SIAM J. Imaging Sciences
Vol. 8, No. 1, pp. 484–512

Optical Flow on Moving Manifolds*

Martin Bauer†, Markus Grasmair‡, and Clemens Kirisits§

Abstract. Optical flow is a powerful tool for the study and analysis of motion in a sequence of images. In this paper we study a Horn–Schunck-type spatio-temporal regularization functional for image sequences that have a non-Euclidean, time varying image domain. To that end we construct a Riemannian metric that describes the deformation and structure of this evolving surface. The resulting functional can be seen as a natural geometric generalization of previous work by Weickert and Schnörr in 2001 and Lefèvre and Baillet in 2008 for static image domains. In this paper we show the existence and well-posedness of the corresponding optical flow problem and derive necessary and sufficient optimality conditions. We demonstrate the functionality of our approach in two experiments using both synthetic and real data.

Key words. optical flow, evolving surfaces, spatio-temporal regularization, variational methods

AMS subject classifications. 58J90, 65M30, 68U10

DOI. 10.1137/140965235

1. Introduction.

Optical flow. Optical flow is a powerful tool for detecting and analyzing motion in a sequence of images. The underlying idea is to depict the displacement of patterns in the image sequence as a vector field—the optical flow vector field—generating the corresponding displacement function. This framework has applications in a variety of areas connected to computer graphics and video analysis, e.g., in video compressing, video surveillance, or vision-based robot navigation.

Variational methods. In their seminal paper [20], Horn and Schunck proposed a *variational ansatz* for the computation of the optical flow vector field. In this approach one minimizes an *energy functional* consisting of a *similarity (data) term* and a *regularity term*:

$$\tilde{u} = \operatorname*{argmin}_{u \in \mathcal{H}} \mathcal{E}(u) = \operatorname*{argmin}_{u \in \mathcal{H}} \left(\mathcal{S}(u) + \mathcal{R}(u)\right).$$

Here $\mathcal{H}$ denotes an admissible space of vector fields, $\mathcal{R}$ denotes the regularity term for the vector field u, and $\mathcal{S}$ denotes the similarity term, which depends on the image sequence $\mathcal{I}$ under consideration. This method turned out to be particularly successful, as the resulting optical

*Received by the editors April 15, 2014; accepted for publication (in revised form) November 17, 2014; published electronically March 3, 2015.
http://www.siam.org/journals/siims/8-1/96523.html

†Faculty of Mathematics, University of Vienna, Vienna 1090, Austria (Bauer.Martin@univie.ac.at). This author's research was supported by FWF project P24625.

‡Department of Mathematical Sciences, NTNU Trondheim, Trondheim 7491, Norway (Markus.Grasmair@math.ntnu.no).

§Computational Science Center, University of Vienna, Vienna 1090, Austria (Clemens.Kirisits@univie.ac.at). This author's research was supported by the Vienna Graduate School in Computational Science (IK I059-N) and FWF project S10505-N20.

flow fields satisfy certain desirable properties governed by the choice of the regularization term $\mathcal{R}$.

In their paper, Horn and Schunck considered the optical flow problem for a sequence of images defined on some domain in $\mathbb{R}^2$. They proposed using the L^2-norm of the first derivative of the vector field u as a regularization term. The *well-posedness* of this ansatz was first shown by Schnörr in [36]. There the author had to impose an additional assumption on the image sequence in order to ensure the coercivity of the functional $\mathcal{E}$; this is mainly caused by the so-called *aperture problem*, which results from the impossibility of detecting or discriminating certain types of motion in a very regular image.

Since the development of the Horn–Schunck functional, several extensions and improvements of the regularization term have been developed; see, e.g., [4, 12, 13, 29, 30, 32]. A survey on variational techniques for optical flow can be found in [40]. In [31], Nagel proposed adding *regularization in time* via smoothing across image discontinuities. *Time smoothing of the flow field*, on the other hand, was introduced by Weickert and Schnörr in [42], where they considered an additional term containing the time derivative of the vector field u in the definition of the regularization functional. This alteration still yields a convex energy functional, and thus the well-posedness of the optical flow problem can be proved by employing methods similar to those used for the original Horn–Schunck functional. While these results have been derived for domains in $\mathbb{R}^2$, the situation of more general—possibly curved—image domains has not been considered in [42]. A first attempt in this direction can be found in [21, 38], where the authors introduced the optical flow functional for images on the round sphere. Finally, the case of an arbitrary compact two-dimensional manifold as image domain has been studied in [26]. There the authors discuss the usage of the Horn–Schunck functional on a manifold and prove a well-posedness result similar to that for the plane.

Time varying image domains. Recently, Kirisits, Lang, and Scherzer studied the optical flow problem on a time varying image domain [24, 25]. The motivation for that was an application in volumetric microscopy, where one studies the early development of a zebra fish embryo. In this setting, almost all movement between consecutive images takes place on the surface of the embryo's yolk cell, which, however, is time-dependent as well. In theory, it would be possible to use the complete volumetric data in order to compute a three-dimensional optical flow field. In practice, however, this is not viable because of the huge amount of data involved. Instead, it makes sense to extract the moving surface in a first step and then compute the flow field on this surface in a second, separate step.

The main mathematical challenge at this point is the correct treatment of a vector field on a moving manifold $\mathcal{M}_t \subset \mathbb{R}^3$, $t \in [0,T]$. We assume in this paper that this manifold is given by a family of parametrizations $f(t,\cdot)\colon M \to \mathbb{R}^3$, where the configuration space M is a fixed compact two-dimensional manifold (possibly with boundary). The image sequence is defined on this moving manifold, and it is assumed that the structure of the manifold has an influence on the deformation of the image sequence. The difficulty is *to capture the structure of the moving manifold in the optical flow field.* Therefore, one has to develop a regularization term that depends on *the induced, changing Riemannian metric.*

At this point, we want to note that it would, in principle, be possible to use some fixed Riemannian metric on M in order to obtain a regularization term as in [26]. Then one would, however, lose all the information about the correct manifold $\mathcal{M}_t$ as well as its movement in space. In the experimental section, section 6, we will compare this naive approach to our geometrical method, and we will see that there is a significant difference in the resulting optical flow fields. Furthermore, we note that similar situations and challenges occur in other applications in biological and medical imaging, for instance, during the reconstruction of the epicardial potential, if one wants to take into account the beating of the heart, and our approach can readily be translated to other application that require spatio-temporal regularization on moving surfaces.

Contributions of the paper. One possibility of a regularization term capturing the structure of a moving manifold has already been given in [24, 25]. In this paper we propose a different term that is induced by a metric $\bar{g}$ on the product manifold $\bar{M} = [0,T] \times M$. This metric $\bar{g}$ is constructed in such a way that it incorporates all available information on the moving image domain:

$$\bar{g}(\cdot,\cdot) = \begin{pmatrix} \alpha^2 & 0 \\ 0 & f^*\langle\cdot,\cdot\rangle_{\mathbb{R}^3} \end{pmatrix}.$$

The constant $\alpha > 0$ is a weighting parameter, and $f^*\langle\cdot,\cdot\rangle_{\mathbb{R}^3}$ denotes the induced surface metric of the parametrization f at time t. Given such a metric, we can use a weighted H^1-norm as regularization term:

$$\mathcal{R}(\bar{u}) = \int_{\bar{M}} \beta\bar{g}(\bar{u},\bar{u}) + \gamma\bar{g}^1_1(\bar{\nabla}\bar{u},\bar{\nabla}\bar{u})\,\mathrm{vol}(\bar{g}).$$

This regularization term is defined for vector fields $\bar{u}$ on the product manifold $\bar{M}$. However, since we do not want to change the time parametrization, we will consider only vector fields with vanishing time component; cf. Remark 7 for a more detailed explanation of this choice. Moreover, $\bar{g}^1_1$ denotes the extension of the metric to 1-1 tensor fields, $\bar{\nabla}$ denotes the covariant derivative of $\bar{g}$, and $\mathrm{vol}(\bar{g})$ is the corresponding volume form. Note that this term enforces spatio-temporal regularity, as it contains derivatives in both time and space. The parameter α that is included in the definition of the metric allows us to penalize regularity in time and space separately. This choice for the regularization term is a natural geometric generalization of the regularization term on the static manifold $[0,T] \times \mathbb{R}^2$ from [42].

If we decide to enforce no regularity in time, then the optical flow problem reduces for each time point t_i to the optical flow problem on the static manifold $\mathcal{M}_{t_i}$. In this case, our regularization term equals the regularization term used in [26].

The similarity term we use in this paper is simply the squared L^2-norm of the defect of the optical flow equation, that is,

$$\mathcal{S}(\bar{u}) = \int_{\bar{M}} (\partial_t I + g(\nabla^g I, u))^2\,\mathrm{vol}(\bar{g})\,.$$

Regarding the well-posedness of this optical flow problem, we obtain the following result.

Theorem 1 (well-posedness of the optical flow problem). *Let $\mathcal{M}_t \subset \mathbb{R}^3$ be a moving two-dimensional compact surface, the movement of which is described by a family of parametrizations $f\colon M \to \mathbb{R}^3$. For all parameters $\beta,\gamma > 0$ and any image sequence $I \in W^{1,\infty}(\bar{M})$, the optical flow functional*

$$\mathcal{E}(\bar{u}) = \mathcal{S}(\bar{u}) + \mathcal{R}(\bar{u})$$

has a unique minimizer in

$$\operatorname{dom}(\mathcal{E}) := \{\bar{u} \in H^1(\bar{M}, T\bar{M}) : \bar{u} = (0\partial_t, u) \textit{ with } u \in H^1(\bar{M}, TM)\}\,.$$

A similar result is also shown under the assumption of partial Dirichlet boundary conditions. The L^2-norm in the regularization term is added to enforce the coercivity of the energy functional. We also discuss under which assumptions we can set the parameter β to zero and still obtain a well-posedness result for our functional. We compare our functional to the functionals introduced by Kirisits, Lang, and Scherzer [24, 25] and discuss the well-posedness of the optical flow problem using the regularization terms that are employed there.

Finally, we demonstrate the functionality of our approach in two experiments using both synthetic and real data. In these experiments we also show the difference between our approach and the straightforward approach, which does not use the actual structure of the moving manifold. In both experiments one can see notably different results in regions of the manifold where either the curvature or the deformation of the manifold is large.

Another important topic is the strong dependence of the optical flow field on the parametrization of the moving manifold. In Appendix A we present a brief discussion on a possible approach to computing realistic parametrizations given an observed moving unparametrized manifold. The long-term goal will be the combination of segmentation and computation of the optical flow, which we hope will lead to more reliable results.

Organization of the paper. In section 2 we recall the differential geometric and functional analytic tools that we will use throughout the paper. Readers who are acquainted with the theory of Sobolev spaces of vector fields on Riemannian manifolds might skip this part and directly start with section 3, which contains the rigorous mathematical formulation of the optical flow problem studied in this paper. In section 4 we construct the regularization term that we employ in this paper, prove the well-posedness of the corresponding functional, and derive the optimality conditions. Up to this point all calculations and results are presented in a coordinate-independent manner. In order to obtain an implementable version, we derive in section 5 a coordinate version of the optimality conditions. This involves rather technical calculations, which are partly postponed to Appendix B. In section 6 we show numerical experiments that demonstrate the functionality of the proposed energy functional. Appendix A contains a discussion on how to compute the parametrization of the moving manifold and the actual calculations of the coordinate version of the optimality conditions.

2. Mathematical preliminaries. In this section we are going to recall the differential geometric and functional analytic tools for Sobolev spaces of vector fields on two-dimensional embedded surfaces, which we will need throughout the paper. A more detailed overview of these topics can be found in, e.g., [10, sect. 3].

Riemannian geometry. We are working on two-dimensional surfaces that are embedded in $\mathbb{R}^3$ and parametrized by a mapping

$$f : M \to \mathbb{R}^3$$

from some *configuration space* M into $\mathbb{R}^3$. We will always assume that M is a compact two-dimensional manifold, possibly with boundary; typical examples are the two-dimensional sphere S^2 or the torus $S^1 \times S^1$. The mapping f is assumed to be smooth (that is, at least

C^2) and injective with injective tangential mapping $Tf\colon TM \to \mathbb{R}^3$ (in other words, f is a smooth embedding).

The embedding f induces in a natural way via pullback a Riemannian metric g on the configuration space M. For tangent vectors $X, Y \in T_xM$, $x \in M$, it is given by

$$g(X,Y) := (f^*\langle\cdot,\cdot\rangle_{\mathbb{R}^3})(X,Y) := \langle T_xf.X, T_xf.Y\rangle_{\mathbb{R}^3}.$$

Here $\langle\cdot,\cdot\rangle_{\mathbb{R}^3}$ denotes the standard scalar product on $\mathbb{R}^3$, and $T_xf.X$ denotes the application of the differential $T_xf : T_xM \to \mathbb{R}^3$ of the embedding f to the tangent vector $X \in T_xM$.

In a chart (V, v) on M the expression of the metric reads as

$$g|_V = \sum_{i,j} g_{ij} dv^i \otimes dv^j = \sum_{i,j} \langle \partial_i f, \partial_j f\rangle_{\mathbb{R}^3} dv^i \otimes dv^j$$

with $\partial_i = \frac{\partial}{\partial v^i}$.

Next we note that the metric induces an isomorphism $\check{g}$ between the tangent bundle and the cotangent bundle defined by

$$\check{g}\colon TM \to T^*M, \qquad X \mapsto g(X,\cdot) := X^\flat,$$

with inverse $\check{g}^{-1}$. Therefore g defines a metric on the cotangent bundle T^*M via

$$g^{-1}(\alpha,\beta) = \alpha(\check{g}^{-1}(\beta)).$$

In this paper we will need the extension of the metric to 1-1 tensor fields. The reason for this is that this type of tensor field occurs as the derivative of a vector field on M, which will be a part of our regularization term. On these tensor fields the metric is given by

$$g_1^1 := g \otimes g^{-1}.$$

Applied to a 1-1 tensor field A, this equals the squared Hilbert–Schmidt norm of A:

$$g_1^1(A, A) = \mathrm{Tr}(A^*A),$$

where the adjoint A^* is computed with respect to the Riemannian metric g. Here we have interpreted A as a linear mapping from T_xM to T_xM.

Sobolev spaces of vector fields. The Riemannian metric g on M induces a unique volume density, which we will denote by $\mathrm{vol}(g)$. In the chart (V, v) its formula reads as

$$\mathrm{vol}(g)|_V = \sqrt{\det(\langle\partial_i f, \partial_j f\rangle_{\mathbb{R}^3})}\;|dv^1 \wedge dv^2|\,.$$

The Levi–Civita covariant derivative of the metric g, which is the unique torsion-free connection preserving the metric g, will be denoted by ∇^g. When it is clear from the context, we omit the g and simply write ∇ instead of ∇^g. Note that ∇^g is just the tangential component of the usual derivative in the ambient space $\mathbb{R}^3$; more precisely, for a vector field $u \in C^\infty(M, TM)$,

$$Tf.\nabla^g u = \mathrm{proj}_{T^{1,1}M}\,\nabla^{\mathbb{R}^3}(Tf.u)\,.$$

We define the Sobolev norms of orders zero and one by

$$\|u\|_{0,g}^2 = \int_M g(u,u)\,\mathrm{vol}(g),$$
$$\|u\|_{1,g}^2 = \int_M g(u,u) + g_1^1(\nabla^g u, \nabla^g u)\,\mathrm{vol}(g).$$

The Sobolev space $H^1(M, TM)$ is then defined as the completion of the space of all vector fields $u \in C^\infty(M, TM)$ with respect to the norm $\|\cdot\|_{1,g}$. On a compact manifold different metrics yield equivalent norms and thus lead to the same Sobolev spaces. We note that, as for Sobolev spaces in $\mathbb{R}^n$, there is an alternative, equivalent definition of $H^1(M, TM)$ as the space of all square integrable vector fields with square integrable weak derivatives.

For the definition of more general Sobolev spaces on manifolds, we refer the reader to [39]; see also [7] for an exposition in notation similar to that used in the current paper. An extension of this theory to noncompact manifolds can be found in the book [14].

3. Problem formulation. We assume that we are given a moving two-dimensional compact surface $\mathcal{M}_t \subset \mathbb{R}^3$, $t \in [0,T]$, the movement of which is described by a family of parametrizations f. For the moment, we will restrict ourselves to compact surfaces without boundary, but we will discuss the situation of manifolds with boundary later. More precisely, we assume that there exist a two-dimensional compact C^2-manifold M and a C^2-mapping

$$f\colon [0,T] \times M \to \mathbb{R}^3$$

such that for every fixed time $t \in [0,T]$ the mapping $f(t,\cdot)$ is an embedding and its image equals $\mathcal{M}_t$. The mapping f defines the movement of the manifold in the sense that the path of a point $y = f(0,x) \in \mathcal{M}_0$ is precisely the curve $t \mapsto f(t,x)$. Or, a point $y_1 \in \mathcal{M}_{t_1}$ corresponds to a point $y_2 \in \mathcal{M}_{t_2}$ if and only if there exists $x \in M$ with $y_1 = f(t_1, x)$ and $y_2 = f(t_2, x)$.

Next we model the movement of an image on this moving surface. For simplicity we will consider only gray-scale images, although the model does not change significantly if we also allow color, that is, vector-valued, images. We stress that in our model the movement of the image is not driven solely by the movement *of* the surface, but that there is also an additional movement *on* the surface, the reconstruction of which is precisely what we are aiming for.

The image sequence we are considering is given by a real-valued function $\mathcal{I}$ on

$$\mathcal{M} := \bigcup_{0 \le t \le T} \{t\} \times \mathcal{M}_t \subset [0,T] \times \mathbb{R}^3\,;$$

for each $t \in [0,T]$, the function $\mathcal{I}(t,\cdot)\colon \mathcal{M}_t \to \mathbb{R}$ is the image at the time t. Moreover, there exists a family of diffeomorphisms $\psi(t,\cdot)\colon \mathcal{M}_0 \to \mathcal{M}_t$ such that

$$\mathcal{I}(0,x) = \mathcal{I}(t, \psi(t,x)).$$

That is, the diffeomorphisms $\psi(t,\cdot)$ generate the movement of the image on the evolving surface.

Next, it is possible to pull back the image and the driving family of diffeomorphisms to the configuration space M. Doing so, we obtain a time-dependent function $I\colon [0,T]\times M \to \mathbb{R}$ defined by

$$I(t,x) = \mathcal{I}(t, f(t,x))$$

and a family of diffeomorphisms $\varphi(t,\cdot)$ of M defined by

$$f(t,\varphi(t,x)) = \psi(t, f(0,x))$$

such that

$$I(0,x) = I(t,\varphi(t,x)) \tag{1}$$

for all $t \in [0,T]$ and $x \in M$.

Furthermore, we assume that the diffeomorphisms $\varphi(t,\cdot)$ are generated by a time-dependent vector field u on M. Then the curves $t \mapsto \varphi(t,x)$ are precisely the integral curves of u; that is,

$$\partial_t \varphi(t,x) = u(t,\varphi(t,x))\,. \tag{2}$$

If the image I is sufficiently smooth, it is possible to compute the time derivative of (1). Using the relation (2) and the fact that $\varphi(t,\cdot)$ is surjective, we then see that the image I and the vector field u satisfy the *optical flow equation*

$$0 = \partial_t I(t,x) + D_x I(t,x) u(t,x) \tag{3}$$

on $[0,T]\times M$. We do note that in (3) all information about the movement of the manifold is suppressed, as all the functions have been pulled back to M. It is, however, possible to reintroduce some knowledge of $\mathcal{M}$ by formulating the optical flow equation not in terms of differentials but rather in terms of gradients. To that end we denote by g the time-dependent Riemannian metric on M that is induced by the family of embeddings $f(t,\cdot)$. Since by definition $\nabla^g I(t,\cdot)^\flat = D_x I(t,\cdot)$, we can rewrite the optical flow equation as

$$0 = \partial_t I(t,x) + g(\nabla^g I(t,x), u(t,x)) \tag{4}$$

for all $(t,x) \in [0,T]\times M$.

Now assume that the model manifold M is a compact manifold with boundary. Then the same model of a moving image on the embedded manifolds $\mathcal{M}_t$ is possible, as long as it is guaranteed that the boundary of the manifold acts as a barrier for the movement of $\mathcal{I}$. That is, the diffeomorphisms $\varphi(t,\cdot)$ satisfy the additional boundary condition $\varphi(t,x) = x$ for $x \in \partial M$. In this case, one arrives at the same optical flow equation (4), but, additionally, one obtains (partial) Dirichlet boundary conditions of the form $u(t,x) = 0$ for all $(t,x) \in [0,T]\times \partial M$.

The situation is different when the image $\mathcal{I}$ actually moves across the boundary of $\mathcal{M}_t$, which can occur if the manifold with boundary $\mathcal{M}_t$ represents the limited field of view on a larger manifold that contains the moving image. Then it is not reasonable to model the movement of the image by a family of global diffeomorphisms $\psi(t,\cdot)$. However, locally it can still be modeled as being generated by a family of local diffeomorphisms, which in turn can be assumed to be generated by a time-dependent vector field on M. With this approach, one arrives, again, at the same optical flow equation (4). The difference from the situations discussed above is that the integral curves of u may be defined only on bounded intervals.

The inverse problem. Now we consider the inverse problem of reconstructing the *movement* of a family of images from the image sequence. We assume that we are given the family of manifolds $\mathcal{M}_t$ together with the parametrizations $f(t,\cdot)$ and the family of images $I(t,\cdot)$ (already pulled back to M). Our task is to find a time-dependent vector field u on M that generates the movement of I; in other words, a vector field u that satisfies the optical flow equation (4).

Solving this equation directly is not sensible, as, in general, the solution, if it exists, will not be unique: The optical flow equation does not "see" a flow that is tangential to the level lines of the image I. Thus, if u is any solution of (4) and the vector field w satisfies

$$g(\nabla^g I(t,x), w(t,x)) = 0,$$

then $u+w$ is also a solution; this is called the *aperture problem* (see [16]). In addition, the whole model fails in the case of noise leading to nondifferentiable data I. In order to be still able to formulate the optical flow equation, it is possible to presmooth the image I, but this will invariably lead to errors in the model, and thus the optical flow equation will be satisfied only approximately by the generating vector field u. For these reasons, it is necessary to introduce some kind of regularization. Note that the main focus here lies in the problem of solution selection.

Intrinsic invariances. Usually, there is no canonical way of choosing the model manifold M and the initial parametrization $f(0,\cdot)$ of $\mathcal{M}_0$. However, the optical flow problem is invariant with respect to reparametrizations in the following sense: If N is another (diffeomorphic) manifold and $\chi\colon N \to M$ a diffeomorphism, then we can pull back everything to N. That is, we obtain parametrizations $f_N(t,\cdot) := f(t,\cdot)\circ\chi$ and images $I_N(t,\cdot) := I(t,\cdot)\circ\chi$. Following the same steps as before, we see that the movement of the image I_N is generated by a time-dependent vector field u_N on N that satisfies the optical flow equation on N,

$$0 = \partial_t I_N(t,x) + D_x I_N(t,x) u_N(t,x). \tag{5}$$

By construction, $\partial_t I_N(t,x) = \partial_t I(t,\chi(x))$, and $D_x I_N(t,x) = D_x I(t,\chi(x)) \circ D\chi(x)$. As a consequence, the vector field $u(t,\chi(x)) := D\chi(x) u_N(t,x)$, which is the pushforward of u_N by means of χ, satisfies the optical flow equation on M. In other words, the solutions of the optical flow equation on M are simply the pushforward of solutions of the optical flow equation on N via the diffeomorphism χ connecting N and M. One main goal of the regularization method we will develop in the following is to retain this invariance also for the inverse problem.

4. Classical variational regularization.

One of the most straightforward regularization methods is the application of Tikhonov regularization, where we try to minimize a functional composed of two terms—a similarity term, which ensures that the equation is almost satisfied, and a regularity term, which ensures the existence of a regularized solution and is responsible for the solution selection.

4.1. Spatial regularization.

If we consider only spatial regularity, the definition of the regularity term is straightforward, using for each time point t the pullback metric

$$g(t)(\cdot,\cdot) = f(t,\cdot)^*\langle\cdot,\cdot\rangle_{\mathbb{R}^3}\,.$$

This leads to the energy functional

$$\boxed{\begin{aligned}\mathcal{E}(u) &:= \mathcal{S}(u) + \mathcal{R}(u)\\ &= \int_0^T \int_M (\partial_t I + g(\nabla^g I, u))^2 + \beta g(u,u) + \gamma g_1^1(\nabla^g u, \nabla^g u)\,\mathrm{vol}(g)\,dt,\end{aligned}}$$

where β and γ are weighting parameters. In this case the problem completely decouples in space and time; i.e., the optimal vector field u has to be minimal for each time point separately. Thus the problem reduces for each time t to the calculation of the optical flow on the (static) Riemannian manifold $(M, g(t))$, which yields for each time point t the energy functional

$$\boxed{\mathcal{E}(u(t,\cdot)) := \int_M (\partial_t I + g(\nabla^g I, u))^2 + \beta g(u,u) + \gamma g_1^1(\nabla^g u, \nabla^g u)\,\mathrm{vol}(g).}$$

This functional is well investigated. For $\beta > 0$ the coercivity of the energy functional is clear, and one can easily deduce the well-posedness of the optical flow problem. In [26, 36] it was shown that one can guarantee the coercivity of the energy functional for $\beta = 0$ by requiring the image sequence to satisfy additional conditions. The conditions in [36] for optical flow in the plane require that the partial derivatives of the image I are linearly independent functions. This is equivalent to the requirement that no nontrivial constant vector field u satisfies the optical flow equation for the given image. Similar requirements are commonly found for Tikhonov regularization with derivative-based regularization terms; see, e.g., [1, 5, 17] and [34, sect. 3.4]. In the case of a nonflat manifold M, the condition translates to the nonexistence of a nontrivial covariantly constant vector field satisfying the optical flow equation. Obviously, this condition is automatically satisfied if the only covariantly constant vector field is $u = 0$, and thus it may be omitted in manifold settings; see [26].

In addition, note that the energy functional defined above does not depend on the parametrization of $\mathcal{M}$. That is, assuming that $\chi\colon N \to M$ is a diffeomorphism between a manifold N and M, we can pull back the optical flow equation to N as in (5). If we then use the same method as above to define an energy functional, then the resulting functional $\mathcal{E}_N$ will satisfy $\mathcal{E}_N(u(t,\cdot)) = \mathcal{E}(D\chi \circ u(t,\cdot))$ whenever u is a time-dependent vector field on N. In particular, this implies that the minimizer of $\mathcal{E}$ on M is the pushforward of the minimizer of $\mathcal{E}_N$ on N, which further means that the resulting vector fields on $\mathcal{M}_t$ are identical. We remark that this property does not hold if one uses an arbitrary Riemannian metric on M that is not inherited from $\mathcal{M}$. For instance, if $M \subset \mathbb{R}^2$ is an open set, it is in principle possible to compute an energy functional based on the Euclidean metric. In this case, a reparametrization of M will in general yield a different flow field.

4.2. Regularization in time and space. In the following we will look for solutions that additionally satisfy a regularity constraint in time t. For the optical flow in the plane $\mathbb{R}^2$, this method has been introduced in [42]. Spatio-temporal regularization of the optical flow on moving manifolds has also been considered in [24], but with a different regularity term from the one we will construct in the following.

In order to construct the regularity term, we consider the product manifold

$$\bar{M} := [0,T] \times M$$

and equip it with the almost product metric

$$\bar{g}(\cdot,\cdot) = \begin{pmatrix} \alpha^2 & 0 \\ 0 & f^*\langle\cdot,\cdot\rangle_{\mathbb{R}^3} \end{pmatrix}.$$

The parameter $\alpha > 0$ is a weighting parameter, which is included in order to be able to penalize spatial regularity and regularity in the time variable differently. This metric is called the almost product metric due to the dependence of the metric $g(\cdot,\cdot) = f^*\langle\cdot,\cdot\rangle_{\mathbb{R}^3}$ on the time t. In order to simplify notation we denote $\nabla^{\bar{g}}$ by $\bar{\nabla}$ from now on.

Remark 2. In the following, we will always indicate by a "bar" $(\bar{\cdot})$ that an object is related to the product manifold $\bar{M}$. For instance, $\bar{g}$ denotes a metric on $\bar{M}$, whereas g denotes a (time-dependent) metric on M. Similarly, $\bar{u}$ will later denote a vector field on $\bar{M}$, whereas u will denote a time-dependent vector field on M.

Remark 3. We could also consider $\mathcal{M}$ as an embedded submanifold of $\mathbb{R} \times \mathbb{R}^3$:

$$\bar{f}\colon \begin{cases} [0,T] \times M & \to \quad \mathbb{R} \times \mathbb{R}^3\,, \\ \quad (t,x) & \mapsto \quad (t, f(t,x))\,. \end{cases}$$

We stress here that the metric $\bar{g}$ is *not* the pullback of the (time-scaled) Euclidean metric on $\mathcal{M}$ by the parametrization $\bar{f}$. Instead, it is constructed in such a way that the paths of the points on M are at each time t orthogonal to the manifold $\mathcal{M}_t$. Moreover, these paths are geodesics with respect to $\bar{g}$. These properties do not, in general, hold for the usual pullback metric.

From now on we will identify a time-dependent vector field u on M with the vector field

$$\bar{u}(t,x) := (0\partial_t, u(t,x)) \in C^\infty(\bar{M}, T\bar{M})$$

and define both the similarity term and the regularization term in terms of $\bar{u}$. Taking the squared L^2-norm with respect to the metric $\bar{g}$ of the right-hand side of the optical flow equation (3), we obtain for the similarity term the functional

$$\begin{aligned}
\mathcal{S}(\bar{u}) &= \|\partial_t I + g(\nabla^g I, u)\|_{0,\bar{g}}^2 \\
&= \int_{\bar{M}} \left(\partial_t I(t,x) + g(\nabla^g I(t,x), u(t,x))\right)^2 \operatorname{vol}(\bar{g}) \\
&= \alpha \int_0^T \int_M \left(\partial_t I(t,x) + g(\nabla^g I(t,x), u(t,x))\right)^2 \operatorname{vol}(g)\, dt.
\end{aligned}$$

Here we used the fact that the volume form on $\bar{M}$ splits into $\operatorname{vol}(\bar{g}) = \alpha \operatorname{vol}(g)\, dt$.

For the regularization term we use a weighted H^1-norm of the vector field $\bar{u}$, that is,

$$\mathcal{R}(\bar{u}) = \beta\|\bar{u}\|_{0,\bar{g}}^2 + \gamma\|\bar{\nabla}\bar{u}\|_{0,\bar{g}}^2\,, \tag{6}$$

where β and γ are weighting parameters. In (6), the term $\|\bar{u}\|_{0,\bar{g}}^2$ denotes the L^2-norm of the vector field $\bar{u}$, and $\|\bar{\nabla}\bar{u}\|_{0,\bar{g}}^2$ denotes the norm of its derivative with respect to the Riemannian metric $\bar{g}$; that is,

$$\|\bar{u}\|_{0,\bar{g}}^2 = \alpha \int_0^T \int_M g(u,u) \operatorname{vol}(g)\, dt, \tag{7}$$

$$\|\bar{\nabla}\bar{u}\|_{0,\bar{g}}^2 = \alpha \int_0^T \int_M \bar{g}_1^1(\bar{\nabla}\bar{u}, \bar{\nabla}\bar{u}) \operatorname{vol}(g)\, dt. \tag{8}$$

To summarize, we propose to solve the optical flow problem on a moving manifold by minimizing the energy functional

$$\boxed{\begin{aligned}\mathcal{E}(\bar{u}) &:= \mathcal{S}(\bar{u}) + \mathcal{R}(\bar{u}) \\ &= \|\partial_t I + g(\nabla^g I, u)\|^2_{0,\bar{g}} + \beta\|\bar{u}\|^2_{0,\bar{g}} + \gamma\|\bar{\nabla}\bar{u}\|^2_{0,\bar{g}}\,,\end{aligned}} \tag{9}$$

which is defined on

$$\mathrm{dom}(\mathcal{E}) = \left\{\bar{u} \in H^1(\bar{M}, T\bar{M}) : \bar{u} = (0\partial_t, u)\right\}.$$

Remark 4. Note that the regularization term depends implicitly on the parameter α as well. However, a large value of α leads to less time regularity, which is in contrast to the influence of the parameters β and γ. Formally, the limit $\alpha \to \infty$ corresponds to no time regularization at all.

Remark 5. We stress the difference between the regularization term proposed in this paper and the one from [25], which is given by

$$\mathcal{R}(u) = \int_0^T \int_{\mathcal{M}_t} \lambda_0 |\mathrm{proj}_{T\mathcal{M}_t}\, \partial_t(Tf.u)|^2 + \lambda_1 \|\mathrm{proj}_{T^{1,1}\mathcal{M}_t}\, \nabla^{\mathbb{R}^3}(Tf.u)\|^2.$$

Even though the latter functional is also a natural generalization of [42]—from an embedded point of view—there is no obvious metric on $\bar{M}$ for which it is a weighted homogeneous H^1-norm.

Remark 6. In the case $\beta = 0$, where $\mathcal{R}$ is the homogeneous Sobolev seminorm, only variations of the movement on the manifold are penalized but not the overall speed of the movement. In contrast, a positive value of β encourages a low speed, which may lead to a systematic underestimation of the magnitude of the computed flow. For this reason the choice $\beta = 0$ is usually preferable. Note, however, that one of the basic assumptions in our model is that most of the movement of the image is driven by the movement of the manifold. Thus, using a positive value of β can be justified and is somehow natural, provided that this assumption holds. In addition, the actual numerical computation of the flow field is easier for $\beta > 0$ because the condition of the resulting linear equation becomes better with increasing β. Still, we have used the parameter choice $\beta = 0$ for our numerical experiments later in the paper.

Remark 7. It is also possible to identify the nonautonomous vector field u on M with the vector field $\hat{u}(t,x) := (1\partial_t, u(t,x))$ on $\bar{M}$, which incorporates the movement of the image in both time and space. If one does so, however, one has to be careful about the regularity term. Simply using the squared (weighted) H^1-norm of $\hat{u}$ has the undesirable effect that the natural movement of the manifold, which is given by the vector field $\hat{u}_0 := (1\partial_t, 0)$, need not be of minimal energy for the regularization term: The vector field $\hat{u}_0$ is in general not covariantly constant. Instead of the norm of the vector field $\hat{u}$ itself, one should therefore penalize the norm of the difference between $\hat{u}$ and $\hat{u}_0$. Doing so, one arrives at the same regularization term (8) as above, although the interpretation is slightly different.

Remark 8. From the construction of the Riemannian metric, it follows that the proposed energy functional is invariant under diffeomorphisms of the model manifold M: If we pull back the optical flow equation to a manifold N by means of a diffeomorphism χ as in (5) and use the same method for constructing an energy functional $\mathcal{E}_N$ on $[0,T] \times N$, then for every vector field $\bar{u} = (0\partial_t, u)$ on $[0,T] \times N$, the equality $\mathcal{E}_N(\bar{u}) = \mathcal{E}((0\partial_t, D\chi \circ u))$ holds.

4.3. Well-posedness. The proof of the well-posedness of our model, that is, the question of whether the proposed energy functional $\mathcal{E}$ attains a unique minimizer in $\mathrm{dom}(\mathcal{E})$, is quite straightforward. In the following result, we denote by $W^{1,\infty}(\bar{M})$ the space of functions on $\bar{M}$ with an essentially bounded weak derivative.

Theorem 9. *Assume that $\alpha, \beta, \gamma > 0$ and that $I \in W^{1,\infty}(\bar{M})$. Then the functional*

$$\mathcal{E}(\bar{u}) = \mathcal{S}(\bar{u}) + \mathcal{R}(\bar{u})$$

defined in (9) *has a unique minimizer in* $\mathrm{dom}(\mathcal{E})$.

Remark 10. We will see in section 4.4 that this optimization problem is in a natural way connected to Neumann boundary conditions. If we want to consider mixed boundary conditions instead—Dirichlet in space and Neumann in time—we have to restrict the domain of the energy functional to

$$\mathrm{dom}_0(\mathcal{E}) := \{\bar{u} \in \mathrm{dom}(\mathcal{E}) : \bar{u} = 0 \text{ on } [0,T] \times \partial M\}.$$

The well-posedness result remains valid on $\mathrm{dom}_0(\mathcal{E})$.

Proof. The condition $I \in W^{1,\infty}(\bar{M})$ guarantees that the similarity term $\mathcal{S}(\bar{u})$ is finite for every square integrable vector field on $\bar{M}$; in particular, it is proper. From the condition $\beta > 0$ we obtain that the regularization term $\mathcal{R}$ and therefore also the energy functional $\mathcal{E}$ are coercive. Thus $\mathcal{E}$ is a proper and coercive, quadratic functional on the Hilbert space $\mathrm{dom}(\mathcal{E})$, which implies the existence of a unique minimizer (cf. [34, sect. 3.4] or [36]). ∎

Remark 11. The condition $\beta > 0$ is not necessary if there is another way of guaranteeing the coercivity of the regularization term. This is possible, for instance, if there exists no nontrivial covariantly constant vector field of the form $\bar{u} = (0\partial_t, u)$ on $\bar{M}$. In that case, the homogeneous Sobolev seminorm $\|\bar{\nabla}\bar{u}\|^2_{0,\bar{g}}$ is, in fact, a norm on $\mathrm{dom}(\mathcal{E})$ that is equivalent to the standard Sobolev norm, and therefore also the parameter choice $\beta = 0$ guarantees the coercivity of $\mathcal{E}$. Note that this condition is independent of the moving image I.

More generally, even if there are nontrivial, admissible, covariantly constant vector fields on $\bar{M}$, the energy function will still be coercive for $\beta = 0$, as long as no such vector field satisfies the optical flow equation $\partial_t I + g(\nabla^g I, u) = 0$. Note, however, that the numerical computation of a minimizer may become difficult because the problem, though still well-posed, may become ill-conditioned as the parameter β approaches zero.

Remark 12. With a similar argumentation one can show that the functionals proposed in [24, 25] are well-posed, provided that they are coercive; cf. Remark 5. Because the regularization terms in these papers penalize only the derivative of the vector field u but not its size, the coercivity (and thus well-posedness) will hold only if one of the conditions in Remark 11 is satisfied.

4.4. The optimality conditions.

Lemma 13. *The L^2-gradient of the optical flow energy functional $\mathcal{E}$ is given by*

$$\mathrm{grad}\,\mathcal{E}(\bar{u}) = 2\Big(\partial_t I + g\big(\nabla^g I, u\big)\Big)(0, \nabla^g I) + 2\beta\bar{u} + 2\gamma\Delta^{\mathrm{B}}\bar{u}\,.$$

- *For $\mathcal{E}$ seen as functional on $\operatorname{dom}(\mathcal{E})$, its domain of definition is the set of all vector fields $\bar{u} \in \operatorname{dom}(\mathcal{E})$ satisfying Neumann boundary conditions, i.e.,*

$$\operatorname{dom}(\operatorname{grad}(\mathcal{E})) = \left\{\bar{u} \in \operatorname{dom}(\mathcal{E}) : \bar{\nabla}_\nu \bar{u}\big|_{\partial \bar{M}} = 0\right\},$$

where ν denotes the normal to the boundary of $\bar{M}$ with respect to $\bar{g}$.
- *For $\mathcal{E}$ restricted to $\operatorname{dom}_0(\mathcal{E})$, its domain of definition is the set of all vector fields $\bar{u} \in \operatorname{dom}_0(\mathcal{E})$ satisfying mixed boundary conditions, more precisely,*

$$\operatorname{dom}_0(\operatorname{grad}(\mathcal{E})) = \left\{\bar{u} \in \operatorname{dom}_0(\mathcal{E}) : \bar{\nabla}_\nu \bar{u}\big|_{\{0,T\}\times M} = 0\right\}.$$

Note here that on $\{0,T\} \times M$ the normal vector ν is given by $\nu = \partial_t$.

Here Δ^{B} denotes the Bochner Laplacian of $\bar{g}$, which is defined via

$$\Delta^{\mathrm{B}} = \bar{\nabla}^* \bar{\nabla},$$

with $\bar{\nabla}^*$ denoting the L^2-adjoint of the covariant derivative. The Bochner Laplacian differs only by a sign from the usual Laplace–Beltrami operator.

Proof. We calculate the gradients for the two terms separately. Using a variation $\overline{\delta u} = (0, \delta u)$, we obtain the following expression for the variation of the similarity term:

$$D\left(\mathcal{S}(\bar{u})\right)(\overline{\delta u}) = 2\alpha \int_0^T \int_M \Big(\partial_t I + g\big(\nabla^g I, u\big)\Big) g\big(\nabla^g I, \delta u\big) \operatorname{vol}(g)\, dt.$$

From this equation one can easily read off the L^2-gradient of the similarity term. It reads as

$$\operatorname{grad}^{L^2}(\mathcal{S}(\bar{u})) = 2\Big(\partial_t I + g\big(\nabla^g I, u\big)\Big)(0, \nabla^g I).$$

The variation of the regularization term is given by

$$\begin{aligned} D\left(\mathcal{R}(\bar{u})\right)(\overline{\delta u}) &= 2\alpha\beta \int_0^T \int_M \bar{g}(\bar{u}, \overline{\delta u}) \operatorname{vol}(g)\, dt + 2\alpha\gamma \int_0^T \int_M \bar{g}^1_1(\bar{\nabla}\bar{u}, \bar{\nabla}\overline{\delta u}) \operatorname{vol}(g)\, dt \\ &= 2\alpha\beta \int_0^T \int_M \bar{g}(\bar{u}, \overline{\delta u}) \operatorname{vol}(g)\, dt + 2\alpha\gamma \int_0^T \int_M \bar{g}(\bar{\nabla}^*\bar{\nabla}\bar{u}, \overline{\delta u}) \operatorname{vol}(g)\, dt \\ &\quad + 2\alpha\gamma \int_0^T \int_{\partial M} \bar{g}(\bar{\nabla}_\nu \bar{u}, \overline{\delta u}) \operatorname{vol}(g)|_{\partial M}\, dt + 2\alpha\gamma \int_M \bar{g}(\bar{\nabla}_{\partial_t}\bar{u}, \overline{\delta u}) \operatorname{vol}(g)\Big|_0^T. \end{aligned}$$

The second step consists of a partial integration using the L^2-adjoint of the covariant derivative, which we denote by $\bar{\nabla}^*$. The last two terms in the above expression are the boundary terms that result from the partial integration. From this we can read off the formula for the gradient of the regularization term. Taking into account that the outer normal vector to the boundary of $\{0,T\} \times M$ is given by $\nu = \partial_t$, this concludes the proof on $\operatorname{dom}(\mathcal{E})$. For the proof on $\operatorname{dom}_0(\mathcal{E})$ the situation is simpler, since the first boundary integral is already zero if $\overline{\delta u} \in \operatorname{dom}_o(\mathcal{E})$. ■

Because of the strict convexity of the energy functional $\mathcal{E}$, a vector field $\bar{u}$ is a minimizer if and only if it is an element of $\operatorname{dom}(\operatorname{grad}(\mathcal{E}))$ and $\operatorname{grad}\mathcal{E}(u) = 0$. Thus we obtain the following result.

Theorem 14. *The minimizer of the energy functional* $\mathcal{E}$ *on* $\mathrm{dom}(\mathcal{E})$ *defined in* (9) *is the unique solution* $\bar{u} = (0\partial_t, u)$ *of the equation*

$$\begin{aligned} \Big(\partial_t I + g\big(\nabla^g I, u\big)\Big)(0, \nabla^g I) + \beta\bar{u} + \gamma\Delta^{\mathrm{B}}\bar{u} &= 0 \quad \textit{in } \bar{M}, \\ \bar{\nabla}_\nu \bar{u} = (0, \nabla_\nu u) &= 0 \quad \textit{in } [0,T]\times\partial M, \\ \bar{\nabla}_{\partial_t}\bar{u} &= 0 \quad \textit{in } \{0,T\}\times M. \end{aligned}$$

If we restrict the energy functional to $\mathrm{dom}_0(\mathcal{E})$*, it is the unique solution of*

$$\begin{aligned} \Big(\partial_t I + g\big(\nabla^g I, u\big)\Big)(0, \nabla^g I) + \beta\bar{u} + \gamma\Delta^{\mathrm{B}}\bar{u} &= 0 \quad \textit{in } \bar{M}, \\ \bar{\nabla}_{\partial_t}\bar{u} &= 0 \quad \textit{in } \{0,T\}\times M, \\ \bar{u} &= 0 \quad \textit{in } [0,T]\times\partial M. \end{aligned}$$

5. The optimality conditions in local coordinates.

The aim of this section is to express the previously derived optimality conditions in a local coordinate chart in order to obtain an implementable version of the previous sections. To simplify the exposition, we will restrict ourselves to the case of mixed boundary conditions. Note that this includes in particular the situation where M is a compact manifold without boundary.

Let (V, v) be a local chart on M with coordinate frame ∂_1, ∂_2. In the following we will use the Einstein summation convention in order to simplify the notation.

The main computational difficulty is the computation of the Bochner Laplacian $\Delta^{\mathrm{B}} = \bar{\nabla}^*\bar{\nabla}$, as it involves the adjoint of the covariant derivative. This is most easily done in an orthonormal frame with respect to the metric $\bar{g}$. We stress here that the natural frame $(\partial_t, \partial_1, \partial_2)$ is, in general, not orthonormal, because $\bar{g}(\partial_1, \partial_2) = g(\partial_1, \partial_2) = \langle \partial_1 f, \partial_2 f\rangle_{\mathbb{R}^3}$ will be different from 0. Note, however, that the construction of the metric implies that $\bar{g}(\partial_t, \partial_i) = 0$ for $i = 1, 2$. We can therefore obtain an orthonormal frame by scaling the vector ∂_t to unit length and, for instance, applying the Gram–Schmidt orthogonalization process to the (time- and space-dependent) vectors ∂_1, ∂_2. Doing so, we obtain an orthonormal frame of the form

$$\bar{X}_0 = \left(\frac{1}{\alpha}\partial_t, 0\right), \qquad \bar{X}_1 = (0, X_1), \qquad \bar{X}_2 = (0, X_2)$$

with space-dependent vector fields $X_1(t,\cdot)$ and $X_2(t,\cdot)$ on $M \cap V$. The (time- and space-dependent) coordinate change matrix between these two bases will be denoted by $\bar{A}$; we have

$$\begin{pmatrix} \bar{X}_0 \\ \bar{X}_1 \\ \bar{X}_2 \end{pmatrix} = \bar{A}\begin{pmatrix} \partial_t \\ \partial_1 \\ \partial_2 \end{pmatrix} = \begin{pmatrix} \frac{1}{\alpha} & 0 & 0 \\ 0 & a_1^1 & a_1^2 \\ 0 & a_2^1 & a_2^2 \end{pmatrix}\begin{pmatrix} \partial_t \\ \partial_1 \\ \partial_2 \end{pmatrix}.$$

Note that the coefficient function a_1^2 will be the constant 0 if the Gram–Schmidt process is used for orthogonalization.

In the orthonormal frame $\{\bar{X}_i\}$, the norm of $\bar{\nabla}\bar{u}$ can be written as

$$\bar{g}_1^1(\bar{\nabla}\bar{u}, \bar{\nabla}\bar{u}) = \sum_i \bar{g}(\bar{\nabla}_{\bar{X}_i}\bar{u}, \bar{\nabla}_{\bar{X}_i}\bar{u}),$$

where $\bar\nabla_{\bar X_i}\bar u$ denotes the covariant derivative of the vector field $\bar u$ along $\bar X_i$. Writing $\bar u$ as $\bar u^j \bar X_j$, the covariant derivative can be computed as

$$\bar\nabla_{\bar X_i}\bar u = \left(\bar X_i \bar u^j + \bar u^k \bar\omega_{ik}^j\right)\bar X_j,$$

where $\bar\omega_{ik}^j$ are the *connection coefficients*. These are defined by the equations

$$\bar\omega_{ik}^m \bar X_m = \bar\nabla_{\bar X_i}\bar X_k.$$

In order to actually compute the connection coefficients, we use the fact that they are related to the Christoffel symbols $\bar\Gamma_{k\ell}^i$ of $\bar g$ with respect to $(\partial_t, \partial_1, \partial_2)$ via

$$\bar\omega_{ik}^j = \left(\bar a_i^\ell \partial_\ell \bar a_k^m + \bar a_i^\ell \bar a_k^n \bar\Gamma_{\ell n}^m\right)\bar a_j^h \bar g_{mh}.$$

Finally, the Christoffel symbols are defined as

$$\bar\Gamma_{k\ell}^i = \frac{1}{2}\bar g^{im}\left(\bar g_{mk,\ell} + \bar g_{m\ell,k} - \bar g_{k\ell,m}\right),$$

where $\{\bar g^{ik}\}$ denote the coefficients of the inverse metric $\bar g^{-1}$, and $\{\bar g_{ik,\ell}\}$ denotes the partial derivatives of the coefficients of the metric $\bar g$ with respect to $(\partial_t, \partial_1, \partial_2)$.

Applying these definitions to our special situation, we arrive at the following explicit formula for the optimality conditions in local coordinates.

Theorem 15 (optimality conditions in local coordinates). *In the chart (V, v), the unique minimizer*

$$\bar u = (0\partial_t, u) = \left(0, u^i X_i\right)$$

of the energy functional $\mathcal{E}$ defined in (9) solves the equation

$$\left(A^j + B_m^j u^m + C_k^{mj}\partial_m u^k + D^{\ell m}\partial_{\ell m} u^j\right) X_j = 0 \quad \text{in } V \times (0, T), \tag{10}$$

$$\left(\frac{1}{\alpha}\partial_t u^j + u^k \bar\omega_{0k}^j\right) X_j = 0 \quad \text{on } V \times \{0, T\}, \tag{11}$$

where

$$\begin{aligned}
A^j &= \partial_t I \partial_i I g^{ik} b_k^j,\\
B_m^j &= \partial_\ell I a_m^\ell \partial_i I g^{ik} b_k^j + \beta\delta_m^j - \gamma \textstyle\sum_i (\bar a_i^\ell \partial_\ell \bar\omega_{im}^j + \bar\omega_{im}^n \bar\omega_{in}^j) - \frac{\gamma}{\alpha}\bar\omega_{0m}^j \bar\Gamma_{n0}^n,\\
C_k^{mj} &= -\gamma \textstyle\sum_i (\delta_k^j \bar a_i^\ell \partial_\ell \bar a_i^m + 2\bar\omega_{ik}^j \bar a_i^m) - \frac{\gamma}{\alpha}\delta_k^j \bar a_0^m \bar\Gamma_{n0}^n,\\
D^{\ell m} &= -\gamma \textstyle\sum_i \bar a_i^\ell \bar a_i^m,
\end{aligned}$$

and b_k^j is the X_j-component of ∂_k.

The connection coefficients $\bar\omega_{ik}^j$ of $\{X_i\}$ are given by

$$\bar\omega_{ik}^j = \begin{cases}
0 & \text{if } j = 0 \text{ and either } i = 0 \text{ or } k = 0,\\
\alpha a_i^\ell a_i^n \bar\Gamma_{\ell n}^0 & \text{if } j = 0 \text{ and } i, k \neq 0,\\
0 & \text{if } j \neq 0 \text{ and } i = k = 0,\\
\frac{1}{\alpha}\left(\partial_t a_k^m + a_k^n \bar\Gamma_{0n}^m\right) a_j^h g_{mh} & \text{if } i = 0 \text{ and } j, k \neq 0,\\
\frac{1}{\alpha} a_i^\ell a_j^h g_{mh} \bar\Gamma_{\ell 0}^m & \text{if } k = 0 \text{ and } j, i \neq 0,\\
\left(a_i^\ell \partial_\ell a_k^m + a_i^\ell a_k^n \Gamma_{\ell n}^m\right) a_j^h g_{mh} & \text{if } j, i, k \neq 0.
\end{cases} \tag{12}$$

Moreover, the Christoffel symbols $\bar{\Gamma}^j_{ik}$ *have the form*

$$\bar{\Gamma}^j_{ik} = \begin{cases} 0 & \text{if } j = 0 \text{ and either } i = 0 \text{ or } k = 0, \\ -\dfrac{\partial_t g_{ik}}{2\alpha^2} & \text{if } j = 0 \text{ and } i, k \neq 0, \\ 0 & \text{if } j \neq 0 \text{ and } i = k = 0, \\ g^{j\ell}\partial_t g_{k\ell} & \text{if } i = 0 \text{ and } j, k \neq 0, \\ g^{j\ell}\partial_t g_{i\ell} & \text{if } k = 0 \text{ and } j, i \neq 0, \\ 2g^{j\ell}\langle \partial_{ik} f, \partial_\ell f\rangle_{\mathbb{R}^3} & \text{if } j, i, k \neq 0. \end{cases} \tag{13}$$

The derivations of this equation and of the coordinate expression of the connection symbols and the Christoffel symbols are postponed to Appendix B.

6. Experimental results.

Numerical implementation. We illustrate the behavior of the proposed model in two experiments. In both, the moving surface $\mathcal{M}_t \subset \mathbb{R}^3$ is parametrized globally by a function $f : [0,T] \times M \to \mathbb{R}^3$, where $M \subset \mathbb{R}^2$ is a rectangular domain. Note that, due to the chosen images, we can always set $\beta = 0$ and still have well-posedness; cf. Remark 11. We also conducted experiments with a positive value of β, which yielded faster convergence of the numerical method. The main difference from the results with $\beta = 0$ was shortened flow fields.

We solve the optimality conditions from Theorem 15 with finite differences on a $k \times m \times n$ grid approximation of $\bar{M}$. Derivatives in all three directions are approximated by central differences, and the resulting sparse linear system is solved with the standard MATLAB implementation of the generalized minimal residual (GMRES) method.

Remark 16. In the two examples that we present in the following, the moving surface is of a relatively simple nature: It is quite smooth, and, moreover, there will (up to corners and periodic boundary conditions) exist a global chart for the surface at each time step. This will allow us to compare the resulting optical flow field obtained with our method to the optical flow field that is computed using the standard metric on the plane. We emphasize, however, that this naive approach of computing the optical flow field on the plane depends highly on the chosen charts, i.e., parametrizations, and a different atlas will yield completely different results. This observation shows the necessity of our intrinsic definition of the optical flow problem. Furthermore, the possibility of computing everything in the plane will fail in the case of more complicated surfaces that do not admit global charts or have large or rapidly varying curvature.

Experiment I, synthetic data. The first image sequence we apply our model to consists of 20 frames of the well-known Hamburg Taxi sequence,[1] scaled to the unit interval. The sequence has a resolution of 255×190, which leads to a total number of $9.7 \cdot 10^5$ grid points.

The surface we consider is a ring torus whose major circle turns into an ellipse while its

[1] The movie can be downloaded from http://i21www.ira.uka.de/image_sequences/.

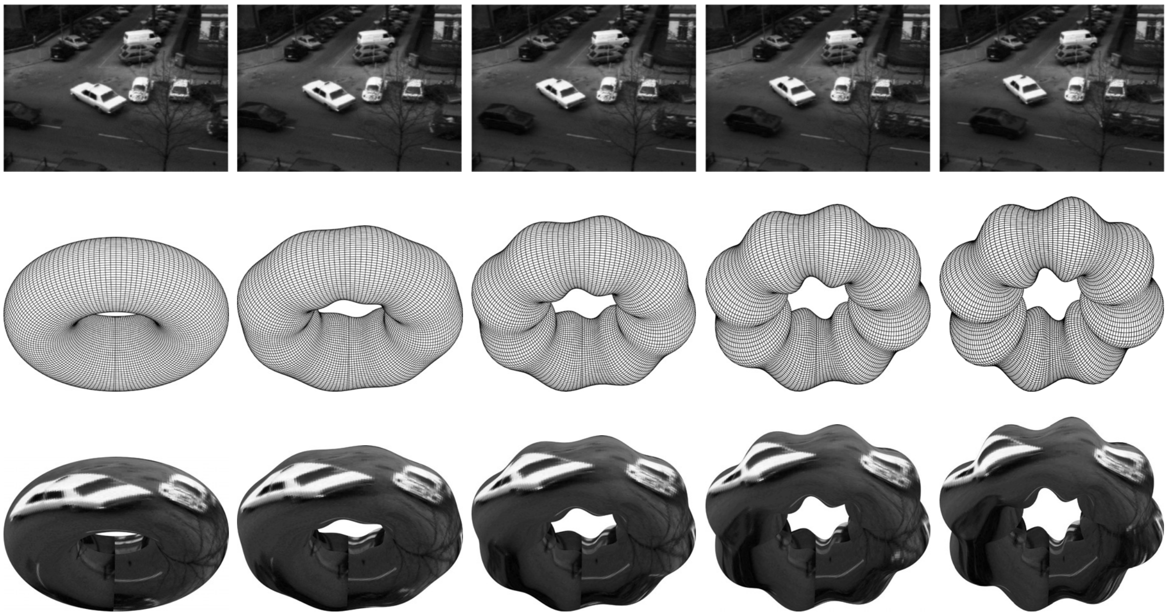

Figure 1. *The data considered in Experiment* I *at frames* 1, 6, 11, 16*, and* 20*. Top row: The pulled back image sequence I. Middle row: The moving surface $\mathcal{M}_t$. Bottom row: The image sequence $\mathcal{I}$.*

tube of uniform thickness grows ripples over time. The corresponding embedding reads as

$$f(t,x_1,x_2) = \begin{pmatrix} (R + \frac{t}{T} + r(t,x_1)\cos x_2)\cos x_1 \\ (R + r(t,x_1)\cos x_2)\sin x_1 \\ r(t,x_1)\sin x_2 \end{pmatrix},$$

where $R = 2$, $r(t,x_1) = 1 + \frac{t}{5T}\sin 8x_1$, and $(x_1,x_2) \in [0,2\pi)^2$. In Figure 1 we show I, $\mathcal{M}_t$, and $\mathcal{I}$.

In Figures 2 and 3 results for the parameter choice $\alpha = \gamma = 1$ and $\beta = 0$ are depicted. The finite difference step size h was set to 1 for all three directions. The GMRES algorithm was terminated after a maximum of 2000 iterations with a restart every 30 iterations. This led to a relative residual of $5.1 \cdot 10^{-3}$. In Figure 3 we use the color coding from [6] to visualize the optical flow. This is done by applying it to the pulled back vector field first, and drawing the resulting color image onto $\mathcal{M}_t$ via f afterwards.

Finally, we illustrate how the moving surface influences the optical flow vector field. To that end we repeat Experiment I on the flat torus with all parameters unchanged. That is, we compute the optical flow from the Hamburg Taxi sequence according to the model of Weickert and Schnörr [42] but with periodic boundary conditions. Note that this is indeed the more realistic model in this case, since this synthetic video sequence is not influenced by the actual structure of the manifold, but is only projected on the moving manifold. The reason that we still decided to conduct this experiment was to demonstrate the existence of a significant difference in these two approaches. However, the advantage of our method is present only if the image sequence is indeed influenced by the structure and movement of the manifold.

In Figure 4 we juxtapose the resulting vector field with the optical flow computed on the

Figure 2. *The optical flow vector field resulting from Experiment* I *at frames* 1, 10, *and* 19. *Top row: The pullback of the vector field. Bottom row: The vector field on the moving surface. The vectors have been scaled for better visibility.*

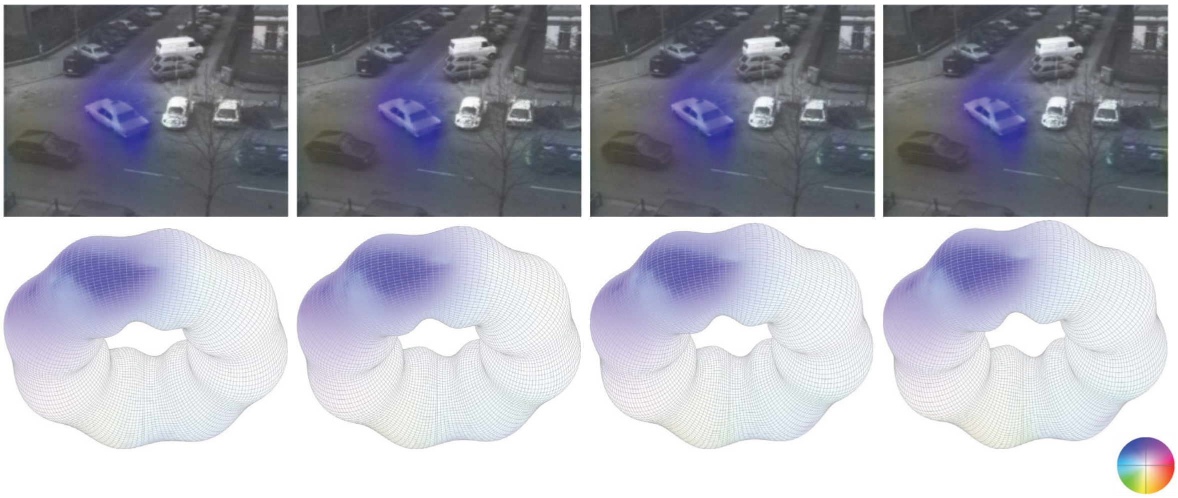

Figure 3. *The color-coded optical flow vector field resulting from Experiment* I *at frames* 10, 11, 12, *and* 13. *First row: Image sequence with pulled back vector field superimposed. Second row: Vector field on the moving surface. The color wheel is shown at the bottom right.*

deforming torus. A common measure for comparing two optical flow vector fields u and v is the angular error

$$\arccos \frac{\langle (1,u),(1,v)\rangle_{\mathbb{R}^3}}{|(1,u)|_{\mathbb{R}^3}|(1,v)|_{\mathbb{R}^3}}.$$

See [6], for example. The main purpose of adding the additional component 1 to both vectors

Figure 4. *Comparison of the optical flow from Experiment* I *with the optical flow computed in the plane. First row: Optical flow from Figure* 3 *(first row), with removed image. Second row: Optical flow computed in the plane with periodic boundaries. The color wheel is shown at the bottom right.*

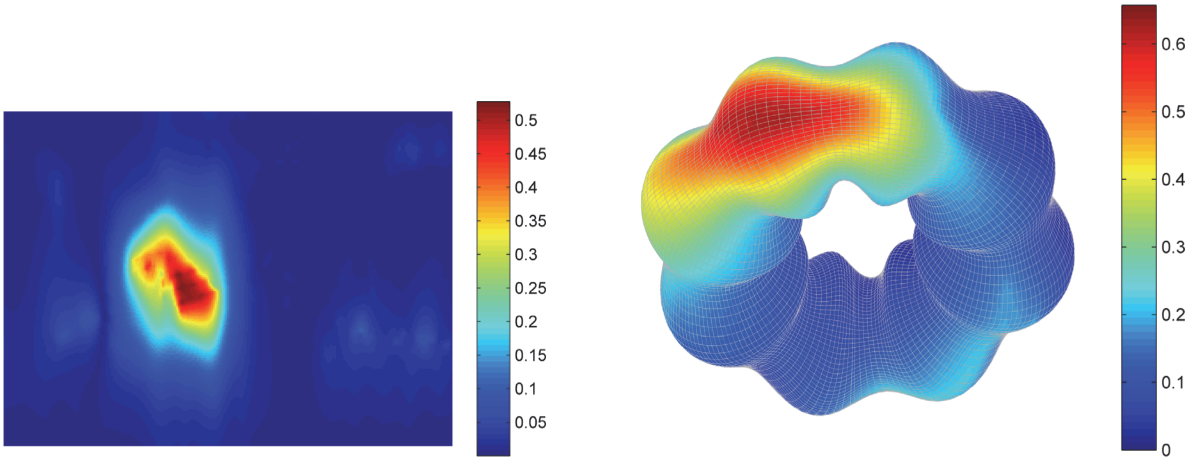

Figure 5. *Angular errors at frame* 19 *between the optical flow vector fields computed on the flat and the deforming torus, respectively. Left:* $\mathbb{R}^2$ *angular error between the pulled back vector fields. Right:* $\mathbb{R}^3$ *angular error between the pushed forward vector fields.*

is to avoid division by zero. Extending the above definition in a straightforward way to vector fields in $\mathbb{R}^3$, we show in Figure 5 the angular error between the optical flow computed on the deforming and the flat torus, respectively, both before and after pushforward to the deforming torus. Note that two unit vectors $u, v \in \mathbb{R}^2$ standing at an angle of $\pi/5$ would have an angular error of approximately 0.44. Another common measure is the so-called endpoint error $|u-v|_{\mathbb{R}^3}$, which also takes into account the lengths of the vectors. However, since vector lengths are typically affected by the choice of regularization parameters and finding comparable values for the flat and the deforming torus is not straightforward, we chose not to visualize the endpoint error.

Experiment II, *microscopy data.* Finally, we test our model on real-world data. The image sequence under consideration in this section depicts a living zebra fish embryo during the gastrula period and has been recorded with a confocal laser-scanning microscope. The only visible feature in this dataset are the embryo's endodermal cells, which, expressing a green fluorescent protein, proliferate on the surface of the embryo's yolk. Understanding and reconstructing cell motion during embryogenesis is a major topic in developmental biology, and optical flow is one way to automatically extract this information [2, 3, 25, 35]. See [23] for a detailed account of the embryonic development of a zebra fish, and [27] for more information on laser-scanning microscopy and fluorescent protein technology.

The considered data do not depict the whole embryo but only a cuboid section of approximately $540 \times 490 \times 340\,\mu m^3$. They have a spatial resolution of $512 \times 512 \times 40$ voxels, and the elapsed time between two consecutive frames is about four minutes. As in [25, 35] we avoid computational challenges by exploiting the fact that during gastrulation endodermal cells form a monolayer. This means that they can be regarded as sitting on a two-dimensional surface. Therefore, by fitting a surface through the cells' positions, we can reduce the spatial dimension of the data by one. We refer the reader to [25] for details on how this surface extraction was done.

In this particular experiment we apply our model to 21 frames of the resulting two-dimensional cell images with a resolution of 373×373 and again scaled to the unit interval. The extracted surface can be conveniently parametrized as the graph of a function $z(t, x_1, x_2)$ describing the height of the surface. That is, f takes the form

$$f(t, x_1, x_2) = (x_1, x_2, z(t, x_1, x_2)).$$

In Figure 6 we show I, $\mathcal{M}_t$, and $\mathcal{I}$. The regularization parameters were set to $\alpha = 10$, $\beta = 0$, $\gamma = 1$, and for the spatial boundaries we chose homogeneous Dirichlet boundary conditions. The GMRES solver converged faster this time and was terminated after the relative residual dropped below 10^{-3}. Results are shown in Figure 7. In Figure 8 we juxtapose the pulled back optical flow with the optical flow computed in the plane with the same parameters. Finally, we again compare the two vector fields by computing their angular error (see Figure 9). This time we do so after pushforward only, since for real-world data we are primarily interested in the vector field on the embedded surface.

7. Conclusion.

Choosing a suitable regularization term is a major challenge in the computation of the optical flow on a moving manifold when using variational methods. The main question is how to incorporate the structure of the manifold and its movement into the regularization. In this paper we have approached this problem from a purely differential geometric point of view. We have constructed a Riemannian metric on the time-dependent manifold in such a way that the paths of points on the manifold are geodesics with respect to this metric. We have then used a Horn–Schunck-type quadratic regularization term with additional time smoothing for the computation of the optical flow. The experiments performed within this setting indicate the viability of this approach and also show that using the manifold structure can have a significant influence on the computed optical flow field. Still, because of the usage of a quadratic regularization term that is not adapted to the image structure, the resulting

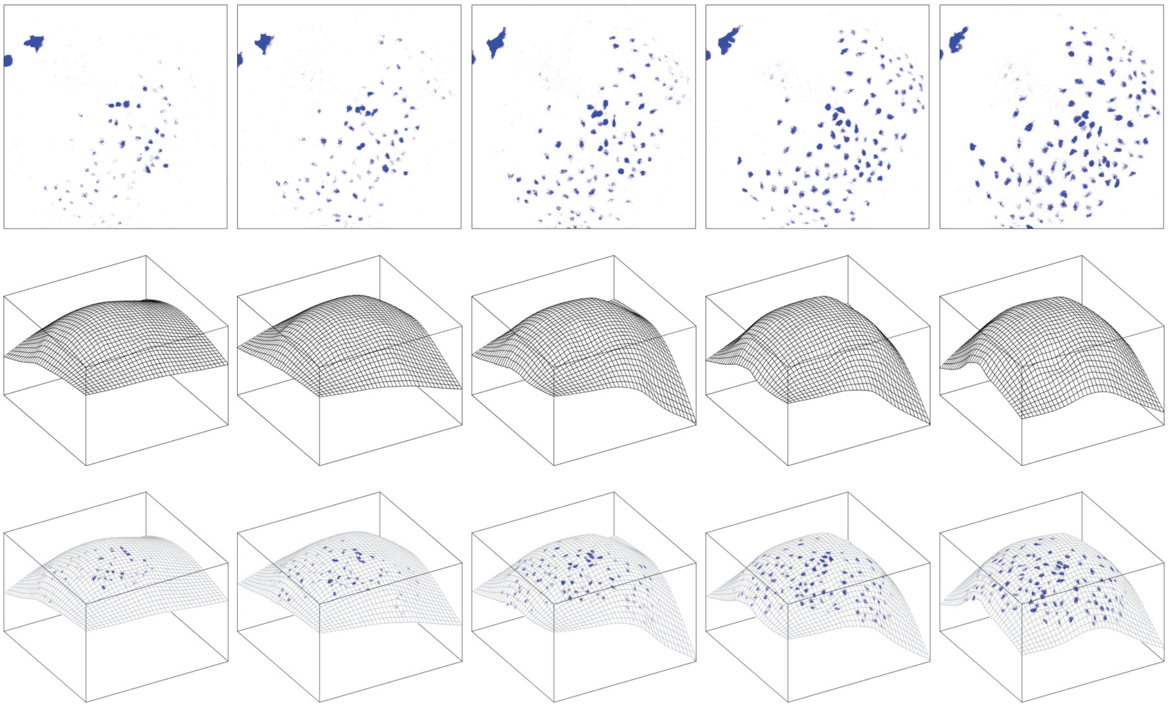

Figure 6. *The data considered in Experiment* II *at frames* 1, 6, 11, 16, *and* 20. *Top row: The pulled back image sequence I. Middle row: The moving surface $\mathcal{M}_t$. Bottom row: The image sequence $\mathcal{I}$.*

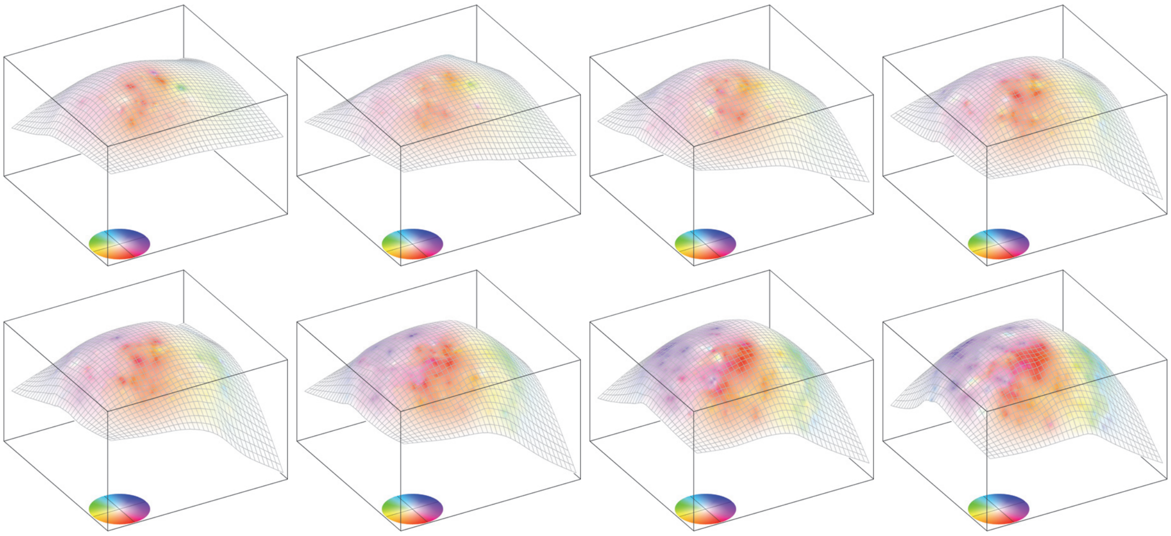

Figure 7. *The color-coded optical flow vector field resulting from Experiment* II *at frames* 1, 4, 7, 10, 11, 14, 17, and 20.

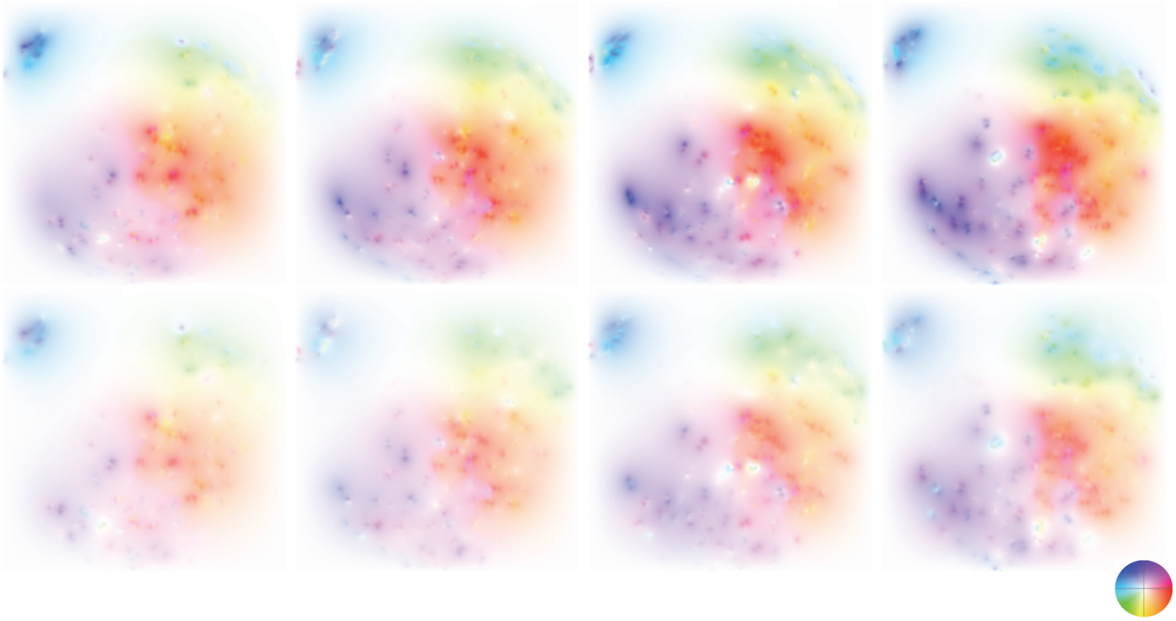

Figure 8. *Comparison of the optical flow from Experiment* II *(frames* 11 14, 17, 20*) with the optical flow computed in the plane. First row: Pullback of the optical flow from Figure* 7 *(second row). Second row: Optical flow computed in the plane. The color wheel is shown at the bottom right.*

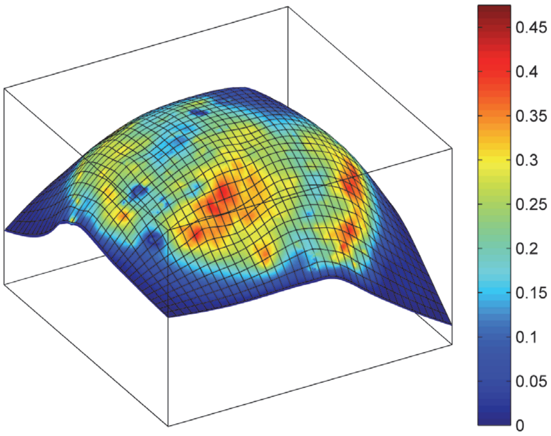

Figure 9. $\mathbb{R}^3$ *angular errors at frame* 20 *between the optical flow computed on the zebra fish surface and the optical flow computed in the plane, but pushed forward to the same surface.*

flow fields tend to be oversmoothed. The next step is therefore the extension to more complicated, anisotropic regularization terms as discussed in [41], which may be more accurate for certain applications of optical flow.

Appendix A. Finding the parametrization. So far we have assumed that we are given the moving surface with a fixed parametrization. In applications this parametrization might be unknown; i.e., one might observe only the shape of the surface but not its actual parametrization. Thus one will need to extract the parametrization from the observed data. In this appendix we will briefly sketch a possible approach to achieve this goal.

Remark 17. Note that the choice of parametrization will have a tremendous influence on the resulting optical flow field. In particular one can choose a parametrization such that the optical flow field is almost zero. To achieve this one can take any fixed parametrization $f(t,\cdot)$ and solve the optical flow problem for this parametrization using small regularization parameters. Then one can use the resulting optical flow field v to generate a path of diffeomorphisms $\varphi(t,\cdot) \in \mathrm{Diff}(M)$. Then the path $\tilde{f}(t,x) = f(t,\varphi(t,x))$ has the desired property.

In the following we will assume that the evolution of the image has no influence on the evolution of the surface—the influence of the surface evolution on the image evolution is taken into account by the nature of the regularization term. Furthermore, we assume that we are given only the shape of the surface at each time point t, but not the actual parametrization, i.e., that we are given a path in the space of unparametrized, embedded surfaces; see [28, 18] for a rigorous mathematical definition of this infinite-dimensional manifold.

At each time point t we can now choose any parametrization of the surfaces yielding a path of embeddings

$$f : [0,T] \times M \mapsto \mathbb{R}^3.$$

Thus we have reduced the problem to finding the path of reparametrizations that best corresponds to the observed shape evolution.

One way to tackle this problem is to define an energy functional on the space of embeddings that incorporates the available information on realistic shape evolutions. In order to be independent of the initial parametrization of the path of surfaces, we require that the energy functional be invariant under the action of the diffeomorphism group, i.e., $E(f(t,\varphi(x)) = E(f(t,x))$ for all $\varphi \in \mathrm{Diff}(M)$. In this case, the energy functional on the space of parametrized surfaces induces an energy functional on the shape space of unparametrized surfaces. Such a functional can be defined using a Riemannian metric, a Finsler-type metric, or some even more general Lagrangian; see, e.g., [9, 11, 37, 22, 33, 19, 8].

For the sake of simplicity, we will focus on the Riemannian case only, i.e.,

$$E(f) = \int_0^T G_f(f_t, f_t)dt,$$

where G is some reparametrization invariant metric on the manifold of all embeddings.

For historical reasons going back to Euler [15], these metrics are often represented via the corresponding inertia operator L:

$$G_f^L(f_t, f_t) := \int_M \langle f_t, L_f f_t \rangle \,\mathrm{vol}(g).$$

The simplest such metric is the reparametrization invariant L^2-metric—or H^0-metric. This metric is induced by the operator $L = \mathrm{Id}$:

$$G_f^0(f_t, f_t) = \int_M \langle f_t, f_t \rangle \,\mathrm{vol}(g).$$

In order to guarantee that the bilinear form G^L really induces a Riemannian metric, we require L to be an elliptic pseudodifferential operator that is symmetric and positive with respect to the L^2-metric. In addition, we assume that L is invariant under the action of the reparametrization group $\operatorname{Diff}(M)$. The invariance of L implies that the induced metric G^L is invariant under the action of $\operatorname{Diff}(M)$, as required. Using the operator L, one can include physical or biological model parameters in the definition of the metric.

Now we want to find the optimal reparametrization of the initial path f with respect to this energy functional. Therefore we have to solve the optimization problem

$$\psi(t,x) = \underset{\varphi \in C^\infty([0,T],\operatorname{Diff}(M))}{\operatorname{argmin}} E(f(t,\varphi(t,x))).$$

Further expanding the energy functional using the invariance of the Riemannian metric yields

$$\begin{aligned} E(f(t,\varphi(t,x))) = \int_0^T & G_f(f_t,f_t) + G_f(f_t, Tf(\varphi_t \circ \varphi^{-1})) \\ & + G_f(Tf(\varphi_t \circ \varphi^{-1}), Tf(\varphi_t \circ \varphi^{-1}))dt. \end{aligned}$$

Remark 18. As an example we want to consider this functional for the L^2-metric. Therefore we decompose f_t for each time point t into a part that is normal to the surface f and a part that is tangential:

$$f_t = Tf.f_t^\top + f_t^\perp.$$

Since these parts are orthogonal to each other—with respect to the L^2-metric—the energy functional reads as

$$E(f(t,\varphi(t,x))) = \int_M \langle f_t^\perp, f_t^\perp \rangle \operatorname{vol}(g) + \int_M g(f_t^\top + \varphi_t \circ \varphi^{-1}, f_t^\top + \varphi_t \circ \varphi^{-1}) \operatorname{vol}(g).$$

This functional is minimal for

$$\varphi_t \circ \varphi^{-1} = -f_t^\top.$$

This, however, corresponds to a reparametrization φ such that $\tilde{f} = f \circ \varphi$ consists only of a deformation in normal direction.

Remark 19. For a more general metric G^L this will no longer hold, since normal and tangential vector fields might not be orthogonal with respect to the G^L-metric. Instead one can show that for the optimal path $\tilde{f}$ we will have that $L\tilde{f}_t$ is normal; cf. [9].

Appendix B. Proof of Theorem 15. In the following we give a sketch of the derivation of the formulas in Theorem 15.

Lemma 20. *The Christoffel symbols of the metric $\bar{g}$ have the form given in* (13).

Proof. This is a straightforward computation using the definition of the Christoffel symbols as

$$\bar{\Gamma}^i_{kl} = \frac{1}{2}\bar{g}^{im}\left(\bar{g}_{mk,l} + \bar{g}_{ml,k} - \bar{g}_{kl,m}\right)$$

and the fact that the metric $\bar{g}$ and its inverse have the forms

$$\bar{g} = \begin{pmatrix} \alpha^2 & 0 & 0 \\ 0 & g_{11} & g_{12} \\ 0 & g_{12} & g_{22} \end{pmatrix} \qquad \text{and} \qquad \bar{g}^{-1} = \begin{pmatrix} \alpha^{-2} & 0 & 0 \\ 0 & g^{11} & g^{12} \\ 0 & g^{12} & g^{22} \end{pmatrix},$$

respectively, and

$$g_{ij} = \langle \partial_i f, \partial_j f\rangle_{\mathbb{R}^3}. \qquad \blacksquare$$

Lemma 21. *The symbols $\bar\omega^j_{ik}$ have the form given in* (12).

Proof. The connection coefficients $\bar\omega^j_{ik}$ are defined as

$$\bar\omega^j_{ik} = \left(\bar a^\ell_i \partial_\ell \bar a^m_k + \bar a^\ell_i \bar a^n_k \bar\Gamma^m_{\ell n}\right) \bar a^h_j \bar g_{mh}\,.$$

Moreover, the coordinates $\bar a^\ell_i$ have the form

$$\bar a^\ell_i = \begin{cases} \alpha^{-1} & \text{if } i=\ell=0,\\ 0 & \text{if } i=0 \text{ and } \ell\neq 0, \text{ or } i\neq 0 \text{ and } \ell = 0,\\ a^\ell_i & \text{if } i,\ell\neq 0.\end{cases}$$

Using these facts and the form of the Christoffel symbols derived in (13), the result follows from a straightforward calculation. $\blacksquare$

Lemma 22. *The L^2-gradient of the similarity term $\mathcal{S}$ in the energy functional $\mathcal{E}$ can be written for $\bar u = (0, u^j X_j)$ as*

$$\operatorname{grad}\mathcal{S}(\bar u) = 2\big(\partial_t I + \partial_\ell I a^\ell_m u^m\big)\partial_k I g^{ik} b^j_k X_j.$$

Proof. As shown in Theorem 13, the gradient of $\mathcal{S}$ has the form

$$\operatorname{grad}\mathcal{S}(\bar u) = 2\big(\partial_t I + g(\nabla^g I, u)\big)(0, \nabla^g I)\,.$$

Denote now by $\tilde u^j$ the coordinates of u with respect to ∂_j, that is, $u = \tilde u^j \partial_j$. Then

$$g(\nabla^g I, \bar u) = (D_x I)\bar u = (\partial_\ell I)\tilde u^\ell.$$

Moreover, we have

$$\tilde u^\ell = a^\ell_m u^m.$$

In addition, the coordinate expression of $\nabla^g I$ is $(\partial_k I) g^{ik}\partial_i$. Therefore we obtain

$$\operatorname{grad}\mathcal{S}(\bar u) = 2\big(\partial_t I + \partial_\ell I a^\ell_m u^m\big)\partial_k I g^{ik}\partial_i.$$

Since $\partial_i = b^j_k X_j$, we obtain the claimed representation. $\blacksquare$

Lemma 23. *In the local coordinate frame $\bar X_0 = \frac{1}{\alpha}\partial_t, \bar X_1, \bar X_2$, the Bochner Laplacian on the Riemannian manifold $(\bar M, \bar g)$ of a vector field $\bar u$ satisfying Neumann boundary conditions*

$$\bar\nabla_\nu \bar u\big|_{\partial\bar M} = \bar\nabla_{\partial_t}\bar u(\cdot, x)\big|_0^T = 0$$

is given by

$$\Delta^B \bar u = \bar\nabla^*\bar\nabla \bar u = -\sum_{i=0}^{2}\bar\nabla^2_{\bar X_i,\bar X_i}\bar u - \frac{\operatorname{Tr}(g^{-1}\partial_t g)}{2\alpha}\bar\nabla_{\bar X_0}\bar u\,. \tag{14}$$

Proof. To calculate the expression of the Laplacian, we have to compute the formula for the L^2-adjoint of the covariant derivative. Taking two vector fields $\bar{u}$, $\bar{v}$, we have

$$\int_0^T \int_M \bar{g}(\Delta^B u, \bar{v}) \operatorname{vol}(g)\, dt = \int_0^T \int_M \bar{g}_1^1(\bar{\nabla}\bar{u}, \bar{\nabla}\bar{v}) \operatorname{vol}(g)\, dt$$
$$= \sum_{i=0}^{2} \int_0^T \int_M \bar{g}(\bar{\nabla}_{\bar{X}_i}\bar{u}, \bar{\nabla}_{\bar{X}_i}\bar{v}) \operatorname{vol}(g)\, dt.$$

Using $(\bar{\nabla}_{\bar{X}_i}\bar{g}) = 0$, we obtain the following expression for the first summand ($i = 0$):

$$\frac{1}{\alpha^2}\int_0^T \int_M \bar{g}(\bar{\nabla}_{\partial_t}\bar{u}, \bar{\nabla}_{\partial_t}\bar{v}) \operatorname{vol}(g)\, dt$$
$$= \frac{1}{\alpha^2}\int_0^T \int_M \partial_t \left(\bar{g}(\bar{\nabla}_{\partial_t}\bar{u}, \bar{v})\right) \operatorname{vol}(g)\, dt - \int_0^T \int_M \frac{1}{\alpha^2}\bar{g}\left(\bar{\nabla}_{\partial_t}\left(\bar{\nabla}_{\partial_t}\bar{u}\right), \bar{v}\right) \operatorname{vol}(g)\, dt$$
$$= \frac{1}{\alpha^2}\int_0^T \partial_t \left(\int_M \bar{g}(\bar{\nabla}_{\partial_t}\bar{u}, \bar{v}) \operatorname{vol}(g)\right) dt - \int_0^T \int_M \frac{1}{\alpha^2}\bar{g}(\bar{\nabla}_{\partial_t}\bar{u}, \bar{v})\partial_t \operatorname{vol}(g)\, dt$$
$$- \int_0^T \int_M \frac{1}{\alpha^2}\bar{g}(\bar{\nabla}^2_{\partial_t,\partial_t}\bar{u}, \bar{v}) \operatorname{vol}(g)\, dt\,.$$

Using the variational formula [10, sect. 4.6]

$$\partial_t \operatorname{vol}(g) = \operatorname{Tr}(g^{-1}\partial_t g) \operatorname{vol}(g)$$

yields

$$\frac{1}{\alpha^2}\int_0^T \int_M \bar{g}(\bar{\nabla}_{\partial_t}\bar{u}, \bar{\nabla}_{\partial_t}\bar{v}) \operatorname{vol}(g)\, dt$$
$$= \frac{1}{\alpha^2}\left(\int_M \bar{g}(\bar{\nabla}_{\partial_t}\bar{u}, \bar{v}) \operatorname{vol}(g)\right)\Big|_0^T$$
$$- \int_0^T \int_M \frac{1}{\alpha^2}\bar{g}(\bar{\nabla}_{\partial_t}\bar{u}, \bar{v}) \operatorname{Tr}(g^{-1}\partial_t g) \operatorname{vol}(g)\, dt$$
$$- \int_0^T \int_M \frac{1}{\alpha^2}\bar{g}(\bar{\nabla}^2_{\partial_t,\partial_t}\bar{u}, \bar{v}) \operatorname{vol}(g)\, dt\,.$$

Note that for Neumann boundary conditions the first term in the above expression vanishes.

Since M has no boundary, the other summands in the formula for Δ^B are similar but simpler:

$$\sum_{i=1}^{2}\int_0^T \int_M \bar{g}(\bar{\nabla}_{\bar{X}_i}\bar{u}, \bar{\nabla}_{\bar{X}_i}\bar{v}) \operatorname{vol}(g)\, dt$$
$$= \sum_{i=1}^{2}\int_0^T \int_M 0 - \bar{g}(\bar{\nabla}_{X_i}\left(\bar{\nabla}_{X_i}\bar{u}\right), \bar{v}) \operatorname{vol}(g)\, dt$$
$$= -\sum_{i=1}^{2}\int_0^T \int_M \bar{g}(\bar{\nabla}^2_{X_i,X_i}\bar{u}, \bar{v}) \operatorname{vol}(g)\, dt\,.$$

Combining these equations, we obtain the desired formula for Δ^B. ■

Proof of Theorem 15. We have already derived the representation of the Christoffel symbols and connection coefficients.

Next, we will derive an explicit representation of the Bochner Laplacian in coordinates. To that end, we treat the two terms in (14) separately. For the second term we note that

$$\bar{\nabla}_{\bar{X}_0}\bar{u} = (\bar{a}_0^m \partial_m u^j + u^m \bar{\omega}_{0m}^j)X_j,$$

and thus we obtain

$$-\frac{1}{2\alpha}\operatorname{Tr}(g^{-1}\partial_t g)\bar{\nabla}_{\bar{X}_0}\bar{u} = -\frac{1}{2\alpha}g^{ik}\partial_t g_{ik} u^m \bar{\omega}_{0m}^j.$$

Moreover, we obtain from (13) that

$$g^{ik}\partial_t g_{ik} = \bar{\Gamma}_{n0}^n.$$

Hence, the first term becomes

$$\frac{\operatorname{Tr}(g^{-1}\partial_t g)}{2\alpha}\bar{\nabla}_{\bar{X}_0}\bar{u} = -\frac{1}{2\alpha}\bar{\Gamma}_{n0}^n(\bar{a}_0^m\partial_m u^j + u^m\bar{\omega}_{0m}^j)X_j. \tag{15}$$

For the second term, we compute

$$\begin{aligned}
\bar{\nabla}_{\bar{X}_i}\left(\bar{\nabla}_{\bar{X}_i}\bar{u}\right) &= \bar{\nabla}_{\bar{X}_i}\left(\bar{a}_i^\ell \partial_\ell u^j + u^k \bar{\omega}_{ik}^j\right)\bar{X}_j \\
&= \Big(\bar{a}_i^m \partial_m\left(\bar{a}_i^\ell\partial_\ell u^j + u^k\bar{\omega}_{ik}^j\right) + \left(\bar{a}_i^\ell \partial_\ell u^m + u^k \bar{\omega}_{ik}^m\right)\bar{\omega}_{im}^j\Big)X_j \\
&= \Big((\bar{a}_i^m\partial_m\bar{\omega}_{ik}^j + \bar{\omega}_{ik}^m\bar{\omega}_{im}^j)u^k + \bar{a}_i^m\partial_m\bar{a}_i^\ell\partial_\ell u^j \\
&\qquad + 2\bar{a}_i^m\bar{\omega}_{ik}^j\partial_m u^k + \bar{a}_i^m\bar{a}_i^\ell\partial_{\ell m}u^j\Big)X_j.
\end{aligned} \tag{16}$$

Combining Lemma 22, equations (15) and (16), and the fact that the gradient of $\|\bar{u}\|_{0,\bar{g}}^2$ is simply $2u^j X_j$, we arrive, after dividing everything by 2, at (10); equations (11) are simply the Neumann boundary conditions in coordinate form. ■

Acknowledgments. We thank Pia Aanstad from the University of Innsbruck for kindly providing the microscopy data used in this paper. We are grateful for helpful comments from and discussions with Lukas Lang, Peter Michor, and Otmar Scherzer. Finally, we want to thank the referees for their careful proofreading and their valuable comments.

REFERENCES

[1] R. Acar and C. R. Vogel, *Analysis of bounded variation penalty methods for ill-posed problems*, Inverse Problems, 10 (1994), pp. 1217–1229.

[2] 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. Methods, 11 (2014), pp. 951-958.

[3] F. Amat, E. W. Myers, and P. J. Keller, *Fast and robust optical flow for time-lapse microscopy using super-voxels*, Bioinformatics, 29 (2013), pp. 373–380.

[4] G. Aubert, R. Deriche, and P. Kornprobst, *Computing optical flow via variational techniques*, SIAM J. Appl. Math., 60 (1999), pp. 156–182.

[5] G. Aubert and L. Vese, *A variational method in image recovery*, SIAM J. Numer. Anal., 34 (1997), pp. 1948–1979.

[6] 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. Vis., 92 (2011), pp. 1–31.

[7] M. Bauer, M. Bruveris, P. Harms, and P. W. Michor, *Geodesic distance for right invariant Sobolev metrics of fractional order on the diffeomorphism group*, Ann. Global Anal. Geom., 44 (2013), pp. 5–21.

[8] M. Bauer, M. Bruveris, and P. W. Michor, *Overview of the geometries of shape spaces and diffeomorphism groups*, J. Math. Imaging Vision, 50 (2014), pp. 60–97.

[9] M. Bauer, P. Harms, and P. W. Michor, *Sobolev metrics on shape space of surfaces*, J. Geom. Mech., 3 (2011), pp. 389–438.

[10] M. Bauer, P. Harms, and P. W. Michor, *Almost local metrics on shape space of hypersurfaces in n-space*, SIAM J. Imaging Sci., 5 (2012), pp. 244–310.

[11] M. Bauer, P. Harms, and P. W. Michor, *Sobolev metrics on shape space,* II: *Weighted Sobolev metrics and almost local metrics*, J. Geom. Mech., 4 (2012), pp. 365–383.

[12] A. Bruhn, J. Weickert, and Ch. Schnörr, *Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods*, Int. J. Comput. Vis., 61 (2005), pp. 211–231.

[13] I. Cohen, *Nonlinear variational method for optical flow computation*, in Proceedings of the Eighth Scandinavian Conference on Image Analysis, Vol. 1, Norway, 1993, pp. 523–530.

[14] J. Eichhorn, *Global Analysis on Open Manifolds*, Nova Science Publishers, New York, 2007.

[15] L. P. Euler, *Du mouvement de rotation des corps solides autour d'un axe variable*, Mémoires de l'Académie des sciences de Berlin, 14 (1765), pp. 154–193.

[16] D. Fleet and Y. Weiss, *Optical flow estimation*, in Handbook of Mathematical Models in Computer Vision, Springer, New York, 2006, pp. 239–257.

[17] M. Grasmair, *Linear convergence rates for Tikhonov regularization with positively homogeneous functionals*, Inverse Problems, 27 (2011), 075014.

[18] R. S. Hamilton, *The inverse function theorem of Nash and Moser*, Bull. Amer. Math. Soc. (N.S.), 7 (1982), pp. 65–222.

[19] B. Heeren, M. Rumpf, M. Wardetzky, and B. Wirth, *Time-discrete geodesics in the space of shells*, Comput. Graph. Forum, 31 (2012), pp. 1755–1764.

[20] B. K. P. Horn and B. G. Schunck, *Determining optical flow*, Artificial Intelligence, 17 (1981), pp. 185–203.

[21] A. Imiya, H. Sugaya, A. Torii, and Y. Mochizuki, *Variational analysis of spherical images*, in Computer Analysis of Images and Patterns, Lecture Notes in Comput. Sci. 3691, A. Gagalowicz and W. Philips, eds., Springer, Berlin, Heidelberg, 2005, pp. 104–111.

[22] I. H. Jermyn, S. Kurtek, E. Klassen, and A. Srivastava, *Elastic shape matching of parameterized surfaces using square root normal fields*, in Proceedings of the 12th European Conference on Computer Vision (ECCV 2012), Part V, Springer-Verlag, Berlin, Heidelberg, 2012, pp. 804–817.

[23] C. B. Kimmel, W. W. Ballard, S. R. Kimmel, B. Ullmann, and T. F. Schilling, *Stages of embryonic development of the zebrafish*, Dev. Dyn., 203 (1995), pp. 253–310.

[24] C. Kirisits, L. F. Lang, and O. Scherzer, *Optical flow on evolving surfaces with an application to the analysis of 4D microscopy data*, in Proceedings of the Fourth International Conference on Scale Space and Variational Methods in Computer Vision (SSVM'13), Lecture Notes in Comput. Sci. 7893, A. Kuijper, K. Bredies, T. Pock, and H. Bischof, eds., Springer-Verlag, Berlin, Heidelberg, 2013, pp. 246–257.

[25] C. Kirisits, L. F. Lang, and O. Scherzer, *Optical flow on evolving surfaces with space and time regularisation*, J. Math. Imaging Vision, to appear.

[26] J. Lefèvre and S. Baillet, *Optical flow and advection on 2-Riemannian manifolds: A common framework*, IEEE Trans. Pattern Anal. Mach. Intell., 30 (2008), pp. 1081–1092.

[27] 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 (2003), pp. 1407–1420.

[28] P. W. Michor, *Manifolds of Differentiable Mappings*, Shiva Publishing, Nantwich, UK, 1980.

[29] H.-H. Nagel, *On the estimation of optical flow: Relations between different approaches and some new results*, Artificial Intelligence, 33 (1987), pp. 299–324.

[30] H.-H. Nagel, *Image sequences—ten (octal) years—from phenomenology towards a theoretical foundation*, Int. J. Patt. Recogn. Artif. Intell., 02 (1988), pp. 459–483.

[31] H.-H. Nagel, *Extending the 'oriented smoothness constraint' into the temporal domain and the estimation of derivatives of optical flow*, in Proceedings of the First European Conference on Computer Vision (ECCV '90), Springer-Verlag, London, 1990, pp. 139–148.

[32] H.-H. Nagel and W. Enkelmann, *An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences*, IEEE Trans. Pattern Anal. Mach. Intell., 8 (1986), pp. 565–593.

[33] M. Rumpf and B. Wirth, *An elasticity-based covariance analysis of shapes*, Int. J. Comput. Vis., 92 (2011), pp. 281–295.

[34] O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen, *Variational Methods in Imaging*, Appl. Math. Sci. 167, Springer, New York, 2009.

[35] B. Schmid, G. Shah, N. Scherf, M. Weber, K. Thierbach, C. Campos Pérez, I. Roeder, P. Aanstad, and J. Huisken, *High-speed panoramic light-sheet microscopy reveals global endodermal cell dynamics*, Nat. Commun., 4 (2013), 2207.

[36] Ch. Schnörr, *Determining optical flow for irregular domains by minimizing quadratic functionals of a certain class*, Int. J. Comput. Vis., 6 (1991), pp. 25–38.

[37] A. Srivastava, E. Klassen, S. H. Joshi, and I. H. Jermyn, *Shape analysis of elastic curves in Euclidean spaces*, IEEE Trans. Pattern Anal. Mach. Intell., 33 (2011), pp. 1415–1428.

[38] A. Torii, A. Imiya, H. Sugaya, and Y. Mochizuki, *Optical flow computation for compound eyes: Variational analysis of omni-directional views*, in Brain, Vision, and Artificial Intelligence, Lecture Notes in Comput. Sci. 3704, M. De Gregorio, V. Di Maio, M. Frucci, and C. Musio, eds., Springer, Berlin, Heidelberg, 2005, pp. 527–536.

[39] H. Triebel, *Theory of Function Spaces*. II, Monogr. Math. 84, Birkhäuser Verlag, Basel, 1992.

[40] J. Weickert, A. Bruhn, T. Brox, and N. Papenberg, *A survey on variational optic flow methods for small displacements*, in Mathematical Models for Registration and Applications to Medical Imaging, Math. Ind. 10, O. Scherzer, ed., Springer, Berlin, Heidelberg, 2006, pp. 103–136.

[41] J. Weickert and Ch. Schnörr, *A theoretical framework for convex regularizers in PDE-based computation of image motion*, Int. J. Comput. Vis., 45 (2001), pp. 245–264.

[42] J. Weickert and Ch. Schnörr, *Variational optic flow computation with a spatio-temporal smoothness constraint*, J. Math. Imaging Vision, 14 (2001), pp. 245–255.