = -\frac{\partial_t G}{|\nabla_{\mathbb{R}^3} G|}.$$

□

In other words, different specifications of a surface can only differ in their respective tangential velocities.

Next we prove the transformation law (5.11), (5.20) for the connection coefficients $\tilde{\Gamma}_{\mu j}^j$.

Lemma 3. *The symbols defined by (5.10) are given by (5.11).*

Proof. Take inner products on both sides of (5.11) with \mathbf{e}_j to get

$$\mathbf{e}_j \cdot \nabla_{\mathbf{e}_i} \mathbf{e}_k = \tilde{\Gamma}_{ik}^j.$$

Next express both terms on the left hand side in the coordinate basis by using $\mathbf{e}_j = \alpha_j^m \partial_m \mathbf{x}$ and formula (5.7). The assertion follows now immediately. □

An analogous calculation yields formula (5.20).

For our implementation the Euler-Lagrange equations (5.29) are needed in the following form

$$\begin{aligned} d^{\nu\sigma} \partial_{\nu\sigma} w^m + c_i^{\sigma m} \partial_{\sigma} w^i + b_i^m w^i &= a^m, & \text{in } D, \\ q^{\nu\sigma} \partial_{\sigma} w^m + p_i^{\nu m} w^i &= 0, & \text{on } \{\xi^{\nu} = 0\} \cup \{\xi^{\nu} = 1\}, \end{aligned} \quad (5.33)$$

where we assumed $D = (0, 1)^3$. As usual the system is to be understood for $m = 1, 2$ and $\nu = 0, 1, 2$. Below we give the exact coefficients.

$$\begin{aligned} a^m &= -\alpha_m^i \partial_i f \partial_t f \\ b_i^m &= \alpha_m^j \alpha_i^k \partial_j f \partial_k f + \sum_{\mu} \lambda_{\mu} \left(\sum_j \tilde{\Gamma}_{\mu m}^j \tilde{\Gamma}_{\mu i}^j - G_{\nu} \alpha_{\mu}^{\nu} \tilde{\Gamma}_{\mu i}^m + \partial_{\nu} (\alpha_{\mu}^{\nu} \tilde{\Gamma}_{\mu i}^m) \right) \\ c_i^{\sigma m} &= \sum_{\mu} \lambda_{\mu} \left(\alpha_{\mu}^{\sigma} \tilde{\Gamma}_{\mu m}^i - \alpha_{\mu}^{\sigma} \tilde{\Gamma}_{\mu i}^m - \delta_{im} (G_{\nu} \alpha_{\mu}^{\nu} \alpha_{\mu}^{\sigma} + \partial_{\nu} (\alpha_{\mu}^{\nu} \alpha_{\mu}^{\sigma})) \right) \\ d^{\nu\sigma} &= -\sum_{\mu} \lambda_{\mu} \alpha_{\mu}^{\nu} \alpha_{\mu}^{\sigma} \\ p_i^{\nu m} &= \sum_{\mu} \lambda_{\mu} \alpha_{\mu}^{\nu} \tilde{\Gamma}_{\mu i}^m \\ q^{\nu\sigma} &= \sum_{\mu} \lambda_{\mu} \alpha_{\mu}^{\nu} \alpha_{\mu}^{\sigma} \end{aligned} \quad (5.34)$$

Here we used the shorthand

$$G_\nu = \frac{\partial_\nu \sqrt{\det g}}{2\sqrt{\det g}}.$$

Recall that the functional without time regularisation (5.25) leads to a sequence of decoupled systems for every instant t . Each of those has the form

$$\begin{aligned} d^{jk} \partial_{jk} w^m + c_i^{km} \partial_k w^i + b_i^m w^i &= a^m, & \text{in } D, \\ q^{jk} \partial_k w^m + p_i^{mj} w^i &= 0, & \text{on } \{\xi^j = 0\} \cup \{\xi^j = 1\}. \end{aligned}$$

Note that, in comparison to system (5.33), we only replaced Greek indices by Latin ones. The coefficients a, b, c, d, p, q of this simpler system can be obtained from the list above by setting $\lambda_0 = 0$.

Chapter 6

Decomposition of Optical Flow on the Sphere

Authors & Contributions The authors are Clemens Kirisits, Lukas Lang, and Otmar Scherzer. The development of this article was a gradual, cooperative process, and each of the authors made significant contributions to every aspect of the paper.

Publication Status Published [51]: C. Kirisits, L. F. Lang, and O. Scherzer. Decomposition of optical flow on the sphere. *GEM. Int. J. Geomath.*, 5(1):117–141, 2014.

The final publication is available at <http://link.springer.com/>.

Decomposition of Optical Flow on the Sphere

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

¹Computational Science Center, University of Vienna,
Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria

²Radon Institute of Computational and Applied Mathematics,
Austrian Academy of Sciences, Altenberger Str. 69, 4040 Linz, Austria

Abstract

We propose a number of variational regularisation methods for the estimation and decomposition of motion fields on the 2-sphere. While motion estimation is based on the optical flow equation, the presented decomposition models are motivated by recent trends in image analysis. In particular we treat $u + v$ decomposition as well as hierarchical decomposition. Helmholtz decomposition of motion fields is obtained as a natural by-product of the chosen numerical method based on vector spherical harmonics. All models are tested on time-lapse microscopy data depicting fluorescently labelled endodermal cells of a zebrafish embryo.

6.1 Introduction

Motion estimation is a fundamental task for the analysis of spatiotemporal data, the prototypical example of which are sequences of images taken by a camera. The term *optical flow* has been coined to designate the apparent motion in such data. Its accurate and efficient estimation has been a major topic in the fields of computer vision and image processing for more than 30 years. However, the applicability of optical flow algorithms is by no means limited to flat two-dimensional projections of real world scenes. The advance of microscopy techniques has led to a particularly promising application of optical flow: cell motion analysis. Reliable optical flow algorithms supplied with microscopy images of sufficiently high spatial and temporal resolution can obviously help understanding cellular dynamics in transparent organisms, see for example [4, 61, 72, 74].

The particular dataset we are working with in this article depicts a living zebrafish embryo during early embryogenesis. Main feature of this dataset are the embryo's endodermal cells, which have been labelled with a fluorescent protein and are known to develop on the surface of the zebrafish's spherical yolk, see Fig. 6.1 and Sec. 6.5. The distribution of these cells can be modelled by a nonnegative function F depending on time t and position x on the 2-sphere, such that the number $F(t, x)$ is directly proportional to the fluorescence




Figure 6.1: Frames no. 57 (left) and 58 (right) of the zebrafish microscopy image sequence. Blue colour indicates fluorescence response. The spherical image is obtained by a radial projection of the unfiltered data onto a fitted sphere. Top and bottom row differ by a rotation of 180 degrees around the z -axis. See Sec. 6.5 for more details on the data and preprocessing.

response of a point x at time t . The models we propose below are motivated, but surely not restricted, to this specific type of data. We argue that the problem of extracting and analysing motion from spherical data is sufficiently general so as to be of potential interest to a wider audience.

Our motion models are based on the optical flow equation

$$\nabla_S F(t, x) \cdot u(t, x) + \partial_t F(t, x) = 0, \quad (6.1)$$

which every vector field u describing the temporal evolution of F should (approximately) satisfy for all x and t . Here ∇_S denotes the surface gradient on the sphere. We derive this equation in more detail in Sec. 6.3. Directly solving the optical flow equation for u is infeasible. We therefore use Tikhonov regularisation to compute an approximate solution to (6.1). Tikhonov regularisation consists in minimising a functional which is a weighted sum of two terms. The first one, usually called data or similarity term, is the squared L^2 norm of the left hand side of (6.1). The second term is a regularising functional $\mathcal{R}(u)$, which in this article will always be a Sobolev H^s norm (Sec. 6.3). These norms are introduced in Sec. 6.2.

Next, we extend the motion estimation method outlined above to two types of decomposition models that are also variational in nature (Sec. 6.3). While the input of those is again F , their outputs are now two or more vector fields capturing different structural parts of the total motion. Both models are adaptations of recently proposed image decomposition techniques to the optical flow setting. The first one is a $u + v$ decomposition. Its idea is to replace u in the data term with a sum $u + v$, and then to add two different regularising

functionals, one for u and one for v . The second one is a hierarchical model. Roughly speaking, one repeatedly minimises an optical flow functional, while the amount of regularisation is constantly decreased. In every iteration, the u in the data term is replaced by a sum $u + \sum_i u_i$, where only u is optimised and the u_i are the results from previous iterations.

Finally, all optimisation problems are solved by projecting them onto finite dimensional spaces spanned by vector spherical harmonics (Sec. 6.4). One advantage of this method is that it automatically yields Helmholtz decompositions of all computed vector fields. Mathematical background on (vector) spherical harmonics is presented in Secs. 6.2 and 6.2. In Sec. 6.5 we provide details of our implementation, give a more detailed account of the used microscopy data, and show experimental results with these data.

To summarise, the main novelties presented in this article are decomposition models for optical flow on the sphere together with their application to microscopy data.

Related Work

The first variational optical flow method is usually attributed to [39], where they used an H^1 seminorm for regularisation. In [75] it was shown that this particular choice leads to a well-posed problem. We refer to [8] for a gentle introduction to optical flow, to [84] for an overview of different optical flow functionals and to [9] for a recent survey and benchmark.

Optical flow algorithms have only recently been extended to data defined on non-Euclidean domains. In [42, 79] images defined on the sphere were treated, whereas in [58] the original functional by Horn and Schunck was generalised to 2-Riemannian manifolds and well-posedness was verified. Most recently, optical flow on evolving manifolds has been considered in [50, 52].

Horn and Schunck [39] numerically solved the variational problem by applying a finite difference scheme to the Euler-Lagrange equations. A similar approach was adopted in [50, 52] after parametrisation of the surface. In [58], however, the problem was solved by finite elements on a surface triangulation. Finally, we mention the work by Schuster and Weickert [76], where they used projection methods to solve the optical flow equation in the plane. Instead of Tikhonov-regularising their solution, they solely relied on regularisation by discretisation. The main reason for choosing a numerical method based on vector spherical harmonics in this article is that H^s -type regularisers, for arbitrary real s , are handled very easily in contrast to most other methods.

The aim of $u + v$ image decomposition models, as pioneered in [62], is to separate the cartoon and texture parts of images. While the cartoon component should capture large-scale structural components and should therefore be piecewise smooth, the texture component is supposed to consist of high-frequency oscillating patterns. Since the original model was promising but hard to implement, a large number of modifications and approximations have been proposed. In some of them the problematic g -norm was approximated by an H^{-1} norm [67, 82]. Recently, $u + v$ models have been extended to the \mathbb{R}^2 optical flow setting [1]. Hierarchical models, originally introduced in [77] for image analysis, have not yet been tried in combination with optical flow. They have, however, the preferable property of producing arbitrarily fine multiscale descriptions of input data. As a concluding remark about vector field decom-

positions, let us remark that Helmholtz-Hodge decompositions of motion fields have enjoyed a certain degree of attention in recent years, not only in the plane [53, 90, 91], but also on surfaces [47].

Applying optical flow algorithms to cell microscopy data has become increasingly popular lately. See for example [4, 50, 52, 61, 74] and the references therein. We highlight the article [74], where also endodermal cells of zebrafish embryos have been analysed. There the authors point out that, although of immense importance for developmental biology, only little is known about the motion behaviour of this type of cells.

6.2 Notation and Background

Scalar Spherical Harmonics

Let

$$\mathcal{S} = \{x \in \mathbb{R}^3 : |x| = 1\}$$

be the two-sphere embedded in \mathbb{R}^3 . For functions $F : \mathcal{S} \rightarrow \mathbb{R}$ we define the Laplace-Beltrami operator by

$$\Delta_{\mathcal{S}} F = -\Delta \bar{F},$$

where Δ is the usual Laplacian of \mathbb{R}^3 and $\bar{F}(x) = F(x/|x|)$ is the radially constant extension of F to $\mathbb{R}^3 \setminus \{0\}$. The eigenvalues of $\Delta_{\mathcal{S}}$ are

$$\lambda_n = n(n+1), \quad n \in \mathbb{N}_0. \quad (6.2)$$

The corresponding eigenspaces Harm_n have dimension $2n+1$ and are mutually orthogonal in $L^2(\mathcal{S})$. Their direct sum equals $L^2(\mathcal{S})$. Every eigenfunction $Y_n \in \text{Harm}_n$ lies in $C^\infty(\mathcal{S})$ and is called (*scalar*) *spherical harmonic of degree n* . Their name derives from the equivalent characterisation of Harm_n as the restriction to \mathcal{S} of the space of harmonic polynomials $P : \mathbb{R}^3 \rightarrow \mathbb{R}$ that are homogeneous of degree n . From now on

$$\{Y_{nj} : n \in \mathbb{N}_0, 1 \leq j \leq 2n+1\} \quad (6.3)$$

always refers to a particular orthonormal basis of $L^2(\mathcal{S})$ consisting of real-valued spherical harmonics. For numerical experiments we use the so-called fully normalised spherical harmonics, see [63, Sec. 5.2] for a detailed construction.

Vector Spherical Harmonics

Denote by ν the outward unit normal to \mathcal{S} and let

$$\nabla_{\mathcal{S}} F = \nabla \bar{F}$$

be the surface gradient of $F : \mathcal{S} \rightarrow \mathbb{R}$.

Definition 1. Let $n \in \mathbb{N}_0$ and $Y_n \in \text{Harm}_n$. Whenever a function $y : \mathcal{S} \rightarrow \mathbb{R}^3$ that does not vanish identically admits one of the following three representations

$$y = \begin{cases} y_n^{(1)} := Y_n \nu, \\ y_n^{(2)} := \nabla_{\mathcal{S}} Y_n, \\ y_n^{(3)} := \nabla_{\mathcal{S}} Y_n \times \nu, \end{cases} \quad (6.4)$$

then $y = y_n^{(i)}$ is called a vector spherical harmonic of degree n and type i . For obvious reasons we refer to types 2 and 3 as tangential vector spherical harmonics.

Note that there is no tangential spherical harmonic of degree 0.

We are mainly interested in the space $L^2(\mathcal{S}, T\mathcal{S})$ of square integrable tangent vector fields on \mathcal{S} endowed with the inner product

$$\langle u, v \rangle = \int_{\mathcal{S}} u \cdot v \, d\mathcal{S},$$

where $d\mathcal{S}$ is the usual surface measure on the sphere. An orthonormal basis of this space is obtained from (6.3) by setting

$$\begin{aligned} y_{nj}^{(2)} &= \lambda_n^{-1/2} \nabla_{\mathcal{S}} Y_{nj}, \\ y_{nj}^{(3)} &= \lambda_n^{-1/2} \nabla_{\mathcal{S}} Y_{nj} \times \nu \end{aligned} \tag{6.5}$$

for all $n \in \mathbb{N}$ and $1 \leq j \leq 2n+1$. In Fig. 6.2 a handful of the elements of both bases (6.3) and (6.5) are depicted. Every $v \in L^2(\mathcal{S}, T\mathcal{S})$ has the following Fourier series representation

$$\sum_{i=2}^3 \sum_{n=1}^{\infty} \sum_{j=1}^{2n+1} \langle v, y_{nj}^{(i)} \rangle y_{nj}^{(i)}.$$

In particular, we have Parseval's identity

$$\|v\|_{L^2(\mathcal{S}, T\mathcal{S})}^2 = \sum_{i,n,j} \langle v, y_{nj}^{(i)} \rangle^2.$$

For a comprehensive and unified treatment of both scalar and vector spherical harmonics we refer to [28].

Sobolev Spaces on the Sphere

For an arbitrary real number s , the space $H^s(\mathcal{S})$ is commonly defined as the domain of $\Delta_{\mathcal{S}}^{s/2}$. See [59, p. 37] or [63, Sec. 6.2] for example. In this section we introduce the spaces $H^s(\mathcal{S}, T\mathcal{S})$ by means of the vectorial counterpart of $\Delta_{\mathcal{S}}$.

For tangent vector fields v we define the Laplace-Beltrami operator by

$$\Delta_{\mathcal{S}} v = P \Delta_{\mathcal{S}} v,$$

where application of $\Delta_{\mathcal{S}}$ to v is understood componentwise and $P = P(x)$ is the orthogonal projector onto the tangent plane $T_x \mathcal{S}$, compare [28, Def. 5.26]. The tangential vector spherical harmonics introduced in Def. 1 are eigenfunctions of this operator to the same eigenvalues as their scalar counterparts: If we let

$$\text{harm}_n = \text{span} \left\{ y_{nj}^{(i)} : 1 \leq j \leq 2n+1, i = 2, 3 \right\},$$

then

$$\Delta_{\mathcal{S}} y_n = \lambda_n y_n$$




Figure 6.2: Fully normalised scalar and corresponding vector spherical harmonics of degree $n = 2$. Scalar spherical harmonics are depicted using a blue to red colour bar. Type 2 vector spherical harmonics are visualised with red arrows, type 3 ones with blue arrows. Note that the length of the vectors has been scaled for better illustration.

for every $y_n \in \text{harm}_n$. The λ_n are as defined in (6.2), the only difference being that now $n > 0$ and therefore the spectrum of $\Delta_{\mathcal{S}}$ is strictly positive. Applying functional calculus, we formally define the s -th power of $\Delta_{\mathcal{S}}$ by

$$\Delta_{\mathcal{S}}^s v = \sum_{i,n,j} \lambda_n^s \langle v, y_{nj}^{(i)} \rangle y_{nj}^{(i)}.$$

Finally, for every $s \in \mathbb{R}$, set

$$\|v\|_{H^s(\mathcal{S}, T\mathcal{S})}^2 := \|\Delta_{\mathcal{S}}^{s/2} v\|_{L^2(\mathcal{S}, T\mathcal{S})}^2 = \sum_{i,n,j} \lambda_n^s \langle v, y_{nj}^{(i)} \rangle^2. \quad (6.6)$$

Note that, in contrast to the scalar setting, this functional is an actual norm, rather than only a seminorm. We therefore define, for every real s , $H^s(\mathcal{S}, T\mathcal{S})$ as the space of all distributions $v \in C^\infty(\mathcal{S}, T\mathcal{S})'$ for which the series in (6.6) is finite. Clearly, if $(\mu_n)_n$ is any sequence satisfying

$$c\mu_n \leq \lambda_n^s \leq C\mu_n \quad (6.7)$$

for two positive constants c, C and for all n , then replacing λ_n^s with μ_n in (6.6) leads to an equivalent norm and thus to the same space. For every sequence of positive numbers μ_n we denote the resulting norm simply by $\|\cdot\|_{\mu_n}$.

6.3 Decomposition Models for Optical Flow

Optical Flow on the Sphere

Let $I = [0, T] \subset \mathbb{R}$ be a time interval. We assume to be given a scalar time-varying (brightness) function

$$F : I \times \mathcal{S} \rightarrow \mathbb{R}.$$

The problem of estimating optical flow consists in tracking the temporal evolution of the data F by means of a time-dependent vector field. Our optical flow model is based on the so-called *brightness constancy assumption*: We assume existence of a function $\phi : I \times \mathcal{S} \rightarrow \mathcal{S}$ satisfying

$$\begin{aligned} F(t, \phi(t, x)) &= F(0, x), \\ \phi(0, x) &= x, \end{aligned} \quad (6.8)$$

for all x and t . Intuitively this means that for every starting point x on the sphere, the function F remains constant along the trajectory $t \mapsto \phi(t, x)$. In addition we require that $\phi(t, \cdot)$ is a diffeomorphism of \mathcal{S} for every t . The first equation in (6.8) implies that

$$\frac{d}{dt} F(t, \phi(t, x)) = \nabla_{\mathcal{S}} F(t, \phi(t, x)) \cdot \partial_t \phi(t, x) + \partial_t F(t, \phi(t, x)) = 0.$$

This equation is typically written in terms of the vector field $u : I \times \mathcal{S} \rightarrow T\mathcal{S}$ whose integral curves are the trajectories $t \mapsto \phi(t, x)$, which is defined by the equation $u(t, \phi(t, x)) = \partial_t \phi(t, x)$. The resulting *optical flow equation* reads

$$\nabla_{\mathcal{S}} F \cdot u + \partial_t F = 0. \quad (6.9)$$

Regularisation

Solving the optical flow equation directly is problematic. In general, a solution to (6.9) need not exist, and if it exists, it cannot be unique. The typical remedy is Tikhonov regularisation, where one minimises a functional of the form

$$\|\nabla_{\mathcal{S}} F \cdot u + \partial_t F\|_{L^2(I \times \mathcal{S})}^2 + \alpha \mathcal{R}(u). \quad (6.10)$$

with \mathcal{R} being a regularising functional that incorporates a-priori knowledge about desirable solutions. The parameter $\alpha > 0$ controls the amount of regularisation. In the context of optical flow one usually tries to enforce spatial (and temporal) smoothness on the solution. A natural candidate for \mathcal{R} would then be the squared Sobolev $H^1(I \times \mathcal{S}, T\mathcal{S})$ (semi-)norm, which penalises first derivatives in space and time equally, compare [52, 86]

Another, and in fact more popular, possibility is to drop time regularisation, in which case minimisation of (6.10) is equivalent to minimising

$$\|\nabla_{\mathcal{S}} F \cdot u + \partial_t F\|_{L^2(\mathcal{S})}^2 + \alpha \|u\|_{H^1(\mathcal{S}, T\mathcal{S})}^2 \quad (6.11)$$

for each instant t separately. This corresponds to the original approach of Horn and Schunck [39]. From now on we denote the above data term by $\mathcal{D}(u, F)$. Instead of (6.11) we consider the more general class of optical flow functionals

$$\mathcal{E}_{\mu_n}(u) := \mathcal{D}(u, F) + \|u\|_{\mu_n}^2. \quad (6.12)$$

The regularisation parameter is omitted, as it can be considered a constant factor in the sequence $(\mu_n)_n$. Functional (6.12) forms the basic optical flow setting of this article. All variational models considered here are extensions of (6.12).

Optical Flow Decomposition

The following two decomposition models are inspired by techniques that have recently been developed in the context of image analysis. The fact that motion estimation based on (6.9) can be viewed as denoising of vector-valued images suggests the translation of said image decomposition models to the optical flow setting [1].

$u + v$ Models

The aim is now not to extract one, but two vector fields u and v in such a way that they capture different structural parts of the total motion $u + v$ of F . The idea is to solve the following variational problem

$$\min_{u, v} \mathcal{E}_{\mu_n, \nu_n}(u, v)$$

where the functional $\mathcal{E}_{\mu_n, \nu_n}$ is defined as

$$\mathcal{E}_{\mu_n, \nu_n}(u, v) := \mathcal{D}(u + v, F) + \|u\|_{\mu_n}^2 + \|v\|_{\nu_n}^2. \quad (6.13)$$

Choosing, for instance, the two regularisers to be $H^1(\mathcal{S}, T\mathcal{S})$ and $H^{-1}(\mathcal{S}, T\mathcal{S})$ norms, respectively, would lead to a model which, in spirit, comes closest to the image decomposition model considered in [82]. Generalising (6.13) to a decomposition into $k \in \mathbb{N}$, instead of two, constituents is possible, but will not be considered here [29].

Hierarchical Models

Hierarchical image decomposition models have been introduced in [77]. There, an original image is decomposed by repeatedly applying denoising steps. The input of one such step is the residual of the previous one. In every step the degree of regularisation is decreased. In contrast to $u + v$ models, hierarchical decomposition models provide multiscale descriptions of the input data.

This iterative procedure can be transferred to the optical flow setting as follows. Let $\|\cdot\|_{\mu_n^{(k)}}^2$, $k \in \mathbb{N}$, be a sequence of norms as defined in Section 6.2, such that

$$\mu_n^{(k+1)} \leq \mu_n^{(k)} \quad (6.14)$$

for all n and k . For every such sequence of sequences we propose the following iterative scheme,

$$u_k = \begin{cases} \arg \min_u \mathcal{E}_{\mu_n^{(1)}}(u), & \text{if } k = 1, \\ \arg \min_u \mathcal{D}(u + \sum_{i=1}^{k-1} u_i, F) + \|u\|_{\mu_n^{(k)}}^2, & \text{if } k > 1. \end{cases} \quad (6.15)$$

The resulting sequence of accumulated solutions

$$\left\{ u^{(k)} := \sum_{i=1}^k u_i : k \in \mathbb{N} \right\}$$

provides a multiscale representation which, with an appropriate choice of sequences $\mu_n^{(k)}$, can be made arbitrarily fine.

The hierarchical model as formulated above is a slight generalisation of the originally proposed one, in the sense that the sequences of regularising functionals considered in [77] are always of the form $\alpha^{-k} \mathcal{R}(\cdot)$. That is, the regularising functional of step k is the same as the one of previous steps save for a smaller regularisation parameter. Requirement (6.14) allows a more general setup.

Helmholtz Decomposition

The Helmholtz decomposition theorem states that every continuously differentiable tangent vector field w on the sphere can be uniquely represented as the sum of a consoidal (curl-free) and a toroidal (divergence-free) vector field. More precisely, there exist two uniquely determined tangent vector fields $w^{(2)}$ and $w^{(3)}$ satisfying

$$\begin{aligned} \nabla_{\mathcal{S}} \cdot (w^{(2)} \times \nu) &= 0, \\ \nabla_{\mathcal{S}} \cdot w^{(3)} &= 0, \\ w^{(2)} + w^{(3)} &= w. \end{aligned}$$

See [28, Sec. 5.3], for example.

The projection method we use to numerically solve the variational problems presented above leads to solutions w that are finite Fourier sums

$$w = \sum_{i,n,j} w_{nj}^i y_{nj}^{(i)},$$

where $w_{nj}^i \in \mathbb{R}$. Now, from the definition of basis (6.5), and the fact that $\nabla_S \cdot (\nabla_S Y \times \nu) = 0$ for all sufficiently smooth functions Y , the Helmholtz decomposition of w is obtained immediately

$$\begin{aligned} w &= \sum_{n,j} w_{nj}^2 y_{nj}^{(2)} + \sum_{n,j} w_{nj}^3 y_{nj}^{(3)} \\ &= \underbrace{\nabla_S \left(\sum_{n,j} w_{nj}^2 \lambda_n^{-1/2} Y_{nj} \right)}_{w^{(2)}} + \underbrace{\nabla_S \left(\sum_{n,j} w_{nj}^3 \lambda_n^{-1/2} Y_{nj} \right)}_{w^{(3)}} \times \nu. \end{aligned}$$

6.4 Numerical Solution

In the first subsection below we describe the numerical optimisation of the optical flow functional (6.12) based on the optical flow equation (6.9) and explain the modifications necessary for the hierarchical decomposition. Subsection 6.4 is dedicated to the $u + v$ decomposition model. Finally, we explain how the resulting spherical integrals are approximated (Sec. 6.4).

For convenience we relabel the orthonormal basis (6.5) using a single index $p \in \mathbb{N}$ and write, for example, $u = \sum_p u_p y_p$ from now on.

Optical Flow

Let $s \in \mathbb{R}$ and let $(\mu_n)_n$ be a sequence comparable to $(\lambda_n^s)_n$ in the sense of (6.7). Then, a minimiser of \mathcal{E}_{μ_n} , if it exists, has to be in $H^s(\mathcal{S}, T\mathcal{S})$. We solve the problem of finding

$$\min_{u \in H^s(\mathcal{S}, T\mathcal{S})} \mathcal{E}_{\mu_n}(u)$$

by a projection method. That is, we let u range only over a finite-dimensional subspace \mathcal{U} of $H^s(\mathcal{S}, T\mathcal{S})$, where

$$\mathcal{U} = \text{span}\{y_p : p \in I_{\mathcal{U}}\}$$

and $I_{\mathcal{U}} \subset \mathbb{N}$ is a finite index set. The unknown vector field can now be written as

$$u = \sum_{p \in I_{\mathcal{U}}} u_p y_p \quad (6.16)$$

and the problem of finding an optimal $u \in H^s(\mathcal{S}, T\mathcal{S})$ simplifies to a minimisation problem over $\mathbb{R}^{|I_{\mathcal{U}}|}$. Plugging (6.16) into the optical flow functional gives

$$\mathcal{E}_{\mu_n}(u) = \int_{\mathcal{S}} \left(\sum_{p \in I_{\mathcal{U}}} u_p (\nabla_S F \cdot y_p) + \partial_t F \right)^2 d\mathcal{S} + \sum_{p \in I_{\mathcal{U}}} \mu_p u_p^2, \quad (6.17)$$

which is minimal, if the optimality conditions $\partial \mathcal{E}_{\mu_n} / \partial u_p = 0$, for all $p \in I_{\mathcal{U}}$, are satisfied. They read

$$\sum_{q \in I_{\mathcal{U}}} u_q \int_{\mathcal{S}} (\nabla_S F \cdot y_p)(\nabla_S F \cdot y_q) d\mathcal{S} + \mu_p u_p = - \int_{\mathcal{S}} \partial_t F \nabla_S F \cdot y_p d\mathcal{S}, \quad p \in I_{\mathcal{U}},$$

or in matrix-vector form

$$(A + D)w = b, \quad (6.18)$$

where $w = (u_1, \dots, u_{|I_{\mathcal{U}}|})^\top$ is the vector of unknown coefficients, the elements of matrix $A = (a_{pq})_{pq}$ read

$$a_{pq} = \int_{\mathcal{S}} (\nabla_{\mathcal{S}} F \cdot y_p)(\nabla_{\mathcal{S}} F \cdot y_q) d\mathcal{S}, \quad (6.19)$$

$D = \text{diag}(\mu_1, \dots, \mu_{|I_{\mathcal{U}}|})$ is a diagonal matrix that corresponds to the regularisation term and the right hand side $b = (b_1, \dots, b_{|I_{\mathcal{U}}|})^\top$ is given by

$$b_p = - \int_{\mathcal{S}} \partial_t F \nabla_{\mathcal{S}} F \cdot y_p d\mathcal{S}. \quad (6.20)$$

With a slight abuse of notation we identified the set $I_{\mathcal{U}}$ with $\{1, \dots, |I_{\mathcal{U}}|\}$ in the definitions of A, D, w, b . We continue to do so below.

Hierarchical Decomposition

The hierarchical model only needs a minor modification for the case $k > 1$. We can rewrite the data term from (6.15) as

$$\mathcal{D}(u + u^{(k-1)}, F) = \|\nabla_{\mathcal{S}} F \cdot u + \partial_t \tilde{F}\|_{L^2(\mathcal{S})}^2,$$

where $\partial_t \tilde{F} = \partial_t F + \nabla_{\mathcal{S}} F \cdot u^{(k-1)}$. Therefore, only the right hand side of the optimality system (6.18) has to be updated in every step. For simplicity we can assume that the approximation space \mathcal{U} is the same in every step, so that $u^{(k-1)}$ has the representation $\sum_{p \in I_{\mathcal{U}}} c_p^{k-1} y_p$, where the coefficients c_p^{k-1} are already known from previous steps. Letting b^k denote the right hand side of the optimality system for step k , we calculate

$$\begin{aligned} b_p^k &= - \int_{\mathcal{S}} \partial_t \tilde{F} \nabla_{\mathcal{S}} F \cdot y_p d\mathcal{S} \\ &= b_p - \sum_{q \in I_{\mathcal{U}}} c_q^{k-1} a_{pq}, \end{aligned}$$

or simply

$$b^k = b - A c^{k-1}.$$

$u + v$ Decomposition

The projection approach explained above is easily adapted to the $u + v$ decomposition problem. Again, let $r \neq s$ be real numbers and choose two sequences $(\mu_n)_n, (\nu_n)_n$ so that $\|\cdot\|_{\mu_n}$ is a norm for $H^r(\mathcal{S}, T\mathcal{S})$ and $\|\cdot\|_{\nu_n}$ is one for $H^s(\mathcal{S}, T\mathcal{S})$. Now, we solve

$$\min_{(u,v) \in \mathcal{U} \times \mathcal{V}} \mathcal{E}_{\mu_n, \nu_n}(u, v),$$

where

$$\begin{aligned} \mathcal{U} &= \text{span}\{y_p : p \in I_{\mathcal{U}}\}, \\ \mathcal{V} &= \text{span}\{y_p : p \in I_{\mathcal{V}}\} \end{aligned}$$

are finite dimensional spaces. Proceeding as in the previous section, we obtain the following optimality conditions

$$\begin{aligned} \sum_{p \in I_{\mathcal{U}}} u_p a_{kp} + \sum_{q \in I_{\mathcal{V}}} v_q a_{kq} + \mu_k u_k &= -b_k, \quad k \in I_{\mathcal{U}}, \\ \sum_{p \in I_{\mathcal{U}}} u_p a_{\ell p} + \sum_{q \in I_{\mathcal{V}}} v_q a_{\ell q} + \nu_{\ell} v_{\ell} &= -b_{\ell}, \quad \ell \in I_{\mathcal{V}}, \end{aligned}$$

where the coefficients a_{pq} and b_p are as defined in (6.19) and (6.20), respectively. Concatenating the two coefficient vectors $(u_p)_p$ and $(v_q)_q$ into a single vector $w \in \mathbb{R}^{|I_{\mathcal{U}}|+|I_{\mathcal{V}}|}$, so that the u_p occupy the first $|I_{\mathcal{U}}|$ entries while the v_q occupy the last $|I_{\mathcal{V}}|$ entries, the linear system reads

$$\tilde{A}w = \tilde{b}.$$

The matrix \tilde{A} is given by

$$\tilde{A} = \begin{pmatrix} U + D_1 & W \\ W^{\top} & V + D_2 \end{pmatrix},$$

where

$$\begin{aligned} U &= (a_{pq})_{p,q \in I_{\mathcal{U}}}, \\ V &= (a_{pq})_{p,q \in I_{\mathcal{V}}}, \\ W &= (a_{pq})_{p \in I_{\mathcal{U}}, q \in I_{\mathcal{V}}}, \\ D_1 &= \text{diag}(\mu_1, \dots, \mu_{|I_{\mathcal{U}}|}), \\ D_2 &= \text{diag}(\nu_1, \dots, \nu_{|I_{\mathcal{V}}|}), \end{aligned}$$

and \tilde{b} is concatenated from two versions of b in the same way as w .

Evaluation of Integrals

It remains to discuss the numerical evaluation of the integrals (6.19), (6.20). First, we approximate the 2-sphere \mathcal{S} with a polyhedron $\hat{\mathcal{S}} = (\mathcal{V}, \mathcal{T})$ defined by a set $\mathcal{V} = \{v_1, \dots, v_m\} \subset \mathcal{S}$ of vertices and a set $\mathcal{T} = \{T_1, \dots, T_n\} \subset \mathcal{V} \times \mathcal{V} \times \mathcal{V}$ of triangular faces. Each triangle $T_i \in \mathcal{T}$ is associated with a tuple (i_1, i_2, i_3) identifying the corresponding vertices $(v_{i_1}, v_{i_2}, v_{i_3})$. How the triangulated sphere is obtained in practice, is explained in Sec. 6.5.

Second, in every experiment data F are given only at the vertices and for two time steps $t = 0$ and $t = 1$. We set $F_0(\cdot) := F(0, \cdot)$ and $F_1(\cdot) := F(1, \cdot)$ and extend both functions to all of $\hat{\mathcal{S}}$ by linear interpolation on every triangle. Thus we obtain two continuous piecewise linear functions \hat{F}_0, \hat{F}_1 . The time derivative of F is approximated by a simple forward difference

$$\partial_t \hat{F} = \hat{F}_1 - \hat{F}_0,$$

which is again piecewise linear on $\hat{\mathcal{S}}$. The surface gradient $\nabla_{\mathcal{S}} F$ is replaced by a vector field $\nabla_{\hat{\mathcal{S}}} \hat{F}$ that is constant on every triangle. It is given by

$$\nabla_{\hat{\mathcal{S}}} \hat{F}|_{T_i} = \left(\hat{F}(v_{i_1}) - \hat{F}(v_{i_2}) \right) \frac{h_{i_2}}{|h_{i_2}|^2} + \left(\hat{F}(v_{i_1}) - \hat{F}(v_{i_3}) \right) \frac{h_{i_3}}{|h_{i_3}|^2},$$




Figure 6.3: Top view of frames no. 57 (left) and 58 (right) of the processed zebrafish images. The embryo's body axis runs from bottom left to top right.

where $h_{i_j} \in \mathbb{R}^3$ is the height vector of the triangle T_i pointing from vertex v_{i_j} to the opposite side, compare [13, Sec. 3.3.3].

Finally, we also replace the fully normalised scalar spherical harmonics Y_{nj} by their piecewise linear approximations \hat{Y}_{nj} defined on $\hat{\mathcal{S}}$. As before we obtain piecewise constant approximations \hat{y}_p of y_p . The resulting approximated integrals read

$$\begin{aligned} \int_{\mathcal{S}} (\nabla_{\mathcal{S}} F \cdot y_p) (\nabla_{\mathcal{S}} F \cdot y_q) d\mathcal{S} &\approx \int_{\hat{\mathcal{S}}} (\nabla_{\hat{\mathcal{S}}} \hat{F} \cdot \hat{y}_p) (\nabla_{\hat{\mathcal{S}}} \hat{F} \cdot \hat{y}_q) d\hat{\mathcal{S}} \\ &= \sum_{i=1}^n (\nabla_{\hat{\mathcal{S}}} \hat{F}|_{T_i} \cdot \hat{y}_p|_{T_i}) (\nabla_{\hat{\mathcal{S}}} \hat{F}|_{T_i} \cdot \hat{y}_q|_{T_i}) A_i, \end{aligned}$$

where A_i denotes the area of T_i , and

$$\begin{aligned} \int_{\mathcal{S}} \partial_t F \nabla_{\mathcal{S}} F \cdot y_p d\mathcal{S} &\approx \int_{\hat{\mathcal{S}}} \partial_t \hat{F} \nabla_{\hat{\mathcal{S}}} \hat{F} \cdot \hat{y}_p d\hat{\mathcal{S}} \\ &= \sum_{i=1}^n (\nabla_{\hat{\mathcal{S}}} \hat{F}|_{T_i} \cdot \hat{y}_p|_{T_i}) \int_{T_i} \partial_t \hat{F} dT_i \\ &= \sum_{i=1}^n (\nabla_{\hat{\mathcal{S}}} \hat{F}|_{T_i} \cdot \hat{y}_p|_{T_i}) \frac{A_i}{3} \sum_{j=1}^3 \partial_t \hat{F}(v_{i_j}). \end{aligned}$$

6.5 Experiments

Description of Microscopy Data

The data which motivated the study of the proposed decomposition models are time-lapse volumetric (4D) images. The obtained sequence depicts a live zebrafish embryo during the gastrula period, approximately five to ten hours after fertilisation. With the help of confocal laser-scanning microscopy, endoderm cells expressing a green fluorescence protein were recorded separately from the

background. For details on the imaging techniques and the fluorescence marker we refer to [60] and [66], respectively.

The sequence obtained by the microscope captures a cuboid region of approximately $860 \times 860 \times 340 \mu\text{m}^3$. The spatial resolution is $512 \times 512 \times 44$ voxels and the intensity is in the range $\{0, \dots, 255\}$. A total number of 77 images were taken, one every 240 s. In the following, the microscopic data will be denoted by

$$F^\delta \in \{0, \dots, 255\}^{77 \times 512 \times 512 \times 44}.$$

During this early stage, endodermal cell proliferation is known to take place on a so-called monolayer [83]. In other words, cells move and divide without stacking and, as the yolk is ball-shaped, admit for the extraction of a spherical image sequence. For further explanations and numerous illustrations of the developmental process of zebrafish embryos we refer to [48].

Acquisition of Spherical Data

We extracted spherical images from the dataset by first fitting a sphere to the approximate cell centres in each pair of consecutive frames. For simplicity we restrict our attention to one such pair of frames which we denote by F_0^δ, F_1^δ . Cell centres are typically characterised by local maxima in intensity and can be found by applying a Gaussian filter and simple thresholding. Without loss of generality the radius of the fitted sphere is assumed to be 1. In a second step, we created a point grid $\mathcal{V} \subset \mathcal{S}$ starting from an icosahedron inscribed in the sphere. In each iteration every triangular face is split into four sub-triangles by connecting the edge midpoints with each other and projecting them onto the sphere. Thus, the total number of faces is $20 \cdot 4^k$, where k is the number of refinements. In our experiments we found that $k = 7$ iterations suffice.

In order to project the volumetric time-lapse data F_j^δ , onto the grid \mathcal{V} , we define

$$\hat{F}_j(v_i) := \max_{c \in [1-\epsilon, 1+\epsilon]} \bar{F}_j^\delta(cv_i),$$

for the said pair of consecutive frames $j = 0, 1$. Here $\bar{F}_j^\delta(x)$ is a piecewise linear extension to \mathbb{R}^3 of F_j^δ and $\epsilon > 0$ is sufficiently large. Deviations of the monolayer from a perfect sphere are thereby corrected. The obtained data are subsequently scaled to the range $[0, 1]$. Note that, in contrast to our previous work [50], here we consider the unfiltered microscopy data for optical flow estimation.

The support of the obtained data is contained in the northern hemisphere. Thus, it suffices to consider only half of the triangulation leading to a total number of around 164000 faces. Figure 6.1 depicts a sample of two frames \hat{F}_j . In Fig. 6.3, a top view of the two frames is shown.

Visualisation of Tangent Vector Fields

In order to visualise our results we will apply the standard flow colour-coding [9] using a colour disk. Figure 6.4 (rightmost image) depicts this colour space. Each vector is assigned a colour determined by its angle and length. However, this colour-coding is defined for planar vector fields only. As a possible remedy we suggest to first project tangential velocities to the plane and then correct




Figure 6.4: The first image shows a tangential velocity field with the adjusted colour-coding. The second image depicts the colour space in the unit circle. The third and fourth images portray the vector field and the colour space, respectively, but in a top view.

Figure	6.5 left	6.5 mid	6.5 right	6.7 left	6.7 mid	6.7 right	6.8	6.9	6.11	6.12
R	0.0081	0.0046	0.0014	0.0110	0.0045	0.0009	0.0204	0.0102	0.0280	0.0267

Table 6.1: Radii R of the colour disks used in the different experiments below.

the length. To this end, let $P_{x_3} : (x_1, x_2, x_3)^\top \mapsto (x_1, x_2, 0)^\top$ be the orthogonal projector of \mathbb{R}^3 onto the x_1 - x_2 -plane. Accordingly, given a tangent vector field v , the planar vector field which we visualise is

$$\frac{|v|}{|P_{x_3} v|} P_{x_3} v.$$

This construction is chosen so that it preserves the length of v . This additional rescaling is different to [50]. The resulting colour image is finally mapped back onto the hemisphere. Figure 6.4 shows a tangent vector field visualised with the proposed approach.¹ From now on we will visualise velocity fields only in top view, as in the right hand side of Fig. 6.4. In addition, for every figure the colour disk's radius R was chosen to be equal to the length of the longest vector under consideration. Specific values of R are given in Table 6.1.

As a second way of illustrating steady velocity fields, we employ streamlines, see e.g. [87]. In all our experiments we consider time as fixed and compute the optical flow v for one pair of frames, cf. Sec. 6.3. Given a tangential velocity field v and a starting point $x_0 \in \mathcal{S}$, a streamline $\gamma(\cdot, x_0)$ on \mathcal{S} is the solution to the ordinary differential equation

$$\begin{aligned} \partial_\tau \gamma(\tau, x_0) &= v(\gamma(\tau, x_0)), \\ \gamma(0, x_0) &= x_0. \end{aligned} \tag{6.21}$$

Numerically, we approximated (6.21) by solving

$$\begin{aligned} \hat{\gamma}(\tau + 1, x_0) &= \hat{\gamma}(\tau, x_0) + h v(\hat{\gamma}(\tau, x_0)), \\ \hat{\gamma}(0, x_0) &= x_0, \end{aligned}$$

where h is a step size, for a number of approximately 1300 initial points $x_0 \in \hat{\mathcal{S}}$ and $\tau = 50$ iterations. The step size was chosen as $h = (10\|v\|_{L^\infty(\mathcal{S}, T\mathcal{S})})^{-1}$.

¹Some figures may appear in colour only in the online version of this article.

This use of integral curves is different from [52], where we computed approximate cell trajectories in a nonsteady velocity field. The visualisation by means of the colour coding is rich in detail and is even capable of indicating individual cell motion. Nevertheless, it fails to deliver intuition about the Helmholtz decomposition. Streamlines provide the anticipated effect.

Experimental Results

We performed numerous experiments and minimised functionals (6.12), (6.13), and (6.15) as outlined in Sec. 6.4 for the two frames shown in Fig. 6.1 and Fig. 6.3. In all experiments the finite-dimensional spaces introduced in Sec. 6.4 were chosen as

$$\mathcal{U} = \mathcal{V} = \text{span}\left\{y_{nj}^{(i)} : 1 \leq n \leq 100, 1 \leq j \leq 2n + 1, i = 2, 3\right\}.$$

All resulting linear systems were solved using the Generalized Minimal Residual Method (GMRES) on an Intel Xeon E5-1620 3.6 GHz workstation with 128 GB RAM. Solutions converged to a relative residual of 0.02 within 100 iterations. The runtime was governed by the evaluation of the integrals, cf. Sec. 6.4, and amounts to approximately five hours for the chosen bases and the chosen triangulation. Nevertheless, once the integrals are computed they can be used in all of the proposed models and the linear systems can be solved in a few seconds for different parameters and different norms. Our Matlab implementation and the data are available on our website.²

Optical Flow

In the first experiment, we minimised functional \mathcal{E}_{μ_n} as defined in (6.12) for $\mu_n = \alpha \lambda_n^s$ and different values of s and α . Figure 6.5, top row, depicts the optical flow field for $s = 1$ and values $\alpha = 1$, $\alpha = 10$, and $\alpha = 100$. The presented results are in accordance with our findings in [50, 52]. As explained in Sec. 6.3, a Helmholtz decomposition $w = w^{(2)} + w^{(3)}$ is obtained immediately. Figure 6.5, middle row, shows $w^{(2)}$ whereas Fig. 6.5, bottom row, shows $w^{(3)}$. Furthermore, in Fig. 6.6, streamlines for the same velocity fields are portrayed and the individual plots are arranged accordingly. In addition, Fig. 6.7 shows the velocity fields for parameters $s = 0.5$, $\alpha = 10$, $\alpha = 10^2$, and $\alpha = 10^3$.

$u + v$ Decomposition

In a next experiment, we computed a minimiser for functional (6.13) in order to obtain a $u + v$ decomposition of the optical flow. The sequences $(\mu_n)_n$ and $(\nu_n)_n$ were chosen as $\mu_n = \alpha \lambda_n^r$ and $\nu_n = \beta \lambda_n^s$, respectively. In Figs. 6.8 and 6.9, the resulting decomposition is shown for two different parameter settings. The motion field in Fig. 6.8 was obtained with parameters $r = 1$, $s = -1$, $\alpha = 10^{-1}$, and $\beta = 10^6$. As anticipated, u and v capture different structural parts of the motion. u is sufficiently smooth whereas v contains spatial oscillations.

The result in Fig. 6.9 was computed by setting $r = 2$, $s = -1$, $\alpha = 10^{-3}$, and $\beta = 10^7$. Expectedly, the velocity field u is, by choice of r , smoother than in the previous setting. In addition, Fig. 6.10 illustrates the characteristics of

²<http://www.csc.univie.ac.at>




Figure 6.5: Minimiser of \mathcal{E}_{μ_n} (top), with $\mu_n = \alpha \lambda_n^s$, $s = 1$, and increasing values $\alpha = 1$, $\alpha = 10$, and $\alpha = 100$ from left to right. The middle row depicts the curl-free component whereas the bottom row depicts the divergence-free component. The embryo's body axis roughly runs from bottom left to top right in all images.

the velocity fields during a cell division in more detail. While u is smooth, v clearly indicates the cell division.

Hierarchical Decomposition

As a final experiment, we computed two types of hierarchical decompositions as proposed in Sec. 6.3. First we chose $\mu_n^{(k)} = 2^{1-k} \alpha \lambda_n^s$ such that α is halved in every iteration. Solutions $(u^{(k)})_{k=1,\dots,16}$ were obtained using parameters $s = 1$ and $\alpha = 1000$. In Fig. 6.11, the subsequence $(u^{(k)})_{k=8,\dots,16}$ is shown. As k increases, the motion field expands on the details. In a second run, $\mu_n^{(k)}$ was set to $\alpha \lambda_n^{s - \frac{k-1}{4}}$, decreasing the exponent of λ_n by 0.25 in every step. We iteratively computed solutions $(u^{(k)})_{k=1,\dots,9}$ with parameters $s = 2$ and $\alpha = 1$, which were kept constant this time. Figure 6.12 depicts the subsequence $(u^{(k)})_{k=2,\dots,7}$. Note that in Figs. 6.11 and 6.12 the colour-coded visualisation is relative to




Figure 6.6: Streamlines illustrating the velocity fields from Fig. 6.5. With increasing parameter τ the streamlines change colours from yellow (bright) to green (dark). The top row shows the total motion whereas middle and bottom rows depict the curl-free and divergence-free parts of the Helmholtz decomposition, respectively. Images are arranged in accordance with Fig. 6.5.

the chosen subsequence. As an exception, here we allowed a maximum number of 1000 iterations instead of 100 for the linear system solve to ensure a relative residual of 0.025 in the first step of the hierarchical decomposition.

6.6 Conclusion

We provided a set of variational methods for the analysis of motion fields. While their applicability is limited to data given on the sphere, the proposed models have great flexibility in terms of possible regularising functionals. In fact, the chosen numerical method based on tangential vector spherical harmonics allows for a straightforward usage of Sobolev H^s norms for any real s . Combined with both $u + v$ and hierarchical decomposition models, which we adapted to the spherical optical flow setting, this flexibility makes it possible to capture different motion characteristics with ease. Feasibility of the proposed models




Figure 6.7: Colour-coded velocity fields for sequence $\mu_n = \alpha\lambda_n^s$ are shown in the top row. Parameters were chosen as $s = 0.5$ and $\alpha = 10$, $\alpha = 10^2$, and $\alpha = 10^3$ and are arranged from left to right. The bottom row depicts the corresponding streamlines.




Figure 6.8: Decomposition of the total motion $u + v$ (left) into structural parts u (middle) and v (right). Sequences were chosen as $\mu_n = \alpha\lambda_n^r$ and $\nu_n = \beta\lambda_n^s$ and parameters were set to $r = 1$, $s = -1$, $\alpha = 10^{-1}$, and $\beta = 10^6$.




Figure 6.9: Decomposition of the total motion $u + v$ (left) into structural parts u (middle) and v (right). Sequences were chosen as $\mu_n = \alpha \lambda_n^r$ and $\nu_n = \beta \lambda_n^s$ and parameters were set to $r = 2$, $s = -1$, $\alpha = 10^{-3}$, and $\beta = 10^7$.




Figure 6.10: Detailed view of a cell division. The same parameters as in Fig. 6.9 are used. The top left image depicts \hat{F}_0 with $u + v$ superimposed. The top right image shows \hat{F}_1 . The bottom left image illustrates u whereas the bottom right image shows v . For better illustration, \hat{F}_0 and \hat{F}_1 have been lightened up and vectors have been scaled.




Figure 6.11: Solutions $(u^{(k)})_{k=8,\dots,16}$ of velocity fields obtained by the hierarchical decomposition. Images are aligned from left to right and top to bottom. At iteration k the sequence was set to $\mu_n^{(k)} = 2^{1-k} \alpha \lambda_n^s$ with parameters $s = 1$ and $\alpha = 1000$.

was verified on a microscopy dataset depicting endodermal cells of a zebrafish embryo.

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).




Figure 6.12: Solutions $(u^{(k)})_{k=2,\dots,7}$ of velocity fields obtained by the hierarchical decomposition. Images are aligned from left to right and top to bottom. At iteration k the sequence was set to $\mu_n^{(k)} = \alpha \lambda_n^{s - \frac{k-1}{4}}$ with parameters $s = 2$ and $\alpha = 1$.

Chapter 7

Optical Flow on Evolving Sphere-Like Surfaces

Authors & Contributions The authors are Lukas Lang and Otmar Scherzer. The development of this article was a gradual, cooperative process, and each of the authors made significant contributions to every aspect of the paper.

Publication Status Submitted to *Inverse Probl. Imaging*. Date of acknowledgement of receipt: June 10, 2015. Preprint available in [55].

Optical Flow on Evolving Sphere-Like Surfaces

Lukas F. Lang¹, Otmar Scherzer^{1,2}

¹Computational Science Center, University of Vienna,
Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria

²Radon Institute of Computational and Applied Mathematics,
Austrian Academy of Sciences, Altenberger Str. 69, 4040 Linz, Austria

Abstract

In this work we consider optical flow on evolving Riemannian 2-manifolds which can be parametrised from the 2-sphere. Our main motivation is to estimate cell motion in time-lapse volumetric microscopy images depicting fluorescently labelled cells of a live zebrafish embryo. We exploit the fact that the recorded cells float on the surface of the embryo and allow for the extraction of an image sequence together with a sphere-like surface. We solve the resulting variational problem by means of a Galerkin method based on vector spherical harmonics and present numerical results computed from the aforementioned microscopy data.

7.1 Introduction

Motion estimation is a fundamental problem in image analysis and computer vision. An important task within is optical flow computation. It is concerned with the inference of a vector field describing the displacements of brightness patterns, such as moving objects, in a sequence of images. Ever since the seminal work of Horn and Schunck [39] a variety of reliable and efficient methods have been proposed and successfully applied in a wide number of fields.

Primarily, optical flow is computed in the plane. However, it is readily generalised to non-Euclidean settings allowing, for instance, for cell motion analysis in time-lapse microscopy data. It has been only recently that high-resolution observations of biological model organisms such as the zebrafish became possible. Despite its importance for tissue and organ formation, little is known about cell migration and proliferation patterns during the zebrafish's early embryonic development [3, 74]. Fluorescence microscopy nowadays allows to record time-lapse images on the scale of single cells, see e.g. [43, 60, 74]. Increasing spatial as well as temporal resolutions result in vast amounts of data, rendering extraction of information through visual inspection carried out by humans impracticable. Automated cell motion estimation therefore is key to large-scale analysis of such data. Optical flow computation delivers necessary




Figure 7.1: Frames 70 (left) and 71 (right) of the volumetric zebrafish microscopy images recorded during early embryogenesis. The sequence contains a total number of 75 frames. Fluorescence response is indicated by blue colour and is proportional to the observed intensity. All dimensions are in micrometer (μm).

quantitative methods and leads to insights into the underlying cellular mechanisms and the dynamic behaviour of cells. See, for example, [4, 61, 72, 74] and the references therein.

The primary biological motivation for this work is the desire to analyse cell motion in a living zebrafish during early embryogenesis. The data at hand depict endodermal cells expressing a green fluorescent protein. By virtue of laser-scanning microscopy, (volumetric time-lapse) 4D images of these labelled cells can be recorded without capturing the background. It is known that endodermal cells float on a so called *monolayer* during early embryonic development meaning that they do not stack on top of each other [83]. Figure 7.1 depicts two frames of the captured sequence, containing only the upper hemisphere of the animal embryo. Observe the salient formation of the cells and the noise present in the images. More precisely, one can see the nuclei of cells forming a round surface in a single layer. For more details on the microscopy data we refer to Sec. 7.5.

We exploit this situation and model this layer as an evolving surface. A natural candidate for a parametrisation of such a zebrafish embryo is a *sphere-like surface*. It is topologically diffeomorphic to the 2-sphere \mathcal{S}^2 and most commonly defined as the set of points

$$\{\tilde{\rho}(x)x : x \in \mathcal{S}^2\}.$$

The function $\tilde{\rho} : \mathcal{S}^2 \rightarrow (0, \infty)$ can be thought of as a radial deformation of \mathcal{S}^2 and will have a dependence on time in the present paper. As a consequence, changes in the embryo's geometry are attributed accordingly, albeit valid only during early stages of its development as cells tend to cluster subsequently. The main intention of this work is to conceive cell motion only on this moving 2-dimensional manifold. As a result we are able to reduce the spatial dimension of the data allowing for more efficient motion estimation in microscopy data. Figure 7.2 depicts two frames of the surface together with images obtained by restriction of the volumetric microscopy data in Fig. 7.1.

In this work we model the data as a time-dependent non-negative function \hat{f} . Its value directly corresponds to the fluorescence response of the observed cells. For a fixed time instant $t \in [0, T]$, the domain of \hat{f} is presumed to be a

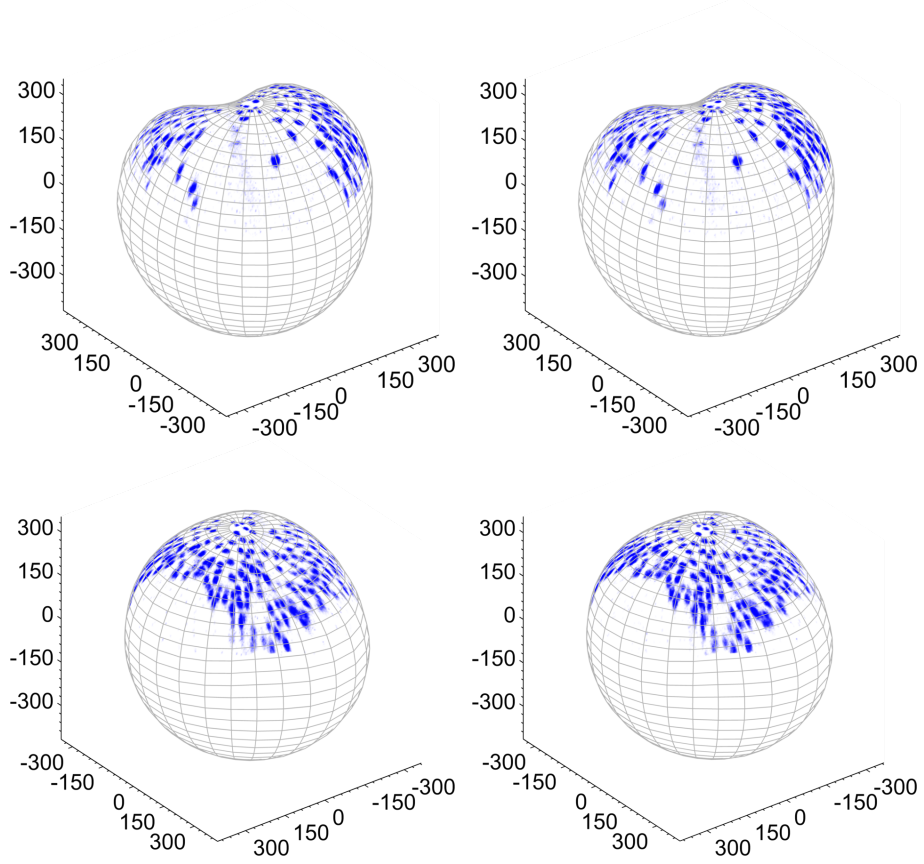


Figure 7.2: Depicted are frames no. 70 (left) and 71 (right) of the processed zebrafish microscopy sequence. Top and bottom row differ by a rotation of 180 degrees around the x_3 -axis. All dimensions are in micrometer (μm).

closed surface $\mathcal{M}_t \subset \mathbb{R}^3$. We assume that this surface can be parametrised by a smooth radial map from the 2-sphere. The temporal evolution of the data \hat{f} can then be tracked by solving an optical flow problem on this moving surface or, more conveniently, an equivalent problem on the round sphere.

Traditionally, the starting point for optical flow is the assumption of constant brightness: a point moving along a trajectory does not change its intensity over time. On a moving domain $\mathcal{M} = \{\mathcal{M}_t\}_t$ one equivalently seeks, for every time $t \in [0, T]$, a tangent vector field $\hat{\mathbf{v}}$ that solves a generalised optical flow equation

$$d_t^{\hat{\mathbf{v}}} \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{v}} = 0 \quad (7.1)$$

at every point $x \in \mathcal{M}$, where \hat{f} is the image sequence living on \mathcal{M} . Here, for a fixed time t , $\nabla_{\mathcal{M}}$ denotes the (spatial) surface gradient, dot the standard inner product, and $d_t^{\hat{\mathbf{v}}} \hat{f}$ an appropriate temporal derivative.

The optical flow problem is ill-posed meaning that equation (7.1) is not uniquely solvable. A common approach to deal with non-uniqueness is Tikhonov

nov regularisation, which consists of computing a minimiser of

$$\mathcal{E}_\alpha(\hat{\mathbf{v}}) = \mathcal{D}(\hat{\mathbf{v}}, \hat{f}) + \alpha \mathcal{R}(\hat{\mathbf{v}}).$$

The first term of the sum is usually the squared L^2 norm of the left-hand side of (7.1) and, in the present article, the second term will be an H^1 Sobolev norm.

Contributions

The primary concern of this article is optical flow computation on evolving 2-dimensional Riemannian manifolds which can be parametrised from the sphere. Motivated by the aforementioned zebrafish microscopy data we consider closed surfaces for which the mapping

$$(t, x) \mapsto \tilde{\rho}(t, x)x, \quad x \in \mathcal{S}^2 \quad (7.2)$$

is a diffeomorphism between the 2-sphere and \mathcal{M}_t for every time $t \in [0, T]$. As a prototypical example we restrict ourselves to radially parametrised surfaces as they suit quite naturally to the given data.

The contributions of this work are as follows. First, we give a variational formulation of optical flow on 2-dimensional closed Riemannian manifolds. We assume a dependence on time and speak of evolving surfaces. The main idea is to solve the problem by a Galerkin method in a finite-dimensional subspace of an appropriate (vectorial) Sobolev space. We take advantage of the fact that tangential vector spherical harmonics form a complete orthonormal system for $L^2(\mathcal{S}^2, T\mathcal{S}^2)$. The sought vector field is thus uniquely determined when expanded in terms of the pushforward—by means of the differential of (7.2)—of these functions. From that we arrive at a minimisation problem over \mathbb{R}^n , where n is the dimension of the finite-dimensional space, and state the optimality conditions. They can be written purely in terms of spherical quantities and solved on the 2-sphere. To this end, we use a standard polyhedral approximation and locally interpolate spherical functions by piecewise quadratic polynomials. For numerical integration we employ appropriate quadrature rules on the approximated sphere.

Second, to obtain the smooth sphere-like surface, which is described by the map (7.2), from the observed microscopy data, we formulate another variational problem on the sphere. The problem is essentially surface interpolation with H^s Sobolev seminorm regularisation. Approximate cell centres serve as sample points of the surface. In particular, our microscopy data are supported only on the upper hemisphere, see Figs. 7.1 and 7.2. Scalar spherical harmonics are the appropriate choice for the numerical solution of the surface fitting problem, as they provide great flexibility with respect to the chosen space H^s .

Finally, we present numerical experiments on the basis of the mentioned cell microscopy data of a live zebrafish. To this end we compute an approximation of the sphere-shaped embryo and obtain a sequence of images living on this moving surface. Eventually, we solve for the optical flow and present the results in a visually adequate manner.

Related Work

The first variational formulation of optical flow is commonly attributed to Horn and Schunck [39]. They attempted to compute a displacement field in \mathbb{R}^2 by minimising a Tikhonov-regularised energy functional. It favours spatially regular vector fields by penalising its squared H^1 Sobolev seminorm. For introductory material on the subject we refer to [6, 8] and to [84] for a survey on various optical flow functionals. Well-posedness of the aforementioned energy was first shown by Schnörr [75]. Moreover, there the problem was extended to irregular planar domains and solved by means of finite elements.

Weickert and Schnörr [86] considered a spatio-temporal model by extending the domain to $\mathbb{R}^2 \times [0, T]$. It additionally favours temporal regularity of the solution by including first derivatives with respect to time. Such models are of particular interest whenever trajectories are to be computed from the optical flow field. A unifying framework including several spatial as well as temporal regularisers was proposed in [85]. For the purpose of evaluation and flow field visualisation a framework was created by Baker et al. [9].

Recently, generalisations to non-Euclidean domains have gained increasing attention. In [42] and [79] optical flow was considered in a spherical setting. Lefèvre and Baillet [58] adapted the Horn-Schunck functional to surfaces embedded in \mathbb{R}^3 . Following Schnörr [75], they proved well-posedness of their formulation and employed a finite element method for solving the discrete problem on a triangle mesh. With an application to cell motion analysis, Kirisits et al. [50, 52] recently considered optical flow on evolving surfaces with boundary. They generalised the spatio-temporal model in [86] to a non-Euclidean and dynamic setting. Eventually, the problem was tackled numerically by solving the corresponding Euler-Lagrange equations in the coordinate domain. Similarly, Bauer et al. [12] studied optical flow on time-varying domains, with and without spatial boundary. They proposed a treatment on surfaces parametrised by product manifolds, constructed an appropriate Riemannian metric, and proved well-posedness of their formulation.

In Kirisits et al. [51], the authors considered various decomposition models for optical flow on the 2-sphere. The proposed functionals were solved by means of projection to a finite-dimensional space spanned by vector spherical harmonics. Concerning projection methods, Schuster and Weickert [76] solved the optical flow problem in \mathbb{R}^2 solely based on regularisation by discretisation.

Regarding sphere-like surfaces and spherical harmonics expansion of closed surfaces we refer to [71] and the references therein.

Finally, let us mention [4, 61, 72, 74], where optical flow was employed for the analysis of cell motion in microscopy data. In particular, in Schmid et al. [74] the embryo of a zebrafish was modelled as a round sphere and motion of endodermal cells computed in map projections.

The remainder of this article is structured as follows. In Sec. 7.2, we formally introduce evolving sphere-like surfaces, recall the definition of vectorial Sobolev spaces on manifolds, and discuss both scalar and vector spherical harmonics on the 2-sphere. Section 7.3 is dedicated to optical flow on evolving surfaces and our variational formulation. In Sec. 7.4 we discuss the numerical solution. In particular, we propose to solve the resulting energy in a finite-dimensional subspace and rewrite the optimality conditions to be defined solely on the 2-sphere. Moreover, we show how to fit a sphere-like surface to the la-

belled cells in the microscopy data. Finally, in Sec. 7.5, we solve for the optical flow field and visualise the results. The appendix contains deferred material.

7.2 Notation and Background

Sphere-Like Surfaces

Let

$$\mathcal{S}^2 = \{x \in \mathbb{R}^3 : \|x\| = 1\}$$

be the 2-sphere embedded in the 3-dimensional Euclidean space. The norm of \mathbb{R}^n , $n = \{2, 3\}$, is denoted by $\|x\| = \sqrt{x \cdot x}$. By

$$\mathbf{x} : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}^3 \quad (7.3)$$

we denote a smooth (local) parametrisation of \mathcal{S}^2 mapping coordinates $\xi = (\xi^1, \xi^2)^\top \in \Omega$ to points $x = (x^1, x^2, x^3)^\top \in \mathcal{S}^2$.

Furthermore, let $I := [0, T] \subset \mathbb{R}$ denote a time interval and let $\mathcal{M} = \{\mathcal{M}_t\}_{t \in I}$ be a family of closed smooth 2-manifolds $\mathcal{M}_t \subset \mathbb{R}^3$. Each \mathcal{M}_t , $t \in I$, is assumed to be regular and oriented by the outward unit normal field $\hat{\mathbf{N}}(t, x) \in \mathbb{R}^3$, $x \in \mathcal{M}_t$. We assume that \mathcal{M} (locally) admits a smooth parametrisation of the form

$$\mathbf{y} : I \times \Omega \rightarrow \mathbb{R}^3, \quad (t, \xi^1, \xi^2)^\top \mapsto \tilde{\rho}(t, \mathbf{x}(\xi^1, \xi^2)) \mathbf{x}(\xi^1, \xi^2) \in \mathcal{M}_t \quad (7.4)$$

and call \mathcal{M} an *evolving sphere-like surface*.

We denote by $\hat{f} : \mathcal{M} \rightarrow \mathbb{R}$ a smooth function on the moving surface. Its coordinate representation $f : I \times \Omega \rightarrow \mathbb{R}$ and its corresponding spherical representation $\tilde{f} : I \times \mathcal{S}^2 \rightarrow \mathbb{R}$ are given by

$$f(t, \xi) = \tilde{f}(t, \mathbf{x}(\xi)) = \hat{f}(t, \mathbf{y}(t, \xi)). \quad (7.5)$$

As a notational convention we indicate functions living on \mathcal{S}^2 with a tilde and functions on \mathcal{M} with a hat, respectively. Their corresponding coordinate version is treated without special indication.

For convenience, we define smooth extensions of \tilde{f} and \hat{f} to $\mathbb{R}^3 \setminus \{0\}$ by

$$\tilde{\tilde{f}}(t, x) := \tilde{f}\left(t, \frac{x}{\|x\|}\right) \quad \text{and} \quad \hat{\hat{f}}(t, x) := \hat{f}\left(t, \tilde{\rho}\left(t, \frac{x}{\|x\|}\right) \frac{x}{\|x\|}\right), \quad (7.6)$$

respectively. Note that, while $\tilde{\tilde{f}}$ is constant in the direction of the surface normal of \mathcal{S}^2 , the extension $\hat{\hat{f}}$ in general is not. We point at Fig. 7.3 illustrating the setting.

Similarly, for vector-valued functions $\tilde{\mathbf{u}} : I \times \mathcal{S}^2 \rightarrow \mathbb{R}^3$ and $\hat{\mathbf{u}} : \mathcal{M} \rightarrow \mathbb{R}^3$ the extensions to $\mathbb{R}^3 \setminus \{0\}$ are defined component-wise and for all times $t \in I$. They are denoted by $\tilde{\tilde{\mathbf{u}}}$ and $\hat{\hat{\mathbf{u}}}$, respectively. As a notational convention, boldface letters are used to denote vector fields. Moreover, we distinguish between lower and upper case boldface letters. The former identify tangent vector fields and their extensions to $\mathbb{R}^3 \setminus \{0\}$ whereas the latter indicate general vector fields in \mathbb{R}^3 .

For a differentiable function $f : I \times \Omega \rightarrow \mathbb{R}$, we write $\partial_i f$ as an abbreviation for the partial derivative of f with respect to ξ^i . That is, $\nabla_{\mathbb{R}^2} f = (\partial_1 f, \partial_2 f)^\top$, where $\nabla_{\mathbb{R}^2}$ is the gradient of \mathbb{R}^2 .




Figure 7.3: Schematic illustration of a cut through the surfaces \mathcal{S}^2 and \mathcal{M}_t intersecting the origin. In addition, a radial line along which the extensions \hat{f} and \tilde{f} are constant is shown. The surface normals are shown in grey.

The tangent plane at a point $\mathbf{y}(t, \xi) \in \mathcal{M}_t$ is denoted by $T_{\mathbf{y}(t, \xi)}\mathcal{M}_t$ and the tangent bundle by $T\mathcal{M}_t = \{\{\mathbf{y}(t, \xi)\} \times T_{\mathbf{y}(t, \xi)}\mathcal{M}_t : \xi \in \Omega\}$. The orthogonal projector onto the tangent plane $T_x\mathcal{M}_t$ at $x \in \mathcal{M}_t$, $t \in I$, is given by

$$P_{\mathcal{M}}(t, x) = \text{Id} - \hat{\mathbf{N}}(t, x)\hat{\mathbf{N}}(t, x)^\top \in \mathbb{R}^{3 \times 3}.$$

In particular if $\mathcal{M}_t = \mathcal{S}^2$, that is $\tilde{\rho}$ in (7.4) is identically one for all $t \in I$, the outward unit normal and the orthogonal projector are given by $\tilde{\mathbf{N}}$ and $P_{\mathcal{S}^2}$, respectively.

In what follows, we define spatial differential operators. As they are identical to those on static surfaces we consider time $t \in I$ arbitrary but fixed. Then, the surface gradient of \hat{f} , as given in (7.5), is defined by

$$\nabla_{\mathcal{M}}\hat{f} := P_{\mathcal{M}}\nabla_{\mathbb{R}^3}\hat{f} \in \mathbb{R}^3, \quad (7.7)$$

where $\nabla_{\mathbb{R}^3}$ denotes the usual gradient of the embedding space. Let us stress that it is independent of the chosen extension, see e.g. [31, p. 389].

We emphasise that, in particular, if $\mathcal{M}_t = \mathcal{S}^2$ for all $t \in I$ it follows that

$$\nabla_{\mathbb{R}^3}\tilde{f} = P_{\mathcal{S}^2}\nabla_{\mathbb{R}^3}\tilde{f} + (\text{Id} - P_{\mathcal{S}^2})\nabla_{\mathbb{R}^3}\tilde{f}.$$

The last term of the sum on the right hand side is the normal derivative of \tilde{f} , which according to the definition of the extension in (7.6) vanishes. Thus,

$$\nabla_{\mathcal{S}^2}\tilde{f} = P_{\mathcal{S}^2}\nabla_{\mathbb{R}^3}\tilde{f} = \nabla_{\mathbb{R}^3}\tilde{f}. \quad (7.8)$$

For convenience let us observe that, by taking $\partial_i f$ in (7.5), we arrive at

$$\partial_i f(t, \xi) = \nabla_{\mathbb{R}^3}\tilde{f}(t, \mathbf{x}(\xi)) \cdot \partial_i \mathbf{x}(\xi) = \nabla_{\mathcal{S}^2}\tilde{f}(t, \mathbf{x}(\xi)) \cdot \partial_i \mathbf{x}(\xi) \quad (7.9)$$

due to the chain rule and the projection onto the tangent plane $T_{\mathbf{x}(\xi)}\mathcal{S}^2$.

Analogously to the surface gradient we define the spherical Laplace-Beltrami of $\tilde{f} : I \times \mathcal{S}^2 \rightarrow \mathbb{R}$ as

$$\Delta_{\mathcal{S}^2} \tilde{f} = -\Delta_{\mathbb{R}^3} \tilde{f}, \quad (7.10)$$

where $\Delta_{\mathbb{R}^3}$ is the standard Laplacian of \mathbb{R}^3 .

The set

$$\{\partial_1 \mathbf{y}(t, \xi), \partial_2 \mathbf{y}(t, \xi)\} \subseteq \mathbb{R}^3, \quad (7.11)$$

where \mathbf{y} is the parametrisation defined in (7.4), forms a basis of the tangent space $T_{\mathbf{y}(t, \xi)} \mathcal{M}_t$ at $\mathbf{y}(t, \xi)$. Its elements form the gradient matrix $D\mathbf{y}$, which is derived as follows.

Let $\tilde{\rho}$ be the extension of $\rho : I \times \mathcal{S}^2 \rightarrow (0, \infty)$ according to (7.6). Then, \mathbf{y} from (7.4) can be rewritten as

$$\mathbf{y}(t, \xi) = \tilde{\rho}(t, \mathbf{x}(\xi)) \mathbf{x}(\xi).$$

By the chain rule,

$$\partial_i \mathbf{y}(t, \xi) = (\nabla_{\mathbb{R}^3} \tilde{\rho}(t, \mathbf{x}(\xi)) \cdot \partial_i \mathbf{x}(\xi)) \mathbf{x}(\xi) + \tilde{\rho}(t, \mathbf{x}(\xi)) \partial_i \mathbf{x}(\xi).$$

Using (7.8) and the fact that $\tilde{\rho}$ equals ρ on \mathcal{S}^2 gives

$$\partial_i \mathbf{y} = (\nabla_{\mathcal{S}^2} \rho \cdot \partial_i \mathbf{x}) \mathbf{x} + \rho \partial_i \mathbf{x},$$

where we have omitted the arguments (t, ξ) and (ξ) for better readability. Whenever convenient and no confusion will arise we will continue to do so.

By applying (7.9) backwards and the fact that $\tilde{\rho}(t, \mathbf{x}(\xi)) = \rho(\xi)$ we have shown

$$\begin{aligned} D\mathbf{y} &= (\partial_1 \mathbf{y} \quad \partial_2 \mathbf{y}) \\ &= ((\partial_1 \rho) \mathbf{x} \quad (\partial_2 \rho) \mathbf{x}) + \rho D\mathbf{x} \in \mathbb{R}^{3 \times 2}, \end{aligned} \quad (7.12)$$

where $D\mathbf{x} = (\partial_1 \mathbf{x}, \partial_2 \mathbf{x})$ is the gradient matrix associated with \mathbf{x} .

As a consequence, we can uniquely represent a tangent vector $\hat{\mathbf{u}} \in T_{\mathbf{y}(t, \xi)} \mathcal{M}_t$ as $\hat{\mathbf{u}} = \sum_{i=1}^2 u^i \partial_i \mathbf{y}$, where $\mathbf{u} = (u^1, u^2)^\top \in \mathbb{R}^2$ is its coordinate representation, see e.g. [57, Prop. 3.15]. We call u^i the components of $\hat{\mathbf{u}}$.

In the sequel we will use Einstein summation convention. We sum over every index letter that appears exactly twice in an expression, once as a sub- and once as a superscript. For instance, we write $\hat{\mathbf{u}} = u^i \partial_i \mathbf{y}$ for the sake of brevity.

We underline that the coordinate basis (7.11) is not orthogonal in general. We will, however, require an orthonormal frame $\{\hat{\mathbf{e}}_1(t, \xi), \hat{\mathbf{e}}_2(t, \xi)\}$ of the tangent space $T_{\mathbf{y}(t, \xi)} \mathcal{M}_t$ from Sec. 7.2 onwards. In the coordinate basis it reads

$$\hat{\mathbf{e}}_i = \alpha_i^j \partial_j \mathbf{y}, \quad (7.13)$$

where $\alpha_i^j : I \times \Omega \rightarrow \mathbb{R}$, $i, j = \{1, 2\}$, are functions obtained from the Gram-Schmidt process.

Combining (7.5) and (7.9) with the expressions derived for $D\mathbf{x}$ and $D\mathbf{y}$ we can conveniently state that

$$\nabla_{\mathbb{R}^2} f = D\mathbf{x}^\top \nabla_{\mathcal{S}^2} \tilde{f} \text{ and } \nabla_{\mathbb{R}^2} f = D\mathbf{y}^\top \nabla_{\mathcal{M}} \hat{f}. \quad (7.14)$$

Let us derive the following useful generalisation of (7.9). For a tangent vector $\tilde{\mathbf{v}} = v^i \partial_i \mathbf{x} \in T_x \mathcal{S}^2$, $x \in \mathcal{S}^2$, the directional derivative of \tilde{f} along $\tilde{\mathbf{v}}$ at x is

$$\nabla_{\mathcal{S}^2} \tilde{f} \cdot \tilde{\mathbf{v}} = \nabla_{\mathcal{S}^2} \tilde{f} \cdot v^i \partial_i \mathbf{x} = (D\mathbf{x}^\top \nabla_{\mathcal{S}^2} \tilde{f}) \cdot \mathbf{v} = \nabla_{\mathbb{R}^2} f \cdot \mathbf{v} = v^i \partial_i f, \quad (7.15)$$

where the third equality follows from the first equation in (7.14). Analogously, for $\hat{\mathbf{v}} = v^i \partial_i \mathbf{y} \in T_x \mathcal{M}_t$, with $x \in \mathcal{M}_t$ and $t \in I$, one can derive

$$\nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{v}} = v^i \partial_i f. \quad (7.16)$$

As soon as we have established the relation between $\hat{\mathbf{v}}$ and $\tilde{\mathbf{v}}$ it will conveniently allow us to switch between (7.15) and (7.16).

Moreover, the coordinate representation of the surface gradient (7.8) is derived as follows. Let us start out with the first equation in (7.14). By writing $\nabla_{\mathcal{S}^2} \tilde{f}$ in the coordinate basis, that is $\nabla_{\mathcal{S}^2} \tilde{f} = D\mathbf{x}\mathbf{u}$ for some \mathbf{u} , we obtain from (7.14)

$$\nabla_{\mathbb{R}^2} f = D\mathbf{x}^\top D\mathbf{x}\mathbf{u}.$$

Multiplying with $(D\mathbf{x}^\top D\mathbf{x})^{-1}$ from the left yields

$$(D\mathbf{x}^\top D\mathbf{x})^{-1} \nabla_{\mathbb{R}^2} f = \mathbf{u}.$$

Thus,

$$\nabla_{\mathcal{S}^2} \tilde{f} = D\mathbf{x}\mathbf{u} = D\mathbf{x}(D\mathbf{x}^\top D\mathbf{x})^{-1} \nabla_{\mathbb{R}^2} f. \quad (7.17)$$

Furthermore, let $t \in I$ be fixed and let $\hat{f}(t, \cdot) : \mathcal{M}_t \rightarrow \mathbb{R}$. The surface integral of \hat{f} is

$$\int_{\mathcal{M}_t} \hat{f} d\mathcal{M}_t = \int_{\Omega} f J\mathbf{y} d\xi, \quad (7.18)$$

where $J\mathbf{y}$ is the Jacobian of \mathbf{y} . According to Theorem 3 in [26, p. 88], it is given by

$$(J\mathbf{y})^2 = \det(D\mathbf{y}^\top D\mathbf{y})$$

and by using (7.12) yields

$$\begin{aligned} (J\mathbf{y})^2 &= \rho^2 ((\partial_1 \rho)^2 \partial_2 \mathbf{x} \cdot \partial_2 \mathbf{x} + (\partial_2 \rho)^2 \partial_1 \mathbf{x} \cdot \partial_1 \mathbf{x} \\ &\quad + \rho^2 (\partial_1 \mathbf{x} \cdot \partial_1 \mathbf{x})(\partial_2 \mathbf{x} \cdot \partial_2 \mathbf{x}) - 2\partial_1 \rho \partial_2 \rho (\partial_1 \mathbf{x} \cdot \partial_2 \mathbf{x}) - \rho^2 (\partial_1 \mathbf{x} \cdot \partial_2 \mathbf{x})^2). \end{aligned} \quad (7.19)$$

Note that $\mathbf{x} \cdot \mathbf{x} = 1$ and thus, terms of the form $\partial_i \mathbf{x} \cdot \mathbf{x}$ vanish. By the differentiability of \mathbf{x} , ones has

$$\partial_i (\mathbf{x} \cdot \mathbf{x}) = 0.$$

Therefore, $\partial_i \mathbf{x} \cdot \mathbf{x} = 0$, meaning that tangential and normal vectors are orthogonal. We emphasise that $D\mathbf{y}^\top D\mathbf{y}$ is commonly referred to as *Riemannian metric*. It is positive definite and thus, $(J\mathbf{y}(t, \xi))^2 > 0$ for all $(t, \xi) \in I \times \Omega$.

The parametrisations \mathbf{x} and \mathbf{y} defined in (7.3) and (7.4), respectively, suggest the straightforward construction of a smooth map $\tilde{\phi}(t, \cdot) : \mathcal{S}^2 \rightarrow \mathcal{M}_t$. It is given by the composition $(\mathbf{y} \circ \mathbf{x}^{-1})(t, \cdot)$, that is

$$\tilde{\phi}(t, x) : x \mapsto \tilde{\rho}(t, x)x.$$

The differential $D\tilde{\phi}(t, x) : T_x\mathcal{S}^2 \rightarrow T_{\tilde{\phi}(t, x)}\mathcal{M}_t$ of $\tilde{\phi}$ is a linear map and is given by

$$D\tilde{\phi}(t, x) = \tilde{\rho}(t, x)\text{Id} + x\nabla_{\mathcal{S}^2}\tilde{\rho}(t, x)^\top \in \mathbb{R}^{3 \times 3}. \quad (7.20)$$

It follows from a direct calculation akin to the derivation of $D\mathbf{y}$ in (7.12).

Let us exhibit the action of $D\tilde{\phi}(t, x)$, for $x = \mathbf{x}(\xi)$ and $t \in I$, onto a tangent vector $\tilde{\mathbf{v}} = v^i \partial_i \mathbf{x} \in T_x\mathcal{S}^2$. We have

$$\begin{aligned} D\tilde{\phi}(t, x)(\tilde{\mathbf{v}}) &= \tilde{\rho}(t, x)\tilde{\mathbf{v}} + x(\nabla_{\mathcal{S}^2}\tilde{\rho}(t, x) \cdot \tilde{\mathbf{v}}) \\ &\stackrel{(7.15)}{=} \tilde{\rho}(t, x)\tilde{\mathbf{v}} + xv^i \partial_i \rho(\xi) \\ &= \tilde{\rho}(t, x)v^i \partial_i \mathbf{x} + xv^i \partial_i \rho(\xi) \\ &= v^i (\tilde{\rho}(t, x)\partial_i \mathbf{x} + x\partial_i \rho(\xi)) \\ &\stackrel{(7.12)}{=} v^i \partial_i \mathbf{y}(\xi). \end{aligned} \quad (7.21)$$

In other words, the components $(v^1, v^2)^\top$ are preserved whenever a tangent vector is mapped from \mathcal{S}^2 to \mathcal{M}_t via the differential (7.20).

As a matter of fact, given a tangent vector field $\tilde{\mathbf{v}} = v^i \partial_i \mathbf{x}$ on \mathcal{S}^2 , the differential $D\tilde{\phi}$ gives rise to a unique tangent vector field $\hat{\mathbf{v}} = v^i \partial_i \mathbf{y}$ on \mathcal{M}_t , see [57, Chapter 8]. Whenever we use $\tilde{\mathbf{v}}$ and $\hat{\mathbf{v}}$ in the sequel we refer to their unique identification via the differential (7.20) and call $\hat{\mathbf{v}}$ the pushforward of $\tilde{\mathbf{v}}$. At this point, the reader might find it helpful to have a look at Fig. 7.5.

With the above definitions at hand we are able to relate the surface integral (7.18) to an integral on \mathcal{S}^2 via a change of variables. The key is to compute a meaningful surface element as $|\det(D\tilde{\phi})|$ is the magnitude of the change of the volume element. The following lemma provides the required form.

Lemma 4. *Let $\mathbf{x} : [0, \pi] \times [0, 2\pi] \rightarrow \mathbb{R}^3$ be the standard parametrisation of \mathcal{S}^2 ,*

$$(\xi^1, \xi^2)^\top \mapsto (\sin \xi^1 \cos \xi^2, \sin \xi^1 \sin \xi^2, \cos \xi^1)^\top,$$

and let $\hat{f} : \mathcal{M} \rightarrow \mathbb{R}$ and $\tilde{\rho} : I \times \mathcal{S}^2 \rightarrow (0, \infty)$ be as above. Then, for $t \in I$,

$$\int_{\mathcal{M}_t} \hat{f} d\mathcal{M}_t = \int_{\mathcal{S}^2} \tilde{f} \tilde{\rho} \sqrt{\|\nabla_{\mathcal{S}^2} \tilde{\rho}\|^2 + \tilde{\rho}^2} d\mathcal{S}^2.$$

Proof. Let us denote by $\tilde{\mathbf{e}}_1(\xi)$ and $\tilde{\mathbf{e}}_2(\xi)$ the orthogonal unit vectors on \mathcal{S}^2 in direction of ξ^1 and ξ^2 , respectively, which are obtained by normalising the coordinate basis $\{\partial_1 \mathbf{x}(\xi), \partial_2 \mathbf{x}(\xi)\}$. That is,

$$\tilde{\mathbf{e}}_1(\xi) = \partial_1 \mathbf{x}(\xi) \text{ and } \tilde{\mathbf{e}}_2(\xi) = \frac{\partial_2 \mathbf{x}(\xi)}{\|\partial_2 \mathbf{x}(\xi)\|}. \quad (7.22)$$

Moreover, a straightforward calculation gives

$$D\mathbf{x}^\top D\mathbf{x} = \begin{pmatrix} 1 & 0 \\ 0 & \sin^2 \xi^1 \end{pmatrix}$$

and thus, the surface gradient of $\tilde{\rho}$ in spherical coordinates (7.17) is given by

$$\begin{aligned} \nabla_{\mathcal{S}^2} \tilde{\rho}(t, \mathbf{x}(\xi)) &= \partial_1 \rho(\xi) \partial_1 \mathbf{x}(\xi) + \frac{1}{\sin^2 \xi^1} \partial_2 \rho(\xi) \partial_2 \mathbf{x}(\xi) \\ &\stackrel{(7.22)}{=} \partial_1 \rho(\xi) \tilde{\mathbf{e}}_1(\xi) + \frac{1}{\sin \xi^1} \partial_2 \rho(\xi) \tilde{\mathbf{e}}_2(\xi), \end{aligned}$$

where we have replaced the coordinate basis with the orthonormal basis.

Using $D\mathbf{x}^\top D\mathbf{x}$ in (7.19), the Jacobian $J\mathbf{y}$ can be written as

$$\begin{aligned} (J\mathbf{y})^2 &= \rho^2((\partial_1\rho)^2 \sin^2 \xi^1 + (\partial_2\rho)^2 + \rho^2 \sin^2 \xi^1) \\ &= \rho^2((\partial_1\rho)^2 + \frac{1}{\sin^2 \xi^1}(\partial_2\rho)^2 + \rho^2) \sin^2 \xi^1 \\ &= \rho^2(\|\nabla_{S^2}\tilde{\rho}\|^2 + \rho^2) \sin^2 \xi^1. \end{aligned}$$

Here, we have omitted the argument $(t, \mathbf{x}(\xi))$ of $\nabla_{S^2}\tilde{\rho}$. Then, the integral turns out to be

$$\begin{aligned} \int_{\mathcal{M}_t} \hat{f} d\mathcal{M}_t &= \int_0^{2\pi} \int_0^\pi f J\mathbf{y} d\xi \\ &= \int_0^{2\pi} \int_0^\pi f \rho \sqrt{\|\nabla_{S^2}\tilde{\rho}\|^2 + \rho^2} \sin \xi^1 d\xi \\ &= \int_{S^2} \tilde{f} \tilde{\rho} \sqrt{\|\nabla_{S^2}\tilde{\rho}\|^2 + \tilde{\rho}^2} dS^2, \end{aligned}$$

where the last equation follows from (7.18) if $\mathcal{M}_t = S^2$, the fact that $\sin \xi^1 \geq 0$, and

$$J\mathbf{x} = \sqrt{\det(D\mathbf{x}^\top D\mathbf{x})} = \sin \xi^1.$$

□

The concepts introduced above, and further properties thereof, may be found in any standard differential geometry book. For instance, in [22, 23, 56, 57].

Vectorial Sobolev Spaces on Manifolds

We briefly introduce the appropriate function spaces required for the variational optical flow formulation on Riemannian manifolds. Again, let us consider time $t \in I$ arbitrary but fixed.

For a tangent vector field $\hat{\mathbf{v}}$ on \mathcal{M}_t we denote by $\nabla_{\hat{\mathbf{u}}}\hat{\mathbf{v}}(x)$ the covariant derivative of $\hat{\mathbf{v}}$ at $x \in \mathcal{M}_t$ along the direction of a tangent vector $\hat{\mathbf{u}} \in T_x\mathcal{M}_t$. We define it as the tangential part of the usual directional derivative of the extension $\hat{\hat{\mathbf{v}}}$ along $\hat{\mathbf{u}}$ in the embedding space, that is,

$$\nabla_{\hat{\mathbf{u}}}\hat{\mathbf{v}} := P_{\mathcal{M}}\nabla_{\mathbb{R}^3}\hat{\hat{\mathbf{v}}}(\hat{\mathbf{u}}). \quad (7.23)$$

It is a linear operator $\nabla\hat{\mathbf{v}}(x) : T_x\mathcal{M}_t \rightarrow T_x\mathcal{M}_t$ and its Hilbert-Schmidt norm is given by

$$\|\nabla\hat{\mathbf{v}}(x)\|_2^2 = \sum_{i=1}^2 \|\nabla_{\hat{\mathbf{e}}_i}\hat{\mathbf{v}}(x)\|^2, \quad (7.24)$$

where $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2\}$ denotes the orthonormal basis of the tangent space $T_x\mathcal{M}_t$, cf. (7.13). We stress that (7.24) is invariant with respect to the chosen parametrisation.

For each $t \in I$, we define the Sobolev space $H^1(\mathcal{M}_t, T\mathcal{M}_t)$ as the completion of the space of vector fields $C^\infty(\mathcal{M}_t, T\mathcal{M}_t)$ with respect to the norm

$$\|\hat{\mathbf{v}}\|_{H^1(\mathcal{M}_t, T\mathcal{M}_t)}^2 := \int_{\mathcal{M}_t} \|\nabla\hat{\mathbf{v}}\|_2^2 d\mathcal{M}_t, \quad (7.25)$$

where the surface integral is defined in (7.18). Let us emphasise that (7.25) is indeed a norm whenever \mathcal{M}_t is diffeomorphic to the 2-sphere. The reason is that, by virtue of the Hairy Ball Theorem, there is no covariantly constant tangent vector field but $\hat{\mathbf{v}} = 0$, see e.g. [37, p. 125].

Alternatively, one can define Sobolev spaces of vector fields such that each component of a vector field originates from a scalar Sobolev space. See, for instance, Lefèvre and Baillet [58]. On the 2-sphere, however, they are typically introduced by means of the spherical Laplace-Beltrami operator, see e.g. [63, Chapter 6.2] and Sec. 7.2 for the scalar counterpart. For a thorough treatment of Sobolev spaces on Riemannian manifolds we refer to the books [33, 80].

Spherical Harmonics

We denote by Harm_n the space of homogeneous harmonic polynomials of degree $n \in \mathbb{N}_0$ with their domain restricted to \mathcal{S}^2 . Its dimension is

$$\dim(\text{Harm}_n) = 2n + 1.$$

An element $\tilde{Y}_n \in \text{Harm}_n$, $n \in \mathbb{N}_0$, is called a (*scalar*) *spherical harmonic*. It is an infinitely often differentiable eigenfunction of the Laplace-Beltrami operator $\Delta_{\mathcal{S}^2}$, defined in (7.10), with corresponding eigenvalue

$$\lambda_n = n(n + 1).$$

We refer to Theorem 5.6 and Lemma 5.8 in [63, Sec. 5.1] for detailed proofs of the previous statements.

The set

$$\left\{ \tilde{Y}_{nj} : n \in \mathbb{N}_0, j = 1, \dots, 2n + 1 \right\} \quad (7.26)$$

is a complete orthonormal system of $L^2(\mathcal{S}^2)$ with respect to the inner product $\langle \cdot, \cdot \rangle_{L^2(\mathcal{S}^2)}$ on \mathcal{S}^2 . In further consequence, for a function $\tilde{f} \in L^2(\mathcal{S}^2)$, we have the Fourier series representation

$$\tilde{f} = \sum_{n=0}^{\infty} \sum_{j=1}^{2n+1} \langle \tilde{f}, \tilde{Y}_{nj} \rangle_{L^2(\mathcal{S}^2)} \tilde{Y}_{nj},$$

Again, we refer to [63, Sec. 5.1] for the proofs, in particular to Theorem 5.25. In the present article we employ *fully normalised spherical harmonics*. For the explicit construction see [63, Sec. 5.2].

Moreover, the norm of $L^2(\mathcal{S}^2)$ is readily stated in terms of the coefficients in the above expansion via Parseval's identity

$$\|\tilde{f}\|_{L^2(\mathcal{S}^2)}^2 = \sum_{n,j} \langle \tilde{f}, \tilde{Y}_{nj} \rangle_{L^2(\mathcal{S}^2)}^2.$$

For an arbitrary real number $s \in \mathbb{R}$, we define the Sobolev space $H^s(\mathcal{S}^2)$ as the completion of all $C^\infty(\mathcal{S}^2)$ functions with respect to the norm

$$\|\tilde{f}\|_{H^s(\mathcal{S}^2)}^2 := \|(\Delta_{\mathcal{S}^2} + 1)^{s/2} \tilde{f}\|_{L^2(\mathcal{S}^2)}^2 = \sum_{n,j} (\lambda_n + 1)^s \langle \tilde{f}, \tilde{Y}_{nj} \rangle_{L^2(\mathcal{S}^2)}^2.$$

We stress that, by (7.10), $\Delta_{\mathcal{S}^2} \tilde{f} = -\Delta_{\mathbb{R}^3} \tilde{f}$ and we have $\lambda_n \geq 0$ for all $n \in \mathbb{N}_0$ yielding a sound definition. Accordingly, for $s \in \mathbb{R}$, we define the H^s seminorm of order s by

$$|\tilde{f}|_{H^s(\mathcal{S}^2)}^2 := \|\Delta_{\mathcal{S}^2}^{s/2} \tilde{f}\|_{L^2(\mathcal{S}^2)}^2 = \sum_{n,j} \lambda_n^s \langle \tilde{f}, \tilde{Y}_{nj} \rangle_{L^2(\mathcal{S}^2)}^2. \quad (7.27)$$

Now that the space $L^2(\mathcal{S}^2)$ is endowed with a basis, we can proceed to define an orthonormal system for square integrable tangent vector fields on the sphere. This will immediately allow us to treat vector-valued problems consistently.

Let $\tilde{Y}_n \in \text{Harm}_n$ be a scalar spherical harmonic of degree $n \in \mathbb{N}_0$. Any vector field $\tilde{\mathbf{y}} : \mathcal{S}^2 \rightarrow \mathbb{R}^3$ that can be written in the form $\tilde{\mathbf{y}} = \tilde{\mathbf{y}}_n^{(i)}$, where

$$\begin{aligned} \tilde{\mathbf{y}}_n^{(1)} &:= \tilde{Y}_n \tilde{\mathbf{N}}, \\ \tilde{\mathbf{y}}_n^{(2)} &:= \nabla_{\mathcal{S}^2} \tilde{Y}_n, \\ \tilde{\mathbf{y}}_n^{(3)} &:= \nabla_{\mathcal{S}^2} \tilde{Y}_n \times \tilde{\mathbf{N}}, \end{aligned}$$

is called a *vector spherical harmonic* of degree n and type i , cf. [28, Definition 5.2]. Recall that $\tilde{\mathbf{N}}$ is the outward unit normal to \mathcal{S}^2 .

By definition, $\tilde{\mathbf{y}}_n^{(1)}$ is a normal field whereas $\tilde{\mathbf{y}}_n^{(2)}$ and $\tilde{\mathbf{y}}_n^{(3)}$ are tangent vector fields. Consequently, the latter are called *tangential vector spherical harmonics*. Note that, by means of the Hairy-Ball Theorem, no tangential vector spherical harmonics of degree zero exist.

In further consequence, let us denote by $L^2(\mathcal{S}^2, T\mathcal{S}^2)$ the space of square integrable tangent vector fields on \mathcal{S}^2 equipped with the inner product

$$\langle \tilde{\mathbf{u}}, \tilde{\mathbf{v}} \rangle_{L^2(\mathcal{S}^2, T\mathcal{S}^2)} = \int_{\mathcal{S}^2} \tilde{\mathbf{u}} \cdot \tilde{\mathbf{v}} \, d\mathcal{S}^2.$$

Here, $d\mathcal{S}^2$ denotes the usual spherical surface measure, see also Lemma 4.

Since (7.26) is an orthonormal set for $L^2(\mathcal{S}^2)$, the set

$$\left\{ \tilde{\mathbf{y}}_{nj}^{(i)} : n \in \mathbb{N}, j = 1, \dots, 2n+1, i = 2, 3 \right\}, \quad (7.28)$$

is an orthonormal system for $L^2(\mathcal{S}^2, T\mathcal{S}^2)$, where we have defined

$$\begin{aligned} \tilde{\mathbf{y}}_{nj}^{(2)} &= \lambda_n^{-1/2} \nabla_{\mathcal{S}^2} \tilde{Y}_{nj}, \\ \tilde{\mathbf{y}}_{nj}^{(3)} &= \lambda_n^{-1/2} \nabla_{\mathcal{S}^2} \tilde{Y}_{nj} \times \tilde{\mathbf{N}}, \end{aligned} \quad (7.29)$$

for orthonormalisation purpose, see [28, Sec. 5.2]. Thus, every vector field $\tilde{\mathbf{v}} \in L^2(\mathcal{S}^2, T\mathcal{S}^2)$ can be written uniquely as

$$\tilde{\mathbf{v}} = \sum_{i=2}^3 \sum_{n=1}^{\infty} \sum_{j=1}^{2n+1} \langle \tilde{\mathbf{v}}, \tilde{\mathbf{y}}_{nj}^{(i)} \rangle_{L^2(\mathcal{S}^2, T\mathcal{S}^2)} \tilde{\mathbf{y}}_{nj}^{(i)}.$$

We refer to the books [28, 63] for further details on the subject. Table 7.1 contains a summary of notation used in the sequel.

Ω	coordinate domain
I	time interval
\mathcal{S}^2	2-sphere
\mathcal{M}	family of sphere-like surfaces \mathcal{M}_t
$T_x \mathcal{S}^2$	tangent plane at $x \in \mathcal{S}^2$
$T_y \mathcal{M}_t$	tangent plane at $y \in \mathcal{M}_t$
$\tilde{\mathbf{N}}, \hat{\mathbf{N}}$	outward unit normals to \mathcal{S}^2 and \mathcal{M}
\mathbf{x}, \mathbf{y}	parametrisations of \mathcal{S}^2 and \mathcal{M}
$D\mathbf{x}, D\mathbf{y}$	gradient matrix of \mathbf{x} and \mathbf{y}
$\{\partial_1 \mathbf{x}, \partial_2 \mathbf{x}\}$	basis for $T\mathcal{S}^2$
$\{\partial_1 \mathbf{y}, \partial_2 \mathbf{y}\}$	basis for $T\mathcal{M}$
$\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2\}$	orthonormal basis for $T\mathcal{M}_t$
$\hat{\mathbf{V}}$	surface velocity of \mathcal{M}
$\tilde{\phi}, D\tilde{\phi}$	smooth map from \mathcal{S}^2 to \mathcal{M} and its differential
\tilde{f}, \hat{f}, f	scalar function on \mathcal{S}^2 , \mathcal{M} , and their coordinate version
$\nabla_{\mathcal{S}^2} \tilde{f}, \nabla_{\mathcal{M}} \hat{f}$	surface gradient on \mathcal{S}^2 and \mathcal{M}_t
$\tilde{\mathbf{v}}, \hat{\mathbf{v}}, \mathbf{v}$	tangent vector fields on \mathcal{S}^2 , \mathcal{M} , and their coordinate version
$\nabla_{\hat{\mathbf{u}}} \hat{\mathbf{v}}$	covariant derivative of $\hat{\mathbf{v}}$ along direction $\hat{\mathbf{u}}$ on \mathcal{M}_t
\tilde{Y}_{nj}	scalar spherical harmonic of degree n and order j
$\tilde{\mathbf{y}}_{nj}^{(i)}$	vector spherical harmonic of degree n , order j , and type i
$\hat{\mathbf{y}}_{nj}^{(i)}$	pushforward of $\tilde{\mathbf{y}}_{nj}^{(i)}$ via the differential $D\tilde{\phi}$

Table 7.1: Summary of notation used throughout the paper.

7.3 Optical Flow on Evolving Surfaces

Generalised Optical Flow

Optical flow models are typically based on the assumption of *constant brightness*. Given a sequence of (planar) images

$$f : I \times \Omega \subset \mathbb{R} \times \mathbb{R}^2 \rightarrow \mathbb{R}$$

such that $f \in C^1(I \times \Omega)$, it assumes that the intensity $f(t, \gamma(t, \xi))$ stays constant over time when moving along a trajectory $\gamma(\cdot, \xi) : I \rightarrow \Omega$ starting at $\xi \in \Omega$. In other words, in the planar setting, we have

$$\frac{d}{dt} f(t, \gamma(t, \xi)) = \partial_t f + \nabla_{\mathbb{R}^2} f \cdot \partial_t \gamma = 0,$$

which is termed *optical flow equation* and must hold for all $\xi \in \Omega$ and all $t \in I$. For the sake of consistency, we denote by ∂_t the partial and by d/dt the total derivative with respect to time.

It is possible to generalise the idea to a non-Euclidean setting where the image lives on a, potentially moving, manifold. To this end, let us be given an evolving surface

$$\mathcal{M} := \bigcup_{t \in I} (\{t\} \times \mathcal{M}_t) \subset \mathbb{R}^4$$

specified by a parametrisation $\mathbf{y} : I \times \Omega \rightarrow \mathbb{R}^3$ as in (7.4) together with a function \hat{f} , its domain being \mathcal{M} . For a time $t \in I$,

$$\hat{f}(t, \cdot) : \mathcal{M}_t \rightarrow \mathbb{R}$$

is then an image on the surface. Adapting the above idea of constant brightness to the new setting requires that, along a smooth trajectory $\gamma(\cdot, x) : t \mapsto \gamma(t, x) \in \mathcal{M}_t$ that starts at $x \in \mathcal{M}_0$ and always stays on the surface, we must have

$$\hat{f}(t, \gamma(t, x)) = \hat{f}(0, x). \quad (7.30)$$

However, in order to proceed as above one needs to define a meaningful derivative with respect to time.

One possibility, which is pursued in [50, 52], is to consider derivatives along trajectories following the moving surface. Let \mathbf{y} be as above and let $\partial_t \mathbf{y} = \hat{\mathbf{V}}$ be the surface velocity, its domain being $\bigcup_{t \in I} (\{t\} \times \mathcal{M}_t) \subset \mathbb{R}^4$. We emphasise that $\hat{\mathbf{V}}$ is in general not tangent to \mathcal{M}_t , $t \in I$, and hence in our notation denoted by a boldface capital letter. Then,

$$d_t^{\hat{\mathbf{V}}} \hat{f}(t_0, x_0) := \left. \frac{d}{dt} \hat{f}(t, \mathbf{y}(t, \xi)) \right|_{t=t_0} \quad (7.31)$$

is the time derivative of \hat{f} at $x_0 = \mathbf{y}(t_0, \xi)$ along the parametrisation $\mathbf{y}(\cdot, \xi)$. As a consequence, one can deduce that

$$d_t^{\hat{\mathbf{V}}} \hat{f} = d_t^{\hat{\mathbf{N}}} \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{V}}$$

holds, where $d_t^{\hat{\mathbf{N}}} \hat{f}(t_0, x_0)$ is the time derivative of \hat{f} in normal direction. It is defined analogously to (7.31) albeit following a trajectory $\psi_{\hat{\mathbf{N}}}$ through $x_0 \in \mathcal{M}_{t_0}$ for which $\partial_t \psi_{\hat{\mathbf{N}}}(t_0, x_0)$ is orthogonal to $T_{x_0} \mathcal{M}_{t_0}$.

From that one can immediately formulate the above idea of constant brightness (7.30) along γ . To this end, we define by $\hat{\mathbf{M}} := \partial_t \gamma$ the velocity of a point moving along the trajectory γ and demand that

$$d_t^{\hat{\mathbf{M}}} \hat{f} = d_t^{\hat{\mathbf{N}}} \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{M}} = 0 \quad (7.32)$$

must hold. Equation (7.32) is a *generalised optical flow equation*. In Fig. 7.4 we sketch the various trajectories through the evolving surface and their corresponding velocities.

Since we are, however, interested in a coordinate representation of γ , we define a family of trajectories $\beta : I \times \Omega \rightarrow \Omega$ such that

$$\gamma(t, \mathbf{y}(0, \xi)) = \mathbf{y}(t, \beta(t, \xi))$$

holds for all $t \in I$ and all $\xi \in \Omega$. In other words, we want the composition of β with \mathbf{y} , and γ to coincide. By taking the total derivative d/dt on both sides of the above equation we get

$$\partial_t \gamma = \partial_t \mathbf{y} + \partial_t \beta^i \partial_i \mathbf{y}.$$

Let us denote $\hat{\mathbf{v}} := \partial_t \beta^i \partial_i \mathbf{y}$ and recall that $\partial_t \mathbf{y} = \hat{\mathbf{V}}$ is the surface velocity. The above relation states that the total velocity $\hat{\mathbf{M}} = \partial_t \gamma$ along a level line of constant intensity is

$$\hat{\mathbf{M}} = \hat{\mathbf{V}} + \hat{\mathbf{v}} \quad (7.33)$$

and $\hat{\mathbf{v}}$ is a tangential velocity relative to the prescribed surface velocity $\hat{\mathbf{V}}$.




Figure 7.4: Illustration of trajectories through the evolving surface. Their corresponding velocities are shown in grey.

Solving the generalised optical flow equation (7.32), however, is inconvenient as $\psi_{\hat{\mathbf{N}}}$ and, in further consequence, $d_t^{\hat{\mathbf{N}}}$ is unknown or hard to estimate. Nevertheless, one can relate (7.32) and (7.33), as shown in [52, Lemma 2], and arrive at the *parametrised optical flow equation*

$$d_t^{\hat{\mathbf{V}}} \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{v}} = 0. \quad (7.34)$$

Solving for the optical flow then means finding a (time-varying) vector field $\hat{\mathbf{v}}$ that is tangent to the surface at all times and satisfies the above equation at every point $x \in \mathcal{M}$ on the moving surface.

Let us conclude this subsection with a remark that, in general, there exist infinitely many parametrisations \mathbf{y} for a given evolving surface. The actual surface velocity however might be unknown or cannot be estimated from the data, as it is the case in this work. As a remedy we impose a surface velocity by choosing a natural surface parametrisation \mathbf{y} . For a further discussion of this matter we refer to [12].

Moreover, we again stress that the sought tangent vector field $\hat{\mathbf{v}}$ depends on the chosen \mathbf{y} , or equivalently on $\hat{\mathbf{V}}$, and should be interpreted with care. The actual trajectories γ though can be reconstructed by finding the integral curves of (7.33). For this precise approach we point the reader to [52].

Variational Formulation

The optical flow equation (7.34) derived above is underdetermined and, in general, a unique solution is not ensured. A common technique to deal with non-uniqueness is Tikhonov regularisation, where one finds a minimiser of

$$\|d_t^{\hat{\mathbf{V}}} \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{v}}\|_{L^2(\mathcal{M})}^2 + \alpha \mathcal{R}(\hat{\mathbf{v}}).$$

Here, $\mathcal{R}(\hat{\mathbf{v}})$ is a regularisation functional and $\alpha > 0$ a regularisation parameter, balancing the two terms. The first term is typically referred to as *data term* whereas the second is called *smoothness term*. The latter enforces uniqueness and incorporates prior knowledge about favoured solutions.

A common choice for $\mathcal{R}(\hat{\mathbf{v}})$ is the squared H^1 Sobolev seminorm, involving first derivatives with respect to space and time. It favours spatial as well

as temporal regularity and is of particular interest when trying to estimate trajectories of objects, albeit computationally more demanding. See, for example [52, 86] and [12].

Alternatively, one can omit temporal regularisation leading to a regulariser of the form

$$\mathcal{R}(\hat{\mathbf{v}}) = \int_I \|\hat{\mathbf{v}}(t, \cdot)\|_{H^1(\mathcal{M}_t, T\mathcal{M}_t)}^2 dt,$$

which is equivalent to solving for each time instant separately. It resembles the original formulation in [39] and its extension to 2-manifolds [58]. In the present article we follow this approach and attempt in finding the unique minimiser $\hat{\mathbf{v}} \in T\mathcal{M}_t$ of the energy

$$\mathcal{E}_\alpha(\hat{\mathbf{v}}) := \|d_t^{\hat{\mathbf{v}}} \hat{f} + \nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{v}}\|_{L^2(\mathcal{M}_t)}^2 + \alpha \|\hat{\mathbf{v}}\|_{H^1(\mathcal{M}_t, T\mathcal{M}_t)}^2 \quad (7.35)$$

for each time instant $t \in I$ separately. Superimposing temporal regularisation however is straightforward, see [52, Sec. 2.2.2], but not considered here.

7.4 Numerical Solution

Finite-dimensional Projection

For the subsequent discussion we let $t \in I$ be arbitrary but fixed and assume to be given a parametrisation $\mathbf{y}(t, \cdot) : \Omega \rightarrow \mathcal{M}_t$ as defined in (7.4). We defer the question of how to find it to Sec. 7.4. Moreover, for notational convenience, we relabel the set of tangential vector spherical harmonics (7.28) using a single index letter $p \in \mathbb{N}$. For instance, for the expansion of a tangent vector field on \mathcal{S}^2 we simply write $\hat{\mathbf{u}} = \sum_p u_p \hat{\mathbf{y}}_p$, where $u_p \in \mathbb{R}$ are the coefficients.

We intend to approximate the solution of the problem

$$\min_{\hat{\mathbf{v}} \in H^1(\mathcal{M}_t, T\mathcal{M}_t)} \mathcal{E}_\alpha(\hat{\mathbf{v}})$$

in a finite-dimensional subspace $\mathcal{U} \subset H^1(\mathcal{M}_t, T\mathcal{M}_t)$, where \mathcal{E}_α is defined in (7.35). We define this space as

$$\mathcal{U} = \text{span}\{\hat{\mathbf{y}}_p : p \in J_{\mathcal{U}}\},$$

where $J_{\mathcal{U}} \subset \mathbb{N}$ is a finite index set and $\hat{\mathbf{y}}_p$ is the pushforward of a particular vector spherical harmonic $\hat{\mathbf{y}}_p$ via the differential $D\phi$, see (7.20). Figure 7.5 gives a descriptive view of the relation between the introduced spaces and tangent vector fields.

The sought vector field is then uniquely expanded as

$$\hat{\mathbf{v}} = \sum_{p \in J_{\mathcal{U}}} v_p \hat{\mathbf{y}}_p, \quad (7.36)$$

with $v_p \in \mathbb{R}$, $p \in J_{\mathcal{U}}$, being the coefficients. Minimisation of functional (7.35) results in a finite-dimensional optimisation problem over $\mathbb{R}^{|J_{\mathcal{U}}|}$. Plugging ansatz (7.36) into (7.35) and writing out definition (7.25) of the Sobolev $H^1(\mathcal{M}_t, T\mathcal{M}_t)$ norm gives

$$\mathcal{E}_\alpha(\hat{\mathbf{v}}) = \int_{\mathcal{M}_t} \left((d_t^{\hat{\mathbf{v}}} \hat{f} + \sum_{p \in J_{\mathcal{U}}} v_p (\nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{y}}_p))^2 + \alpha \left\| \sum_{p \in J_{\mathcal{U}}} v_p \nabla \hat{\mathbf{y}}_p \right\|_2^2 \right) d\mathcal{M}_t. \quad (7.37)$$




Figure 7.5: Commutative diagram relating spaces Ω , \mathcal{S}^2 , and \mathcal{M}_t , and tangent vector fields. We highlight that \mathbf{y}_p is the coordinate representation, see Sec. 7.2, of a particular tangential vector spherical harmonic $\tilde{\mathbf{y}}_p$ and $\hat{\mathbf{y}}_p$ is its uniquely identified tangent vector field on \mathcal{M}_t .

By applying the definition of the Hilbert-Schmidt norm (7.24), using linearity of the covariant derivative $\nabla_{\hat{\mathbf{u}}} \hat{\mathbf{v}}$ with respect to $\hat{\mathbf{v}}$, and the definition of the norm of \mathbb{R}^3 we obtain the representation

$$\begin{aligned}
 \|\nabla \sum_{p \in J_{\mathcal{U}}} v_p \hat{\mathbf{y}}_p\|_2^2 &= \sum_{i=1}^2 \left\| \sum_{p \in J_{\mathcal{U}}} v_p \nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{y}}_p \right\|^2 \\
 &= \sum_{i=1}^2 \left(\sum_{p \in J_{\mathcal{U}}} v_p \nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{y}}_p \cdot \sum_{q \in J_{\mathcal{U}}} v_q \nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{y}}_q \right) \\
 &= \sum_{i=1}^2 \sum_{p, q \in J_{\mathcal{U}}} v_p v_q (\nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{y}}_p \cdot \nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{y}}_q)
 \end{aligned}$$

for the regularisation term.

The optimality conditions for the discrete minimisation problem (7.37) are obtained by taking $\partial \mathcal{E}_{\alpha} / \partial v_p = 0$, for all $p \in J_{\mathcal{U}}$, and are given by

$$\begin{aligned}
 \sum_{q \in J_{\mathcal{U}}} v_q \int_{\mathcal{M}_t} \left((\nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{y}}_p) (\nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{y}}_q) + \alpha \sum_{i=1}^2 (\nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{y}}_p \cdot \nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{y}}_q) \right) d\mathcal{M}_t \\
 = - \int_{\mathcal{M}_t} d_t^{\hat{\mathbf{v}}} \hat{f} (\nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{y}}_p) d\mathcal{M}_t, \quad p \in J_{\mathcal{U}}.
 \end{aligned} \tag{7.38}$$

In matrix form they read

$$(A + \alpha D)v = b,$$

where $v = (v_1, \dots, v_{|J_{\mathcal{U}}|})^{\top} \in \mathbb{R}^{|J_{\mathcal{U}}|}$ is the vector of unknowns. The entries of the matrix $A = (a_{pq})_{pq}$ are

$$a_{pq} = \int_{\mathcal{M}_t} (\nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{y}}_p) (\nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{y}}_q) d\mathcal{M}_t,$$

the entries of the matrix $D = (d_{pq})_{pq}$ associated with the regularisation term are given by

$$d_{pq} = \int_{\mathcal{M}_t} \sum_{i=1}^2 (\nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{y}}_p \cdot \nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{y}}_q) d\mathcal{M}_t,$$

and the entries of the vector $b = (b_p)_p$ are

$$b_p = - \int_{\mathcal{M}_t} d_t^{\hat{\mathbf{y}}} \hat{f} (\nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{y}}_p) d\mathcal{M}_t.$$

Rewriting the Optimality Conditions

Even though directly solving the derived optimality conditions (7.38) is perfectly legitimate, we take a different approach. The goal of this section is to rewrite the optimality conditions in terms of quantities defined on the 2-sphere, thereby allowing a more general treatment. On the one hand, we want to deal with all surfaces \mathcal{M}_t for all $t \in I$ in a unified manner and, on the other hand, we aim at evaluating (7.38) numerically on the (approximated) sphere without having to deal with multiple charts, see e.g. [22].

The following is a straight-forward generalisation of [52, Lemma 2].

Lemma 5. *Consider time $t \in I$ arbitrary but fixed. Let $\tilde{\mathbf{v}} = v^i \partial_i \mathbf{x}$ and $\hat{\mathbf{v}} = v^i \partial_i \mathbf{y}$ be two tangent vector fields on \mathcal{S}^2 and \mathcal{M}_t , respectively, such that they are related via the differential (7.20). Then, the parametrised optical flow equation (7.34) is equivalent to*

$$\partial_t \tilde{f} + \nabla_{\mathcal{S}^2} \tilde{f} \cdot \tilde{\mathbf{v}} = 0.$$

Proof. According to the definitions (7.31) and (7.5), we have

$$\begin{aligned} d_t^{\hat{\mathbf{y}}} \hat{f}(t, \mathbf{y}(t, \xi)) &= \frac{d}{dt} \hat{f}(t, \mathbf{y}(t, \xi)) \\ &= \frac{d}{dt} \tilde{f}(t, \mathbf{x}(\xi)) \\ &= \partial_t \tilde{f}(t, \mathbf{x}(\xi)) \end{aligned}$$

and it remains to show the identity

$$\nabla_{\mathcal{M}} \hat{f} \cdot \hat{\mathbf{v}} = \nabla_{\mathcal{S}^2} \tilde{f} \cdot \tilde{\mathbf{v}},$$

where we have omitted the arguments $(t, \mathbf{y}(t, \xi))$ on the left and $(t, \mathbf{x}(\xi))$ on the right hand side, respectively. It follows directly from the coordinate representation of the directional derivatives (7.15) and (7.16). \square

In order to give coordinate expressions for the terms in (7.38) arising from the regularisation term we locally choose an orthonormal frame $\{\hat{\mathbf{e}}_1(t, \xi), \hat{\mathbf{e}}_2(t, \xi)\}$ of the tangent space, see (7.13). As a consequence, the sought tangent vector field $\hat{\mathbf{v}}$ can be written as

$$\hat{\mathbf{v}} = w^i \hat{\mathbf{e}}_i \tag{7.39}$$

for some components $(w^1, w^2)^\top$. The reason for expressing the unknown in an orthonormal frame, rather than the coordinate frame, is to simplify matters with regard to the Hilbert-Schmidt norm (7.24) of the covariant derivative.

However, the chosen Galerkin method expands the unknown $\hat{\mathbf{v}}$ in terms of the pushforward of vector fields which are defined on the 2-sphere, cf. (7.36). We necessarily need to establish the relation between the intended form (7.39) and the expression in terms of the coordinate frame.

Lemma 6. *Again, let $t \in I$ be arbitrary but fixed and let $\hat{\mathbf{u}} = u^i \partial_i \mathbf{x}$ be a tangent vector field on \mathcal{S}^2 . Then, for a tangent vector field $\hat{\mathbf{v}} = w^i \hat{\mathbf{e}}_i$ on \mathcal{M}_t , we have $\hat{\mathbf{v}} = D\tilde{\phi}(\hat{\mathbf{u}})$ if and only if $w^i = (\alpha^{-1})_\ell^i u^\ell$.*

Proof. (\Leftarrow) First, note that $\alpha_i^j (\alpha^{-1})_\ell^i = \delta_{j\ell}$. Expanding $\hat{\mathbf{v}}$ gives

$$\hat{\mathbf{v}} = w^i \hat{\mathbf{e}}_i = w^i \alpha_i^j \partial_j \mathbf{y} = (\alpha^{-1})_\ell^i u^\ell \alpha_i^j \partial_j \mathbf{y} = u^j \partial_j \mathbf{y} = D\tilde{\phi}(\hat{\mathbf{u}}),$$

where we have used (7.21), cf. also Fig. 7.5.

(\Rightarrow) Suppose $\hat{\mathbf{v}} = D\tilde{\phi}(\hat{\mathbf{u}})$. Let us take the inner product with $\hat{\mathbf{e}}_i$ on both sides. For the left hand side we have

$$\hat{\mathbf{v}} \cdot \hat{\mathbf{e}}_i = w^j \hat{\mathbf{e}}_j \cdot \hat{\mathbf{e}}_i = w^j \delta_{ji} = w^i.$$

For the right hand side we first observe that, by inversion of the matrix α in (7.13), it holds that $\partial_j \mathbf{y} = (\alpha^{-1})_j^\ell \hat{\mathbf{e}}_\ell$. Then,

$$\begin{aligned} D\tilde{\phi}(\hat{\mathbf{u}}) \cdot \hat{\mathbf{e}}_i &= u^\ell \partial_\ell \mathbf{y} \cdot \hat{\mathbf{e}}_i \\ &= u^\ell (\alpha^{-1})_\ell^j \hat{\mathbf{e}}_j \cdot \hat{\mathbf{e}}_i \\ &= u^\ell (\alpha^{-1})_\ell^j \delta_{ji} \\ &= u^\ell (\alpha^{-1})_\ell^i \end{aligned}$$

and we conclude that $w^i = u^\ell (\alpha^{-1})_\ell^i$ as required. \square

With the above relation at hand we obtain the following form.

Lemma 7. *Let $t \in I$ and let $\hat{\mathbf{u}} = u^i \partial_i \mathbf{y}$ and $\hat{\mathbf{v}} = v^i \partial_i \mathbf{y}$ be two tangent vector fields on \mathcal{M}_t . Then, we have*

$$\nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{u}} \cdot \nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{v}} = \sum_{j=1}^2 D_i u^j D_i v^j,$$

where

$$D_i u^j := \alpha_i^k \partial_k ((\alpha^{-1})_\ell^j u^\ell) + (\alpha^{-1})_\ell^k u^\ell \hat{\Gamma}_{ik}^j, \quad i, j = \{1, 2\},$$

and $D_i v^j$ are defined accordingly.

$\hat{\Gamma}_{ik}^j$ denote the Christoffel symbols with regard to the orthonormal frame $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2\}$ and are defined as

$$\nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{e}}_k = \hat{\Gamma}_{ik}^j \hat{\mathbf{e}}_j. \quad (7.40)$$

We refer to [52, Lemma 3] for their derivation.

Proof. First let us show that, for $\hat{\mathbf{u}} = w^j \hat{\mathbf{e}}_j$ as in (7.39), it holds that

$$\nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{u}} = D_i w^j \hat{\mathbf{e}}_j.$$

By the product rule for the covariant derivative (7.23),

$$\nabla_{\hat{\mathbf{e}}_i} w^j \hat{\mathbf{e}}_j = \hat{\mathbf{e}}_j \nabla_{\hat{\mathbf{e}}_i} w^j + w^j \nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{e}}_j. \quad (7.41)$$

Consider the first term of the sum and let $\hat{\mathbf{e}}_i$ be represented in the coordinate basis as in (7.13). Then,

$$\nabla_{\hat{\mathbf{e}}_i} w^j = \nabla_{\alpha_i^k \partial_k \mathbf{y}} w^j.$$

Linearity of the lower argument of the covariant derivative with respect to $C^\infty(\mathcal{M}_t)$ functions, cf. (7.16), yields

$$\nabla_{\alpha_i^k \partial_k \mathbf{y}} w^j = \alpha_i^k \nabla_{\partial_k \mathbf{y}} w^j$$

and by realising that $\nabla_{\partial_k \mathbf{y}} w^j$ is just the directional derivative (7.16) along $\partial_k \mathbf{y}$ we obtain

$$\alpha_i^k \nabla_{\partial_k \mathbf{y}} w^j = \alpha_i^k \partial_k w^j.$$

Moreover, in the second term of the sum in (7.41) we use definition (7.40). Thus, by summing up all terms in (7.41) we obtain

$$\nabla_{\hat{\mathbf{e}}_i} w^j \hat{\mathbf{e}}_j = (\alpha_i^k \partial_k w^j + w^j \hat{\Gamma}_{ik}^j) \hat{\mathbf{e}}_j.$$

Applying the previous lemma gives coefficients $D_i u^j$ and $D_i v^j$ in the intended form. Finally, it remains to observe that

$$\begin{aligned} \nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{u}} \cdot \nabla_{\hat{\mathbf{e}}_i} \hat{\mathbf{v}} &= D_i u^j \hat{\mathbf{e}}_j \cdot D_i v^j \hat{\mathbf{e}}_j \\ &= \sum_{j=1}^2 D_i u^j D_i v^j, \end{aligned}$$

since by definition $\hat{\mathbf{e}}_i \cdot \hat{\mathbf{e}}_j = \delta_{ij}$. \square

Finally, by combining Lemmas 4, 5, and 7 we are able to express the optimality conditions (7.38) in terms of integrals on the 2-sphere. Thus, we arrive at the optimality conditions

$$\begin{aligned} \sum_{q \in J_{\mathcal{U}}} v_q \int_{S^2} \left((\nabla_{S^2} \tilde{f} \cdot \tilde{\mathbf{y}}_p) (\nabla_{S^2} \tilde{f} \cdot \tilde{\mathbf{y}}_q) + \alpha \sum_{i,j=1}^2 D_i y_p^j D_i y_q^j \right) \tilde{\rho} \sqrt{\|\nabla_{S^2} \tilde{\rho}\|^2 + \tilde{\rho}^2} dS^2 \\ = - \int_{S^2} \partial_t \tilde{f} (\nabla_{S^2} \tilde{f} \cdot \tilde{\mathbf{y}}_p) \tilde{\rho} \sqrt{\|\nabla_{S^2} \tilde{\rho}\|^2 + \tilde{\rho}^2} dS^2, \quad p \in J_{\mathcal{U}}, \end{aligned} \quad (7.42)$$

where $\tilde{\rho} \sqrt{\|\nabla_{S^2} \tilde{\rho}\|^2 + \tilde{\rho}^2}$ arises from the Jacobian (7.19), see also Lemma 4.

The entries of the matrices A , D and of the vector b , respectively, are then given by

$$a_{pq} = \int_{S^2} (\nabla_{S^2} \tilde{f} \cdot \tilde{\mathbf{y}}_p) (\nabla_{S^2} \tilde{f} \cdot \tilde{\mathbf{y}}_q) \tilde{\rho} \sqrt{\|\nabla_{S^2} \tilde{\rho}\|^2 + \tilde{\rho}^2} dS^2, \quad (7.43)$$

$$d_{pq} = \int_{\mathcal{S}^2} \sum_{i,j=1}^2 D_i y_p^j D_i y_q^j \tilde{\rho} \sqrt{\|\nabla_{\mathcal{S}^2} \tilde{\rho}\|^2 + \tilde{\rho}^2} d\mathcal{S}^2, \quad (7.44)$$

and

$$b_p = - \int_{\mathcal{S}^2} \partial_t \tilde{f} (\nabla_{\mathcal{S}^2} \tilde{f} \cdot \tilde{\mathbf{y}}_p) \tilde{\rho} \sqrt{\|\nabla_{\mathcal{S}^2} \tilde{\rho}\|^2 + \tilde{\rho}^2} d\mathcal{S}^2. \quad (7.45)$$

Surface Parametrisation

In order to actually compute the above optimality conditions it remains to determine the radius $\tilde{\rho} : I \times \mathcal{S}^2 \rightarrow (0, \infty)$ in the presumed parametrisation (7.4). Again, we continue the discussion for one particular but fixed time $t \in I$ and drop the argument whenever convenient.

Estimating $\tilde{\rho}(t, \cdot) : \mathcal{S}^2 \rightarrow (0, \infty)$ is closely related to surface interpolation from scattered data. Given noisy data $\tilde{\rho}^\delta$ and a parameter $\beta > 0$, it amounts to finding the unique minimiser of the functional

$$\mathcal{F}_\beta(\tilde{\rho}) := \|\tilde{\rho} - \tilde{\rho}^\delta\|_{L^2(\mathcal{S}^2)}^2 + \beta \|\tilde{\rho}\|_{H^s(\mathcal{S}^2)}^2, \quad (7.46)$$

where $s > 0$ is a sufficiently large real number, cf. definition (7.27). The first term penalises deviation from the observed data whereas the second term enforces spatial regularity of the solution.

In practice, however, $N > 0$ evaluations $\{\tilde{\rho}^\delta(x_i) : x_i \in \mathcal{S}^2\}_{i=1}^N$ are given at pairwise distinct points on the 2-sphere. In our particular application these correspond to taking the norm in \mathbb{R}^3 of pairwise distinct sampling points lying on the sphere-like surface \mathcal{M}_t :

$$\tilde{\rho}^\delta(\bar{x}_i) = \|x_i\|, \quad x_i \in \mathbb{R}^3 \setminus \{0\}, \quad i = 1, \dots, N, \quad (7.47)$$

where $\bar{x}_i = x_i/\|x_i\|$ is the radial projection onto \mathcal{S}^2 . We again point the reader to Fig. 7.3.

Furthermore, before turning to the numerical solution of (7.46), let us briefly discuss the regularity requirements. In [12], the authors demand twice continuous differentiability for both the manifold \mathcal{M}_t and the map $\mathbf{y}(t, \cdot)$ to obtain well-posedness of the optical flow problem. By definition of the parametrisation (7.4) we require that $\tilde{\rho}(t, \cdot) \in C^2(\mathcal{S}^2)$. As a consequence of Theorem 2.7 in [34, Chapter 2.6] regarding Sobolev embeddings, the space $H^s(\mathcal{S}^2)$ for $s > 3$ is the appropriate choice, i.e. $H^s(\mathcal{S}^2) \subset C^2(\mathcal{S}^2)$.

Numerically, we approximate the solution of the problem

$$\min_{\tilde{\rho} \in H^s(\mathcal{S}^2)} \mathcal{F}_\beta(\tilde{\rho})$$

by considering a finite-dimensional subspace $\mathcal{Q} \subset H^s(\mathcal{S}^2)$ and point evaluations (7.47). In contrast to above, the space

$$\mathcal{Q} = \text{span}\{\tilde{Y}_p : p \in J_{\mathcal{Q}}\},$$

where $J_{\mathcal{Q}} \subset \mathbb{N}_0$ again is an index set, is spanned by scalar spherical harmonics. The sought function is expanded as

$$\tilde{\rho} = \sum_{p \in J_{\mathcal{Q}}} \rho_p \tilde{Y}_p,$$

where the unknowns are the coefficients $\rho_p \in \mathbb{R}$, for $p \in J_Q$. Plugging into (7.46), applying definition (7.27), and taking $\partial\mathcal{F}/\partial\rho_p$, for all $p \in J_Q$, gives the optimality conditions

$$\sum_{q \in J_Q} \rho_q \left(\sum_{i=1}^N \tilde{Y}_p(\bar{x}_i) \tilde{Y}_q(\bar{x}_i) \right) + \beta \lambda_p^s \rho_p = \sum_{i=1}^N \|x_i\| \tilde{Y}_p(\bar{x}_i), \quad p \in J_Q. \quad (7.48)$$

Denoting by $\varrho = (\rho_1, \dots, \rho_{|J_Q|})^\top \in \mathbb{R}^{|J_Q|}$ the vector of unknown coefficients, the equations (7.48) can be written in matrix-vector form as

$$(L + \beta M)\varrho = c,$$

The entries of the matrix $L = (l_{pq})_{pq}$ are

$$l_{pq} = \sum_{i=1}^N \tilde{Y}_p(\bar{x}_i) \tilde{Y}_q(\bar{x}_i),$$

the matrix $M = \text{diag}(\lambda_1^s, \dots, \lambda_{|J_Q|}^s)$ is a diagonal matrix, and

$$c_p = \sum_{i=1}^N \|x_i\| \tilde{Y}_p(\bar{x}_i).$$

Numerical Approximation

Let us finally discuss the numerical solution of the optimality conditions (7.42). In particular, one needs to (approximately) evaluate the integrals (7.43), (7.44), and (7.45). Even though integrals on the 2-sphere can be computed exactly and quadrature rules exist up to a certain degree, see e.g. [5, 35], we instead prefer to use a triangulation together with an appropriate quadrature. The reason is that numerical quadrature on the sphere would have to be of rather high degree to reproduce small details and features of the data, contrary to the chosen quadrature, which can easily be refined up to the desired precision. Finally let us mention that, for a more accurate evaluation of the integrals, one can introduce an intermediate (radial) map from the polyhedron to geodesic triangles. See e.g. [35, Sec. 7.2].

We use a polyhedral approximation $\mathcal{S}_h^2 = (\mathcal{V}, \mathcal{T})$ of the 2-sphere \mathcal{S}^2 . It is defined by a set $\mathcal{V} = \{v_1, \dots, v_n\} \subset \mathcal{S}^2$ of vertices and a set $\mathcal{T} = \{T_1, \dots, T_m\} \subset \mathcal{V} \times \mathcal{V} \times \mathcal{V}$ of triangular faces. Each triangle is most easily parametrised using barycentric coordinates, see e.g. [13, Chapter 5]. We associate with each triangle $T_i \in \mathcal{T}$ a tuple (i_1, i_2, i_3) identifying the corresponding vertices $(v_{i_1}, v_{i_2}, v_{i_3})$, which are arranged in clockwise order. The parametrisation (7.3) then reads

$$\mathbf{x}_i(\xi) = v_{i_1} + \xi^1(v_{i_3} - v_{i_1}) + \xi^2(v_{i_2} - v_{i_1})$$

with

$$\Omega = \{\xi \in \mathbb{R}^2 : \xi^1 \in [0, 1] \text{ and } \xi^2 \in [0, 1 - \xi^1]\},$$

which is referred to as the reference triangle. The gradient matrix of T_i is then simply

$$D\mathbf{x}_i = \begin{pmatrix} \partial_1 \mathbf{x}_i & \partial_2 \mathbf{x}_i \end{pmatrix} = \begin{pmatrix} v_{i_3} - v_{i_1} & v_{i_2} - v_{i_1} \end{pmatrix}.$$




Figure 7.6: Illustration of a triangular face (filled gray) intersecting the sphere \mathcal{S}^2 at the vertices (hollow circles). The six nodal points consist of the vertices of the triangle together the edge midpoints (filled black dots). The approximated sphere-like surface is shown by the hatched gray area. A radial line passing through the vertex v_i is shown. The hollow circle indicates the intersection with \mathcal{S}^2 at which $\tilde{f}(v_i)$ in (7.49) is taken. \tilde{f} itself, as described in Sec. 7.5, is assigned by taking the maximum image intensity along the drawn radial line between the two cross marks.

The surface normal is constant on T_i and is denoted by $\tilde{\mathbf{N}}_i$.

We approximate all functions on \mathcal{S}^2 by corresponding functions on the polyhedron \mathcal{S}_h^2 . A continuous function $\tilde{f} : \mathcal{S}^2 \rightarrow \mathbb{R}$ is replaced by its piecewise polynomial interpolation $\tilde{f}_h : \mathcal{S}_h^2 \rightarrow \mathbb{R}$ on \mathcal{S}_h^2 . We define it as

$$\tilde{f}_h(\cdot) := \sum_{j=1}^{N_h} \tilde{f}(v_j) \tilde{\varphi}_j(\cdot). \quad (7.49)$$

Here, $\{\tilde{\varphi}_j\}$ are $N_h = 6$ quadratic shape functions forming a nodal basis together with nodal points $\{v_j\} \subset \mathcal{S}_h^2$ and \tilde{f} is the usual radially constant extension, cf. (7.6) in Sec. 7.2. In other words, \tilde{f}_h is both a radial projection from the 2-sphere to the polyhedron \mathcal{S}_h^2 and to piecewise quadratic functions. Note that the shape functions are defined on the triangular faces T_i . Whenever a function \tilde{f} has a dependence on time we simply compute its approximation \tilde{f}_h separately for all times $t \in I$. We point the reader to Fig. 7.6 for a figurative illustration.

In further consequence, the fully normalised scalar spherical harmonics, which were introduced in (7.26), are substituted with their corresponding ap-

proximations on \mathcal{S}_h^2 . For $\tilde{Y} \in \text{Harm}_n$, $n \in \mathbb{N}_0$, we have

$$\tilde{Y}_h(\cdot) = \sum_{j=1}^{N_h} \tilde{Y}(v_j) \tilde{\varphi}_j(\cdot). \quad (7.50)$$

We chose piecewise quadratic approximations for \tilde{Y} so that we can adequately apply $\nabla_{\mathcal{S}_h^2}$ and obtain piecewise linear vector fields. Accordingly, we define approximations of the vector spherical harmonics, introduced in (7.29), as follows.

Proposition 2. *Let $\tilde{Y} \in \text{Harm}_n$, $n \in \mathbb{N}$. The piecewise linear interpolations of the corresponding tangential vector spherical harmonics on a triangular face $T_i \in \mathcal{T}$ are*

$$\tilde{\mathbf{y}}_h^{(2)}(\mathbf{x}_i(\xi)) = \lambda_n^{-1/2} \sum_{j=1}^{N_h} \tilde{Y}(v_j) \nabla_{\mathcal{S}_h^2} \tilde{\varphi}_j(\mathbf{x}_i(\xi)), \quad (7.51)$$

$$\tilde{\mathbf{y}}_h^{(3)}(\mathbf{x}_i(\xi)) = \frac{\lambda_n^{-1/2}}{2|T_i|} \sum_{j=1}^{N_h} \tilde{Y}(v_j) (\partial_2 \varphi_j(\xi) \partial_1 \mathbf{x}_i(\xi) - \partial_1 \varphi_j(\xi) \partial_2 \mathbf{x}_i(\xi)). \quad (7.52)$$

Their derivation is deferred to the appendix.

Without loss of generality, let $\tilde{f}_h(0, \cdot)$ and $\tilde{f}_h(1, \cdot)$ be the approximations of the data \tilde{f} at two subsequent frames. We define the derivative with respect to time by the forward difference

$$\partial_t \tilde{f}_h(\cdot) := \tilde{f}_h(1, \cdot) - \tilde{f}_h(0, \cdot).$$

Moreover, we replace the surface gradient $\nabla_{\mathcal{S}^2} \tilde{f}$ of a function on \mathcal{S}^2 with its counterpart $\nabla_{\mathcal{S}_h^2} \tilde{f}_h$ on \mathcal{S}_h^2 , which is computed according to (7.17). The function $\tilde{\rho}$ is obtained by solving (7.48) and, for numerical computations, is further replaced with its piecewise quadratic interpolation $\tilde{\rho}_h$ as in (7.49). Coefficients α_i^j are computed by the Gram-Schmidt process at the nodal points. For numerical computations piecewise quadratic approximations, as defined in (7.49), are used.

Finally, for the calculation of the integrals we employ the standard quadrature on triangulated spheres, see e.g. [5, 35]. Let $\xi_c = (1/3, 1/3)^\top$ be the centroid of the reference triangle Ω . Then, we approximate the spherical integral over a function $\tilde{f} : \mathcal{S}^2 \rightarrow \mathbb{R}$ on the 2-sphere by

$$\int_{\mathcal{S}^2} \tilde{f} d\mathcal{S}^2 \approx \int_{\mathcal{S}_h^2} \tilde{f}_h d\mathcal{S}_h^2 \approx \sum_{i=1}^m |T_i| \tilde{f}_h(\mathbf{x}_i(\xi_c)).$$

7.5 Experiments

Microscopy Data

The present data consist of volumetric time-lapse (4-dimensional) images of a live zebrafish embryo during the gastrula period. These videos were recorded

approximately five to ten hours after fertilisation by means of confocal laser-scanning microscopy and feature endodermal cells expressing a green fluorescence protein. As a consequence, these labelled cells are recorded without background and allow for a separate treatment. We refer the reader to [48] for many illustrations and a detailed discussion of the zebrafish's developmental process. Regarding the imaging techniques used during data acquisition we refer to [60] and for the treatment of the specimen we point the reader to [66].

The crucial feature of endodermal cells is the fact that they form a so-called monolayer during early morphogenesis, see [83]. Essentially, it means that the labelled cells do not sit on top of each other but float side by side forming an artificial sphere-shaped layer. It can be regarded as a surface and allows for the straightforward extraction of an image sequence. Clearly, this surface is subject to geometric approximations. For instance, in [51, 74] it is assumed an ideal sphere, whereas in [12] and [52] only a fraction of the data is considered and modelled as a moving manifold and a height field, respectively, both possessing a boundary.

The recorded data features a cuboid region of approximately $860 \times 860 \times 320 \mu\text{m}^3$ of the animal hemisphere. The spatial resolution is $512 \times 512 \times 44$ voxels and the recorded image intensities are in the range $\{0, \dots, 255\}$. Our sequence contains 75 images with a temporal interval of 240 s. For the further discussion, we denote the data by

$$f^\delta \in \{0, \dots, 255\}^{75 \times 512 \times 512 \times 44}.$$

Preprocessing and Surface Data Acquisition

Let us briefly discuss the preprocessing steps required to obtain an image sequence together with the evolving surface. We limit our consideration to two consecutive frames and denote the respective volumetric data by f_0^δ and f_1^δ .

For each frame, the approximate surface is found by minimising the functional (7.46) with approximate cell centres acting as sample points. They appear as local maxima in image intensity and are readily located by Gaussian filtering followed by plain thresholding. However, beforehand the points are centred around the origin by first fitting a sphere and subsequently subtracting the spherical centre.

The triangle mesh \mathcal{S}_h^2 is obtained by iterative refinement of an icosahedron that is inscribed in the 2-sphere, see e.g. [13, Chapter 1.3.3]. Every refinement step halves the edge lengths by connecting the edge midpoints and projecting them to the unit sphere. Consequentially, every triangular face is split into four smaller triangles and the total number of faces after $k \in \mathbb{N}_0$ subdivisions is $20 \cdot 4^k$. In our case, $k = 7$ refinements are required to resolve the data adequately.

It remains to discuss the acquisition of the approximations $\tilde{f}_h(0, \cdot)$ and $\tilde{f}_h(1, \cdot)$ on the polyhedron. For a frame $j \in \{0, 1\}$, we define the value at a nodal point $v_i \in \mathcal{S}_h^2$ in (7.49) via the projection

$$\tilde{f}(j, v_i) := \max_{c \in [1-\varepsilon, 1+\varepsilon]} f_j^\delta(c \tilde{\rho}_h(j, v_i) v_i),$$

where $\varepsilon > 0$ is chosen sufficiently large. \mathring{f}_j^δ denotes the piecewise linear extension of f_j^δ to \mathbb{R}^3 , which is necessary for gridded data. The above projection

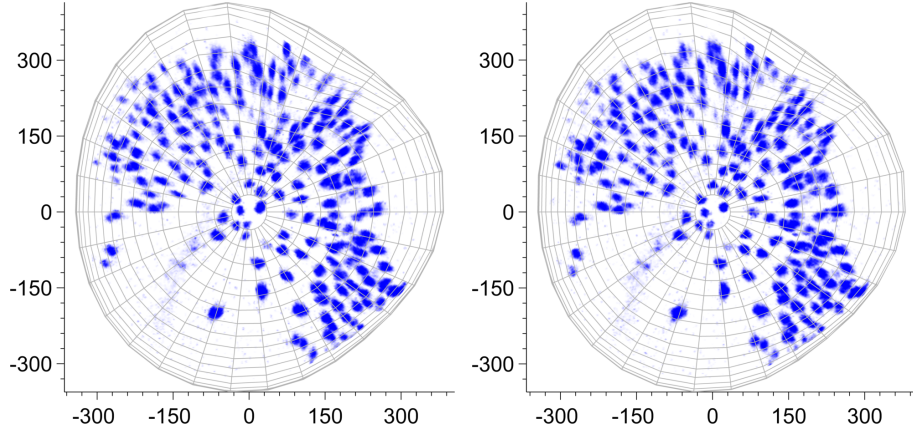


Figure 7.7: Frames no. 70 (left) and 71 (right) of the processed image sequence in a top view. The embryo's body axis is oriented from bottom left to top right.

within the narrow band

$$[(1 - \varepsilon)\tilde{\rho}_h(j, v_i)v_i, (1 + \varepsilon)\tilde{\rho}_h(j, v_i)v_i]$$

corrects for small deviations of the cells from the fitted surface. Again, we refer to Fig. 7.6 for illustration. Finally, all intensities are scaled to the interval $[0, 1]$. Figure 7.2 shows two frames of the extracted image sequence defined on the sphere-like evolving surface. Figure 7.7 depicts the same matter but in a top view. For better illustration we have added an artificial mesh. Its radius has been widened by one percent.

Visualisation of Results

We employ the standard flow colour-coding [9] for the visualisation of the computed vector fields. Its purpose is to create a colour image by assigning every vector a colour from a pre-defined colour disk. The colour associated is determined by a vector's angle and its length.

However, it was originally defined for planar vector fields and requires adaptation to our particular purpose of tangent vector field visualisation. To this end, we follow the idea developed in [51] by first projecting each vector to the plane and then rescaling its length. Let us denote by $P_{x_3} : (x_1, x_2, x_3)^\top \mapsto (x_1, x_2, 0)^\top$ the orthogonal projector of \mathbb{R}^3 onto the x_1 - x_2 -plane. For a tangent vector field $\hat{\mathbf{v}}$ we apply the colour-coding to the planar vector field

$$\frac{\|\hat{\mathbf{v}}\|}{\|P_{x_3}\hat{\mathbf{v}}\|} P_{x_3}\hat{\mathbf{v}}.$$

It is constructed so that the length of individual vectors is preserved. Subsequently, the obtained colour image is mapped back onto the surface. Clearly, in the above construction, one has to distinguish the cases where $x_3 \geq 0$ and $x_3 < 0$. Moreover, P_{x_3} is required to be injective in either case.

The radius R of the colour disk is chosen to be equal to the longest vector in the respective vector field we attempt to visualise. Table 7.2 lists all values of

Figure	7.10	7.11 (a)	7.11 (b)	7.11 (c)	7.11 (d)	7.12 (a)	7.12 (b)	7.12 (c)	7.12 (d)
R	5.18	9.86	5.18	3.64	2.52	9.86	5.18	3.64	2.52

Table 7.2: Radii R of the colour disks used for colour-coded visualisation of tangent vector fields.

R for the different figures in this section. In Fig. 7.10 we show a colour-coded tangent vector field together with the colour disk.

For simplicity reasons, for image functions as well as surfaces we plot their piecewise linear approximations. Moreover, the visualised vector fields are evaluated at the centroids and result in piecewise constant colour-coded images.

Results

We performed several experiments on said zebrafish microscopy data. In order to obtain an approximation of the evolving surface, we minimised functional (7.46) by solving the optimality conditions (7.48). As mentioned in Sec. 7.5, approximate cell centres serve as input. The parameter of the Sobolev space $H^s(\mathcal{S})$ was chosen as $s = 3 + \epsilon$, where $\epsilon = 2.2204 \cdot 10^{-16}$ is the machine precision, cf. also the discussion regarding theoretical requirements in Sec. 7.4. The regularisation parameter was set to $\beta = 10^{-4}$ and the finite-dimensional subspace was chosen as

$$\mathcal{Q} = \text{span}\{\tilde{Y}_{nj} : n = 0, \dots, 30, j = 1, \dots, 2n + 1\}.$$

In the second step, we computed a minimiser of functional (7.35) as outlined in Sec. 7.4. Here, the finite-dimensional subspace was chosen as

$$\mathcal{U} = \text{span}\{\hat{\mathbf{y}}_{nj}^{(i)} : n = 1, \dots, 50, j = 1, \dots, 2n + 1, i = 2, 3\}.$$

The linear systems resulting from optimality conditions (7.42) and (7.48) were solved by means of the General Minimal Residual Method (GMRES) using an Intel Xeon E5-1620 3.6 GHz workstation equipped with 128 GB RAM. Solutions to (7.42) and (7.48) converged within 1000 and 100 iterations, respectively, to a relative residual of 10^{-2} . The overall runtime was dominated by the evaluation of the integrals (7.43), (7.44), and (7.45). In our Matlab implementation it amounts to several hours. However, the resulting linear system can typically be solved within seconds. Both implementation and data are available on our website.¹

Figures 7.8 and 7.9 portray a minimising function of \mathcal{F}_β for frames 70 and 71 of the image sequence. The resulting surface is depicted in Fig. 7.2 and in Fig. 7.7. Clearly, it reflects the geometry appropriately and contains the desired cell features, cf. also the unprocessed microscopy data in Fig. 7.1.

In a second step we solved for minimisers of \mathcal{E}_α for different values of the regularisation parameter α . Figure 7.10 depicts the optical flow field for $\alpha = 10^{-1}$. The tangent vector field is visualised as discussed in Sec. 7.5. Note that in all figures the colour disk has been scaled for better illustration. In Fig. 7.11, we illustrate tangent vector fields by means of the colour-coding obtained for $\alpha = 10^{-2}$, $\alpha = 10^{-1}$, $\alpha = 1$, and $\alpha = 10$. Finally, in Fig. 7.12 we show the same results but in a top view.

¹<http://www.csc.univie.ac.at>




Figure 7.8: Function $\tilde{\rho}_h$ obtained by minimising \mathcal{F}_β for frames 70 (left) and 71 (right). Colour corresponds to the radius of the fitted surface. \mathcal{S}_h^2 is depicted in a top view.

7.6 Conclusion

With the goal of efficient cell motion analysis we considered optical flow on evolving surfaces. As a prototypical example we restricted ourselves to surfaces parametrised from the round sphere and showed that 4D microscopy data of a living zebrafish embryo can be faithfully represented in this way. In contrast to previous works, where only a section of the embryo or a spherical approximation was considered, our approach fully attributes the geometry and models the embryo as a closed surface of genus zero. The resulting energy functional was solved by means of a Galerkin method based on vector spherical harmonics. Moreover, the parametrisation of the moving sphere-like surface was obtained from the data by solving a surface interpolation problem. Scalar spherical harmonics expansion allows to easily meet the smoothness requirements of the surface. Finally, we conducted several experiments based on said microscopy data. Our results show that cell motion can be indicated reasonably well by the proposed approach.

Acknowledgements We thank Pia Aanstad from the University of Innsbruck for sharing her biological insight and for kindly providing the microscopy data. Moreover, we are grateful to Peter Elbau and Clemens Kirsits for their helpful comments, and to José A. Iglesias for carefully proofreading an earlier version of this article and providing valuable feedback. 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 network “Geometry + Simulation” (project S11704, Variational Methods for Imaging on Manifolds).

Appendix

It remains to give the calculations regarding the piecewise linear approximations of vector spherical harmonics on \mathcal{S}_h^2 . Both equations in Prop. 2 follow




Figure 7.9: Illustration of the function $\tilde{\rho}_h$ for frames 70 (left column) and 71 (right column). The bottom row differs from the top row by a rotation of 180 degrees around the x_3 -axis.

directly by expanding the definitions of the tangential vector spherical harmonics (7.29). For the first identity, that is (7.51), we have

$$\tilde{\mathbf{y}}_h^{(2)} = \lambda_n^{-1/2} \nabla_{S_h^2} \tilde{Y}_h = \lambda_n^{-1/2} \sum_{j=1}^{N_h} \tilde{Y}(v_j) \nabla_{S_h^2} \tilde{\varphi}_j.$$

The second identity, that is (7.52), follows by the fact that

$$2|T_i| = |\partial_1 \mathbf{x}_i \times \partial_2 \mathbf{x}_i|, \quad \tilde{\mathbf{N}}_i = \frac{\partial_1 \mathbf{x}_i \times \partial_2 \mathbf{x}_i}{|\partial_1 \mathbf{x}_i \times \partial_2 \mathbf{x}_i|},$$




Figure 7.10: Tangent vector field minimising \mathcal{E}_α . Depicted is the colour-coded optical flow field computed between frames 70 and 71. The right image differs from the left by a rotation of 180 degrees around the x_3 -axis.




Figure 7.11: Visualisation of the optical flow field obtained for different values of α . The bottom row differs from the top view by a rotation of 180 degrees around the x_3 -axis. From left to right: a) $\alpha = 10^{-2}$, b) $\alpha = 10^{-1}$, c) $\alpha = 1$, and d) $\alpha = 10$.

and by application of the vector triple product rule, yielding

$$\begin{aligned}
 \tilde{\mathbf{y}}_h^{(3)} &= \lambda_n^{-1/2} \nabla_{S_h^2} \tilde{Y}_h \times \tilde{\mathbf{N}}_i \\
 &= \lambda_n^{-1/2} \nabla_{S_h^2} \tilde{Y}_h \times \frac{\partial_1 \mathbf{x}_i \times \partial_2 \mathbf{x}_i}{|\partial_1 \mathbf{x}_i \times \partial_2 \mathbf{x}_i|} \\
 &= \frac{\lambda_n^{-1/2}}{2|T_i|} ((\nabla_{S_h^2} \tilde{Y}_h \cdot \partial_2 \mathbf{x}_i) \partial_1 \mathbf{x}_i - (\nabla_{S_h^2} \tilde{Y}_h \cdot \partial_1 \mathbf{x}_i) \partial_2 \mathbf{x}_i),
 \end{aligned}$$

where the last equality results from the definition of the interpolation (7.50) of




Figure 7.12: Top view of the optical flow field computed for different values of α . From left to right, top to bottom: a) $\alpha = 10^{-2}$, b) $\alpha = 10^{-1}$, c) $\alpha = 1$, and d) $\alpha = 10$.

\tilde{Y}_h and the directional derivative (7.9) on the triangular face, that is

$$\nabla_{S_h^2} \tilde{Y}_h \cdot \partial_k \mathbf{x}_i = \sum_{j=1}^{N_h} \tilde{Y}(v_j) \partial_k \varphi_j.$$

Part III

Appendix

Bibliography

- [1] J. Abhau, Z. Belhachmi, and O. Scherzer. On a decomposition model for optical flow. In *Energy Minimization Methods in Computer Vision and Pattern Recognition*, volume 5681 of *Lecture Notes in Computer Science*, pages 126–139. Springer-Verlag, Berlin, Heidelberg, 2009.
- [2] M. D. Abràmoff and M. A. Viergever. Computation and visualization of three-dimensional soft tissue motion in the orbit. *IEEE Trans. Med. Imag.*, 21(4):296–304, 2002.
- [3] F. Amat, W. Lemon, D. P. Mossing, K. McDole, Y. Wan, K. Branson, E. W. Myers, and P. J. Keller. Fast, accurate reconstruction of cell lineages from large-scale fluorescence microscopy data. *Nat. Meth.*, 11(9):951–958, 2014.
- [4] F. Amat, E. W. Myers, and P. J. Keller. Fast and robust optical flow for time-lapse microscopy using super-voxels. *Bioinformatics*, 29(3):373–380, 2013.
- [5] K. Atkinson and W. Han. *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction*, volume 2044 of *Lecture Notes in Mathematics*. Springer, Heidelberg, 2012.
- [6] G. Aubert, R. Deriche, and P. Kornprobst. Computing optical flow via variational techniques. *SIAM J. Appl. Math.*, 60:156–182, 1999.
- [7] G. Aubert and P. Kornprobst. A mathematical study of the relaxed optical flow problem in the space $BV(\Omega)$. *SIAM J. Math. Anal.*, 30(6):1282–1308, 1999.
- [8] G. Aubert and P. Kornprobst. *Mathematical problems in image processing*, volume 147 of *Applied Mathematical Sciences*. Springer, New York, 2 edition, 2006. Partial differential equations and the calculus of variations, With a foreword by Olivier Faugeras.
- [9] 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.
- [10] J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *Int. J. Comput. Vision*, 12(1):43–77, 1994.
- [11] G. K. Batchelor. *An introduction to fluid dynamics*. Cambridge Mathematical Library. Cambridge University Press, Cambridge, 1999.

- [12] M. Bauer, M. Grasmair, and C. Kirisits. Optical flow on moving manifolds. *SIAM J. Imaging Sciences*, 8(1):484–512, 2015.
- [13] M. Botsch, L. Kobbelt, M. Pauly, P. Alliez, and B. Lévy. *Polygon Mesh Processing*. A K Peters, 2010.
- [14] M. Buibas, D. Yu, K. Nizar, and G. A. Silva. Mapping the spatiotemporal dynamics of calcium signaling in cellular neural networks using optical flow. *Ann. Biomed. Eng.*, 38(8):2520–2531, 2010.
- [15] 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.
- [16] M. Chalfie, Y. Tu, G. Euskirchen, W. W. Ward, and D. C. Prasher. Green fluorescent protein as a marker for gene expression. *Science*, 263(5148):802–805, 1994.
- [17] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Imaging Vision*, 40(1):120–145, 2011.
- [18] T. Corpetti, É. Mémin, and P. Pérez. Dense estimation of fluid flows. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(3):365–380, 2002.
- [19] R. Courant and D. Hilbert. *Methods of mathematical physics. Vol. I*. Interscience Publishers, Inc., New York, N.Y., 1953.
- [20] A. Del Bimbo, P. Nesi, and J. L. C. Sanz. Analysis of optical flow constraints. *IEEE Trans. Image Process.*, 4(4):460–469, 1995.
- [21] J. Delpiano, J. Jara, J. Scheer, O. A. Ramírez, J. Ruiz-del Solar, and S. Härtel. Performance of optical flow techniques for motion analysis of fluorescent point signals in confocal microscopy. *Mach. Vis. Appl.*, 23(4):675–689, 2012.
- [22] M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, 1976.
- [23] M. P. do Carmo. *Riemannian Geometry*. Birkhäuser, 1992.
- [24] G. Dziuk and C. M. Elliott. Finite element methods for surface PDEs. *Acta Numer.*, 22:289–396, 5 2013.
- [25] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of inverse problems*, volume 375 of *Mathematics and its Applications*. Kluwer Academic Publishers Group, Dordrecht, 1996.
- [26] L. C. Evans and R. F. Gariepy. *Measure theory and fine properties of functions*. Studies in Advanced Mathematics. CRC Press, Boca Raton, FL, 1992.
- [27] D. J. Fleet and Y. Weiss. Optical flow estimation. In Paragios et al. [70], chapter 15, pages 239–257.

- [28] W. Freeden and M. Schreiner. *Spherical functions of Mathematical Geosciences. A scalar, vectorial, and tensorial setup*. Springer Berlin Heidelberg, 2009.
- [29] F. Frühauf, C. Pontow, and O. Scherzer. Texture enhancing based on variational image decomposition. In M. Bergounioux, editor, *Mathematical Image Processing*, volume 5 of *Springer Proceedings in Mathematics*, pages 127–140, Berlin Heidelberg, 2011. Springer.
- [30] I. M. Gelfand and S. V. Fomin. *Calculus of variations*. Revised English edition translated and edited by Richard A. Silverman. Prentice-Hall Inc., Englewood Cliffs, N. J., 1963.
- [31] D. Gilbarg and N. Trudinger. *Elliptic Partial Differential Equations of Second Order*. Classics in Mathematics. Springer Verlag, Berlin, 2001. Reprint of the 1998 edition.
- [32] H.W. Haussecker and D.J. Fleet. Computing optical flow with physical models of brightness variation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(6):661–673, June 2001.
- [33] E. Hebey. *Sobolev Spaces on Riemannian Manifolds*, volume 1635 of *Lecture Notes in Mathematics*. SV, Berlin, 1996.
- [34] E. Hebey. *Nonlinear analysis on manifolds: Sobolev spaces and inequalities*. Courant Lecture Notes in Mathematics. New York University, Courant Institute of Mathematical Sciences, New York; American Mathematical Society, Providence, RI, 1999.
- [35] K. Hesse, I. H. Sloan, and R. S. Womersley. Numerical integration on the sphere. In W. Freeden, M. Z. Nashed, and T. Sonar, editors, *Handbook of Geomathematics*, pages 1187–1219. Springer, 2010.
- [36] W. Hinterberger, O. Scherzer, C. Schnörr, and J. Weickert. Analysis of optical flow models in the framework of calculus of variations. *Numer. Funct. Anal. Optim.*, 23(1-2):69–89, 2002.
- [37] M. W. Hirsch. *Differential topology*, volume 33 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994.
- [38] B. K. P. Horn. *Robot Vision*. MIT electrical engineering and computer science series. MIT Press, 1986.
- [39] B. K. P. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
- [40] J. Hubený. *Applications of PDE-Based Image Processing in Fluorescence Microscopy*. PhD thesis, Masaryk University, Brno, Czech Republic, 2007.
- [41] J. Hubený, V. Ulman, and P. Matula. Estimating large local motion in live-cell imaging using variational optical flow. In *VISAPP: Proc. of the Second International Conference on Computer Vision Theory and Applications*, pages 542–548. INSTICC, 2007.

- [42] A. Imiya, H. Sugaya, A. Torii, and Y. Mochizuki. Variational analysis of spherical images. In A. Galalowicz 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.
- [43] P. J. Keller. Imaging morphogenesis: Technological advances and biological insights. *Science*, 340(6137), 2013.
- [44] P. J. Keller, A. D. Schmidt, A. Santella, K. Khairy, Z. Bao, J. Wittbrodt, and E. H. K. Stelzer. Fast, high-contrast imaging of animal development with scanned light sheet-based structured-illumination microscopy. *Nat. Meth.*, 7(8):637–642, 2010.
- [45] P. J. Keller, A. D. Schmidt, J. Wittbrodt, and E. H. K. Stelzer. Reconstruction of zebrafish early embryonic development by scanned light sheet microscopy. *Science*, 322(5904):1065–1069, 2008.
- [46] K. Khairy and P. J. Keller. Reconstructing embryonic development. *Genesis*, 49(7):488–513, 2011.
- [47] S. Khan, J. Lefèvre, H. Ammari, and S. Baillet. Feature detection and tracking in optical flow on non-flat manifolds. *Pattern Recogn. Lett.*, 32(15):2047–2052, 2011.
- [48] 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.
- [49] C. Kirisits. *Optical Flow on Evolving Surfaces*. PhD thesis, University of Vienna, Austria, Vienna, June 2014.
- [50] C. Kirisits, L. F. Lang, and O. Scherzer. Optical flow on evolving surfaces with an application to the analysis of 4D microscopy data. In A. Kuijper, K. Bredies, T. Pock, and H. Bischof, editors, *SSVM’13: Proceedings of the fourth International Conference on Scale Space and Variational Methods in Computer Vision*, volume 7893 of *Lecture Notes in Computer Science*, pages 246–257, Berlin, Heidelberg, 2013. Springer-Verlag.
- [51] C. Kirisits, L. F. Lang, and O. Scherzer. Decomposition of optical flow on the sphere. *GEM. Int. J. Geomath.*, 5(1):117–141, 2014.
- [52] C. Kirisits, L. F. Lang, and O. Scherzer. Optical flow on evolving surfaces with space and time regularisation. *J. Math. Imaging Vision*, 52(1):55–70, 2015.
- [53] T. Kohlberger, E. Memin, and Ch. Schnörr. Variational dense motion estimation using the helmholtz decomposition. In L. D. Griffin and M. Lillholm, editors, *Scale Space Methods in Computer Vision*, volume 2695 of *Lecture Notes in Computer Science*, pages 432–448. Springer, Berlin, 2003.
- [54] W. Kühnel. *Differential geometry*, volume 16 of *Student Mathematical Library*. American Mathematical Society, Providence, RI, 2 edition, 2006. Curves—surfaces—manifolds, Translated from the 2003 German original by Bruce Hunt.

- [55] L. F. Lang and O. Scherzer. Optical flow on evolving sphere-like surfaces. Preprint on ArXiv arXiv:1506.03358, University of Vienna, Austria, 2015.
- [56] J. M. Lee. *Riemannian manifolds*, volume 176 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997. An introduction to curvature.
- [57] J. M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, New York, 2 edition, 2013.
- [58] 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.
- [59] J.-L. Lions and E. Magenes. *Non-Homogeneous Boundary Value Problems and Applications I*, volume 181 of *Die Grundlehren der Mathematischen Wissenschaften*. Springer Verlag, New York, 1972.
- [60] 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.
- [61] C. Melani, M. Campana, B. Lombardot, B. Rizzi, F. Veronesi, C. Zanella, P. Bourguine, K. Mikula, N. Peyri  ras, 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.
- [62] Y. Meyer. *Oscillating patterns in image processing and nonlinear evolution equations*, volume 22 of *University Lecture Series*. American Mathematical Society, Providence, RI, 2001. The fifteenth Dean Jacqueline B. Lewis memorial lectures.
- [63] V. Michel. *Lectures on constructive approximation. Fourier, spline, and wavelet methods on the real line, the sphere, and the ball*. Birkh  user Basel, 2013.
- [64] R. Mikut, T. Dickmeis, W. Driever, P. Geurts, F. A. Hamprecht, B. X. Kausler, M. J. Ledesma-Carbayo, R. Mar  e, K. Mikula, P. Pantazis, O. Ronneberger, A. Santos, R. Stotzka, u. Str  hle, and N. Peyri  ras. Automated processing of zebrafish imaging data: A survey. *Zebrafish*, 10:401–421, 2013.
- [65] 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.
- [66] 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.
- [67] S. Osher, A. Sol  , and L. Vese. Image decomposition and restoration using total variation minimization and the H^{-1} -norm. *Multiscale Model. Simul.*, 1(3):349–370, 2003.

- [68] S. W. Paddock. Confocal laser scanning microscopy. *BioTechniques*, 27:992–1007, 1999.
- [69] S. W. Paddock. Principles and practices of laser scanning confocal microscopy. *Mol. Biotechnol.*, 16(2):127–149, 2000.
- [70] N. Paragios, Y. Chen, and O. Faugeras, editors. *Handbook of mathematical models in computer vision*. Springer, New York, 2006.
- [71] M. A. Penna and K. A. Dines. A simple method for fitting sphere-like surfaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(9):1673–1678, September 2007.
- [72] 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.
- [73] O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen. *Variational methods in imaging*, volume 167 of *Applied Mathematical Sciences*. Springer, New York, 2009.
- [74] B. Schmid, G. Shah, N. Scherf, M. Weber, K. Thierbach, C. Campos Pérez, I. Roeder, P. Aanstad, and J. Huiskens. High-speed panoramic light-sheet microscopy reveals global endodermal cell dynamics. *Nat. Commun.*, 4:2207, 2013.
- [75] Ch. Schnörr. Determining optical flow for irregular domains by minimizing quadratic functionals of a certain class. *Int. J. Comput. Vision*, 6:25–38, 1991.
- [76] T. Schuster and J. Weickert. On the application of projection methods for computing optical flow fields. *Inverse Probl. Imaging*, 1(4):673–690, 2007.
- [77] E. Tadmor, S. Nezzar, and L. Vese. A multiscale image representation using hierarchical (BV, L^2) decompositions. *Multiscale Model. Simul.*, 2(4):554–579 (electronic), 2004.
- [78] A. N. Tikhonov and V. Y. Arsenin. *Solutions of Ill-Posed Problems*. John Wiley & Sons, Washington, D.C., 1977.
- [79] 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.
- [80] H. Triebel. *Theory of function spaces. II*, volume 84 of *Monographs in Mathematics*. Birkhäuser Verlag, Basel, 1992.
- [81] A. Verri and T. Poggio. Motion field and optical flow: qualitative properties. *IEEE Trans. Pattern Anal. Mach. Intell.*, 11(5):490–498, 1989.

- [82] L. Vese and S. Osher. Modeling textures with total variation minimization and oscillating patterns in image processing. *J. Sci. Comput.*, 19(1–3):553–572, 2003. Special issue in honor of the sixtieth birthday of Stanley Osher.
- [83] R. M. Warga and C. Nüsslein-Volhard. Origin and development of the zebrafish endoderm. *Development*, 126(4):827–838, February 1999.
- [84] J. Weickert, A. Bruhn, T. Brox, and N. Papenberg. A survey on variational optic flow methods for small displacements. In O. Scherzer, editor, *Mathematical Models for Registration and Applications to Medical Imaging*, volume 10 of *Mathematics in Industry*, pages 103–136. Springer, Berlin Heidelberg, 2006.
- [85] J. Weickert and Ch. Schnörr. A theoretical framework for convex regularizers in PDE-based computation of image motion. *Int. J. Comput. Vision*, 45(3):245–264, 2001.
- [86] J. Weickert and Ch. Schnörr. Variational optic flow computation with a spatio-temporal smoothness constraint. *J. Math. Imaging Vision*, 14:245–255, 2001.
- [87] D. Weiskopf and G. Erlebacher. Overview of flow visualization. In C. D. Hansen and C. R. Johnson, editors, *The Visualization Handbook*, pages 261–278. Elsevier, Amsterdam, 2005.
- [88] R. P. Wildes, A. M. Amabile, M. J. and Lanzillotto, and T.S. Leu. Recovering estimates of fluid flow from image sequence data. *Comput. Vision Image Understanding*, 80(2):246–266, 2000.
- [89] L. Wolpert, C. Tickle, and A. M. Arias. *Principles of development*. Oxford University Press, 2 edition, 2002.
- [90] J. Yuan, C. Schnörr, and G. Steidl. Simultaneous higher-order optical flow estimation and decomposition. *SIAM J. Sci. Comput.*, 29(6):2283–2304, 2007.
- [91] J. Yuan, C. Schnörr, and G. Steidl. Convex hodge decomposition and regularization of image flows. *J. Math. Imaging Vision*, 33(2):169–177, 2009.

Zusammenfassung

Die Bestimmung von Bewegungen in Videos ist ein allgegenwärtiges Ziel im Gebiet der mathematischen Bildanalyse und der Computer Vision. Ein zentrales Anliegen stellt die Berechnung des optischen Flusses in einer Bildsequenz dar. Dessen erklärtes Ziel ist es, aus Fluktuationen der Bildintensität ein Vektorfeld zu bestimmen welches die Bewegungen von Objekten in einer aufgenommenen Szene erfasst.

Üblicherweise wird der optische Fluss in planaren Bildsequenzen bestimmt. Generalisierungen auf nichteuklidische Räume ermöglichen beispielsweise die Analyse von Zellbewegungen in Zeitraffer-Mikroskopaufnahmen.

Fluoreszenzmikroskopie erlaubt heutzutage hochauflösende Beobachtungen von biologischen Modellorganismen wie zum Beispiel dem Zebrafisch. Trotz wesentlicher Bedeutung für die Organ- und Gewebebildung ist nur wenig bekannt hinsichtlich der Bewegungs- und Proliferationsmuster von Zellen während dessen frühen Entwicklungsstadiums. Die Beantwortung vieler Fragen bezüglich dieser Aspekte basiert auf dem Wissen um Zellbewegungen.

Steigende räumliche und zeitliche Auflösung von bildgebenden Verfahren und die dabei entstehenden beachtlichen Datenmengen machen die manuelle, von Menschen durchgeführte Bildanalyse, impraktikabel. Automatische Auswertung ist unausweichlich um Zellbewegungen in den oben erwähnten Mikroskopdaten zu verfolgen. Die Bestimmung des optischen Flusses trägt zur Erforschung von Zellmechanismen und dem dynamischen Verhalten von Zellen bei und liefert das dazu notwendige quantitative Verfahren.

Die primäre biologische Motivation für diese Dissertation ist der Wunsch, Zellbewegungen in einem lebenden Zebrafischembryo während dessen Embryogenese zu analysieren. Die vorliegenden Bilddaten zeigen endoderme Zellen, welche mit einem grün fluoreszierenden Protein markiert wurden. Anhand eines Laser-Scanning-Mikroskops lassen sich vierdimensionale Bildsequenzen von den markierten Zellen aufnehmen.

Während des frühen Entwicklungsstadiums des Zebrafisches formen endoderme Zellen einen sogenannten *Monolayer*. Dies bedeutet, dass Zellen dieses Typs nebeneinander auf einer gekrümmten Oberfläche angeordnet sind. Die wesentliche Idee dieser Dissertation ist es, diesen Umstand zu nützen und diese Schicht als bewegte zweidimensionale Oberfläche zu modellieren und Zellbewegungen nur bezüglich dieser sich deformierenden Mannigfaltigkeit zu betrachten. Da sich die räumliche Dimension verringert, folgt daraus als Konsequenz eine Reduzierung der zu untersuchenden Datenmenge und ermöglicht so eine effiziente Bewegungsanalyse in den obengenannten Bildsequenzen.

In dieser Dissertation formulieren wir die Bestimmung von Zellbewegungen als Variationsproblem und behandeln den optischen Fluss auf einer bewegten

zweidimensionalen Mannigfaltigkeit. Je nach Wahl der Geometrie liegen verschiedene Approximationen zugrunde.

Im ersten Teil der Arbeit nehmen wir die Oberfläche des Embryos als veränderlich an und übertragen das (Tikhonov-regularisierte) Horn-Schunck-Funktional und dessen raumzeitliche Generalisierung von Weickert und Schnörr auf dieses nichteuklidische und dynamische Szenario.

Im zweiten Teil widmen wir uns der Topologie des Embryos. Zuerst nehmen wir an, es sei eine statische Kugel und untersuchen verschiedene Zerlegungsfunktionale für Vektorfelder. Wir orientieren uns an aktuellen Trends in der Bildzerlegung und ergründen $u + v$ und hierarchische Zerlegungsmodelle für den optischen Fluss. Die gewählte numerische Optimierungsmethode approximiert die Lösung in einem endlich-dimensionalen Raum aufgespannt durch Vektor Spherical Harmonics. Daraus ergibt sich der Vorteil großer Flexibilität bezüglich der Regularisierungsfunktionale und eine automatische Helmholtz-Zerlegung des Flussfelds.

Ferner widmen wir uns der genaueren geometrischen Modellierung des Zebrafischembryos. Wir betrachten diesen als sphärenähnliche Oberfläche, die sich im Laufe der Entwicklung des Embryos verformt. Zu diesem Zwecke erweitern wir das Optische-Fluss-Funktional von Lefèvre und Baillet um das Szenario einer veränderlichen sphärenähnlichen Mannigfaltigkeit. Das Variationsproblem lösen wir anhand einer Galerkinmethode basierend auf tangentialer Vektor Spherical Harmonics. Um die sphärenähnliche Oberfläche aus den Mikroskopiedaten zu bestimmen, formulieren wir dieses Problem ebenfalls als Variationsproblem und lösen dieses mittels Expansion in skalaren Spherical Harmonics.

Schlussendlich präsentieren wir numerische Ergebnisse basierend auf den obengenannten Zellmikroskopiedaten eines lebenden Zebrafischembryos und veranschaulichen diese mit adäquaten Visualisierungsmethoden.

Curriculum Vitae

University Education

- since 2012 Doctoral studies in *Computer Science* (Informatik, IK: Computational Science), University of Vienna, Austria. Thesis: *Optical Flow on Evolving Manifolds with an Application to the Analysis of Fluorescence Microscopy Data* (in progress), supervised by Univ.-Prof. Dr. Otmar Scherzer and Univ.-Prof. Dr. Monika Henzinger.
- 2010-2012 Master studies (Dipl.-Ing.) in *Computational Intelligence*, Vienna University of Technology, Austria; graduated with distinction. Thesis: *Image Denoising with Variational Methods via Graph Cuts*, supervised by Univ.-Prof. Dr. Otmar Scherzer and Univ.-Prof. Dr. Monika Henzinger.
- 2005-2010 Bachelor studies (BSc.) in *Software & Information Engineering*, Vienna University of Technology, Austria. Thesis: *Metaheuristics for the multicast tree scheduling problem*, supervised by Priv.-Doz. Dr. Nysret Musliu and Priv.-Doz. Dr. Stanford Bessler

Publications

- C. Kirisits, L. F. Lang, and O. Scherzer. Optical flow on evolving surfaces with space and time regularisation. *J. Math. Imaging Vision*, 52(1):55–70, 2015.
- C. Kirisits, L. F. Lang, and O. Scherzer. Decomposition of optical flow on the sphere. *GEM. Int. J. Geomath.*, 5(1):117–141, 2014.
- C. Kirisits, L. F. Lang, and O. Scherzer. Optical flow on evolving surfaces with an application to the analysis of 4D microscopy data. In A. Kuijper, K. Bredies, T. Pock, and H. Bischof, editors, *SSVM'13: Proceedings of the fourth International Conference on Scale Space and Variational Methods in Computer Vision*, volume 7893 of *Lecture Notes in Computer Science*, pages 246–257, Berlin, Heidelberg, 2013. Springer-Verlag.
- S. Bessler, L. F. Lang, and N. Musliu. Scheduling content multicast trees with minimal repetitions. In M. Pióro, K. Szczypiorski, J. Rak, and O. Gonzalez-Soto, editors, *Proc. 14th International Telecommunications Network Strategy and Planning Symposium (NETWORKS)*, 2010.

Preprints

L. F. Lang and O. Scherzer. Optical flow on evolving sphere-like surfaces. Preprint on ArXiv arXiv:1506.03358, University of Vienna, Austria, 2015.

Conference Talks

- 04/2014 *Decomposition of Optical Flow on the Sphere.* 4nd NFN seminar: Internal workshop of the FSP S117, Universitätszentrum Obergurgl, Obergurgl, Austria.
- 06/2013 *Optical Flow on Evolving Surfaces: Numerical Aspects.* Joint Mathematical Conference of the Catalan Mathematical Society, Slovenian Mathematical Society, Austrian Mathematical Society, Slovak Mathematical Society, and Czech Mathematical Society (CSASC), University of Primorska, Koper, Slovenia.
- 07/2012 *Time-lapse cellular analysis of zebrafish embryogenesis.* European Conference on Mathematics for Industry (ECMI), University of Lund, Lund, Sweden.

Other Activities

- 10/2014 Attended *Workshop on Variational methods in imaging, Special Semester on New Trends in Calculus of Variations* at Johann Radon Institute for Computational and Applied Mathematics (RICAM), Linz, Austria.
- 10/2014 Attended *School on Imaging, Special Semester on New Trends in Calculus of Variations* at Johann Radon Institute for Computational and Applied Mathematics (RICAM), Linz, Austria.
- 06/2013 Attended *Fourth International Conference on Scale Space Methods and Variational Methods in Computer Vision (SSVM)* at Schloss Seggau, Seggau, Austria.
- 09/2012 Attended *Young researchers' workshop on Statistics, Learning and Variational Methods in Imaging* at DAMTP, University of Cambridge, UK.