

NFN - Nationales Forschungsnetzwerk

Geometry + Simulation

<http://www.gs.jku.at>



**Simultaneous optical flow and
source estimation: space-time
discretization and
preconditioning**

R. Andreev, O. Scherzer, W.
Zulehner

G+S Report No. 20

November 2014

FWF

Der Wissenschaftsfonds.



JKU
JOHANNES KEPLER
UNIVERSITY LINZ

Simultaneous optical flow and source estimation: space-time discretization and preconditioning

Roman Andreev¹, Otmar Scherzer^{1,2}, and Walter Zulehner³

¹ Johann Radon Institute for Computational and Applied Mathematics (RICAM)

Austrian Academy of Sciences
Altenberger Str. 69, 4040 Linz, Austria

² Computational Science Center
University of Vienna

Oskar-Morgenstern Platz 1, 1090 Wien, Austria

³ Institute of Computational Mathematics

Johannes Kepler University
Altenberger Str. 69, 4040 Linz, Austria

Abstract. We consider the simultaneous estimation of an optical flow field and an illumination source term in a movie sequence. The particular optical flow equation is obtained by assuming that the image intensity is a conserved quantity up to possible sources and sinks which represent varying illumination. We formulate this problem as an energy minimization problem and propose a space-time simultaneous discretization for the optimality system in saddle-point form. We investigate a preconditioning strategy that renders the discrete system well-conditioned uniformly in the discretization resolution. Numerical experiments complement the theory.

November 27, 2014

1 Introduction

Optical flow is the apparent motion of objects, surfaces, and edges in a visual scene. In Computer Vision the optical flow is calculated as the flow field registering pixels of a movie, which we understand as a sequence of ordered images. The classical optical flow equation is based on the following assumptions:

1. **Brightness constancy** along each characteristic of the flow.
2. A **slowly varying** image sequence.

The first assumption implies that changes of intensities caused by varying illumination (for instance shades) influence and distort the recovered flow. In this paper we admit violations of this assumption, and aim at estimating global changes of illumination and the effect of them onto the flow field. Indeed, the problem we consider here consists in simultaneously determining the optical flow, the brightness, *and* its possible sources and sinks in space and time. While standard optical flow algorithms process two successive frames of a movie sequence at a time [16,17,18,3,13,6,1], we work with a continuous space-time simultaneous

formulation similarly to [20]. An brief overview of gradient based optical flow models and algorithms may be found in [14].

Related to our work is [21]. Therein, the standard optical flow constraint was replaced by the continuity equation $\partial_t \rho + \nabla \cdot (\rho v) = 0$, motivated by mass conservation in fluid flow. We relax the continuity equation by allowing distributed *sources* g , and replace the momentum ρv by one unknown j , called the *flux*; we thus obtain the constraint $\partial_t \rho + \nabla \cdot j = g$.

The work [15], see also references therein, suggested incorporating physical models to account for illumination changes, such as the diffusion of the heat observed by infrared cameras. The parameters governing the physical model are then estimated in parallel with the optical flow. Our ansatz is more naive in that we treat the possible change in illumination as an additional unknown, but a physical description can in principle be used instead.

As a further difference to the bulk of optical flow literature we treat the density ρ as an unknown. We thus estimate (ρ, j, g) in parallel by minimizing a functional \mathcal{J} which consists of a discrepancy term for ρ and regularization terms for j and g . This principle similar to [7]; however, using the flux j instead of the velocity v results in *linear* optimality conditions characterizing the minimizer of the functional \mathcal{J} . While this is convenient for computation, it leads to somewhat unexpected flow patterns (see Sections 5–6).

It is important to note that these optimality conditions form a set of equations that are coupled in space-time. This is typical for optimal control problems with spatio-temporal constraints due to the coupling of the original problem forward in time and the adjoint problem backward in time. The main contribution of this work is to propose and test a preconditioner for the discretized space-time system in saddle-point form that is robust in the discretization parameters.

The outline of the paper is as follows. In Section 2 we introduce our optical flow model, which consists in minimization of a quadratic cost functional over a space of space-time dependent functions subject to an optical flow constraint. In Section 3 the functional analytic framework is established. In Section 4 we rewrite the optimality conditions in a saddle-point form and discuss its discretization and preconditioning. In Section 5 we report on our numerical experiments. Section 6 concludes the paper.

In the following, $\langle \cdot, \cdot \rangle$ denotes the duality pairing, while $\langle \cdot, \cdot \rangle_X$ and $\|\cdot\|_X$ denote the scalar product and the norm of a Hilbert space X . The usual Lebesgue and Sobolev spaces on a domain D are written as $L_p(D)$ and $H^s(D)$. We abbreviate $\|\cdot\|_{L_2(D)}$ as $\|\cdot\|_D$, and similarly for other L_2 spaces.

2 Optical flow model

We consider a movie sequence of time length $T > 0$, with each frame defined on a rectangular domain $D \subset \mathbb{R}^2$. We write $J := (0, T)$ for the temporal interval. Our optical flow model is based on the scalar transport equation

$$\boxed{\partial_t \rho + \nabla \cdot j = g \quad \text{in } J \times D.} \quad (1)$$

Here, ρ denotes the *density* which is the intensity of the image considered as a piecewise constant function; j denotes the *flux* which is the optical flow weighted by the density; and g is a function that models varying spatial illumination. The divergence $\nabla \cdot$ acts on the spatial variable only. The functions ρ and g are scalar-valued, while j is vector-valued on $\bar{J} \times D$. Integrating (1) over a subdomain $D' \subset D$ and using the divergence theorem for the j term, one sees that the density ρ is either transported in or out of D' over its boundary, or is created/annihilated by means of the source term g .

In this paper we consider the problem of identifying the flux j and the source g from a finite number of frames of a movie, indexed by $\mathcal{T} \subseteq \bar{J}$:

$$\rho_\tau \in L_2(D), \quad \tau \in \mathcal{T}. \quad (2)$$

We aim at minimizing the data fidelity functional

$$F_{\mathcal{T}}(\rho) := \frac{1}{2} \frac{1}{T} \sum_{\tau \in \mathcal{T}} \|\rho_\tau - \rho(\tau)\|_D^2 \quad (3)$$

subject to the transport equation (1) and further constraints on the flux j and the source g as discussed below. Here, the symbol $\sum_{\tau \in \mathcal{T}}$ denotes the averaged sum $\frac{1}{\#\mathcal{T}} \sum_{\tau \in \mathcal{T}}$.

For comparison purposes we recall the standard optical flow equation, which reads as follows:

$$\partial_t \rho + \Phi \cdot \nabla \rho = 0 \quad \text{in } J \times D. \quad (4)$$

Equation (4) can be formally derived from (1) by identifying j and $\rho \Phi$ and neglecting small terms: Indeed, from (1) we get

$$\partial_t \rho + \Phi \cdot \nabla \rho + \rho \nabla \cdot \Phi = 0 \quad \text{in } J \times D. \quad (5)$$

Hence, if the term $\rho \nabla \cdot \Phi$ is negligible, and if there are no changes in illumination, the equations are identical.

3 Functional analytic framework

We now introduce function spaces for which the transport equation (1) is well-defined. We shall work with image intensity $\rho \in H^{1,0} := H^1(J; L_2(D))$, flux $j \in H^{0,\text{div}} := L_2(J; H^{\text{div}}(D))$, and source $g \in H^{0,0} := L_2(J; L_2(D))$. The superscripts indicate the Sobolev smoothness in time and space, respectively. Here, $H^{\text{div}}(D)$ is the space of vector-valued functions in $[L_2(D)]^2$ with distributional divergence in $L_2(D)$. As a consequence of the Fubini–Tonelli theorem [22, Section 0.3], the space $H^{0,0}$ is isometrically isomorphic to $L_2(J \times D)$. The norms on $H^1(J)$ and $H^{\text{div}}(D)$ are defined by

$$\begin{aligned} \|f\|_{H^1(J)}^2 &:= T^{-2} \|f\|_J^2 + \|f'\|_J^2, \quad f \in H^1(J), \\ \|v\|_{H^{\text{div}}(D)}^2 &:= \text{diam}(D)^{-2} \|v\|_D^2 + \|\nabla \cdot v\|_D^2, \quad v \in H^{\text{div}}(D), \end{aligned}$$

where the scaling is so as to match the units. We introduce the product space

$$X := H^{1,0} \times H^{0,\text{div}} \times H^{0,0}$$

and associate it with the norm

$$\|(\rho, j, g)\|_X^2 := \|\rho\|_{H^{1,0}}^2 + \|j\|_{H^{0,\text{div}}}^2 + \|g\|_{H^{0,0}}^2. \quad (6)$$

The transport equation constraint (1) is reformulated via the linear operator

$$G : X \rightarrow H^{0,0}, \quad G(\rho, j, g) := \partial_t \rho + \nabla \cdot j - g. \quad (7)$$

Since G is continuous, the preimages of closed subsets are closed; in particular, its kernel

$$\Gamma := G^{-1}(0) \subset X \quad (8)$$

is a closed linear subspace. It contains precisely the triples $(\rho, j, g) \in X$ that satisfy the transport equation (1).

Given two regularization parameters $\alpha_j > 0$ and $\alpha_g > 0$, we define the penalization functional

$$R(j, g) := \frac{\alpha_j}{2} \|j\|_{H^{0,\text{div}}}^2 + \frac{\alpha_g}{2} \|g\|_{H^{0,0}}^2, \quad (j, g) \in H^{0,\text{div}} \times H^{0,0}, \quad (9)$$

and, recalling $F_{\mathcal{T}}$ from (3), the cost functional

$$\mathcal{J}(\rho, j, g) := F_{\mathcal{T}}(\rho) + R(j, g), \quad (\rho, j, g) \in X. \quad (10)$$

The parameters α_j and α_g are dimensionless. The following lemma is straightforward.

Lemma 1. *The functional \mathcal{J} is Gâteaux differentiable on X . With ρ_τ from (2), its Gâteaux derivative at $(\rho, j, g) \in X$ is the continuous linear functional $\mathcal{J}'(\rho, j, g) = A(\rho, j, g) - \ell$, where $A : X \rightarrow X'$ and $\ell \in X'$ are given by*

$$\langle A(\rho, j, g), (\tilde{\rho}, \tilde{j}, \tilde{g}) \rangle = \frac{1}{T} \sum_{\tau \in \mathcal{T}} \langle \rho(\tau), \tilde{\rho}(\tau) \rangle_D + \alpha_j \langle j, \tilde{j} \rangle_{H^{0,\text{div}}} + \alpha_g \langle g, \tilde{g} \rangle_{H^{0,0}}, \quad (11)$$

$$\ell(\tilde{\rho}, \tilde{j}, \tilde{g}) = \frac{1}{T} \sum_{\tau \in \mathcal{T}} \langle \rho_\tau, \tilde{\rho}(\tau) \rangle_D, \quad \forall (\tilde{\rho}, \tilde{j}, \tilde{g}) \in X. \quad (12)$$

Moreover, A is continuous, self-adjoint, and Γ -elliptic,

$$\exists \alpha > 0 : \quad \langle Av, v \rangle \geq \alpha \|v\|_X^2 \quad \forall v \in \Gamma. \quad (13)$$

Further, A induces a seminorm $\|\cdot\|$ on X given by

$$\|(\rho, j, g)\|^2 := \frac{1}{T} \sum_{\tau \in \mathcal{T}} \|\rho(\tau)\|_D^2 + \alpha_j \|j\|_{H^{0,\text{div}}}^2 + \alpha_g \|g\|_{H^{0,0}}^2. \quad (14)$$

This seminorm is a norm $\Gamma = G^{-1}(0)$, where it is equivalent to $\|\cdot\|_X$.

With the help of this lemma and standard arguments from variational calculus, one can show strict convexity and coercivity properties of the functional \mathcal{J} :

Lemma 2. *The functional \mathcal{J} defined in (10) is continuous and convex on X . Moreover, it is strictly convex on Γ ,*

$$\mathcal{J}(\lambda v_+ + (1 - \lambda)v_-) < \lambda \mathcal{J}(v_+) + (1 - \lambda)\mathcal{J}(v_-) \quad \forall v_{\pm} \in \Gamma, \lambda \in (0, 1), \quad (15)$$

$v_+ \neq v_-$, and coercive on Γ ,

$$\exists \alpha > 0, \beta \geq 0: \quad \mathcal{J}(v) \geq \alpha \|v\|_X^2 - \beta \quad \forall v \in \Gamma. \quad (16)$$

Our optical flow problem now reads as follows:

Definition 1. *For given regularization parameters $\alpha = (\alpha_j, \alpha_g)$, the optical flow u_α is the unique minimizer of the functional \mathcal{J} over Γ .*

Existence of the minimizer is due to standard arguments of the calculus of variations [12], because \mathcal{J} is non-negative and proper ($\mathcal{J} \neq \infty$) on Γ . Uniqueness follows from the strict convexity (15) of \mathcal{J} on Γ .

The minimizer u_α of \mathcal{J} is equivalently characterized in terms of first order optimality conditions. By [12, Theorems 1.3–1.4 in §3.1.3], we have $\mathcal{J}'(u_\alpha) = 0$, where $\mathcal{J}'(u) \in \Gamma'$ is the Gateaux derivative of \mathcal{J} at $u \in \Gamma$. To simplify the notation we will omit the dependence on α . Using Lemma 1, the requirement $\mathcal{J}'(u) = 0 \in \Gamma'$ is equivalent to the variational problem

$$\boxed{\langle Au - \ell, v \rangle = 0 \quad \forall v \in \Gamma.} \quad (17)$$

The fact that this variational problem is posed on the implicitly defined subspace $\Gamma = G^{-1}(0)$ leads to the saddle-point problem introduced in the next section.

4 Numerical solution

In order to minimize \mathcal{J} we solve the equivalent variational problem (17). The constraint $u \in \Gamma$ is implemented using a Lagrange multiplier λ to complement (17) to a saddle-point problem: Find $(u, \lambda) \in X \times H^{0,0}$ such that

$$\langle Au, v \rangle + \langle Gu, \mu \rangle + \langle Gv, \lambda \rangle = \langle \ell, v \rangle \quad \forall (v, \mu) \in X \times H^{0,0}. \quad (18)$$

4.1 Discretization

In order to discretize (18) we introduce finite dimensional subspaces $X_h \subset X$ and $M_h \subset H^{0,0}$, as specified below, and consider the discrete saddle-point system: Find $(u_h, \lambda_h) \in X_h \times M_h$ such that

$$\langle Au_h, v \rangle + \langle Gu_h, \mu \rangle + \langle Gv, \lambda_h \rangle = \langle \ell, v \rangle \quad \forall (v, \mu) \in X_h \times M_h. \quad (19)$$

With $A_h : X_h \rightarrow X'_h$ and $G_h : X_h \rightarrow M'_h$, defined by $A_h v := (Av)|_{X_h}$ and $G_h v := (Gv)|_{M_h}$, $v \in X_h$, as well as $\ell_h := \ell|_{X_h}$, the system (19) is equivalent to

$$\mathcal{A}_h \begin{pmatrix} u_h \\ \lambda_h \end{pmatrix} := \begin{pmatrix} A_h & G'_h \\ G_h & 0 \end{pmatrix} \begin{pmatrix} u_h \\ \lambda_h \end{pmatrix} = \begin{pmatrix} \ell_h \\ 0 \end{pmatrix}, \quad (20)$$

where $G'_h : M_h \rightarrow X'_h$ is the $H^{0,0}$ -adjoint.

Analogously to the continuous case, we set

$$\Gamma_h := \{u_h \in X_h : \langle Gu_h, \mu \rangle = 0 \ \forall \mu \in M_h\}. \quad (21)$$

The Brezzi equivalence theorem, see [8, Satz III.4.3], states that the left-hand-side of (19) defines an isomorphism $X_h \times M_h \rightarrow X'_h \times M'_h$ if and only if the following two conditions are fulfilled:

1. A is Γ_h -elliptic,

$$\exists \alpha_h > 0 : \quad \langle Av, v \rangle \geq \alpha_h \|v\|_X^2 \quad \forall v \in \Gamma_h. \quad (22)$$

2. The discrete inf-sup constant is positive,

$$\beta_h := \inf_{\lambda \in M_h \setminus \{0\}} \sup_{v \in X_h \setminus \{0\}} \frac{\langle Gv, \lambda \rangle}{\|v\|_X \|\lambda\|_{H^{0,0}}} > 0. \quad (23)$$

We shall employ a discretization $X_h \times M_h \subset X \times H^{0,0}$ for which the two rather restrictive conformity conditions

$$\Gamma_h \subset \Gamma \quad \text{and} \quad M_h \subset GX_h \quad (24)$$

hold. The first, together with Γ -ellipticity (13) of A , implies Γ_h -ellipticity (22) of A , and the second immediately implies the discrete inf-sup condition (23) with $\beta_h \geq 1$.

The square domain D is partitioned into non-overlapping open rectangles and/or triangles, collected in D_h . Geometric compatibility conditions are imposed following [11, p.51]: each edge of any geometric element $K \in D_h$ is either part of the boundary ∂D or is also an edge of some other geometric element $K' \in D_h$. By D_h^∂ we mean the set of edges E of all geometric elements in D_h (shared edges occur only once). The temporal interval $J = (0, T)$ is partitioned into open subintervals collected in J_h , such that $\bar{J} = \bigcup_{I \in J_h} \bar{I}$. Quantities in $H^{0,0}$, such as $\partial_t \rho$, $\nabla \cdot j$, g , and λ , are discretized as piecewise constants on each space-time geometric element $I \times K$, $(I, K) \in J_h \times D_h$. The conserved quantity $\rho \in H^{1,0}$ is discretized by continuous functions $J \rightarrow L_2(D)$ that are affine on each $I \in J_h$ and have values in the space of piecewise constant functions on D_h . The flux $j \in H^{0,\text{div}}$ is discretized by assigning a flux density j_{IE} to each pair $(I, E) \in J_h \times D_h^\partial$, such that $j_{IE} = \frac{1}{|I||E|} \int_I \int_E j(t, x) \cdot n(x) d\sigma(x) dt$, and interpolated linearly into the inner of each geometric element $I \times K \in J_h \times D_h$. This is trivially possible on each rectangle, and corresponds to Raviart–Thomas interpolation on triangles [8, p.141]. It is then clear that the conformity conditions (24), and therefore (22) and (23), are satisfied.

4.2 Preconditioning

Several classes of preconditioners have been developed for saddle-point systems of the form (20), see the survey [4]. Preconditioners based on approximations of A_h^{-1} and of the Schur complement here $S_h = G_h A_h^{-1} G_h'$ are most widely used. However, our A_h is not, in general, invertible on the whole finite-dimensional space X_h . Indeed, recall from (11) that A has the block structure

$$\langle A_h v, v \rangle = \frac{1}{T} \sum_{\tau \in \mathcal{T}} \|\rho(\tau)\|_D^2 + \alpha_j \|j\|_{H^{0,\text{div}}}^2 + \alpha_g \|g\|_{H^{0,0}}^2, \quad (25)$$

where $v = (\rho, j, g)$, so that the first block is not positive definite on $H^{1,0}$.

In [5] a similar saddle-point system for a related problem from image registration was obtained, for which block-triangular preconditioners were proposed. We replace the triangular preconditioner by the symmetric and indefinite preconditioner $\mathcal{P}_h : X_h \times M_h \rightarrow X_h' \times M_h'$, given by

$$\mathcal{P}_h := \begin{pmatrix} \widehat{A}_h & 0 \\ G_h & I \end{pmatrix} \begin{pmatrix} \widehat{A}_h^{-1} & 0 \\ 0 & -\widehat{S}_h \end{pmatrix} \begin{pmatrix} \widehat{A}_h & G_h' \\ 0 & I \end{pmatrix}, \quad (26)$$

see [2]. It was shown in [19] that this allows to transform the indefinite saddle-point system into a symmetric positive definite one for a new inner product (see (27) below), so that the conjugate gradient (CG) method can be applied. This is not possible for the preconditioned system of [5] due to nonreal spectrum. Another well-known method to transform the indefinite saddle-point system into a symmetric positive definite one was proposed in [10]. There it is also required that A_h be positive definite on X_h , which is not true in our case.

It was shown in [19, Theorem 2.1] that if

$$\mathcal{N}_h := \begin{pmatrix} \widehat{A}_h - A_h & 0 \\ 0 & G_h \widehat{A}_h^{-1} G_h' - \widehat{S}_h \end{pmatrix} \quad (27)$$

is positive definite on $X_h \times M_h$ then $\mathcal{P}_h^{-1} \mathcal{A}_h$ is symmetric positive definite with respect to the scalar product defined by \mathcal{N}_h . To that end, assume:

1. $\widehat{A}_h > A_h$ and $G_h \widehat{A}_h^{-1} G_h' > \widehat{S}_h$ to assert that $\mathcal{N}_h > 0$, as well as
2. $\widehat{A}_h \leq M_0 A_h$ on $\ker G_h$, for some real constant $M_0 \geq 1$,
3. $G_h \widehat{A}_h^{-1} G_h' \leq M_1 \widehat{S}_h$ for some real constant $M_1 \geq 1$.

Under those assumptions, the spectrum of $\mathcal{P}_h^{-1} \mathcal{A}_h$ is positive and is contained in an interval determined by the constants M_0 and M_1 [19, Theorem 2.2]. Importantly, $\widehat{A}_h \leq M_0 A_h$ in the second condition is only required to hold on $\ker G_h = \Gamma_h$, where A_h is positive definite by (22), and not on all of X_h .

We now specify our choice of \widehat{A}_h and \widehat{S}_h . For \widehat{A}_h we take the block operator (25) with the first block replaced by

$$2C_\rho \left((\#\mathcal{T})^{-2} \|\partial_t \rho\|_{H^{0,0}}^2 + T^{-2} \|\rho\|_{H^{0,0}}^2 \right), \quad (28)$$

while the second and the third nonzero blocks are multiplied by two. Here, $C_\rho > 0$ is computed to satisfy

$$T^{-1} \sum_{\tau \in \mathcal{T}} |f(\tau)|^2 \leq C_\rho \left((\#\mathcal{T})^{-2} \|f'\|_J^2 + T^{-2} \|f\|_J^2 \right) \quad (29)$$

for all $f \in H^1(J)$. This entails $\widehat{A}_h \geq 2A_h$, so that $\widehat{A}_h > A_h$ is fulfilled in the first assumption. The second assumption follows from the norm equivalence stated in Lemma 1. The constant C_ρ can be chosen independently of T and \mathcal{T} , as long as the subintervals defined by \mathcal{T} are of comparable length.

For \widehat{S}_h we take the C_λ -fold of the Riesz isomorphism on $H^{0,0}$, with, for instance, $0 < C_\lambda < \beta_h^2 \|\widehat{A}_h\|^{-1}$, where $\beta_h > 0$ is the discrete inf-sup constant (23). Then $G_h \widehat{A}_h^{-1} G_h' > \widehat{S}_h$ is fulfilled in the first assumption. The third assumption follows by continuity of $G_h \widehat{A}_h^{-1} G_h'$ on $H^{0,0}$. In the numerical experiments below we set $C_\lambda = \frac{1}{2}$.

It can be seen from the factorization (26) the application of the inverse of the preconditioner only requires the application of \widehat{A}_h^{-1} and \widehat{S}_h^{-1} (as well as G_h and G_h'). The application of \widehat{S}_h^{-1} is trivial since we work with piecewise constant basis functions. In order to apply the block-diagonal operator \widehat{A}_h^{-1} , it is useful to observe that each block of \widehat{A} has the Kronecker product form $T \otimes X$. The application of $T \otimes X$ and $(T \otimes X)^{-1} = T^{-1} \otimes X^{-1}$ is done by means of the matrix identity $(\mathbf{T} \otimes \mathbf{X})^{\pm 1} \text{Vec}(\mathbf{u}) = \text{Vec}(\mathbf{X}^{\pm 1} \mathbf{u} \mathbf{T}^{\pm \text{T}})$, where Vec stacks the columns of a matrix one after another into one long vector. In our Matlab implementation, we use this identity and apply the direct solver to compute $\mathbf{X}^{-1} \mathbf{u}$ column-wise, $\mathbf{u} \mathbf{T}^{-\text{T}}$ row-wise, etc. This technology is an important feature of space-time tensor product discretizations.

5 Numerical examples

5.1 Robustness of preconditioner

We investigate the robustness of the discretization and the preconditioner with respect to the spatial discretization. The image domain is $D := (0, 3) \times (0, 4)$ and the temporal interval is $J := (0, 2)$. A sequence of partitions $D_{2^{-k}}$, $k = 0, 1, 2, \dots, 7$, of D is obtained by subdividing D into $(3 \times 2^k) \times (4 \times 2^k)$ equal rectangles. In each case, the partition J_h of J consists of 100 subintervals of equal length (this choice is arbitrary). The input data (2) consists of two frames ρ_0 and ρ_2 , hence $\mathcal{T} = \{0, 2\}$, as shown in Figure 1. The parameter choices in the penalization functional (9) are $\alpha_j = 10^{-4}$ and $\alpha_g = 1$, which promotes the flux and puts a heavy penalty on the source term. We performed 20 CG iterations with the preconditioner (26) and the energy norm (27) as described in Section 4.2. The residual in the energy norm is reduced by the iteration by a factor $\approx 10^6$, for each choice of the spatial resolution level k . The results are depicted in Figure 1. In Figure 2 (left) we report on the value of the cost functional in the course of the CG iteration, for different spatial resolutions levels k . We observe

that the iteration is robust in k . In Figure 2 (right), the convergence of the time-averaged flux over the horizontal midline (for the final CG iterate) as a function of the spatial resolution is shown. First order convergence in terms of the total number of degrees of freedom is observed.




Fig. 1. *Top:* Input data ρ_τ for $\tau = 0$ (left) and $\tau = 2$ (right). *Middle, left to right:* The divergence of the computed flux, $\nabla \cdot j$, at $t \approx 0$, $t \approx 1$ and $t \approx 2$. For the purpose of visualization, the flux is interpolated from the edges to the midpoint of each geometric element. *Bottom:* The computed source term g at $t = 0, 1, 2$. Note the small scale of 10^{-3} in accordance with the choice of $\alpha_g/\alpha_j = 10^4$.

5.2 Vienna Prater movie

In this example we apply the method to a movie showing the Riesenrad at the Vienna Prater park turning clockwise. A vertical dark strip moving from left to right across the image has been superimposed artificially simulating shading, see Figure 3. The movie consists of 200 frames of 256×216 pixels each. This data is mapped to the domain $D = (0, 4) \times (0, 3)$ and the temporal interval $J = (0, 2)$. The parameter choices in (9) is again $\alpha_j = 10^{-4}$ and $\alpha_g = 1$. We look at frames 91 and 92 and zoom in on two regions highlighted in Figure 3. The computed




Fig. 2. *Left:* The value of the cost functional \mathcal{J} in the course of the conjugate gradient iteration. The curves for different spatial resolutions overlap. *Right:* First order convergence of the time-averaged flux over the horizontal midline as the spatial resolution is increased. See Section 5.1.

discrete flux between these two frames in those two regions is shown in Figure 3. It is interesting to note that the model generates a strong, almost uniform flux j across the superimposed dark strip in the opposite direction of the movement of the strip in order to transport the bright intensity from its bow to its stern. One observes finer flow pattern within the strip away from its boundary caused by the rotation of the Riesenrad, as the graycolor displaying $\nabla \cdot j$ in Figure 3 suggests. On the other hand, the moving bright spokes of the Riesenrad cause the intensities to flow across them in the same direction (as one might expect from the optical flow model).

6 Conclusions

We have considered a version of the in optical flow equations in which the image brightness evolves as a conserved quantity up to possible sources (or sinks). We have formulated the problem of estimating the optical flow and the source as an energy minimization problem. We have investigated a space-time discretization and preconditioning strategy for the resulting saddle-point equations. The discretization was shown to be stable in the Galerkin sense, and the preconditioner to be robust in the discretization resolution. In particular, applying the conjugate gradient method (with a suitable scalar product for which the system matrix is symmetric and positive definite) allows to solve the complete space-time problem within a few iterations. We have found that the transport equation model with the chosen penalization functional produces rather global flows; we do not expect this to happen if the “kinetic energy” $\|j\|^2/\rho$ is penalized instead of the divergence of the flux j and wish to investigate this type of penalization next. We note, however, that our model is typical for control problems with a quadratic cost functional subject to a linear time-dependent PDE constraint, and we therefore expect the results to be applicable elsewhere. It would further be of interest to investigate robustness of the proposed preconditioner in the penalization parameters.

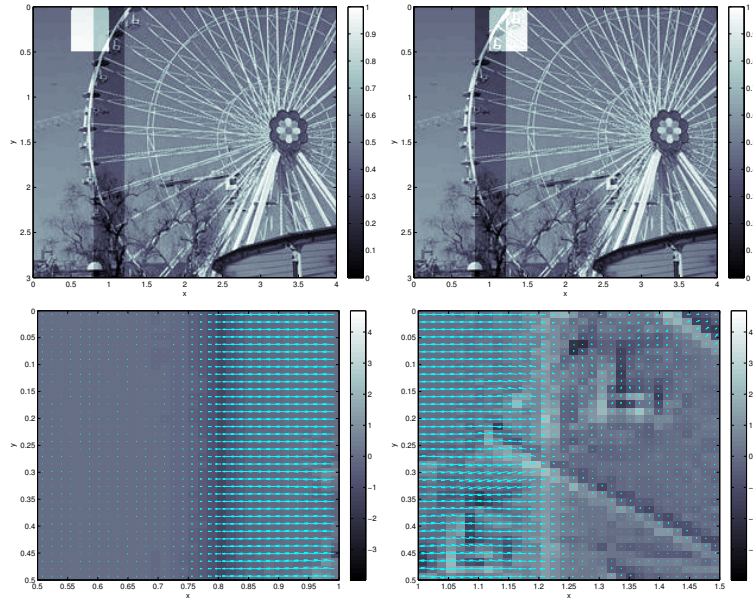


Fig. 3. *Top:* Two consecutive frames ρ_τ for $\tau \approx 0.91$ in the example in Section 5.2. Highlighted in white are the regions of $[33, 64] \times [1, 36]$ pixels (left) and $[65, 96] \times [1, 36]$ pixels (right). The Riesenrad rotates clockwise and the superimposed dark vertical strip moves eastwards (\rightarrow). *Bottom:* Computed discrete flux for the example in Section 5.2. Left: The region of $[33, 64] \times [1, 36]$ pixels. Right: The region of $[65, 96] \times [1, 36]$ pixels. In both cases, there is a strong flux westward (\leftarrow) across the superimposed dark strip, which moves eastwards (\rightarrow). There is a less distinct flux directed north-east (\nearrow) across the bright spokes of the Riesenrad. The background color shows the divergence $\nabla \cdot j$.

Acknowledgements

OS and WZ were supported by the Austrian Science Fund (FWF) within the NFN Geometry and Simulation, projects S11702 and S11704. OS was also supported by the project *Modeling Visual Attention as a Key Factor in Visual Recognition and Quality of Experience* funded by the Wiener Wissenschafts und Technologie Funds – WWTF. The Vienna Prater movie was provided by Ch. Valuch, Cognitive Science Research Platform, Cognitive Psychology, University of Vienna, and preprocessed by A. Patrone, University of Vienna.

References

1. Gilles Aubert, Rachid Deriche, and Pierre Kornprobst. Computing optical flow via variational techniques. *SIAM J. Appl. Math.*, 60(1):156–182, 2000.

2. Randolph E. Bank, Bruno D. Welfert, and Harry Yserentant. A class of iterative methods for solving saddle point problems. *Numer. Math.*, 56(7):645–666, 1990.
3. Steven S. Beauchemin and John L. Barron. The computation of optical flow. *ACM Comput. Surv.*, 27(3):433–466, 1995.
4. Michele Benzi, Gene H. Golub, and Jörg Liesen. Numerical solution of saddle point problems. *Acta Numer.*, 14:1–137, 2005.
5. Michele Benzi, Eldad Haber, and Lauren Taralli. A preconditioning technique for a class of PDE-constrained optimization problems. *Adv. Comput. Math.*, 35(2-4):149–173, 2011.
6. Michael J. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Comput. Vis. Image Und.*, 63:75–104, 1996.
7. Alfio Borzi, Kazufumi Ito, and Karl Kunisch. Optimal control formulation for determining optical flow. *SIAM J. Sci. Comput.*, 24(3):818–847, 2002.
8. Dietrich Braess. *Finite Elemente*. Springer, 4th edition, 2007.
9. James H. Bramble and Joseph E. Pasciak. Corrigenda: “A preconditioning technique for indefinite systems resulting from mixed approximations of elliptic problems”. *Math. Comp.*, 51(183):387–388, 1988.
10. James H. Bramble and Joseph E. Pasciak. A preconditioning technique for indefinite systems resulting from mixed approximations of elliptic problems. *Math. Comp.*, 50(181):1–17, 1988.
11. Philippe G. Ciarlet. *The finite element method for elliptic problems*. North-Holland, 1978.
12. Bernard Dacorogna. *Direct methods in the calculus of variations*. Springer, 1989.
13. Alberto Del Bimbo, Paolo Nesi, and Jorge L. C. Sanz. Analysis of optical flow constraints. *IEEE T. Image Process.*, 4(4):460–469, 1995.
14. David J. Fleet and Yair Weiss. Optical flow estimation. In Nikos Paragios, Yunmei Chen, and Olivier D. Faugeras, editors, *Handbook of Mathematical Models in Computer Vision*, chapter 15. Springer, 2006.
15. Horst W. Haussecker and David J. Fleet. Computing optical flow with physical models of brightness variation. *IEEE T. Pattern Anal.*, 23(6):661–673, 2001.
16. Berthold K. P. Horn and Brian G. Schunck. Determining optical flow. *Artif. Intell.*, 17:185–203, 1981.
17. Bruce D. Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th international joint conference on Artificial intelligence*, volume 2, pages 674–679. Morgan Kaufmann Publishers Inc. San Francisco, CA, USA, 1981.
18. Nans-Helmut Nagel and Wilfried Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE T. Pattern Anal.*, 8:565–593, 1986.
19. Joachim Schöberl and Walter Zulehner. Symmetric indefinite preconditioners for saddle point problems with applications to PDE-constrained optimization problems. *SIAM J. Matrix Anal. Appl.*, 29(3):752–773 (electronic), 2007.
20. Joachim Weickert and Christoph Schnörr. Variational optic flow computation with a spatio-temporal smoothness constraint. *J. Math. Imaging Vis.*, 14(3):245–255, 2001.
21. Richard P. Wildes, Michael J. Amabile, Ann-Marie Lanzillotto, and Tzong-Shyng Leu. Recovering estimates of fluid flow from image sequence data. *Comput. Vis. Image Und.*, 80(2):246–266, 2000.
22. Kôsaku Yosida. *Functional analysis*. Springer, 1980.