Computer Vision, Graphics, and Pattern Recognition Group
Department of Mathematics and Computer Science
University of Mannheim
68131 Mannheim, Germany

Reihe Informatik
13/2000

A Theoretical Framework for Convex Regularizers in PDE-Based Computation of Image Motion

Joachim Weickert and Christoph Schnörr

Technical Report 13/2000
Computer Science Series
June 2000

A Theoretical Framework for Convex Regularizers in PDE-Based Computation of Image Motion

Joachim Weickert and Christoph Schnörr

Computer Vision, Graphics, and Pattern Recognition Group
Department of Mathematics and Computer Science
University of Mannheim, 68131 Mannheim, Germany

{Joachim.Weickert,Christoph.Schnoerr}@uni-mannheim.de
http://www.cvgpr.uni-mannheim.de

Abstract

Many differential methods for the recovery of the optic flow field from an image sequence can be expressed in terms of a variational problem where the optic flow minimizes some energy. Typically, these energy functionals consist of two terms: a data term, which requires e.g. that a brightness constancy assumption holds, and a regularizer that encourages global or piecewise smoothness of the flow field. In this paper we present a systematic classification of rotation invariant convex regularizers by exploring their connection to diffusion filters for multichannel images. This taxonomy provides a unifying framework for data-driven and flow-driven, isotropic and anisotropic, as well as spatial and spatio-temporal regularizers. While some of these techniques are classic methods from the literature, others are derived here for the first time. We prove that all these methods are well-posed: they posses a unique solution that depends in a continuous way on the initial data. An interesting structural relation between isotropic and anisotropic flow-driven regularizers is identified, and a design criterion is proposed for constructing anisotropic flow-driven regularizers in a simple and direct way from isotropic ones. Its use is illustrated by several examples.

Keywords: optic flow, differential methods, regularization, diffusion filtering, well-posedness

1998 ACM Computing Classification: I.4.8, I.4.3, I.4.4.

1 Introduction

Even after two decades of intensive research, robust motion estimation continues to be a key problem in computer vision. Motion is linked to the notion of optic flow, the displacement field of corresponding pixels in subsequent frames of an image sequence. Optic flow provides information that is important for many applications, ranging from the estimation of motion parameters for robot navigation to the design of second generation video coding algorithms. Surveys of the state-of-the-art in motion computation can be found in papers by Mitiche and Bouthemy [32], and Stiller and Konrad [50]. For a performance evaluation of some of the most popular algorithms we refer to Barron *et al.* [5] and Galvin *et al.* [20].

Bertero *et al.* [6] pointed out that, depending on its formulation, optic flow calculations may be ill-conditioned or even ill-posed. It is therefore common to use implicit or explicit smoothing steps in order to stabilize or regularize the process.

Implicit smoothing steps appear for instance in the robust calculation of image derivatives, where one usually applies some amount of spatial or temporal smoothing (averaging over several frames). It is not rare that these steps are only described as algorithmic details, but indeed they are often very crucial for the quality of the algorithm.

Thus, it would be consequent to make the role of smoothing more explicit by incorporating it already in a continuous problem formulation. This way has been pioneered by Horn and Schunck [25] and improved by Nagel [34] and many others. Approaches of this type calculate optic flow as the minimizer of an energy functional, which consists of a data term and a smoothness term. Formulations in terms of energy functionals allow a conceptually clear formalism without any hidden model assumptions, and several evaluations have shown that these methods perform well [5, 20].

The data term in the energy functional involves optic flow constraints such as the assumption that corresponding pixels in different frames should reveal the same grey value. The smoothness term usually requires that the optic flow field should vary smoothly in space [25]. Such a term may be modified in an *image-driven* way in order to suppress smoothing at or across image boundaries [1, 34]. As an alternative, *flow-driven* modifications have been proposed which reduce smoothing across flow discontinuities [8, 12, 14, 29, 40, 43, 54]. Most smoothness terms require only *spatial* smoothness. *Spatio-temporal* smoothness terms have been considered to a much smaller extent [7, 33, 36, 56]. Since smoothness terms fill in information from regions where reliable flow estimates exist to regions where no estimates are possible, they create dense flow fields. In many applications, this is a desirable quality which distinguishes regularization methods from other optic flow algorithms. The latter ones create non-dense flow fields, that have to be postprocessed by interpolation, if 100 % density is required.

Modeling the optic flow recovery problem in terms of *continuous* energy functionals offers the advantage of having a formulation that is as independent of the pixel grid as possible. A correct continuous model can be rotation invariant, and the use of well-established numerical methods shows how this rotation invariance can be approximated in a mathematically consistent way.

From both a theoretical and practical point of view, it can be attractive to use energy functionals that are *convex*. They have a unique minimum, and this global minimum can be found in a stable way by using standard techniques from convex optimization, for instance gradient descent methods. Having a unique minimum allows to use globally convergent algorithms, where every arbitrary flow initialization leads to the same solution: the global minimum of the functional. This property is an important quality of a robust algorithm. *Nonconvex* energy functionals, on the other hand, may have many local minima, and it is difficult to find algorithms that are both efficient and converge to a global minimum. Typical algorithms which converge to a global minimum (such as simulated annealing [30]) are computationally very expensive, while methods which are more efficient (such as graduated non-convexity algorithms [9]) may get trapped in a local minimum.

Minimizing continuous energy functionals leads in a natural way to partial differential equations (PDEs): applying gradient descent, for instance, yields a system of coupled diffusion–reactions equations for the two flow components. The fastly emerging use of PDE-based image restoration methods [22, 39], such as nonlinear diffusion filtering and total variation denoising, has motivated many researchers to apply similar ideas to estimate optic flow [1, 4, 12, 14, 24, 29, 38, 40, 43, 54]. A systematic framework that links the diffusion and optic flow paradigms, however, has not been studied so far. Furthermore, from the framework of diffusion filtering it is also well-known that anisotropic filters with a diffusion tensor have more degrees of freedom than isotropic ones with scalar-valued diffusivities. These additional degrees of freedom can be used to obtain better results in specific situations [53]. However, similar nonlinear anisotropic regularizers have not been considered in the optic flow literature so far.

The goal of the present paper is to address these issues. We present a theoretical framework for a broad class of regularization methods for optic flow estimation. For the reasons explained above, we focus on models that allow a formulation in terms of convex and rotation invariant continuous energy functionals. We consider image-driven and flow-driven models, isotropic and anisotropic ones, as well as models with spatial and spatio-temporal smoothing terms. We prove that all these approaches are well-posed in the sense of Hadamard: they have a unique solution that depends in a continuous (and therefore predictable) way on the input data.

We shall see that our taxonomy includes not only many existing models, but also interesting novel ones. In particular, we will derive novel regularization functionals for optic flow estimation that are flow-driven and anisotropic. They are the optic flow analogues of anisotropic diffusion filters with a diffusion tensor. Many of the spatio-temporal methods have not been proposed before as well. With the increased computational possibilities of modern computers it is likely that they will become more important in the future. In the present paper we also focus on interesting relations between isotropic and anisotropic flow-driven methods. They allow us to formulate a general design principle which explains how one can create anisotropic optic flow regularizers from isotropic ones.

Our paper is organized as follows. In Section 2 we first review and classify existing image-driven and isotropic flow-driven models, before we derive a novel energy functional leading to anisotropic flow driven models. Then we show how one has to modify all models with a spatial smoothness term in order to obtain methods with spatio-temporal regularization. A unifying energy functional is derived that incorporates the previous models as well as novel ones. Its well-posedness is established in Section 3. In Section 4 we take advantage of structural similarities between isotropic and anisotropic approaches in order to formulate a design principle for anisotropic optic flow regularizers. The paper is concluded with a summary in Section 5.

2 A Framework for Convex Regularizers

2.1 Spatial regularizers

2.1.1 Basic structure

In order to formalize the optic flow estimation problem, let us consider a real-valued image sequence $f(x, y, \theta)$, where (x, y) denotes the location within the image domain $\Omega \in \mathbb{R}^2$, and the time parameter $\theta \in [0, T]$ specifies the frame. The optic flow field $(u_1(x, y, \theta), u_2(x, y, \theta))$ describes the displacement between two subsequent frames θ and $\theta + 1$, i.e. $f(x, y, \theta)$ and $f(x + u_1(x, y, \theta), y + u_2(x, y, \theta), \theta + 1)$ should depict the same image detail. Frequently it is assumed that image objects keep their grey value over time:

$$f(x, y, \theta) - f(x + u_1(x, y, \theta), y + u_2(x, y, \theta), \theta + 1) = 0. \tag{1}$$

Such a model assumes that illumination changes do not appear, and that occlusions or disocclusions do not happen. Numerous generalizations to multiple constraint equations and/or different "conserved quantities" (replacing intensity) exist; see e.g. [18, 51]. However, since the goal of the present paper is to study different regularizers, we restrict ourselves to (1). If the spatial and temporal sampling is sufficiently fine, we may replace (1) by its first order Taylor approximation

$$f_x u_1 + f_y u_2 + f_\theta = 0, \tag{2}$$

where the subscripts x, y and θ denote partial derivatives. This so-called *optic flow constraint (OFC)* forms the basis of many differential methods for estimating the optic flow. Evidently such a single equation is not sufficient to determine the two unknown functions u_1 and u_2 uniquely. In order to recover a unique flow field, we need an additional assumption. Regularization-based optic flow methods use as additional assumption the requirement that the optic flow field should be smooth (or at least piecewise smooth). The basic idea is to recover the optic flow as a minimizer of some energy functional of type

$$E(u_1, u_2) := \int_\Omega \Big((f_x u_1 + f_y u_2 + f_\theta)^2 \; + \; \alpha\, V(\nabla f, \nabla u_1, \nabla u_2) \Big)\, dx\, dy \tag{3}$$

where $\nabla := (\partial_x, \partial_y)^T$ denotes the spatial nabla operator, and $u := (u_1, u_2)^T$. The first term in the energy functional is a data term requiring that the OFC be fulfilled, while the second term penalizes deviations from (piecewise) smoothness. The smoothness term $V(\nabla f, \nabla u_1, \nabla u_2)$ is called *regularizer*, and the positive smoothness weight α is the *regularization parameter*. One would expect that the specific choice of the regularizer has a strong influence on the result. Therefore, let us discuss different classes of convex regularizers next.

2.1.2 Homogeneous regularization

In 1981 Horn and Schunck [25] pioneered the field of regularization methods for optic flow computations. They used the regularizer

$$V_H(\nabla f, \nabla u_1, \nabla u_2) \;:=\; |\nabla u_1|^2 + |\nabla u_2|^2. \tag{4}$$

It is a classic result from the calculus of variations [13, 16] that – under mild regularity conditions – a minimizer (u_1, u_2) of some energy functional

$$E(u_1, u_2) := \int_\Omega G(x, y, u_1, u_2, \nabla u_1, \nabla u_2)\, dx\, dy \tag{5}$$

satisfies necessarily the so-called *Euler–Lagrange equations*

$$\partial_x G_{u_{1x}} + \partial_y G_{u_{1y}} - G_{u_1} \;=\; 0, \tag{6}$$
$$\partial_x G_{u_{2x}} + \partial_y G_{u_{2y}} - G_{u_2} \;=\; 0 \tag{7}$$

with homogeneous Neumann boundary conditions:

$$\partial_n u_1 = 0 \qquad \text{on } \partial\Omega, \tag{8}$$
$$\partial_n u_2 = 0 \qquad \text{on } \partial\Omega. \tag{9}$$

Hereby, n is a vector normal to the image boundary $\partial\Omega$.

Applying this framework to the minimization of the Horn and Schunck functional leads to the PDEs

$$\Delta u_1 - \tfrac{1}{\alpha} f_x (f_x u_1 + f_y u_2 + f_\theta) \;=\; 0, \tag{10}$$
$$\Delta u_2 - \tfrac{1}{\alpha} f_y (f_x u_1 + f_y u_2 + f_\theta) \;=\; 0, \tag{11}$$

where $\Delta := \partial_{xx} + \partial_{yy}$ denotes the Laplace operator. These equations can be regarded as the steady state ($t \to \infty$) of the diffusion–reaction system

$$\partial_t u_1 \;=\; \Delta u_1 - \tfrac{1}{\alpha} f_x (f_x u_1 + f_y u_2 + f_\theta), \tag{12}$$
$$\partial_t u_2 \;=\; \Delta u_2 - \tfrac{1}{\alpha} f_y (f_x u_1 + f_y u_2 + f_\theta), \tag{13}$$

where t denotes an artificial evolution parameter that should not be mixed up with the time θ of the image sequence. These equations also arise when minimizing the Horn and Schunck functional using steepest descent. Schnörr [41] has established well-posedness by showing that this functional has a unique minimizer that depends continuously on the input data f. Recently, Hinterberger [24] proved similar well-posedness results for a related model with a different data term.

We observe that the underlying diffusion process in the Horn and Schunck approach is the linear diffusion equation

$$\partial_t u_i = \Delta u_i = \operatorname{div}(g\, \nabla u_i) \tag{14}$$

with $g := 1$ and $i = 1, 2$. This equation is well-known for its regularizing properties and has been extensively used in the context of Gaussian scale-space; see [48] and the references therein. It smoothes, however, in a completely homogeneous way, since its diffusivity g equals 1 everywhere. As a consequence, it also blurs across semantically important flow discontinuities. This is the reason why the Horn and Schunck approach creates rather blurry optic flow fields. The regularizers described in the sequel are attempts to overcome this limitation.

2.1.3 Isotropic image-driven regularization

It seems plausible that motion boundaries are a subset of the image boundaries. Thus, a simple way to prevent smoothing at motion boundaries consists of introducing a weight function into the Horn and Schunck regularizers that becomes small at image edges. This modification yields the regularizer

$$V_{II}(\nabla f, \nabla u_1, \nabla u_2) \;:=\; g(|\nabla f|^2)\, \left(|\nabla u_1|^2 + |\nabla u_2|^2\right), \tag{15}$$

where g is a decreasing, strictly positive function. This regularizer has been proposed and theoretically analysed by Alvarez *et al.* [1]. The corresponding diffusion–reaction equations are given by

$$\partial_t u_1 \;=\; \operatorname{div}\, \left(g(|\nabla f|^2)\, \nabla u_1\right) - \tfrac{1}{\alpha} f_x (f_x u_1 + f_y u_2 + f_\theta), \tag{16}$$
$$\partial_t u_2 \;=\; \operatorname{div}\, \left(g(|\nabla f|^2)\, \nabla u_2\right) - \tfrac{1}{\alpha} f_y (f_x u_1 + f_y u_2 + f_\theta). \tag{17}$$

The underlying diffusion process is

$$\partial_t u_i = \operatorname{div}\, \left(g(|\nabla f|^2)\, \nabla u_i\right) \qquad (i = 1, 2). \tag{18}$$

It uses a scalar-valued diffusivity g that depends on the image gradient. Such a method can therefore be classified as inhomogeneous, isotropic and image-driven. Isotropic refers to the fact that a scalar-valued diffusivity guarantees a direction-independent smoothing behaviour, while inhomogeneous means that this behaviour may be space-dependent. Since the diffusivity does not depend on the flow itself, the diffusion process is linear. For more details on this terminology and diffusion filtering in image processing, we refer to [53]. Homogeneous regularization arises as a special case of (15) when $g(|\nabla f|^2) := 1$ is considered.

2.1.4 Anisotropic image-driven regularization

An early anisotropic modification of the Horn and Schunck functional is due to Nagel [34]; see also [2, 17, 35, 37, 41, 42, 47]. The basic idea is to reduce smoothing across image boundaries, while encouraging smoothing along image boundaries. This is achieved by considering the regularizer

$$V_{AI}(\nabla f, \nabla u_1, \nabla u_2) \;:=\; \nabla u_1^T D(\nabla f) \nabla u_1 + \nabla u_2^T D(\nabla f) \nabla u_2. \tag{19}$$

$D(\nabla f)$ is a regularized projection matrix perpendicular to ∇f:

$$D(\nabla f) := \frac{1}{|\nabla f|^2 + 2\lambda^2} \left(\nabla f^{\perp} \nabla f^{\perp T} + \lambda^2 I \right), \tag{20}$$

where I denotes the unit matrix. This methods leads to the diffusion–reaction equations

$$\partial_t u_1 \;=\; \text{div}\,\left(D(\nabla f)\,\nabla u_1\right) - \tfrac{1}{\alpha} f_x (f_x u_1 + f_y u_2 + f_\theta), \tag{21}$$

$$\partial_t u_2 \;=\; \text{div}\,\left(D(\nabla f)\,\nabla u_2\right) - \tfrac{1}{\alpha} f_y (f_x u_1 + f_y u_2 + f_\theta). \tag{22}$$

The usage of a diffusion tensor $D(\nabla f)$ instead of a scalar-valued diffusivity allows a direction-dependent smoothing behaviour. This method can therefore be classified as anisotropic. Since the diffusion tensor depends on the image f but not on the unknown flow, it is a purely image-driven process that is linear in its diffusion part. Well-posedness for this model has been established by Schnörr [41].

The eigenvectors of D are $v_1 := \nabla f$, $v_2 := \nabla f^{\perp}$, and the corresponding eigenvalues are given by

$$\lambda_1(|\nabla f|) = \frac{\lambda^2}{|\nabla f|^2 + 2\lambda^2}, \tag{23}$$

$$\lambda_2(|\nabla f|) = \frac{|\nabla f|^2 + \lambda^2}{|\nabla f|^2 + 2\lambda^2}. \tag{24}$$

In the interior of objects we have $|\nabla f| \to 0$, and therefore $\lambda_1 \to 1/2$ and $\lambda_2 \to 1/2$. At ideal edges where $|\nabla f| \to \infty$, we obtain $\lambda_1 \to 0$ and $\lambda_2 \to 1$. Thus, we have isotropic behaviour within regions, and at image boundaries the process smoothes anisotropically along the edge. This behaviour is very similar to edge-enhancing anisotropic diffusion filtering [53]. In contrast to edge-enhancing anisotropic diffusion, however, Nagel's optic flow technique is linear. It is interesting to note that only recently it has been pointed out that the Nagel method may be regarded as an early predecessor of anisotropic diffusion filtering [2].

Homogeneous and isotropic image-driven regularizers are special cases of (19), where $D(\nabla f) := I$ and $D(\nabla f) := g(|\nabla f|^2)I$ are chosen.

2.1.5 Isotropic flow-driven regularization

Image-driven regularization methods may create oversegmentations for strongly textured objects: in this case we have much more image boundaries than motion boundaries. In order to reduce smoothing only at motion boundaries, one may consider using a purely flow-driven regularizer. This, however, is at the expense of refraining from quadratic optimization problems. In earlier work [43, 54], the authors considered regularizers of type

$$V_{IF}(\nabla f, \nabla u_1, \nabla u_2) \;:=\; \Psi\left(|\nabla u_1|^2 + |\nabla u_2|^2\right), \tag{25}$$

where $\Psi(s^2)$ is a differentiable and increasing function that is convex in s, for instance

$$\Psi(s^2) \;:=\; \varepsilon s^2 + (1-\varepsilon)\lambda^2\sqrt{1 + s^2/\lambda^2} \qquad (0 < \varepsilon \ll 1,\ \lambda > 0). \tag{26}$$

Regularizer of type (25) lead to the diffusion–reaction system

$$\partial_t u_1 \;=\; \operatorname{div}\,\left(\Psi'(|\nabla u_1|^2+|\nabla u_2|^2)\,\nabla u_1\right)-\tfrac{1}{\alpha}f_x(f_x u_1+f_y u_2+f_\theta), \tag{27}$$
$$\partial_t u_2 \;=\; \operatorname{div}\,\left(\Psi'(|\nabla u_1|^2+|\nabla u_2|^2)\,\nabla u_2\right)-\tfrac{1}{\alpha}f_y(f_x u_1+f_y u_2+f_\theta), \tag{28}$$

where Ψ' denotes the derivative of Ψ with respect to its argument. The scalar-valued diffusivity $\Psi'(|\nabla u_1|^2+|\nabla u_2|^2)$ shows that this model is isotropic and flow-driven. In general, the diffusion process is nonlinear now. For the specific regularizer (26), for instance, the diffusivity is given by

$$\Psi'(s^2) \;=\; \varepsilon+\frac{1-\varepsilon}{\sqrt{1+s^2/\lambda^2}} \tag{29}$$

Since this nonlinear diffusivity is decreasing in its argument, smoothing at flow discontinuities is inhibited. For the specific choice $\Psi(s^2) := s^2$, however, homogeneous regularization with diffusivity $\Psi'(s^2)=1$ is recovered again.

The preceding diffusion–reaction system uses a common diffusivity for both channels. This avoids that edges are formed at different locations in each channel. The same coupling also appears in isotropic nonlinear diffusion filters for vector-valued images as considered by Gerig *et al.* [21], and Whitaker and Gerig [57]. Nonlinear flow-driven regularizers with different diffusivities for each channel are discussed in Section 4.

2.1.6 Anisotropic flow-driven regularization

We have seen that there exist isotropic and anisotropic image-driven regularizers as well as isotropic flow-driven ones. Thus, our taxonomy would be incomplete without having discussed anisotropic flow-driven regularizers. In the context of nonlinear diffusion filtering, anisotropic models with a diffusion tensor instead of a scalar-valued diffusivity offer advantages for images with noisy edges or interrupted structures [55].

How can one construct related optic flow methods? Let us first have a look at diffusion filtering of multichannel images. In the nonlinear anisotropic case, Weickert [52, 55] and Kimmel *et al.* [26] proposed to filter a multichannel image by using a joint diffusion tensor that depends on the gradients of all image channels. Our goal is thus to find an optic flow regularizer that leads to a coupled diffusion–reaction system where the same flux-dependent diffusion tensor $D(\nabla u_1, \nabla u_2)$ is used in each equation.

In order to derive this novel class of regularizers, we have to introduce some definitions first. As in the previous section, we consider an increasing smooth function $\Psi(s^2)$ that is convex in s. Let us assume that A is some symmetric $n \times n$ matrix with orthonormal eigenvectors w_1,...,w_n and corresponding eigenvalues σ_1,...,σ_n. Then we may formally extend the scalar-valued function $\Psi(z)$ to a matrix-valued function $\Psi(A)$ by defining $\Psi(A)$ as the matrix with eigenvectors w_1,...,w_n and eigenvalues $\Psi(\sigma_1)$,...,$\Psi(\sigma_n)$:

$$\Psi(A) \;:=\; \sum_i \Psi(\sigma_i)\, w_i w_i^T. \tag{30}$$

This definition can be motivated from the case where $\Psi(z)$ is represented by a power series $\sum_{k=0}^{\infty} c_k\, z^k$. Then it is easy to see that the corresponding matrix-valued power series $\sum_{k=0}^{\infty} c_k A^k$ has the eigenvectors w_1,...,w_n and eigenvalues $\Psi(\sigma_1)$,...,$\Psi(\sigma_n)$. Another definition that is useful for our considerations below is the trace of a quadratic matrix $A = (a_{ij})$. It is the sum of its diagonal elements, or – equivalently – the sum of its eigenvalues:

$$\mathrm{tr}\,(A) := \sum_i a_{ii} = \sum_i \sigma_i. \tag{31}$$

With these notations we consider the regularizer

$$V_{AF}(\nabla f, \nabla u_1, \nabla u_2) \;:=\; \mathrm{tr}\,\Psi\left(\nabla u_1 \nabla u_1^T + \nabla u_2 \nabla u_2^T\right). \tag{32}$$

Its argument

$$J := \nabla u_1 \nabla u_1^T + \nabla u_2 \nabla u_2^T \tag{33}$$

is a symmetric and positive semidefinite 2×2 matrix. Hence, there exist two orthonormal eigenvectors v_1, v_2 with corresponding nonnegative eigenvalues μ_1, μ_2. These eigenvalues specify the contrast of the vector-valued image (u_1, u_2) in the directions v_1 and v_2, respectively. This concept has been introduced by Di Zenzo for edge analysis of multichannel images [15]. It can be regarded as a generalization of the structure tensor [19], and it is related to the first fundamental form in differential geometry [28].

Our result below states that the regularizer (32) leads to the desired nonlinear anisotropic diffusion-reaction system.

Proposition 1 (Anisotropic Flow-Driven Regularization)
For the energy functional (3) with the regularizer (32), the corresponding steepest descent diffusion–reaction system is given by

$$\begin{aligned}\partial_t u_1 &= \mathrm{div}\,\left(D(\nabla u_1, \nabla u_2)\,\nabla u_1\right) - \tfrac{1}{\alpha} f_x (f_x u_1 + f_y u_2 + f_\theta), \qquad (34)\\ \partial_t u_2 &= \mathrm{div}\,\left(D(\nabla u_1, \nabla u_2)\,\nabla u_2\right) - \tfrac{1}{\alpha} f_y (f_x u_1 + f_y u_2 + f_\theta), \qquad (35)\end{aligned}$$

where the diffusion tensor satisfies

$$D(\nabla u_1, \nabla u_2) \;:=\; \Psi'\left(\nabla u_1 \nabla u_1^T + \nabla u_2 \nabla u_2^T\right). \tag{36}$$

Proof. The Euler–Lagrange equations for minimizing the energy

$$E(u_1, u_2) \;:=\; \int_\Omega \Big((f_x u_1 + f_y u_2 + f_\theta)^2 \;+\; \alpha\,\mathrm{tr}\,\Psi\left(\nabla u_1 \nabla u_1^T + \nabla u_2 \nabla u_2^T\right) \Big)\, dx\, dy \tag{37}$$

are given by

$$\begin{aligned}\partial_x \partial_{u_{1x}} (\mathrm{tr}\,\Psi(J)) + \partial_y \partial_{u_{1y}} (\mathrm{tr}\,\Psi(J)) - \tfrac{2}{\alpha} f_x (f_x u_1 + f_y u_2 + f_\theta) &= 0, \qquad (38)\\ \partial_x \partial_{u_{2x}} (\mathrm{tr}\,\Psi(J)) + \partial_y \partial_{u_{2y}} (\mathrm{tr}\,\Psi(J)) - \tfrac{2}{\alpha} f_y (f_x u_1 + f_y u_2 + f_\theta) &= 0. \qquad (39)\end{aligned}$$

In order to simplify the evaluation of the first and second summand in both equations, we replace (x, y) by (x_1, x_2), and denote by e_i the unit vector in x_i direction. Together with the identities

$$
\begin{aligned}
\Psi'(J) &= \Psi'(\mu_1)\, v_1 v_1^T + \Psi'(\mu_2)\, v_2 v_2^T, && (40)\\
\mathrm{tr}\,(ab^T) &= a^T b, && (41)\\
\mathrm{div}\,(a) &= \sum_i \partial_{x_i}(e_i^T a) && (42)
\end{aligned}
$$

it follows that

$$
\begin{aligned}
\sum_i \partial_{x_i}\partial_{u_{k x_i}}(\mathrm{tr}\,\Psi(J)) &= \sum_i \partial_{x_i}\mathrm{tr}\,\left(\Psi'(J)\,\partial_{u_{k x_i}} J\right)\\
&\overset{(33)}{=} \sum_i \partial_{x_i}\mathrm{tr}\,\left(\Psi'(J)\,(e_i \nabla u_k^T + \nabla u_k e_i^T)\right)\\
&\overset{(40)}{=} \sum_i \partial_{x_i}\mathrm{tr}\,\sum_j \Psi'(\mu_j)\, v_j v_j^T (e_i \nabla u_k^T + \nabla u_k e_i^T)\\
&= \sum_i \partial_{x_i}\mathrm{tr}\,\sum_j \Psi'(\mu_j)\,\left((v_j^T e_i)(v_j \nabla u_k^T) + (v_j^T \nabla u_k)(v_j e_i^T)\right)\\
&\overset{(41)}{=} 2\sum_i \partial_{x_i} \sum_j \Psi'(\mu_j)(e_i^T v_j)(v_j^T \nabla u_k)\\
&= 2\sum_i \partial_{x_i}\left(e_i^T \Psi'(J) \nabla u_k\right)\\
&\overset{(42)}{=} 2\ \mathrm{div}\ (\Psi'(J)\,\nabla u_k) \qquad (k = 1, 2). && (43)
\end{aligned}
$$

Plugging this result into the Euler–Lagrange equations concludes the proof. □

It should be noted that, in general, the eigenvalues $\Psi'(\mu_1)$ and $\Psi'(\mu_2)$ of the diffusion tensor are not equal. Therefore, we have a real anisotropic diffusion process with different behaviour in different directions. Homogeneous regularization is a special case of the regularizer (32), if $\Psi(s^2) := s^2$.

An interesting similarity between the isotropic regularizer (25) and its anisotropic counterpart (32) becomes explicit when writing (25) as

$$
V_{IF}(\nabla f, \nabla u_1, \nabla u_2) \;=\; \Psi\left(\mathrm{tr}\,(\nabla u_1 \nabla u_1^T + \nabla u_2 \nabla u_2^T)\right). \qquad (44)
$$

This shows that it is sufficient to exchange the role of the trace operator and the penalty function Ψ to switch between both regularization techniques. Another structural similarity will be discussed in Section 4.

2.2 A unifying framework

Let us now make a synthesis of all previously discussed models. Table 1 gives an overview of the smoothness terms that we have investigated so far.

Table 1: Classification of regularizers for optic flow models.

	isotropic	anisotropic
image-driven	$g(\|\nabla f\|^2) \sum_{i=1}^{2} \|\nabla u_i\|^2$	$\sum_{i=1}^{2} \nabla u_i^T D(\nabla f) \nabla u_i$
flow-driven	$\Psi\left(\sum_{i=1}^{2} \|\nabla u_i\|^2\right)$	$\mathrm{tr}\ \Psi\left(\sum_{i=1}^{2} \nabla u_i \nabla u_i^T\right)$

One may regard these regularizers as special cases of two more general models. Using the compact notation $\nabla u := (\nabla u_1, \nabla u_2)$, the first model has the structure

$$V_1(\nabla f, \nabla u) \;:=\; \Psi\left(\mathrm{tr}\, \nabla u^T D(\nabla f) \nabla u\right). \tag{45}$$

For $\Psi(s^2) := s^2$, this model comprises pure image-driven models, regardless whether they are isotropic ($D(\nabla f) := g(|\nabla f|^2) I$) or anisotropic. Isotropic flow-driven models arise for $D := I$. In the general case, the model may be both image-driven and flow-driven.

The second model can be written as

$$V_2(\nabla f, \nabla u) \;:=\; \mathrm{tr}\, \Psi\left(\nabla u\, D(\nabla f) \nabla u^T\right). \tag{46}$$

It comprises the anisotropic flow-driven case and its combinations with image-driven approaches. Note the large structural similarities between (45) and (46).

Both models can be assembled to the regularizer

$$V(\nabla f, \nabla u) \;:=\; (1-\beta)\, \Psi\left(\mathrm{tr}\, \nabla u^T D(\nabla f) \nabla u\right) \;+\; \beta\, \mathrm{tr}\, \Psi\left(\nabla u\, D(\nabla f) \nabla u^T\right) \tag{47}$$

where the paramter $\beta \in [0,1]$ determines the anisotropy. This regularizer is embedded into the general optic flow functional

$$E(u) \;=\; \int_{\Omega} \Big((f_x u_1 + f_y u_2 + f_\theta)^2 + \alpha\, V(\nabla f, \nabla u)\Big)\, dx\, dy. \tag{48}$$

2.3 Spatio-temporal regularizers

All regularizers that we have discussed so far use only *spatial* smoothness constraints. Thus, it would be natural to impose some amount of (piecewise) temporal smoothness as well. Using our results from the previous section it is straightforward to extend the smoothness constraint into the temporal domain. Instead of calculating the optic flow (u_1, u_2) as the minimizer of the two-dimensional integral (48) for each time frame θ, we now minimize a single three-dimensional integral whose solution is the optic flow for *all* frames $\theta \in [0, T]$:

$$E(u) \;:=\; \int_{\Omega \times [0,T]} \Big((f_x u_1 + f_y u_2 + f_\theta)^2 \;+\; \alpha\, V(\nabla_\theta f, \nabla_\theta u)\Big)\, dx\, dy\, d\theta \tag{49}$$

where $\nabla_\theta := (\partial_x, \partial_y, \partial_\theta)^T$ denotes the spatio-temporal nabla operator.

The corresponding diffusion–reaction systems of spatio–temporal energy functionals have the same structure as the pure spatial ones that we investigated so far. The only difference is that the spatial nabla operator ∇ has to be replaced by its spatio-temporal analogue ∇_θ. Thus, one has to solve 3D diffusion–reaction systems instead of 2D ones.

Not many spatio-temporal regularizers have been studied in the literature so far. To the best of our knowledge, there have been no attempts to investigate rotation invariant spatio-temporal models that use homogeneous, isotropic image-driven, or anisotropic flow-driven regularizers.

Nagel [36] suggested an extension of his anisotropic image-driven smoothness constraint, where the diffusion tensor (20) is replaced by

$$D(\nabla_\theta f) \;:=\; \frac{1}{2|\nabla_\theta f|^2 + 3\lambda^2}\,\left(\nabla_\theta f^\perp \nabla_\theta f^{\perp T} + \lambda^2 I\right), \tag{50}$$

Its eigenvalues are given by

$$\lambda_1(|\nabla_\theta f|) \;=\; \frac{\lambda^2}{2|\nabla_\theta f|^2 + 3\lambda^2}, \tag{51}$$

$$\lambda_2(|\nabla_\theta f|) \;=\; \frac{|\nabla_\theta f|^2 + \lambda^2}{2|\nabla_\theta f|^2 + 3\lambda^2} \;=\; \lambda_3(|\nabla_\theta f|). \tag{52}$$

Isotropic flow-driven spatio-temporal regularizers have been studied by the authors in [56]. They showed that it outperforms a corresponding spatial regularizer at low additional computing time, if an entire image stack is to be processed.

It appears that the limited memory of previous computer architectures prevented many researchers from studying approaches with spatio-temporal regularizers, since they require to keep the entire image stack in the computer memory. On contemporary PCs or workstations, however, this is no longer a problem, if typical stack sizes are used (e.g. 32 frames with 256×256 pixels). It is thus likely that spatio-temporal regularizers will become more important in the future.

3 Well-Posedness Properties

In this section we shall prove that the energy functionals (48) and (49), respectively, admit a unique solution that continuously depends on the initial data. These favourable properties are the consequence of embedding the optic flow constraint (2) into a convex regularization approach.

From the perspective of regularization, Table 1 reveals another useful classification in this context: while image–driven models correspond to the class of *quadratic* regularizers [6], flow–driven models belong to the more general class of *non–quadratic convex* regularizers. This latter class has been suggested in [11, 45, 49] for generalizing the well-known quadratic regularization approaches (cf. [6]) used for early computational vision.

3.1 Assumptions

In the following, we do not distinguish between the approaches (48) and (49) since with $\Omega \subset \mathbb{R}^n$, our results hold true for arbitrary n. Furthermore, we assume that the function $\Psi(s^2), s \in \mathbb{R}$, is strictly convex with respect to s, and there exist constants $c_1, c_2 > 0$ such that

$$c_1 s^2 \leq \Psi(s^2) \leq c_2 s^2 \;, \quad \forall s \;. \tag{53}$$

We consider only matrices $D(\nabla f)$ that are symmetric and positive definite. We define as the space of admissible optic flow fields the set

$$\mathcal{H} := \left\{ u = (u_1, u_2)^T \;\middle|\; u_i,\, \partial_{x_j} u_i \in L^2(\Omega),\ \forall i, j \right\} \;, \tag{54}$$

endowed with the scalar product

$$(u, v)_{\mathcal{H}} := \int_\Omega \left(u^T v + \operatorname{tr} \nabla u \nabla v^T \right) \, dx_1 \cdots dx_n \tag{55}$$

and its induced norm

$$\|u\|_{\mathcal{H}} := (u, u)_{\mathcal{H}}^{1/2} \;. \tag{56}$$

In what follows, $\langle f, u \rangle$ denotes the action of some linear continuous functional $f \in \mathcal{H}^*$, i.e. some element of the dual space $\mathcal{H}^*$, on some vector field $u \in \mathcal{H}$.

3.2 Convexity

We wish to show that the functional $E(u)$ is strictly convex over $\mathcal{H}$. To this end, we may disregard linear and constant terms in $E(u)$ and consider the functional $F(u)$ defined by

$$\begin{aligned} F(u) \;&:=\; \int_\Omega \Big(\left(\nabla f^T u \right)^2 + \alpha V(\nabla f, \nabla u) \Big) \, dx_1 \cdots dx_n \\ &=\; E(u) + \langle b, u \rangle + c \end{aligned} \tag{57}$$

where

$$\langle b, u \rangle \;:=\; -2 \int_\Omega f_\theta (\nabla f^T u) \, dx_1 \cdots dx_n, \tag{58}$$

$$c \;:=\; - \int_\Omega f_\theta^2 \, dx_1 \cdots dx_n. \tag{59}$$

Strict convexity is a crucial property for the existence of a unique global minimizing optical flow field u of $E(u)$ determined as the root of the equation

$$F'(u) = b \tag{60}$$

for any linear functional $b \in \mathcal{H}^*$. We proceed in several steps. First, we consider the smoothness terms $V_1(\nabla f, \nabla u)$ and $V_2(\nabla f, \nabla u)$ separately. This can be done because the sum of convex functions is again convex. Then we consider all terms together, that is the functional $F(u)$.

The term

$$V_1(\nabla f, \nabla u) \;:=\; \Psi\left(\operatorname{tr} \nabla u^T D(\nabla f) \nabla u\right) \tag{61}$$

belongs to the class of smoothness terms which were considered in earlier work on isotropic nonlinear diffusion of multichannel images (e.g. [44]). To see this, let

$$\operatorname{vec}(\nabla u) := \begin{pmatrix} \nabla u_1 \\ \nabla u_2 \end{pmatrix} \tag{62}$$

denote the vector obtained by stacking the columns of ∇u one upon the other, and let $|\cdot|_D$ denote the norm induced by the scalar product

$$\left(\operatorname{vec}(\nabla u), \operatorname{vec}(\nabla v)\right)_D \;:=\; \operatorname{vec}(\nabla u)^T \begin{pmatrix} D(\nabla f) & 0 \\ 0 & D(\nabla f) \end{pmatrix} \operatorname{vec}(\nabla v) \tag{63}$$

Then V_1 can be rewritten as

$$\Psi\left(\operatorname{tr} \nabla u^T D(\nabla f) \nabla u\right) = \Psi(|\operatorname{vec}(\nabla u)|_D^2), \tag{64}$$

and the framework in [44] is applicable.

The second anisotropic and flow–driven smoothness term

$$V_2(\nabla f, \nabla u) \;:=\; \operatorname{tr} \Psi\left(\nabla u\, D(\nabla f) \nabla u^T\right) \tag{65}$$

is new in the context of optical flow computation. Note that by contrast to term V_1, the function Ψ is matrix-valued. The strict convexity of V_2 is stated in

Proposition 2 (Matrix-Valued Convexity)
Let $\Psi : \mathbb{R} \to \mathbb{R}$ be strictly convex, A and B two positive semidefinite symmetric matrices with $A \neq B$, and $\tau \in (0,1)$. Then

$$\operatorname{tr} \Psi((1-\tau)A + \tau B) \;<\; (1-\tau) \operatorname{tr} \Psi(A) + \tau \operatorname{tr} \Psi(B). \tag{66}$$

Proof. Put $C := (1-\tau)A + \tau B$. Since A, B, C are symmetric, there are orthonormal systems of eigenvectors $\{u_i\}, \{v_i\}, \{w_i\}$ and real-valued eigenvalues $\{\alpha_i\}$, $\{\beta_i\}$, $\{\gamma_i\}$ such that

$$A = \sum_i \alpha_i u_i u_i^T\,, \quad B = \sum_i \beta_i v_i v_i^T\,, \quad C = \sum_i \gamma_i w_i w_i^T\,. \tag{67}$$

Expanding the vectors u_i, v_i with respect to the system $\{w_i\}$ gives

$$u_i = \sum_j (u_i^T w_j) w_j \ , \quad v_i = \sum_j (v_i^T w_j) w_j. \tag{68}$$

With this we obtain

$$\begin{aligned}
\sum_j \gamma_j w_j w_j^T &= C = (1-\tau)A + \tau B \\
&= (1-\tau) \sum_i \alpha_i u_i u_i^T + \tau \sum_i \beta_i v_i v_i^T \\
&= (1-\tau) \sum_i \alpha_i \sum_j (u_i^T w_j)^2 w_j w_j^T + \tau \sum_i \beta_i \sum_j (v_i^T w_j)^2 w_j w_j^T \\
&= \sum_j \Big((1-\tau) \sum_i (u_i^T w_j)^2 \alpha_i + \tau \sum_i (v_i^T w_j)^2 \beta_i \Big) w_j w_j^T. \qquad (69)
\end{aligned}$$

Comparing the coefficients shows that γ_j is a convex combination of $\{\alpha_i\}$ and $\{\beta_i\}$:

$$\gamma_j = (1-\tau) \sum_i (u_i^T w_j)^2 \alpha_i + \tau \sum_i (v_i^T w_j)^2 \beta_i \ , \quad \forall j \ . \tag{70}$$

Since $\Psi : \mathbb{R} \to \mathbb{R}$ is strictly convex and $\sum_j (u_i^T w_j)^2 = \sum_j (v_i^T w_j)^2 = 1$ for all i, we obtain

$$\begin{aligned}
\operatorname{tr} \Psi(C) &= \sum_j \Psi(\gamma_j) \qquad (71) \\
&= \sum_j \Psi\Big((1-\tau) \sum_i (u_i^T w_j)^2 \alpha_i + \tau \sum_i (v_i^T w_j)^2 \beta_i \Big) \\
&< \sum_j \Big((1-\tau) \sum_i (u_i^T w_j)^2 \Psi(\alpha_i) + \tau \sum_i (v_i^T w_j)^2 \Psi(\beta_i) \Big) \\
&= (1-\tau) \sum_i \sum_j (u_i^T w_j)^2 \Psi(\alpha_i) + \tau \sum_i \sum_j (v_i^T w_j)^2 \Psi(\beta_i) \\
&= (1-\tau) \operatorname{tr} \Psi(A) + \tau \operatorname{tr} \Psi(B) \qquad (72)
\end{aligned}$$

This concludes the convexity proof. □

So far we have shown the convexity of the smoothness term $V(\nabla f, \nabla u)$ in (57). To show that $F(u)$ is strictly convex, we may use the equivalent condition that $F'(u)$ is strongly monotone [58]:

$$\exists\, c_m > 0 \ : \quad \langle F'(u) - F'(v), u - v \rangle \ \geq \ c_m \|u - v\|_{\mathcal{H}}^2 \ , \quad \forall\, u, v \in \mathcal{H}. \tag{73}$$

Note that the smoothness term fulfills this condition because it is convex, as we have just shown. Concerning the remaining first term in (57), we have to cope with the small technical difficulty that the vector field u is multiplied with ∇f which may vanish in homogeneous image regions. In this context, we refer to in [41] where this problem has been dealt with.

3.3 Existence, uniqueness, and continuous dependence on the data

It is a well-established result (see, e.g., [58]) that property (73) together with the Lipschitz continuity of the operator F' (which holds true under mild conditions with respect to the data $\nabla f, f_\theta$, cf. [41, 46]) ensure the existence of a unique and globally minimizing optical vector field u that continuously depends on the data. To understand the latter property, suppose we are given two image sequences and corresponding functionals b_1, b_2 (cf. (57)) and minimizers u_1, u_2:

$$
\begin{aligned}
F'(u_1) &= b_1, && (74)\\
F'(u_2) &= b_2. && (75)
\end{aligned}
$$

By virtue of (73) we have

$$
\begin{aligned}
c_m \|u_1 - u_2\|^2_{\mathcal{H}} &\leq \langle F'(u_1) - F'(u_2),\, u_1 - u_2 \rangle \\
&\leq \|F'(u_1) - F'(u_2)\|_{\mathcal{H}^*} \, \|u_1 - u_2\|_{\mathcal{H}}. \qquad (76)
\end{aligned}
$$

Thus,

$$
\|u_1 - u_2\|_{\mathcal{H}} \leq \frac{1}{c_m} \|b_1 - b_2\|_{\mathcal{H}^*}\,, \quad \forall\, b_1, b_2\,. \qquad (77)
$$

This equation states that, for a slight change of the image sequence data, the corresponding optical flow field cannot arbitrarily jump but gradually changes, too. It is therefore an important robustness property.

4 Extensions

All regularizers that we have discussed so far can be motivated from existing nonlinear diffusion methods for multichannel images, where a joint diffusivity or diffusion tensor for all channels is used. As one might expect, this is not the only way to construct useful optic flow regularizers. In particular, there exists a more general design principle for anisotropic flow-driven regularizers which we will discuss next.

Our key observation for deriving this principle is an interesting relation between anisotropic flow-driven regularizers and isotropic flow-driven ones: the anisotropic regularizer $\operatorname{tr}\Psi(J)$ can be expressed by means of the eigenvalues μ_1, μ_2 of J as

$$
V_{AF}(\nabla f, \nabla u_1, \nabla u_2) = \Psi(\mu_1) + \Psi(\mu_2), \qquad (78)
$$

while its isotropic counterpart $\Psi(\operatorname{tr} J)$ can be written as

$$
V_{IF}(\nabla f, \nabla u_1, \nabla u_2) = \Psi(\mu_1 + \mu_2). \qquad (79)
$$

This observation motivates us to formulate the following design principle for rotationally invariant anisotropic flow-driven regularizers:

Design Principle (Rotationally Invariant Anisotropic Regularizers)
Assume that we are given some isotropic regularizer $\Psi(\sum_i |\nabla u_i|^2)$ *with a nonquadratic function* Ψ*, and a decomposition of its argument*

$$\sum_i |\nabla u_i|^2 = \sum_j \rho_j, \tag{80}$$

where the ρ_j *are rotationally invariant expressions. Then the regularizer* $\sum_j \Psi(\rho_j)$ *is rotationally invariant and anisotropic.*

Examples

1. The decomposion that has been used in (78) and (79) to transit from an isotropic to an anisotropic model was the trace identity

$$|\nabla u_1|^2 + |\nabla u_2|^2 \;=\; \mu_1 + \mu_2, \tag{81}$$

where μ_1 and μ_2 are the eigenvalues of $J = \nabla u_1 \nabla u_1^T + \nabla u_2 \nabla u_2^T$.

2. Schnörr [43] proposed the regularizer

$$V_{AFS}(\nabla f, \nabla u_1, \nabla u_2) \;:=\; \Psi\left(\operatorname{div}^2 u\right) + \Psi\left(\operatorname{rot}^2 u\right) + \Psi\left(\operatorname{sh}^2 u\right) \tag{82}$$

with $u := (u_1, u_2)^T$, $\operatorname{rot} u := u_{2x} - u_{1y}$, and $\operatorname{sh} u := \sqrt{(u_{2y} - u_{1x})^2 + (u_{1y} + u_{2x})^2}$. Applying the design principle, one can derive this expression from the identity [27]

$$|\nabla u_1|^2 + |\nabla u_2|^2 = \tfrac{1}{2}\left(\operatorname{div}^2 u + \operatorname{rot}^2 u + \operatorname{sh}^2 u\right). \tag{83}$$

Using the regularizer (82) in the functional (3) leads to the highly anisotropic diffusion–reaction system

$$\begin{aligned}
\partial_t u_1 \;=\;& \partial_x \Big(\big(\Psi'(\operatorname{div}^2 u) + \Psi'(\operatorname{sh}^2 u)\big)\, u_{1x} + \big(\Psi'(\operatorname{div}^2 u) - \Psi'(\operatorname{sh}^2 u)\big)\, u_{2y} \Big) \\
+\;& \partial_y \Big(\big(\Psi'(\operatorname{sh}^2 u) + \Psi'(\operatorname{rot}^2 u)\big)\, u_{1y} + \big(\Psi'(\operatorname{sh}^2 u) - \Psi'(\operatorname{rot}^2 u)\big)\, u_{2x} \Big) \\
-\;& \tfrac{1}{\alpha} f_x (f_x u_1 + f_y u_2 + f_\theta), \qquad (84)\\
\partial_t u_2 \;=\;& \partial_x \Big(\big(\Psi'(\operatorname{sh}^2 u) - \Psi'(\operatorname{rot}^2 u)\big)\, u_{1y} + \big(\Psi'(\operatorname{sh}^2 u) + \Psi'(\operatorname{rot}^2 u)\big)\, u_{2x} \Big) \\
+\;& \partial_y \Big(\big(\Psi'(\operatorname{div}^2 u) - \Psi'(\operatorname{sh}^2 u)\big)\, u_{1x} + \big(\Psi'(\operatorname{div}^2 u) + \Psi'(\operatorname{sh}^2 u)\big)\, u_{2y} \Big) \\
-\;& \tfrac{1}{\alpha} f_y (f_x u_1 + f_y u_2 + f_\theta). \qquad (85)
\end{aligned}$$

Note that now the coupling between both equations is more complicated than in the previous cases, where a joint diffusivity or a joint diffusion tensor has been used. We are not aware of similar diffusion filters for multichannel images. Well-posedness properties and experimental results for this optic flow method are presented in [43, 46].

3. Requiring that the ρ_j in (80) be rotationally invariant ensures the rotation invariance of the anisotropic regularizer. If we dispense with rotation invariance, the

design principle can still be used. As an example, let us study the flow-driven regularization methods that are considered in [4, 12, 14, 29]. They use a regularizer of type

$$V_C(\nabla f, \nabla u_1, \nabla u_2) \;:=\; \Psi(|\nabla u_1|^2) + \Psi(|\nabla u_2|^2). \tag{86}$$

According to our design principle, we may regard this regularizer as an anisotropic version of the isotropic regularizer (25). However, the decomposition of its argument into $|\nabla u_1|^2$ and $|\nabla u_2|^2$ is not rotationally invariant. The corresponding diffusion–reaction system is given by

$$\partial_t u_1 \;=\; \operatorname{div}\,\big(\Psi'(|\nabla u_1|^2)\,\nabla u_1\big) - \tfrac{1}{\alpha} f_x(f_x u_1 + f_y u_2 + f_\theta), \tag{87}$$
$$\partial_t u_2 \;=\; \operatorname{div}\,\big(\Psi'(|\nabla u_2|^2)\,\nabla u_2\big) - \tfrac{1}{\alpha} f_y(f_x u_1 + f_y u_2 + f_\theta), \tag{88}$$

which shows that both systems are completely decoupled in their diffusion terms. Thus, flow discontinuities may be created at different locations for each channel. The same decoupling appears also for some other PDE-based optic flow methods such as [40].

While each of the two diffusion processes is isotropic, the overall process reveals some anisotropy: in general, the two diffusivities $\Psi'(|\nabla u_1|^2)$ and $\Psi'(|\nabla u_2|^2)$ are not identical. Well-posedness results for this approach with a modified data term have been established by Aubert *et al.* [4].

There is also a number of related stochastic methods that lead to discrete models which are not consistent approximations to rotation invariant processes [7, 8, 10, 23, 31, 33]. Nonconvex regularizers are typically used in these approaches. Discrete spatio-temporal versions of the regularizer (86) are investigated in [7, 33].

It is a challenging open question whether there exist more useful rotation invariant convex regularizers than the ones we have just discussed. This is one of our current research topics.

5 Summary and Conclusions

The goal of this paper was to derive a diffusion theory for optic flow functionals. Minimizing optic flow functionals by steepest descent leads to a set of two coupled diffusion–reaction systems. Since similar equations appear for diffusion filtering of multi-channel images, the question arises whether there are optic flow analogues to the various kinds of diffusion filters.

We saw that image-driven optic flow regularizers correspond to linear diffusion filters, while flow-driven regularizers create nonlinear diffusion processes. Pure spatial regularizers can be expressed as 2D diffusion–reaction processes, and spatio-temporal regularizers may be regarded as generalizations to the 3D case. This taxonomy helped us not only to classify existing methods within a unifying framework, but also to identify gaps, where no models are available in the current literature. We filled these gaps by deriving suitable methods with the specified properties, and we proved well-posedness for the class of convex diffusion-based optic flow regularization methods.

One important novelty along these lines was the derivation of regularizers that can be related to anisotropic diffusion filters with a matrix-valued diffusion tensor. This also enabled us to propose a design principle for anisotropic regularizers, and we discovered an interesting structural similarity between isotropic and anisotropic models: it is sufficient to exchange the role of the trace operator and the penalty function in order to switch between the two models.

We are convinced that these relations are only the starting point for many more fruitful interactions between the theories of diffusion filtering and variational optic flow methods. Diffusion filtering has progressed very much in recent years, and so it appears appealing to incorporate recent results from this area into optic flow methods. Conversely, it is clear that novel optic flow regularizers can also be regarded as energy functionals for suitable diffusion filters.

We hope that our systematic taxonomy provides a unifying platform for algorithms for the entire class of convex variational optic flow methods. Our future plans are to use such a platform for a detailed performance evaluation of the different methods in this paper, and for a systematic comparison of different numerical algorithms. Another point on our agenda is an investigation of alternative rotation-invariant decompositions that can be applied to construct useful anisotropic regularizers.

Acknowledgement. CS completed his doctoral thesis under the supervision of Prof. Nagel in 1991. He is grateful to Prof. Nagel who introduced him to the field of computer vision.

References

[1] L. Alvarez, J. Esclarín, M. Lefébure and J. Sánchez, *A PDE model for computing the optical flow*, Proc. XVI Congreso de Ecuaciones Diferenciales y Aplicaciones (C.E.D.Y.A. XVI, Las Palmas de Gran Canaria, Sept. 21–24, 1999), 1349–1356, 1999.

[2] L. Alvarez, J. Weickert, J. Sánchez, *Reliable estimation of dense optical flow fields with large displacements*, Report No. 2, Instituto Universitario de Ciencias y Tecnologías Cibernéticas, University of Las Palmas de Gran Canaria, Campus de Tafira, 35017 Las Palmas, Spain, 1999. To appear in International Journal of Computer Vision.

[3] P. Anandan, *A computational framework and an algorithm for the measurement of visual motion*, Int. J. Comput. Vision, Vol. 2, 283–310, 1989.

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

[5] J.L. Barron, D.J. Fleet, S.S. Beauchemin, *Performance of optical flow techniques*, Int. J. Comput. Vision, Vol. 12, 43–77, 1994.

[6] M. Bertero, T.A. Poggio, V. Torre, *Ill-posed problems in early vision*, Proc. IEEE, Vol. 76, 869–889, 1988.

[7] M.J. Black, P. Anandan, *Robust dynamic motion estimation over time*, Proc. IEEE Comp. Soc. Conf. on Computer Vision and Pattern Recognition (CVPR '91, Maui, June 3–6, 1991), IEEE Computer Society Press, Los Alamitos, 292–302, 1991.

[8] M.J. Black, P. Anandan, *The robust estimation of multiple motions: Parametric and piecewise smooth flow fields*, Computer Vision and Image Understanding, Vol. 63, 75–104, 1996.

[9] A. Blake, A. Zisserman, *Visual reconstruction*, MIT Press, Cambridge (Mass.), 1987.

[10] L. Blanc–Féraud, M. Barlaud, T. Gaidon, *Motion estimation involving discontinuities in a multiresolution scheme*, Optical Engineering, Vol. 32, No. 7, 1475–1482, 1993.

[11] P. Charbonnier, L. Blanc–Féraud, G. Aubert, M. Barlaud, *Two deterministic half-quadratic regularization algorithms for computed imaging*, Proc. IEEE Int. Conf. Image Processing (ICIP–94, Austin, Nov. 13–16, 1994), Vol. 2, IEEE Computer Society Press, Los Alamitos, 168–172, 1994.

[12] I. Cohen, *Nonlinear variational method for optical flow computation*, Proc. Eighth Scandinavian Conf. on Image Analysis (SCIA '93, Tromsø, May 25–28, 1993), Vol. 1, 523–530, 1993.

[13] R. Courant, D. Hilbert, *Methods of mathematical physics*, Vol. 1, Interscience, New York, 1953.

[14] R. Deriche, P. Kornprobst, G. Aubert, *Optical-flow estimation while preserving its discontinuities: A variational approach*, Proc. Second Asian Conf. Computer Vision (ACCV '95, Singapore, December 5–8, 1995), Vol. 2, 290–295, 1995.

[15] S. Di Zenzo, *A note on the gradient of a multi-image*, Computer Vision, Graphics, and Image Processing, Vol. 33, 116–125, 1986.

[16] L.E. Elsgolc, *Calculus of variations*, Pergamon, Oxford, 1961.

[17] W. Enkelmann, *Investigation of multigrid algorithms for the estimation of optical flow fields in image sequences*, Computer Vision, Graphics and Image Processing, Vol. 43, 150–177, 1988.

[18] D.J. Fleet, A.D. Jepson, *Computation of component image velocity from local phase information*, Int. J. Comput. Vision, Vol. 5, 77–104, 1990.

[19] W. Förstner, E. Gülch, *A fast operator for detection and precise location of distinct points, corners and centres of circular features*, Proc. ISPRS Intercommission Conf. on Fast Processing of Photogrammetric Data (Interlaken, June 2–4, 1987), 281–305, 1987.

[20] B. Galvin, B. McCane, K. Novins, *Recovering motion fields: An analysis of eight optical flow algorithms*, Proc. 1998 British Machine Vision Conference (BMVC '98, Southampton, September 14–17, 1998).

[21] G. Gerig, O. Kübler, R. Kikinis, F.A. Jolesz, *Nonlinear anisotropic filtering of MRI data*, IEEE Trans. Medical Imaging, Vol. 11, 221–232, 1992.

[22] B. ter Haar Romeny, L. Florack, J. Koenderink, M. Viergever (Eds.), *Scale-space theory in computer vision*, Lecture Notes in Computer Science, Vol. 1252, Springer, Berlin, 1997.

[23] F. Heitz, P. Bouthemy, *Multimodal estimation of discontinuous optical flow using Markov random fields*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 15, 1217–1232, 1993.

[24] W. Hinterberger, *Generierung eines Films zwischen zwei Bildern mit Hilfe des optischen Flusses*, M.Sc. thesis, Industrial Mathematics Institute, University of Linz, Austria, 1999.

[25] B. Horn, B. Schunck, *Determining optical flow*, Artificial Intelligence, Vol. 17, 185–203, 1981.

[26] R. Kimmel, R. Malladi, N. Sochen, *Images as embedded maps and minimal surfaces: movies, color, texture, and volumetric medical images*, Int. J. Comput. Vision, to appear.

[27] J.J. Koenderink, *Invariant properties of the motion parallax field due to the movement of rigid bodies relative to an observer*, Optica Acta, Vol. 22, 773–791, 1975.

[28] E. Kreyszig, *Differential geometry*, University of Toronto Press, Toronto, 1959.

[29] A. Kumar, A.R. Tannenbaum, G.J. Balas, *Optic flow: a curve evolution approach*, IEEE Trans. Image Proc., Vol. 5, 598–610, 1996.

[30] P.J.M. van Laarhoven, E.H.L. Aarts, *Simulated annealing: theory and applications*, Reidel, Dordrecht, 1988.

[31] E. Mémin, P. Pérez, *Dense estimation and object-based segmentation of the optical flow with robust techniques*, IEEE Trans. Image Proc, Vol. 7, 703–719, 1998.

[32] A. Mitiche, P. Bouthemy, *Computation and analysis of image motion: a synopsis of current problems and methods*, Int. J. Comput. Vision, Vol. 19, 29–55, 1996.

[33] D.W. Murray, B.F. Buxton, *Scene segmentation from visual motion using global optimization*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 9, 220–228, 1987.

[34] H.H. Nagel, *Constraints for the estimation of displacement vector fields from image sequences*, Proc. Eighth Int. Joint Conf. on Artificial Intelligence (IJCAI '83, Karlsruhe, August 8–12, 1983), 945–951, 1983.

[35] H.-H. Nagel, *On the estimation of optical flow: relations between different approaches and some new results*, Artificial Intelligence, Vol. 33, 299–324, 1987.

[36] H.-H. Nagel, *Extending the 'oriented smoothness constraint' into the temporal domain and the estimation of derivatives of optical flow*, O. Faugeras (Ed.), Computer vision – ECCV '90, Lecture Notes in Computer Science, Vol. 427, Springer, Berlin, 139–148, 1990.

[37] H.-H. Nagel, W. Enkelmann, *An investigation of smoothness constraints for the estimation of displacement vector fields from images sequences*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 8, 565–593, 1986.

[38] P. Nesi, *Variational approach to optical flow estimation managing discontinuities*, Image and Vision Computing, Vol. 11, 419–439, 1993.

[39] M. Nielsen, P. Johansen, O.F. Olsen, J. Weickert (Eds.), *Scale-space theories in computer vision*, Lecture Notes in Computer Science, Springer, Berlin, Vol. 1682, 1999.

[40] M. Proesmans, L. Van Gool, E. Pauwels, A. Oosterlinck, *Determination of optical flow and its discontinuities using non-linear diffusion*, J.-O. Eklundh (Ed.), Computer vision – ECCV '94, Lecture Notes in Computer Science, Vol. 801, Springer, Berlin, 295–304, 1994.

[41] C. Schnörr, *Determining optical flow for irregular domains by minimizing quadratic functionals of a certain class*, Int. J. Comput. Vision, Vol. 6, 25–38, 1991.

[42] C. Schnörr, *On functionals with greyvalue-controlled smoothness terms for determining optical flow*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 15, 1074–1079, 1993.

[43] C. Schnörr, *Segmentation of visual motion by minimizing convex non-quadratic functionals*, Proc. 12th Int. Conf. Pattern Recognition (ICPR 12, Jerusalem, Oct. 9–13, 1994), Vol. A, IEEE Computer Society Press, Los Alamitos, 661–663, 1994.

[44] C. Schnörr, *Convex variational segmentation of multi-channel images*, M.-O. Berger, R. Deriche, I. Herlin, J. Jaffré, J.-M. Morel (Eds.), ICAOS '96: Images, wavelets and PDEs, Lecture Notes in Control and Information Sciences, Vol. 219, Springer, London, 201–207, 1996.

[45] C. Schnörr, R. Sprengel, *A nonlinear regularization approach to early vision*, Biol. Cybernetics, Vol. 72, 141–149, 1994.

[46] C. Schnörr, R. Sprengel, B. Neumann, *A variational approach to the design of early vision algorithms*, Computing, Suppl. 11, 149–165, 1996.

[47] M.A. Snyder, *On the mathematical foundations of smoothness constraints for the determination of optical flow and for surface reconstruction*, IEEE Trans. Pattern Anal. Mach. Intell., Vol. 13, 1105–1114, 1991.

[48] J. Sporring, M. Nielsen, L. Florack, P. Johansen (Eds.), *Gaussian scale-space theory*, Kluwer, Dordrecht, 1997.

[49] R.L. Stevenson, B.E. Schmitz, E.J. Delp, *Discontinuity preserving regularization of inverse visual problems*, IEEE Trans. Systems, Man and Cybernetics, Vol. 24, 455–469, 1994.

[50] C. Stiller, J. Konrad, *Estimating motion in image sequences*, IEEE Signal Proc. Magazine, Vol. 16, 70–91, 1999.

[51] J. Weber, J. Malik, *Robust computation of optical flow in a multi-scale differential framework*, Int. J. Comput. Vision, Vol. 14, 67–81, 1995.

[52] J. Weickert, *Scale-space properties of nonlinear diffusion filtering with a diffusion tensor*, Report No. 110, Laboratory of Technomathematics, University of Kaiserslautern, P.O. Box 3049, 67653 Kaiserslautern, Germany, 1994.

[53] J. Weickert, *Anisotropic diffusion in image processing*, Teubner, Stuttgart, 1998.

[54] J. Weickert, *On discontinuity-preserving optic flow*, S. Orphanoudakis, P. Trahanias, J. Crowley, N. Katevas (Eds.), Proc. Computer Vision and Mobile Robotics Workshop (CVMR '98, Santorini, Sept. 17–18, 1998), 115–122, 1998.

[55] J. Weickert, *Coherence-enhancing diffusion of colour images*, Image and Vision Computing, Vol. 17, 199–210, 1999.

[56] J. Weickert, C. Schnörr, *Räumlich–zeitliche Berechnung des optischen Flusses mit nichtlinearen flußabhängigen Glattheitstermen*, W. Förstner, J.M. Buhmann, A. Faber, P. Faber (Eds.), Mustererkennung 1999, Springer, Berlin, 317–324, 1999.

[57] R. Whitaker, G. Gerig, *Vector-valued diffusion*, B.M. ter Haar Romeny (Ed.), Geometry-driven diffusion in computer vision, Kluwer, Dordrecht, 93–134, 1994.

[58] E. Zeidler, *Nonlinear functional analysis and its applications*, Vol. IIb, Springer, Berlin, 1990.