

Spatially-Variant Kernel for Optical Flow Under Low Signal-to-Noise Ratios: Application to Microscopy

Denis Fortun
Center for Biomedical Imaging
Ecole Polytechnique Federale de Lausanne
denis.fortun@epfl.ch

Noémi Debroux
Laboratoire de Mathématiques
INSA de Rouen
noemie.debroux@insa-rouen.fr

Charles Kervrann
Inria Rennes – Bretagne Atlantique
charles.kervrann@inria.fr

Abstract

Local and global approaches can be identified as the two main classes of optical flow estimation methods. In this paper, we propose a framework to combine the advantages of these two principles, namely robustness to noise of the local approach and discontinuity preservation of the global approach. This is particularly crucial in biological imaging, where the noise produced by microscopes is one of the main issues for optical flow estimation. The idea is to adapt spatially the local support of the local parametric constraint in the combined local-global model [6]. To this end, we jointly estimate the motion field and the parameters of the spatial support. We apply our approach to the case of Gaussian filtering, and we derive efficient minimization schemes for usual data terms. The estimation of a spatially varying standard deviation map prevents from the smoothing of motion discontinuities, while ensuring robustness to noise. We validate our method for a standard model and demonstrate how a baseline approach with pixel-wise data term can be improved when integrated in our framework. The method is evaluated on the Middlebury benchmark with ground truth and on real fluorescence microscopy data.

1. Introduction

Accurate optical flow estimation still remains a challenge for various applications [15]. In particular, image sequences acquired in microscopy are often corrupted by a high level of noise, which is one of the main sources of estimation errors. For example, in time-lapse fluorescence microscopy, the power of excitation lasers often has to be reduced to limit phototoxicity and improve temporal resolution of long sequences. This is at the cost of increasing the

amount of noise. Optical flow estimation is also required for the alignment of micrographs with extremely low signal-to-noise ratios in electron microscopy [1].

Optical flow estimation is based on the assumption of conservation of the image intensity of each pixel during its displacement. Due to its pixel-wise nature, this data fitting constraint is poorly descriptive and gives no information about the spatial context around each pixel. This fundamental uncertainty is called *aperture problem* and states that only the normal component of the motion field can be unambiguously determined. The role of regularization is then to overcome this under-determination by adding additional constraints imposing smoothness of the motion field. Two approaches of regularization have been designed for optical flow and are usually referred to as *local* and *global*.

Global methods penalize the norm of the gradient of the motion field to favor piecewise smoothness [18]. Variational minimization of global models has become the dominant approach for optical flow estimation and is the basis framework of current state-of-the-art methods [23, 26]. However the errors generated by the pixel-wise data term are often compensated by an undesirable over-smoothing. A way to deal with this effect is to design more sophisticated data terms as proposed in [21, 11, 14].

The idea of the local approach is to fit parametric motion models in a neighborhood of each pixel [19]. The motion field is estimated from a region-based data term defined by the sum of data constraints of the pixels in the region. Despite their superior robustness to noise [6], these methods are usually unable to compete with global methods in terms of accuracy at low levels of noise. The reason is the difficulty to determine regions where motion can be accurately approximated by a parametric motion model. However, it has been shown that when appropriate regions are chosen, local parametric models can yield excellent performances

[27, 16].

Our aim is to combine the potential of the region-based constraint of local methods with global regularization. A first investigation in that direction has been performed by Bruhn et al. [6]. The locally constant flow assumption was integrated as a data term in a global variational approach. The spatial support was defined as a Gaussian filter with fixed standard deviation in the whole image domain. As a consequence, the counterpart of robustness to noise was an over-smoothing of discontinuities due to the spatial invariance of the Gaussian kernel.

In this paper, we propose to take benefit from local and global methods with spatial adaptation of the support for the local parametric constraint. We propose a framework for joint estimation of the motion field and the parameters of the spatial support. We derive an efficient minimization scheme in the case of the isotropic Gaussian kernel for several data terms. As a result, we achieve at the same time robustness to noise and preservation of the discontinuities. We demonstrate on the Middlebury benchmark how the accuracy of a baseline pixel-wise method can be improved when integrated in our framework. We also show how real fluorescence microscopy problems where noise is one of the main issues can be addressed with our approach.

1.1. Related work

A common strategy to cope with noise in the optical flow estimation process is to apply a pre-filtering operation to the input images. A Gaussian filtering is the standard choice. However, Gaussian smoothing also blurs image discontinuities and thus affects the recovery of discontinuities of the motion field, as analyzed in [6]. Another usual way to tackle noise is to increase the regularization parameter, which also tends to over-smooth the motion field. Indeed, the role of motion regularization is to model the *a priori* assumption on the motion field, regardless of the nature of data. Noise in the data should therefore be taken into account in the data term.

To reduce the over-smoothing produced by the shift-invariant filtering of the data term in [6], the authors of [13] replaced the Gaussian filter by a bilateral filter, and [22] exploited tensor voting. These latter approaches rely on image measurements to specify filters. Similarly to image-based regularization [25, 24], they are particularly sensitive to noise and produce over-segmentation of the motion field. In [3], the covariance of a Gaussian weighting function is locally estimated jointly with the velocity field, but only for cross-correlation based motion estimation. Estimation of shift-variant filters has been investigated for other problems like deconvolution [12] or denoising [2].

2. Combinations of local and global approaches

2.1. Variational optical flow and energy functional

In the global approach, the flow field is assumed to be piecewise smooth and the strategy is to minimize a global energy $E(\mathbf{w})$ that explicitly combines a potential $M(\cdot)$ which penalizes deviations from the brightness constancy equation with a regularization potential $R(\cdot)$ which penalizes high values of the norm of the gradient of the velocity field. Let $f : \Omega \times [0, T] \rightarrow \mathbb{R}$ be an image sequence of T frames, $\mathbf{w} : \Omega \rightarrow \mathbb{R}^2$ the motion field, Ω the image domain and $\mathbf{x} = (x, y)^\top \in \Omega$ the pixel location. The energy functional is defined as

$$E(\mathbf{w}) = \int_{\Omega} M(\mathbf{x}, \mathbf{w}) + \lambda R(\mathbf{x}, \nabla_2 \mathbf{w}) dx \quad (1)$$

where $\mathbf{w}(\mathbf{x}) = (u(\mathbf{x}), v(\mathbf{x}))^\top$ denotes the motion vector at pixel $\mathbf{x} \in \Omega$, $\nabla_2 := (\partial_x, \partial_y)^\top$ denotes the spatial gradient operator and $\lambda > 0$ is a regularization weight.

The standard quadratic pixel-wise data fidelity potential obtained from the brightness constancy assumption writes

$$M_0(\mathbf{x}, \mathbf{w}) = \mathbf{w}^+(\mathbf{x})^\top \mathbf{J}_0 \mathbf{w}^+(\mathbf{x}) \quad (2)$$

with the tensor $\mathbf{J}_0 = \nabla_3 f \nabla_3^\top f$, where $\nabla_3 := (\partial_x, \partial_y, \partial_t)^\top$ denotes the spatio-temporal gradient operator and $\mathbf{w}^+(\mathbf{x}) := (u(\mathbf{x}), v(\mathbf{x}), 1)^\top$. To make the data term robust to additive illumination changes, it is recommended in [5, 28] to consider the following data term, which combines the brightness and gradient constancy assumptions as follows:

$$M_1(\mathbf{x}, \mathbf{w}) = \mathbf{w}^+(\mathbf{x})^\top \mathbf{J}_1 \mathbf{w}^+(\mathbf{x}) + \gamma \mathbf{w}^+(\mathbf{x})^\top \bar{\mathbf{J}}_1 \mathbf{w}^+(\mathbf{x}) \quad (3)$$

where $\mathbf{J}_1 = c \mathbf{J}_0$ and $\bar{\mathbf{J}}_1 = c_x \nabla_3 f_x \nabla_3^\top f_x + c_y \nabla_3 f_y \nabla_3^\top f_y$ where $\gamma > 0$ balances the brightness and gradient assumptions and the normalization factors are defined as: $c = (f_x^2 + f_y^2 + \epsilon^2)^{-1}$, $c_x = (f_{xx}^2 + f_{xy}^2 + \epsilon^2)^{-1}$, $c_y = (f_{yx}^2 + f_{yy}^2 + \epsilon^2)^{-1}$ where $\epsilon > 0$ avoids division by zero.

2.2. Spatial filtering of data terms

The combination of local and global approaches in a single model has been investigated by Bruhn et al. [6]. The so-called ‘‘Combined Local-Global’’ method (CLG) is based on the local filtering of the data term, implying a neighborhood-wise data constancy constraint, less sensitive to noise than pixel-wise measures. The following data term is considered in [6]:

$$M_{0,\sigma}(\mathbf{x}, \mathbf{w}) = \mathbf{w}^+(\mathbf{x})^\top \mathbf{J}_{0,\sigma} \mathbf{w}^+(\mathbf{x}) \quad (4)$$

where $\mathbf{J}_{0,\sigma} := G_\sigma * \mathbf{J}_0$, G_σ is a Gaussian filter of standard deviation σ , and $*$ is the convolution operator. In (4),

the Gaussian filtering is applied only on the image variables f_t, f_x, f_y and not on \mathbf{w} which is therefore assumed to be locally constant. Experiments in [6] demonstrated the increased robustness to noise of the results obtained by Gaussian filtering of the data term. The method has been applied in several fields where acquisition conditions induce noise in the image [9, 10].

In [28], a similar modeling is proposed when the data term $M_1(\mathbf{x}, \mathbf{w})$ is considered:

$$M_{1,\sigma}(\mathbf{x}, \mathbf{w}) = \mathbf{w}^+(\mathbf{x})^\top \mathbf{J}_{1,\sigma} \mathbf{w}^+(\mathbf{x}) + \gamma \mathbf{w}^+(\mathbf{x})^\top \bar{\mathbf{J}}_{1,\sigma} \mathbf{w}^+(\mathbf{x})$$

where $\mathbf{J}_{1,\sigma} := G_\sigma * \mathbf{J}_1$ and $\bar{\mathbf{J}}_{1,\sigma} := G_\sigma * \bar{\mathbf{J}}_1$. Nevertheless, while this strategy focuses on robustness to noise, it does not address the problem of estimating the optimal filter support σ at each location. Gaussian smoothing with fixed standard deviation tends to over-smooth the resulting flow field at motion discontinuities.

3. Spatially adaptive filtering and optimization

3.1. Novel optical flow energy model

Over-smoothing occurs with [6], as for local methods, when the support of the local filtering contains multiple motions, that is, at motion discontinuities. The aim is then to restrict the spatial support to coherently moving regions by replacing Gaussian filtering with fixed standard deviation by an adaptive filtering. Considering the case of isotropic Gaussian smoothing, the parameter to optimize jointly with \mathbf{w} is the standard deviation of the Gaussian filter, now defined as a dense field $\sigma : \Omega \rightarrow \mathbb{R}^+$. We define the optimization problem

$$(\hat{\mathbf{w}}, \hat{\sigma}) = \arg \min_{\mathbf{w}, \sigma} E(\mathbf{w}, \sigma) \quad (5)$$

where the energy $E(\mathbf{w}, \sigma)$ is defined as

$$E(\mathbf{w}, \sigma) = \int_{\Omega} \rho(M_{\nu,\sigma}(\mathbf{x}, \mathbf{w})) d\mathbf{x} + \lambda \int_{\Omega} \phi(\|\nabla \mathbf{w}(\mathbf{x})\|^2) d\mathbf{x} + \beta \int_{\Omega} \psi(|\nabla \sigma(\mathbf{x})|^2) d\mathbf{x} + \mu \int_{\Omega} \frac{1}{\sigma(\mathbf{x})} d\mathbf{x} \quad (6)$$

and $\nu \in \{0, 1\}$, $\lambda, \beta > 0$ are weights that balance the contributions of the data and regularization terms and $\rho(\cdot)$ is a robust penalization function. The first term is a generalization of the data potential (4) and the second and third terms are regularizations on \mathbf{w} and σ . To preserve discontinuities in the regularized flow field, robust penalizers $R(\mathbf{x}, \nabla_2 \mathbf{w}) = \phi(\|\nabla_2 \mathbf{w}(\mathbf{x})\|^2)$ and $\psi(|\sigma(\mathbf{x})|^2)$ are considered. The last term is an inverse barrier that prevents σ from being negative if it is initialized with a positive value, and μ is interpreted as parameter that controls the severity of the barrier.

The minimization (6) w.r.t. \mathbf{w} amounts to the method of [6]. Minimizing w.r.t. σ adapts spatially the standard deviation of the convolution $\sigma(\mathbf{x})$ at each point \mathbf{x} . The aim is to reduce σ at motion discontinuities, where Gaussian smoothing tends to blur the estimated motion field. If a discontinuity is contained in the Gaussian support defined by σ , the locally constant motion assumption will be violated and lower values σ will be favoured. Rather than being adapted to the image content as in [13, 22], σ is guided now by the data term and variations follow motion discontinuities rather than image discontinuities. Note that the inverse barrier term tends to encourage large values of σ , which is desirable for large regions with coherent motion, except at motion discontinuities.

3.2. Energy minimization

The optimization is performed alternatively w.r.t. \mathbf{w} and σ . Minimization of $E(\mathbf{w}, \sigma)$ w.r.t. \mathbf{w} with σ fixed amounts to a standard optimization problem in optical flow. We derive the associated non-linear Euler-Lagrange equations as [5]:

$$\begin{cases} \frac{\partial \rho(M_{\nu,\sigma}(\mathbf{x}, \mathbf{w}))}{\partial u(\mathbf{x})} - \lambda \operatorname{div}(\phi'(\|\nabla \mathbf{w}(\mathbf{x})\|^2) \nabla u(\mathbf{x})) = 0 \\ \frac{\partial \rho(M_{\nu,\sigma}(\mathbf{x}, \mathbf{w}))}{\partial v(\mathbf{x})} - \lambda \operatorname{div}(\phi'(\|\nabla \mathbf{w}(\mathbf{x})\|^2) \nabla v(\mathbf{x})) = 0, \end{cases} \quad (7)$$

where $\operatorname{div}(\cdot)$ is the divergence. To solve the non-linear system of equations (7) we use the fixed point scheme detailed in [4].

To minimize (6) w.r.t. σ with \mathbf{w} fixed, we adopt a gradient-based minimization approach using the quasi-Newton method L-BFGS [20]. It requires the computation the derivative of $E(\mathbf{w}, \sigma)$ w.r.t. $\sigma(\mathbf{x})$. We first we consider the data term $\rho(M_{0,\sigma}(\mathbf{x}, \mathbf{w}))$. It follows that

$$\begin{aligned} \frac{\partial E(\mathbf{w}, \mathbf{x})}{\partial \sigma(\mathbf{x})} &= (\mathbf{w}^+(\mathbf{x})^\top (g_{\sigma(\mathbf{x})} * \mathbf{J}_0) \mathbf{w}^+(\mathbf{x})) \rho'(M_{0,\sigma(\mathbf{x})}(\mathbf{x}, \mathbf{w})) \\ &\quad - 2\beta \operatorname{div}(\nabla \sigma(\mathbf{x}) \psi'(|\nabla \sigma(\mathbf{x})|^2)) - \frac{\mu}{\sigma^2(\mathbf{x})} \end{aligned}$$

where

$$g_{\sigma(\mathbf{x})}(\mathbf{x} - \cdot) := \frac{\partial G_{\sigma(\mathbf{x})}(\mathbf{x} - \cdot)}{\partial \sigma(\mathbf{x})} = \frac{G_{\sigma(\mathbf{x})}(\mathbf{x} - \cdot)}{\sigma(\mathbf{x})} \left(\frac{\|\mathbf{x} - \cdot\|^2}{\sigma^2(\mathbf{x})} - 1 \right)$$

is a filter analogous to the Gaussian filter $G_{\sigma(\mathbf{x})}$. The derivative of $E(\mathbf{w}, \sigma)$ w.r.t. $\sigma(\mathbf{x})$ when considering the data term $\rho(M_{1,\sigma})$ is as follows:

$$\begin{aligned} \frac{\partial E(\mathbf{w}, \mathbf{x})}{\partial \sigma(\mathbf{x})} &= \quad (8) \\ &(\mathbf{w}^+(\mathbf{x})^\top (g_{\sigma(\mathbf{x})} * \mathbf{J}_1) \mathbf{w}^+(\mathbf{x})) \rho'(\mathbf{w}^+(\mathbf{x})^\top \mathbf{J}_{1,\sigma} \mathbf{w}^+(\mathbf{x})) \\ &+ \gamma (\mathbf{w}^+(\mathbf{x})^\top (g_{\sigma(\mathbf{x})} * \bar{\mathbf{J}}_1) \mathbf{w}^+(\mathbf{x})) \rho'(\mathbf{w}^+(\mathbf{x})^\top \bar{\mathbf{J}}_{1,\sigma} \mathbf{w}^+(\mathbf{x})) \\ &- 2\beta \operatorname{div}(\nabla \sigma(\mathbf{x}) \psi'(|\nabla \sigma(\mathbf{x})|^2)) - \frac{\mu}{\sigma^2(\mathbf{x})} \end{aligned}$$




Figure 1. Visual results obtained with CLG-A, CLG₀ and CLG on the *Grove3* sequence corrupted with Gaussian noise with standard deviation 40.

3.3. Implementation

In our implementation, we embedded the estimation in a coarse-to-fine scheme to cope with large displacements. The alternate optimization of w and σ is performed at each level of the pyramid. In our experiments, 3 alternated minimization iterations are sufficient to reach convergence.

As in [5, 28], we choose a TV- l_1 model by taking $\rho(z^2) = \phi(z^2) = \psi(z^2) = \sqrt{z^2 + \varepsilon}$ with a small regularization value to approximate the l_1 norm $\varepsilon = 0.001$, and we set $\gamma = 3$.

4. Experimental results

We evaluate our method with the model $M_{1,\sigma}$ by comparing our spatially variant estimation of σ , that we call CLG-A, with the standard pixel-wise approach without filtering [5], denoted CLG₀, and the method of [6] using a fixed value for σ , denoted CLG. We focus on these methods to isolate the impact of our adaptive filtering approach inside a fixed estimation framework. Our approach could be integrated in more recent motion estimation methods to further improve results. We fix the value of σ to 3 for CLG. We first evaluate our method on the sequences of the Middlebury benchmark with ground truth, and we use the endpoint error (EPE) as an error measure, which is a standard evaluation procedure. Note that the Middlebury database is recommended to evaluate variational methods which focus on

accurate estimation of small displacements. Recent datasets [7, 17] are mostly dedicated to large displacements, occlusions and intensity changes, which should be addressed by specific techniques that could be integrated in our framework. The regularization parameter on the motion field λ has been optimized for each method and each sequence. The computational time of CLG-A is about 5 times higher than the one of CLG when the number of alternate optimization steps is 3. For a C++ implementation, CLG takes about 20s for the *Grove3* sequence.

Figure 1 shows visual results on the *Grove3* sequence that we corrupted with additive Gaussian noise of standard deviation equal to 40. The map of σ estimated with CLG-A is also displayed, encoding σ values by the image intensity (dark regions correspond to small value of σ and bright regions to large value of σ). The range of estimated values of σ is $[0.1, 3.1]$. In Table 1, the EPE obtained on the Middlebury benchmark are reported for several levels of noise. In Fig. 2, we compare the evolution of the EPE with the noise level in the input images for CLG-A, CLG₀, and CLG. In Fig. 1, we observe that to cope with noise in the input image, CLG₀ has to smooth motion details and large regions of the motion field. CLG tends to smooth motion discontinuities because of the fixed size of the Gaussian kernel. In contrast our CLG-A results are both robust to noise and discontinuity-preserving. The estimated σ map shows that large values are estimated in regions of homogeneous

Table 1. Comparison of EPE obtained with CLG_0 , CLG, and our proposed method CLG-A on the sequences of the Middlebury benchmark. We consider three levels of Gaussian noise standard deviation (“noise std”) added to the input images.

Noise std	Method	Dimetrodon	Grove2	Grove3	Hydrangea	RubberWhale	Urban2	Urban3	Venus
20	CLG_0	0.253	0.211	0.801	0.229	0.215	0.721	1.305	0.608
	CLG	0.360	0.381	1.111	0.416	0.340	0.816	1.013	0.619
	CLG-A	0.278	0.244	0.809	0.296	0.239	0.717	0.886	0.592
30	CLG_0	0.313	0.271	0.927	0.418	0.328	0.818	1.470	0.666
	CLG	0.398	0.412	1.154	0.425	0.351	0.821	1.051	0.629
	CLG-A	0.325	0.285	0.909	0.332	0.272	0.761	0.925	0.625
40	CLG_0	0.397	0.323	1.120	0.602	0.352	0.961	1.578	0.678
	CLG	0.412	0.418	1.174	0.433	0.356	0.833	1.079	0.633
	CLG-A	0.373	0.318	1.006	0.379	0.296	0.798	0.960	0.654




Figure 2. Influence of the standard deviation of the noise on the endpoint error obtained with CLG-A, CLG, and CLG_0 , for the sequence *Grove3*.

motion to cope with noise, whereas the filtering is almost canceled at motion discontinuities to avoid over-smoothing. The spatial variations of σ follow motion discontinuities and not image discontinuities. The results of Table 1 shows that our spatial adaptation of σ yields a significant improvement over pixel-wise and fixed σ approaches for noise levels above a variance of 30, and stays competitive for lower noise levels. Figure 2 confirms that CLG-A combines the performances of CLG for high levels of noise with those of CLG_0 for low levels of noise.

We also tested our method on several real fluorescence microscopy sequences obtained in widefield imaging representing the growth of *CElegans* embryos, which is a classical research topic in biology. Figure 3 visually compares the results obtained with the three methods on an image pair extracted from a typical sequence, for several values of λ .

The results of CLG_0 show that it is clearly not adapted to the high level of noise that commonly occurs in fluorescence microscopy. Compared to CLG, CLG-A preserves sharper cell discontinuities. Since the Gaussian kernel also acts as a regularizer, CLG-A is less affected by the change of λ , which can be a convenient feature for users and biologist experts.

5. Conclusion - Perspectives

In this paper, we have proposed a framework for the spatially adaptive integration of a local parametric assumption in the data term of a global optical flow model. In the case of Gaussian filtering, we have derived efficient quasi-Newton optimization schemes for joint estimation of motion and standard deviation of the filter. Our experimental results demonstrated that our method combines robustness to noise with sharpness of motion discontinuities. It yields significant improvements compared to the non-adaptive version for all levels of noise. It is also more accurate in fluorescence microscopy, where the usual pixel-wise approach fails due to the high level of noise.

A natural and potentially powerful extension could be to define the local spatial support with anisotropic Gaussian weights to fit more accurately the possibly complex shapes of motion discontinuities. Another direction of research would be to consider more advanced local anisotropic models [8]. We integrated our model in a variational framework adapted to small displacements. Other approaches to deal with large displacements and occlusions could be considered in the future [16, 23, 26].

Acknowledgement

This work was realized as part of the Quaero program, funded by OSEO, French State agency for innovation. It was also partly supported by the France-BioImaging project




Figure 3. Visual results obtained with CLG-A, CLG_0 and CLG on the *CElegans* sequence, for several values of λ .

granted by the Investissement d’Avenir program, and by UMR 144 CNRS of Institut Curie.

References

- [1] V. Abrishami, J. Vargas, X. Li, Y. Cheng, R. Marabini, C. Ó. S. Sorzano, and J. M. Carazo. Alignment of direct detection device micrographs using a robust optical flow approach. *Journal of structural biology*, 189(3):163–176, 2015.
- [2] N. Azzabou, N. Paragios, F. Guichard, and F. Cao. Variable bandwidth image denoising using image-based noise models. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1–7, 2007.
- [3] F. Becker, B. Wienneke, S. Petra, A. Schroder, and C. Schnorr. Variational adaptive correlation method for flow estimation. *Image Processing*, 21(6):3053–3065, 2012.
- [4] T. Brox. *From pixels to regions: partial differential equations in image analysis*. PhD dissertation, Saarland University, 2005.
- [5] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a the-

- ory for warping. In *European Conference on Computer Vision (ECCV)*, pages 25–36, Prague, Czech Republic, May 2004.
- [6] A. Bruhn, J. Weickert, and C. Schnörr. Lucas/kanade meets horn/schunck: combining local and global optical flow methods. *International Journal of Computer Vision*, 61(3):211–231, 2005.
- [7] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *European Conference on Computer Vision (ECCV)*, pages 611–625. Springer-Verlag, 2012.
- [8] T. Corpetti and E. Mémin. Stochastic uncertainty models for the luminance consistency assumption. *IEEE Transactions on Image Processing*, 21(2):481–493, 2012.
- [9] M. Dawood, F. Buther, X. Jiang, and K. P. Schafers. Respiratory motion correction in 3-d pet data with advanced optical flow algorithms. *IEEE Transactions on Medical Imaging*, 27(8):1164–1175, 2008.
- [10] J. Delpiano, J. Jara, J. Scheer, O. A. Ramírez, J. Ruizdel Solar, and S. Härtel. Performance of optical flow techniques for motion analysis of fluorescent point signals in confocal microscopy. *Machine Vision and Applications*, 23(4):675–689, 2012.
- [11] O. Demetz, D. Hafner, and J. Weickert. Morphologically invariant matching of structures with the complete rank transform. *International Journal of Computer Vision*, 113(3):220–232, 2015.
- [12] L. Denis, E. Thiébaud, F. Soulez, J.-M. Becker, and R. Mourya. Fast approximations of shift-variant blur. *International Journal of Computer Vision*, 115(3):253–278, 2015.
- [13] M. Drulea and S. Nedevschi. Total variation regularization of local-global optical flow. In *Intelligent Transportation Systems Conference (ITSC)*, pages 318–323, 2011.
- [14] M. Drulea and S. Nedevschi. Motion estimation using the correlation transform. *IEEE Transactions on Image Processing*, 22(8):3260–3270, 2013.
- [15] D. Fortun, P. Bouthemy, and C. Kervrann. Optical flow modeling and computation: A survey. *Computer Vision and Image Understanding*, 134:1–21, 2015.
- [16] D. Fortun, P. Bouthemy, and C. Kervrann. Aggregation of local parametric candidates with exemplar-based occlusion handling for optical flow. *Computer Vision and Image Understanding*, 145:81–94, 2016.
- [17] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Computer Vision and Pattern Recognition (CVPR)*, pages 3354–3361, 2012.
- [18] B. Horn and B. Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981.
- [19] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *International Joint Conference on Artificial Intelligence*, pages 674–679, 1981.
- [20] J. Nocedal. Updating quasi-newton matrices with limited storage. *Mathematics of computation*, 35(151):773–782, 1980.
- [21] R. Ranftl, S. Gehrig, T. Pock, and H. Bischof. Pushing the limits of stereo using variational stereo estimation. In *Intelligent Vehicles Symposium (IV)*, pages 401–407, 2012.
- [22] H. A. Rashwan, M. A. García, and D. Puig. Variational optical flow estimation based on stick tensor voting. *IEEE Transactions on Image Processing*, 22(7):2589–2599, 2013.
- [23] J. Revaud, P. Weinzaepfel, and C. Harchoui, Z. Schmid. Epicflow: Edge-preserving interpolation of correspondences for optical flow. In *IEEE Conf. Computer Vision and Pattern Recognition (CVPR’15)*, Boston, MA, 2015.
- [24] A. Wedel, D. Cremers, T. Pock, and H. Bischof. Structure-and motion-adaptive regularization for high accuracy optic flow. In *International Conference on Computer Vision (ICCV)*, pages 1663–1668, Kyoto, Japan, October 2009.
- [25] M. Werlberger, T. Pock, and H. Bischof. Motion estimation with non-local total variation regularization. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2464–2471, San Fransisco, June 2010.
- [26] L. Xu, J. Jia, and Y. Matsushita. Motion detail preserving optical flow estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(9):1744–1757, 2012.
- [27] J. Yang and H. Li. Dense, accurate optical flow estimation with piecewise parametric model. In *IEEE Conf. Computer Vision and Pattern Recognition (CVPR’15)*, Boston, MA, 2015.
- [28] H. Zimmer, A. Bruhn, and J. Weickert. Optic flow in harmony. *International Journal of Computer Vision*, 93(3):1–21, 2011.