Optic Flow Scale Space

Oliver Demetz[1], Joachim Weickert[1], Andrés Bruhn[2], and Henning Zimmer[1]

[1] Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science, Building E1.1
Saarland University, 66041 Saarbrücken, Germany
{demetz,weickert,zimmer}@mia.uni-saarland.de

[2] Vision and Image Processing Group
Cluster of Excellence Multimodal Computing and Interaction,
Saarland University, 66041 Saarbrücken, Germany
bruhn@mmci.uni-saarland.de

Abstract. While image scale spaces are well understood, it is undeniable that the regularisation parameter in variational optic flow methods serves a similar role as the scale parameter in scale space evolutions. However, no thorough analysis of this optic flow scale-space exists to date. Our paper closes this gap by interpreting variational optic flow methods as Whittaker-Tikhonov regularisations of the normal flow, evaluated in a constraint-specific norm. The transition from this regularisation framework to an optic flow evolution creates novel vector-valued scale-spaces that are not in divergence form and act in a highly anisotropic way. From a practical viewpoint, the deep structure in optic flow scale space allows the automatic selection of the most accurate scale by means of an optimal prediction principle. Moreover, we show that our general class of optic flow scale-spaces incorporates novel methods that outperform classical variational approaches.

1 Introduction

Starting with Iijima's pioneering work on Gaussian scale-space and its use in optical character recognition many decades ago [12,13], scale-spaces have become versatile tools for analysing and understanding the multiscale structure of images; see e.g. the monographs [8,15,21,23] and the references therein. While partial differential equations (PDEs) of evolution type provide a natural framework for most scale-space concepts [2], it has also been shown that variational regularisation methods create scale-spaces where the regularisation parameter acts as scale [18]. Such variational methods, however, offer much broader application fields than classical data smoothing. In computer vision, for example, they are widely used for solving correspondence problems. For specific applications such as optic flow computation in image sequences, variational methods have become highly sophisticated and constitute the most accurate methods to date [3]. In contrast to classical image regularisation methods, however, these optic flow methods do not regularise image data, but constraint equations. For instance, in the classical method of Horn and Schunck, a grey value constancy assumption replaces the role of a data fidelity term [11]. Therefore, it is unclear in which sense one may interpret variational optic flow methods as scale-spaces. However, in view of the fact

that the unreliable normal flow is essentially the only information that can be extracted directly from the data in an image sequence, it is astonishing that modern variational optic flow methods are capable of achieving results of such a high quality. Thus, there is a clear need to understand their scale-space behaviour.

The goal of our paper is to address this problem. Our contributions are fourfold:

1. We interpret the classical variational methods of Horn and Schunck [11] and of Nagel and Enkelmann [17] as Whittaker-Tikhonov regularisations of the normal flow. We show that this requires to replace the Euclidean norm by a space-variant matrix-induced norm that respects the data constraints.
2. We generalise this framework to a broader class of methods that also allows to come up with new models that have not been considered before. We show that they can offer better performance than the classical variational methods.
3. Going from the regularisation framework to a scale-space representation, we come up with the novel concept of *optic flow scale-spaces*. They are parabolic evolutions of vector-valued data with the regularisation parameter as scale and the normal flow as initial state. However, in contrast to many image scale-spaces they are not of divergence type and hence do not preserve the average value of the initial data. Moreover, they turn out to be highly anisotropic. Both properties are essential for the remarkable performance of optic flow methods.
4. Having a scale-space evolution, it is natural to explore its deep structure. To this end, we show that the optic flow scale-space provides an efficient framework for automatically selecting the best scale that gives the most accurate optic flow field. As a parameter-free scale selection principle we employ the *Optimal Prediction Principle*, which is specifically tailored to the needs of optic flow estimation [25].

This paper is not intended to present a new high accuracy optic flow method. Rather, we want to investigate the scale space behaviour of variational optic flow methods. Thus, we intentionally do not consider algorithmic sophistications such as robust penalisation strategies, constancy assumptions without linearisation, multiscale warping strategies, and nonlocal search methods. Since the focus of our work lies on the fundamental concepts behind optic flow scale space, the discussion of additional side effects would obliterate the core message of the paper.

Related work. The transformation we apply is related to a proposal by Schnörr [19], who did not pursue this concept further. With respect to the interpretation of variational methods in terms of specific norms, vector spaces, and higher order manifolds, there is a huge amount of literature available; see e.g. Sochen *et al.* [20] and the references therein. Particularly interesting in this context is the work of Ben-Ari and Sochen [5] who derive a class of smoothness terms based on spatially varying norms induced by suitable embeddings of the flow field into higher dimensional vector spaces. Scale selection is a classical issue in Gaussian scale-space theory [15]. More specifically, choosing optimal smoothness parameters is an enduring problem for almost all classes of scale-space and variational methods. In our context, the works by Krajsek *et al.* [14], Mrázek and Navarra [16] and the recent ideas by Zimmer *et al.* [25] are most relevant. While there has been some research on scale spaces for image sequences [7,10], to our knowledge the concept of optic flow scale space has not been considered before.

Organisation of our paper. In Section 2, we describe the transformation of classical optic flow energies into a regularisation framework for the normal flow in a suitable constraint-induced norm that is further generalised in Section 3. The fourth section introduces our scale space framework for optic flow. An efficient numerical scheme is described in Section 5, while the automatic scale selection is discussed in Section 6. Section 7 presents computational results, and we conclude with a summary and an outlook in Section 8.

2 Variational Optic Flow as Whittaker-Tikhonov Regularisation

As starting point of our derivations of an optic flow scale space we consider the classic method of Horn and Schunck [11]. This variational method estimates the optic flow field $\boldsymbol{w} := (u, v)^\top = (u(x, y, z), v(x, y, z))^\top$ as the minimiser of the energy functional

$$E(\boldsymbol{w}) = \int_\Omega \Big((f_x u + f_y v + f_z)^2 + \alpha \left(\|\boldsymbol{\nabla} u\|^2 + \|\boldsymbol{\nabla} v\|^2 \right) \Big) \, \mathrm{d}x \, \mathrm{d}y \, , \tag{1}$$

where $\Omega \subset \mathbb{R}^2$ represents the spatial image domain, $f : \Omega \times [0, \infty) \to \mathbb{R}$ the image sequence, $\|\cdot\|$ stands for the Euclidean norm, subscripts denote partial derivatives, and $\boldsymbol{\nabla} = (\partial_x, \partial_y)^\top$ is the spatial gradient operator.

The first term of this functional is called data term and models the linearised assumption that corresponding pixels in subsequent frames have similar grey value. Since it depends on two unknown functions u and v that describe the horizontal and vertical displacement field, respectively, its solution is under-determined. Evidently, in order to find a unique solution, additional assumptions on u and v are needed. This is realised by the second term – the so-called smoothness term. It penalises variations of the solution and is weighted by the positive regularisation parameter α.

2.1 Towards Regularisation in a Spatially Varying Norm

In the following, we consider a slightly modified version of the energy (1) with the additional terms $\epsilon^2(u^2 + v^2) + c$, where ϵ is a small positive constant and $c(x, y, z) = -\epsilon f_z^2 / (|\boldsymbol{\nabla} f|^2 + \epsilon^2)$ is a function which does not depend on the unknown and thus does not play a role in the actual minimisation. Later on, these terms will be useful for theoretical reasons. The modified energy then reads as

$$E(\boldsymbol{w}) = \int_\Omega \Big((f_x u + f_y v + f_z)^2 + \epsilon^2(u^2 + v^2) + c + \alpha \left(\|\boldsymbol{\nabla} u\|^2 + \|\boldsymbol{\nabla} v\|^2 \right) \Big) \, \mathrm{d}x \, \mathrm{d}y \, . \tag{2}$$

In order to obtain a different much more intuitive understanding of the underlying variational model, we seek to reformulate the latter energy in an image regularisation framework. Since a smoothness term is already present, the main task is now to derive a suitable similarity term. To this end, we make use of the following result (see also [1]): Let $\boldsymbol{w}_n$ be the regularised normal flow, given by

$$\boldsymbol{w}_n = \frac{-f_z \boldsymbol{\nabla} f}{|\boldsymbol{\nabla} f|^2 + \epsilon^2} \, , \tag{3}$$

and let $\boldsymbol{A} : \Omega \to \mathbb{R}^{2\times 2}$ be a symmetric positive definite matrix in every point of the image domain defined as

$$\boldsymbol{A}^2 = \boldsymbol{\nabla} f \boldsymbol{\nabla} f^\top + \epsilon^2 \boldsymbol{I} \,, \tag{4}$$

where $\boldsymbol{I}$ denotes the unit matrix. Then the following equivalence holds:

$$(f_x u + f_y v + f_z)^2 + \epsilon^2 (u^2 + v^2) + c = (\boldsymbol{w} - \boldsymbol{w}_n)^\top \boldsymbol{A}^2 (\boldsymbol{w} - \boldsymbol{w}_n) \,. \tag{5}$$

This can easily be verified by straightforward calculations. Furthermore, we use the latter quadratic form and the concept of matrix-weighted norms to rewrite the modified functional (2) into an image regularisation-like energy [6]

$$E(\boldsymbol{w}) = \int_\Omega \Big(\|\boldsymbol{w} - \boldsymbol{w}_n\|^2_{\boldsymbol{A}^2} + \alpha \left(\|\boldsymbol{\nabla} u\|^2 + \|\boldsymbol{\nabla} v\|^2 \right) \Big) \, \mathrm{d}x \, \mathrm{d}y \,. \tag{6}$$

In this context, for a symmetric positive definite matrix $\boldsymbol{M}$ the corresponding matrix-weighted norm is given by $\|\boldsymbol{x}\|^2_{\boldsymbol{M}} := \langle \boldsymbol{x}, \boldsymbol{x} \rangle_{\boldsymbol{M}} = \boldsymbol{x}^\top \boldsymbol{M} \boldsymbol{x}$. Note that due to the additional term in (2), the matrix $\boldsymbol{A}$ fulfils these requirements by construction.

Having performed the previous rewritings, the following insight becomes explicit: Essentially, the seminal variational optic flow method of Horn and Schunck can be interpreted as Whittaker-Tikhonov regularisation of the normal flow in a matrix-weighted spatially varying norm.

2.2 Analysing the Matrix-Weighted Norm

By analysing the obvious eigenstructure of the constraint matrix $\boldsymbol{A}^2$, we can gain a deeper understanding of the introduced matrix-weighted norm and thus of the data term. Using the eigendecomposition of $\boldsymbol{A}^2$ given by

$$\boldsymbol{A}^2 = \left(|\boldsymbol{\nabla} f|^2 + \epsilon^2 \right) \frac{\boldsymbol{\nabla} f}{|\boldsymbol{\nabla} f|} \frac{\boldsymbol{\nabla} f^\top}{|\boldsymbol{\nabla} f|} \;+\; \epsilon^2 \frac{\boldsymbol{\nabla} f^\perp}{|\boldsymbol{\nabla} f|} \frac{\boldsymbol{\nabla} f^{\perp\top}}{|\boldsymbol{\nabla} f|} \;, \tag{7}$$

we can express the data term as

$$\|\boldsymbol{w} - \boldsymbol{w}_n\|^2_{\boldsymbol{A}^2} = \left(|\boldsymbol{\nabla} f|^2 + \epsilon^2 \right) \left\langle \boldsymbol{w} - \boldsymbol{w}_n, \frac{\boldsymbol{\nabla} f}{|\boldsymbol{\nabla} f|} \right\rangle^2 + \epsilon^2 \left\langle \boldsymbol{w} - \boldsymbol{w}_n, \frac{\boldsymbol{\nabla} f^\perp}{|\boldsymbol{\nabla} f|} \right\rangle^2 \,. \tag{8}$$

This shows that the central quantity being under consideration is the normal flow $\boldsymbol{w}_n$, or rather the difference between the actual solution and the normal flow $\boldsymbol{w} - \boldsymbol{w}_n$. This difference vector is then projected into the local eigensystem of $\boldsymbol{A}^2$. There, its component perpendicular to the image gradient is basically negligible (since ϵ^2 is small), while its component along the image gradient is the part that actually contributes.

This also confirms the classic explanation of the linearised grey value constancy assumption as a constraint line: The expression $f_x u + f_y v + f_t = 0$ defines a line perpendicular to the image gradient with distance $f_t / |\boldsymbol{\nabla} f|$ from the origin [11]. Also the findings in [25] are in accordance with this interpretation: There, the authors rewrite the assumption into a projection of the difference vector $\boldsymbol{w} - \boldsymbol{w}_n$ onto the image gradient, which exactly comes down to our locally adapted norm for $\epsilon = 0$.

3 Generalisation of the Matrix-Weighted Norm

Up to now, we have mainly reformulated the method of Horn and Schunck and identified the matrix-induced norm $\| \cdot \|_{\boldsymbol{A}^2}$ to play a central role. Consequently, we now propose to generalise this idea to a class of norms with varying anisotropy. To this end, we modify the exponent of the constraint matrix inducing the norm. This yields a data term of the general form

$$\|\boldsymbol{w} - \boldsymbol{w}_n\|^2_{\boldsymbol{A}^{2-\beta}} \tag{9}$$

with $0 \leq \beta \leq 2$. Larger choices for β do not make sense, since they would invert the anisotropy, i.e. swap the penalisation directions. The parametrisation of the norm has been chosen in such a way that for $\beta = 0$ the original model of Horn and Schunck, and for $\beta = 2$ a pure decoupled vector-valued regularisation of the normal flow in the Euclidean norm is obtained (since the constraint matrix collapses to the identity matrix).

To establish a consistently extended model, we also equip the smoothness term with the same spatially varying norm. This leads to the regulariser

$$\|\boldsymbol{\nabla} u\|^2_{\boldsymbol{A}^{-\gamma}} + \|\boldsymbol{\nabla} v\|^2_{\boldsymbol{A}^{-\gamma}} \tag{10}$$

where $\gamma \geq 0$. This generally anisotropic image-driven regulariser allows variations of the flow field across image edges but not along them. In the special case of $\gamma = \beta = 0$ the model corresponds to Whittaker-Tikhonov regularisation [24,22] as used by Horn and Schunck [11]. Another special case of our generalised model is obtained for $\gamma = 2$ and $\beta = 0$: Then, our method resembles the method of Nagel and Enkelmann [17]. Recall that the proposed spatially varying norm naturally arises from the linearised constancy assumption in the data term. In this way our approach differs significantly from the work in [5], which derives such a norm by embedding the flow in a higher dimensional vector space and thus disregards the data term throughout the derivation. Incorporating both generalisations, we finally consider the energy functional

$$E(\boldsymbol{w}) = \int_\Omega \Big(\| \boldsymbol{w} - \boldsymbol{w}_n \|^2_{\boldsymbol{A}^{2-\beta}} + \alpha \left(\|\boldsymbol{\nabla} u\|^2_{\boldsymbol{A}^{-\gamma}} + \|\boldsymbol{\nabla} v\|^2_{\boldsymbol{A}^{-\gamma}} \right) \Big) \, \mathrm{d}x \, \mathrm{d}y \,, \tag{11}$$

with α, β and γ as defined before. This energy functional forms the basis for our optic flow scale space introduced in the next section.

4 Optic Flow Scale Space

From the calculus of variations it follows that any minimiser of the functional (11) has to fulfil the associated Euler-Lagrange equation, which reads

$$\frac{\boldsymbol{w} - \boldsymbol{w}_n}{\alpha} = \boldsymbol{A}^{\beta-2} \begin{pmatrix} \operatorname{div} \left(\boldsymbol{A}^{-\gamma} \, \boldsymbol{\nabla} u \right) \\ \operatorname{div} \left(\boldsymbol{A}^{-\gamma} \, \boldsymbol{\nabla} v \right) \end{pmatrix} \,, \tag{12}$$

with reflecting Neumann boundary conditions $\boldsymbol{n}^\top \boldsymbol{A}^{-\gamma} \boldsymbol{\nabla} u = 0$ and $\boldsymbol{n}^\top \boldsymbol{A}^{-\gamma} \boldsymbol{\nabla} v = 0$. The reader should keep in mind that the argument of the energy (11) as well as the latter

PDE is vector-valued, since $\boldsymbol{w} = (u, v)^\top$. Equation (12) can be seen as a fully implicit time discretisation of the filter

$$\boxed{\partial_t \boldsymbol{w} = \boldsymbol{A}^{\beta-2} \begin{pmatrix} \operatorname{div}(\boldsymbol{A}^{-\gamma}\, \boldsymbol{\nabla} u) \\ \operatorname{div}(\boldsymbol{A}^{-\gamma}\, \boldsymbol{\nabla} v) \end{pmatrix}, \qquad \boldsymbol{w}(\,\cdot\,, 0) = \boldsymbol{w}_n\,.} \tag{13}$$

with a single time step of size α, and the normal flow $\boldsymbol{w}_n$ as initial state at time $t = 0$.

Obviously this temporal evolution constitutes a scale space, whose evolution time t coincides with the regularisation parameter α of the associated energy. Interestingly, the initial state of our *optic flow scale space* is the regularised normal flow, which is the only component of the flow field that can be directly extracted from the image data.

Note that we have transformed the *regularisation-like* energy functional (11) into a *diffusion-like* coupled system of parabolic PDEs (13). In the context of image filtering, relations between such methods have been investigated by Scherzer and Weickert [18].

5 Numerical Realisation

For solving the parabolic problem in (13) we use an explicit scheme. To this end we discretise the two flow components u and v by sampling them on a regular grid and stacking all rows in single vectors $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^N$, where N denotes the number of pixels. Using this single-index notation, we discretise the matrix $\boldsymbol{A}^{\beta-2}$ in pixel i by

$$\boldsymbol{A}_i^{\beta-2} = \begin{pmatrix} a_i & b_i \\ b_i & c_i \end{pmatrix}, \quad i = 1, \ldots, N\,. \tag{14}$$

Accordingly, we discretise the diffusive terms using finite differences and obtain a nona-diagonal diffusion matrix $\boldsymbol{D} \in \mathbb{R}^{N \times N}$ (similar to [23]). This leads to the following explicit scheme:

$$\begin{pmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{pmatrix}^{k+1} = \left(\boldsymbol{I} + \tau \left(\begin{array}{ccc|ccc} a_1 & & & b_1 & & \\ & \ddots & & & \ddots & \\ & & a_N & & & b_N \\ \hline b_1 & & & c_1 & & \\ & \ddots & & & \ddots & \\ & & b_N & & & c_N \end{array} \right) \left(\begin{array}{ccc|ccc} & & & & & \\ & \boldsymbol{D} & & & \boldsymbol{0} & \\ & & & & & \\ \hline & & & & & \\ & \boldsymbol{0} & & & \boldsymbol{D} & \\ & & & & & \end{array} \right) \right) \begin{pmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{pmatrix}^k, \tag{15}$$

where every iteration advances the evolution by the time step size $\tau > 0$. Thus, after k iterations $(\boldsymbol{u}^k, \boldsymbol{v}^k)^\top$ contains the flow field at scale $\alpha = k \cdot \tau$. As a consequence, this explicit scheme inherently samples the whole scale space up to the stopping time α in intervals of size τ. Thereby, the whole iteration matrix remains constant for all iterations, since all terms are exclusively image-driven.

In order to accelerate this explicit scheme, we made use of the recently introduced *fast explicit diffusion (FED)* strategy [9], which performs cycles of explicit iterations with varying time step sizes. In particular, up to 50% of the step sizes can exceed the

stability limit significantly, while the overall process remains provably stable. By that, the order of the smoothing time reached in n steps can be increased from $O(n)$ to $O(n^2)$. In our case, this is very beneficial, since the maximal stable time step size of the explicit scheme can decrease drastically with increasing anisotropy or small choices of the parameter β.

6 Scale Selection

In the previous sections, we have set up a novel class of optic flow scale spaces which all evolve in the regularisation parameter α. Evidently, in our optic flow setting there exists one scale within each scale space that provides the flow with the highest accuracy. Since the optic flow on all scales is available by construction, we have access to the deep structure of this scale space and can exploit this information to perform an automatic scale selection. To this end, we adapt the *Optimal Prediction Principle* (OPP) recently developed by Zimmer *et al.* [25]. In short, this principle suggests to rate the quality of an optic flow field between the first and second frame according to its extrapolation quality from the first to the third frame. The underlying assumption is that the velocity of objects (or the camera) remains constant over time. It is shown in [25] that this simple assumption works very well for the automatic estimation of the smoothness weight α.

In our case, for a given flow field $\boldsymbol{w} = (u, v)^\top$ between the frames at time z and $z{+}1$, we assess the extrapolation quality by evaluating the *Average Data Constancy Error* (*ADCE*), which is based on the grey value constancy assumption without linearisation:

$$ADCE_{1,3}(\boldsymbol{w}) = \frac{1}{|\Omega|} \int_\Omega \big(f(x+2u, y+2v, z+2) - f(x,y,z) \big)^2 \mathrm{d}x \, \mathrm{d}y \; . \qquad (16)$$

It is obvious that if the model assumptions hold, a *good* flow field will lead to *small* values of this error measure. Note that in contrast to the optimisation strategy in [25], we can exploit the following advantageous property of our numerical scheme: It explicitly evolves in the parameter α, hence after each iteration the flow field at cumulated time α is available, and the *ADCE* can be evaluated. This on-the-fly computation of the quality estimate is not possible for most other optic flow methods, because they typically require to solve a new system of equations for each value of α.

Besides the OPP, we also tried other schemes for automatic scale estimation. In particular, we investigated the performance of the decorrelation method by Mrázek and Navarra [16]. However, experiments indicated that the underlying assumptions do not hold for our optic flow scale space.

7 Experiments

In order to investigate the behaviour of our optic flow scale space we perform experiments on several image sequences that are publicly available and for which the ground truth flow field is known. In particular, we use the *Yosemite* sequence with and without clouds [4], the *New Marble*[1] sequence as well as the *Rubberwhale* sequence [3].

[1] available from `http://i21www.ira.uka.de/image_sequences`

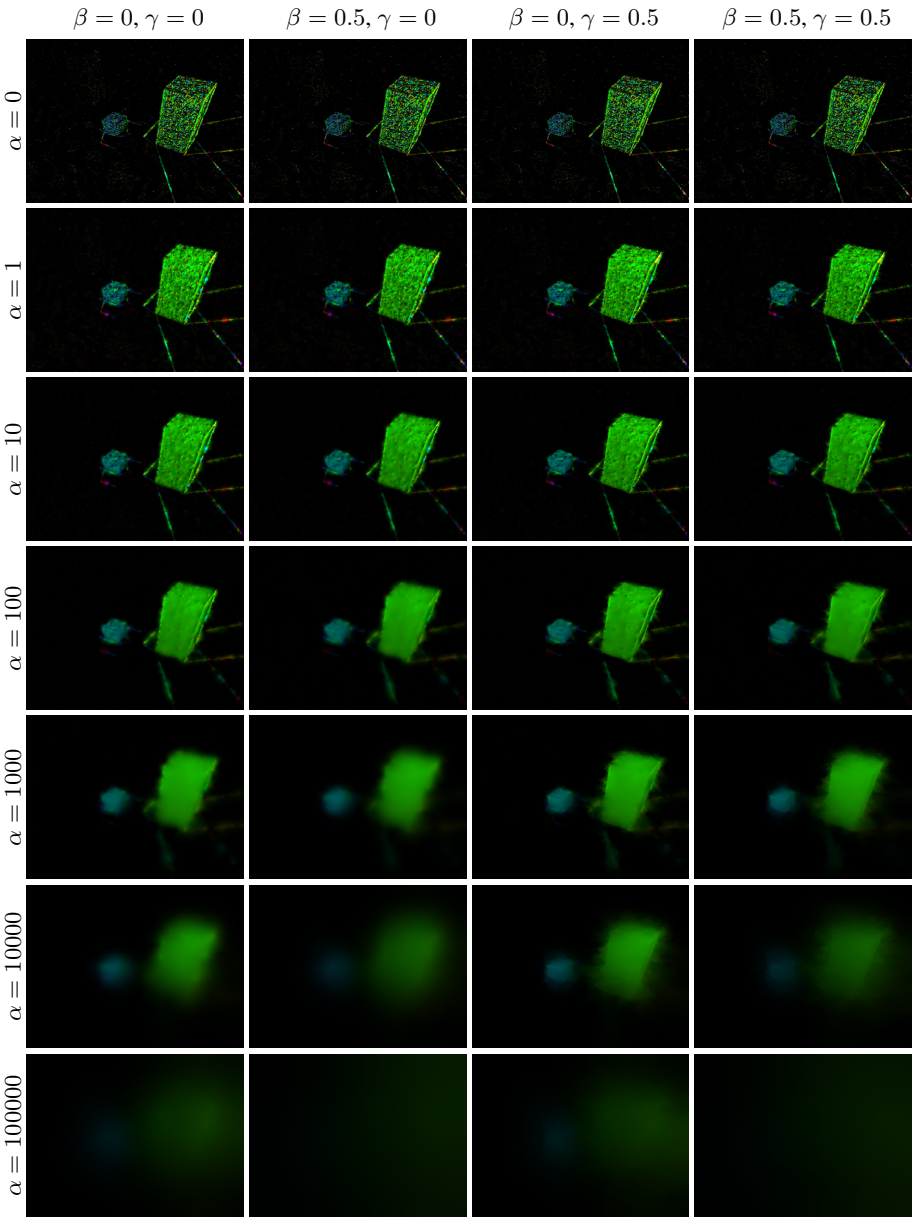

Fig. 1. Scale space at different stopping times for varying β and γ.

In our first experiment we compute samples of the scale space for the New Marble sequence at different evolution times and for several choices of β and γ. Figure 1 shows the corresponding flow fields, where colour encodes the direction and brightness indicates the magnitude of the displacements. Here, one can clearly see the scale space behaviour of the proposed diffusion-like optic flow process: Independently of β and γ, the initial state of all these scale spaces ($\alpha = 0$) is given by the noisy normal flow, while for larger values of α the flow fields become successively smoother. In this context, we make two observations: On the one hand, for $\gamma > 0$ discontinuities are preserved for a longer time, since the regulariser is then of image-driven anisotropic nature. On the other hand, results for $\beta > 0$ are slightly less noisy, since the larger eigenvalue of $\boldsymbol{A}^{2-\beta}$ – the one that depends on the magnitude of the image gradient – is now subject to a smaller exponent, cf. Equation (8). This becomes particularly visible in the magnifications shown in Figure 2.

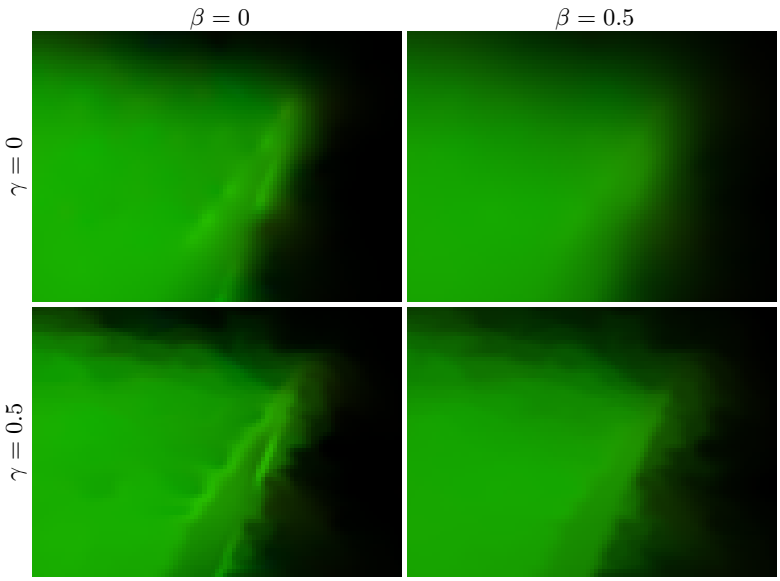

Fig. 2. Zoom into the optic flow fields at scale $\alpha = 1000$ from Figure 1.

In a second experiment, we compare the accuracy of the proposed scheme against the two special cases in our framework: Horn and Schunck ($\beta = \gamma = 0$) and Nagel and Enkelmann ($\beta = 0$, $\gamma = 2$). This is done for the aforementioned image sequences by means of the *Average Angular Error* (*AAE*) [4]. Please note that we keep the presmoothing scale fixed at $\sigma = 1$ throughout all experiments, since its impact is not in the focus of our contribution. Table 1 demonstrates that our method consistently leads to improved results. In particular, it shows that the additional degrees of freedom β and γ provided by our general class of scale spaces can be beneficial.

Table 1. Quantitative error measurements in terms of the *AAE* on different image sequences. For our method, β and γ have been optimised. The actual choices are given in brackets.

Image sequence	Horn / Schunck	Nagel / Enkelmann	Our method
New Marble	2.65	2.77	2.53 ($\beta=1.0$, $\gamma=0.4$)
Rubberwhale	10.58	9.27	9.04 ($\beta=0.5$, $\gamma=1.0$)
Yosemite	7.51	6.86	6.41 ($\beta=0.4$, $\gamma=0.9$)
Yosemite no clouds	2.82	3.63	2.76 ($\beta=0.2$, $\gamma=0.1$)

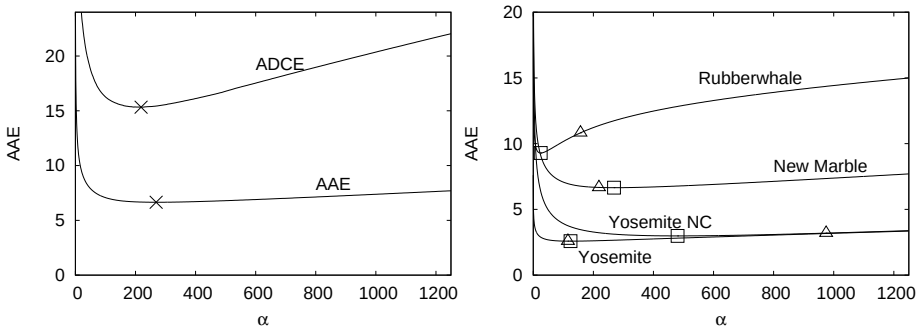

Fig. 3. Automatic scale selection using the Optimal Prediction Principle. **(a) Left:** Graphs of the *ADCE* and *AAE* side-by side. Crosses denote the minimal value of the graph. **(b) Right:** Estimation results of the selected scale α for different image sequences. Triangles indicate the estimated value and squares denote the optimal choice.

In our third experiment, we investigate the automatic selection of the scale parameter α of our model using the OPP. To this end, we first juxtapose the graph of the estimated quality in terms of the *ADCE* with the graph of the measured accuracy given by the *AAE* in Figure 3 (a). This is done for the Yosemite sequence with clouds with $\beta=\gamma=0.5$. One can see that both graphs have a similar and well aligned shape. In particular, the minima of both curves are attained at almost the same position. Secondly, we compare the estimated values for the regularisation parameter α against those that are optimal with respect to the *AAE*. This is done for all four image sequences with $\beta=\gamma=0.5$ fixed. As one can see from Figure 3 (b), the OPP works very well in practice: In all cases, the *AAE* at the estimated scale is close to the one of the optimal scale.

In our final experiment we analyse how the two generalisation parameters β and γ influence the accuracy of the estimation. To this end, we have computed the *AAE* for $\beta, \gamma \in [0, 2]$ using the Yosemite sequence with clouds. As in the previous experiments, the selection of the stopping time α within each scale space has been performed automatically using the OPP. The resulting graphs in Figure 4 (a) and (b) show that for both parameters values larger than zero consistently improve the accuracy.

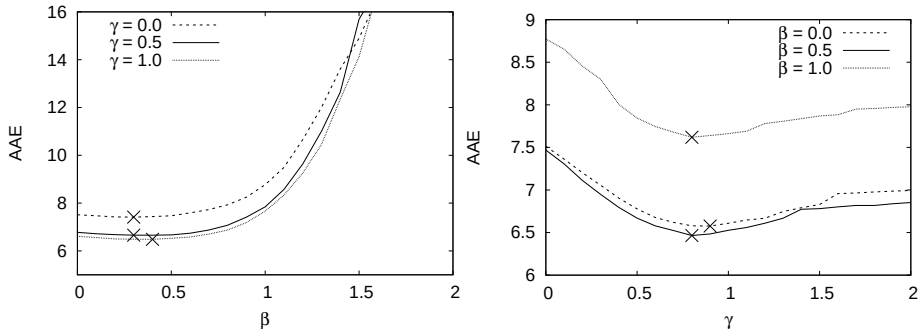

Fig. 4. Influence of the parameters β and γ on the accuracy. **(a) Left:** Behaviour under variations of β for fixed γ. **(b) Right:** Ditto for varying γ and fixed β. In both cases, crosses indicate the minimum of the graphs.

8 Conclusion

It is surprising that in spite of many years of scale space research the scale space character of variational optic flow methods has not been analysed so far. We have shown that such an analysis can be a very worthwhile endeavour: It provides interesting insights in the regularisation behaviour of variational optic flow methods by interpreting them as classical regularisation in a problem-specific norm. Linking this to novel parabolic scale space evolutions that are highly anisotropic and not of diffusion type shows that modern scale space research has only discovered a glimpse of the entire fascinating world of scale space concepts in image processing and computer vision. It would be nice, if our paper serves as a starting point for further research in this direction.

Acknowledgements

Our work has partially been funded by the Cluster of Excellence *Multimodal Computing and Interaction* within the Excellence Initiative of the German Federal Government.

References

1. J. Abhau, Z. Belhachmi, and O. Scherzer. On a decomposition model for optical flow. In D. Cremers, Y. Boykov, A. Blake, and F. R. Schmidt, editors, *Energy Minimization Methods in Computer Vision and Pattern Recognition*, volume 5681 of *Lecture Notes in Computer Science*, pages 126–139. Springer, Berlin, 2009.
2. L. Alvarez, F. Guichard, P.-L. Lions, and J.-M. Morel. Axioms and fundamental equations in image processing. *Archive for Rational Mechanics and Analysis*, 123:199–257, 1993.
3. S. Baker, D. Scharstein, J. P. Lewis, S. Roth, M. J. Black, and R. Szeliski. A database and evaluation methodology for optical flow. Technical Report MSR-TR-2009-179, Microsoft Research, Redmond, WA, December 2009.
4. J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, February 1994.

5. R. Ben-Ari and N. Sochen. A geometric framework and a new criterion in optical flow modeling. *Journal of Mathematical Imaging and Vision*, 33:178–194, February 2009.
6. M. Bertero, T. A. Poggio, and V. Torre. Ill-posed problems in early vision. *Proceedings of the IEEE*, 76(8):869–889, August 1988.
7. D. Fagerström. Spatio-temporal scale-spaces. In F. Sgallari, F. Murli, and N. Paragios, editors, *Scale Space and Variational Methods in Computer Vision*, volume 4485 of *Lecture Notes in Computer Science*, pages 326–337. Springer, Berlin, 2007.
8. L. Florack. *Image Structure*, volume 10 of *Computational Imaging and Vision*. Kluwer, Dordrecht, 1997.
9. S. Grewenig, J. Weickert, and A. Bruhn. From box filtering to fast explicit diffusion. In M. Goesele, S. Roth, A. Kuijper, B. Schiele, and K. Schindler, editors, *Pattern Recognition, Proceedings of the 32nd DAGM*, volume 6376 of *Lecture Notes in Computer Science*, pages 543–552, Berlin, 2010. Springer.
10. F. Guichard. A morphological, affine, and Galilean invariant scale-space for movies. *IEEE Transactions on Image Processing*, 7(3):444–456, March 1998.
11. B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
12. T. Iijima. Theory of pattern recognition. *Electronics and Communications in Japan*, pages 123–134, November 1963. In English.
13. T. Iijima, H. Genchi, and K. Mori. A theory of character recognition by pattern matching method. In *Proc. First International Joint Conference on Pattern Recognition*, pages 50–56, Washington, DC, October 1973. In English.
14. K. Krajsek and R. Mester. A maximum likelihood estimator for choosing the regularization parameters in global optical flow methods. In *Proc. 2006 IEEE International Conference on Image Processing*, pages 1081–1084, Atlanta, GA, 2006.
15. T. Lindeberg. *Scale-Space Theory in Computer Vision*. Kluwer, Boston, 1994.
16. P. Mrázek and M. Navara. Selection of optimal stopping time for nonlinear diffusion filtering. *International Journal of Computer Vision*, 52(2-3):189–203, 2003.
17. H.-H. Nagel and W. Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:565–593, 1986.
18. O. Scherzer and J. Weickert. Relations between regularization and diffusion filtering. *Journal of Mathematical Imaging and Vision*, 12(1):43–63, February 2000.
19. C. Schnörr. On functionals with greyvalue-controlled smoothness terms for determining optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(10):1074–1079, 1993.
20. N. Sochen, R. Kimmel, and F. Bruckstein. Diffusions and confusions in signal and image processing. *Journal of Mathematical Imaging and Vision*, 14(3):195–210, May 2001.
21. J. Sporring, M. Nielsen, L. Florack, and P. Johansen, editors. *Gaussian Scale-Space Theory*, volume 8 of *Computational Imaging and Vision*. Kluwer, Dordrecht, 1997.
22. A. N. Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Soviet Mathematics Doklady*, 4:1035–1038, 1963.
23. J. Weickert. *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart, 1998.
24. E. T. Whittaker. A new method of graduation. *Proceedings of the Edinburgh Mathematical Society*, 41:65–75, 1923.
25. H. Zimmer, A. Bruhn, and J. Weickert. Optic flow in harmony. *International Journal of Computer Vision*, January 2011. Online First.