Universität des Saarlandes

Fachrichtung 6.1 – Mathematik

Preprint Nr. 106

Confidence Measures for Variational Optic Flow Methods

Andrés Bruhn and Joachim Weickert

Saarbrücken 2004

Fachrichtung 6.1 – Mathematik
Universität des Saarlandes

Preprint No. 106
submitted: Juni 17, 2004

Confidence Measures for Variational Optic Flow Methods

Andrés Bruhn

Mathematical Image Analysis Group,
Faculty of Mathematics and Computer Science,
Saarland University, Building 27.1,
66041 Saarbrücken, Germany
bruhn@mia.uni-saarland.de

Joachim Weickert

Mathematical Image Analysis Group,
Faculty of Mathematics and Computer Science,
Saarland University, Building 27,
66041 Saarbrücken, Germany
weickert@mia.uni-saarland.de

Edited by
FR 6.1 – Mathematik
Universität des Saarlandes
Postfach 15 11 50
66041 Saarbrücken
Germany

Fax: + 49 681 302 4443
e-Mail: preprint@math.uni-sb.de
WWW: http://www.math.uni-sb.de/

Abstract

In this paper we investigate the usefulness of confidence measures for variational optic flow computation. To this end we discuss the frequently used sparsification strategy based on the image gradient. Its drawbacks motivate us to propose a novel, energy-based confidence measure that is parameter-free and applicable to the entire class of energy minimising optic flow techniques. Experimental evaluations show that this confidence measure leads to excellent results, independently of the image sequence or the underlying variational approach.

AMS 2000 Subject Classification: 68T45, 49K20, 65K10, 35J60, 65N06
Key Words: optic flow, confidence measures, differential techniques, variational methods, partial differential equations, performance evaluation.

Contents

1 Introduction

The recovery of motion fields from image sequences is one of the key problems in computer vision. Given two consecutive frames of an image sequence, one is interested in finding the projection of the 3-D motion onto the image plane: the so-called optic flow field. In order to estimate this displacement field, optic flow methods often use a constancy assumptions on image features such as the grey value.

At that point, the handling of incomplete data plays a very important role: In general, these constancy assumptions cannot provide sufficient data to determine a unique solution of the optic flow problem. For instance, in the case of the grey value constancy assumption this incompleteness manifests itself in the aperture problem. In this case not more than the flow component parallel to the image gradient can be calculated. At locations where the gradient is zero, not even this component is computable and no estimation is possible.

In order to cope with these situations, variational methods regularise the problem by assuming smoothness or piecewise smoothness of the resulting flow field. At locations, where the problem of incomplete data occurs, this regularisation fills in information from the neighbourhood and thus allows the estimation of a 100 % dense flow field.

However, it is clear that these estimates cannot have the same reliability at all locations. It would therefore be interesting to find a confidence measure that allows to assess the reliability of a dense optic flow field. In particular, such a measure would allow to identify locations where the problem of incomplete data has been solved successfully. Therefore, it is not surprising that [1] have identified the absence of such a good measure as one of the main drawbacks of variational optic flow techniques.

In our paper we address this problem. By discussing the frequently used confidence measure based on the image gradient, we show why this method is not appropriate for sparsifying dense flow fields from variational methods. As a remedy, we propose a novel energy-based confidence measure that offers several advantages and works well over a large range of densities.

Related work. In spite of the fact that there exists a very large number of publications on variational optic flow methods (see e.g. [11, 15, 21, 3]) and on confidence measures for local optic flow approaches (see e.g. [2, 19]), there has been remarkably little work devoted to the application of confidence measures in the context of variational optic flow computation. First approaches go back to [1] who used the magnitude of the image gradient to decide on the local reliability or a flow estimate. More recently, [10] proposed a general classification of confidence measures. However, they did not introduce any novel confidence measure for variational methods.

Organisation of the paper. Our paper is organised as follows. In Section 2 we give a review on variational optic flow computation and discuss different types of regularisation strategies. The gradient-based confidence measure by [1] that is often used to sparsify the resulting flow fields is discussed in Section 3. Based on the results of this discussion we propose a novel energy-based confidence measure in Section 4 and perform a systematic experimental evaluation in Section 5. Finally, Section 6 concludes this paper with a summary.

2 Variational Optic Flow Computation

Let us consider some image sequence $f(x, y, t)$, where (x, y) denotes the location within a rectangular image domain Ω, and $t \in [0, T]$ denotes time. In order to retrieve objects in subsequent frames of this image sequence, many optic flow methods assume that corresponding pixels have the same grey value, i.e. that the grey value of objects remains constant over time. If we denote the movement of such an object by $(x(t), y(t))$ this assumption can be formulated as

$$0 = \frac{df(x(t), y(t), t)}{dt}. \tag{1}$$

By applying the chain rule this leads to the following *optic flow constraint (OFC)*:

$$0 = f_x u + f_y v + f_t, \tag{2}$$

where subscripts denote partial derivatives and the optic flow field satisfies $(u, v)^\top = (\partial_t x, \partial_t y)^\top$.

Evidently, this single equation is not sufficient to uniquely determine the two unknowns u and v. In particular at locations where the image gradient is zero, no estimation of the optic flow is possible. In all other cases, only the flow component parallel to $\nabla f := (f_x, f_y)^\top$ can be computed, the so-called *normal flow*:

$$w_n = -\frac{f_t}{|\nabla f|} \frac{\nabla f}{|\nabla f|}. \tag{3}$$

In the literature, this ambiguity is referred to as the *aperture problem.*

2.1 General Structure

Variational methods overcome the aperture problem by imposing an additional constraint on the solution: They assume that the resulting flow field is smooth or piecewise smooth. Then the optic flow can be computed as minimiser of a global energy functional, where both deviations from the data and deviations from the smoothness constraint are penalised. Let $\nabla_3 := (\partial_x, \partial_y, \partial_t)^\top$ denote the spatiotemporal gradient, let $D(u, v, \nabla_3 f)$ stand for a data term (e.g. the squared OFC) and let $S(\nabla_3 u, \nabla_3 v, \nabla_3 f)$ represent a constraint on the smoothness of the resulting flow field. Then the corresponding energy functional is given by

$$E(u,v) = \int_{\Omega\times[0,T]} \Big(\underbrace{D(u,v,\nabla_3 f)}_{\text{Data term}} + \alpha \underbrace{S(\nabla_3 u, \nabla_3 v, \nabla_3 f)}_{\text{Smoothness term}} \Big) \, dx\, dy\, dt, \tag{4}$$

where α serves as regularisation parameter that steers the smoothness of the estimated flow field.

2.2 Prototypes for Variational Methods

Let us now take a closer look at the different regularisation strategies that may serve as smoothness constraints. As classified in [21] there are basically three different types of regularisation: *Homogeneous regularisation* that assumes overall smoothness, *image-driven regularisation* that assumes piecewise smoothness and respects discontinuities in the image data, and *flow-driven regularisation* that assumes piecewise smoothness and respects discontinuities in the flow field. Moreover, when considering image and flow-driven regularisation, one can distinguish between *isotropic* and *anisotropic* smoothness terms. While isotropic regularisers do not impose any smoothness at discontinuities, anisotropic ones permit smoothing along the discontinuity but not across it.
In order to demonstrate the different regularisation concepts and to allow for a systematic experimental evaluation, we have chosen three prototypes of variational methods that cover all types of regularisation.

2.2.1 The Combined Local–Global Method

As prototype for the class of optic flow techniques with *homogeneous regularisation* we consider the so-called combined local-global (CLG) method [4, 5]. This technique combines the dense flow fields of the global approach of [11] with the high noise robustness of the local method of [13].
Let $\mathbf{w} = (u, v, 1)^\top$ denote the spatiotemporal extended flow vector. Then the energy functional of the CLG method is given by

$$E(u,v) = \int_{\Omega\times[0,T]} \left(\mathbf{w}^\top J_\rho(\nabla_3 f)\, \mathbf{w} + \alpha \left(|\nabla_3 u|^2 + |\nabla_3 v|^2 \right) \right) \, dx\, dy\, dt. \tag{5}$$

where the matrix $J_\rho(\nabla_3 f)$ is the so-called structure tensor [2, 8, 17] given by $K_\rho * (\nabla_3 f \, \nabla_3 f^\top)$, the symbol $*$ denotes convolution in each matrix component, and K_ρ is a Gaussian with standard deviation ρ. One should note that for $\rho \to 0$ the spatial variant of the CLG approach comes down to the Horn and Schunck method, and for $\alpha \to 0$ it becomes the Lucas–Kanade algorithm.

2.2.2 The Method of Nagel and Enkelmann

For the class of optic flow methods with *image-driven* regularisation we consider the *anisotropic* technique of [15]. This method assumes the flow field to be smooth everywhere except across discontinuities in the image data. This can be realised by penalising only the projection of the flow on the plane orthogonal to the image gradient. The corresponding energy functional for the spatiotemporal variant of the Nagel-Enkelmann algorithm is given by [14]

$$
\begin{aligned}
E(u,v) = \int_{\Omega \times [0,T]} & \Big((f_x u + f_y v + f_t)^2 + \alpha \, (\nabla_3 u^\top D(\nabla_3 f) \nabla_3 u) \\
& + \alpha \, (\nabla_3 v^\top D(\nabla_3 f) \nabla_3 v) \Big) \; dx \, dy \, dt,
\end{aligned}
\tag{6}
$$

with the regularised projection matrix

$$
D(\nabla_3 f) = \frac{1}{2|\nabla_3 f|^2 + 3\epsilon^2} \begin{pmatrix} f_y^2 + f_z^2 + \epsilon^2 & -f_x f_y & -f_x f_z \\ -f_x f_y & f_x^2 + f_z^2 + \epsilon^2 & -f_y f_z \\ -f_x f_z & -f_y f_z & f_x^2 + f_y^2 + \epsilon^2 \end{pmatrix}
\tag{7}
$$

perpendicular to $\nabla_3 f$, where ϵ serves as regularisation parameter.

2.2.3 The TV-based Regularisation Method

In contrast to image-driven regularisation methods, *flow-driven* techniques preserve discontinuities at those locations where edges in the flow field occur during computation. Our representative for this third class of variational optic flow techniques is an *isotropic* method that penalises deviations from the smoothness assumption with the L_1 norm of the flow gradient. This corresponds to *total variation* (TV) regularisation [18]. The associated energy functional is given by

$$
\begin{aligned}
E(u,v) = \int_{\Omega \times [0,T]} & \Big((f_x u + f_y v + f_t)^2 \\
& + \alpha \sqrt{|\nabla_3 u|^2 + |\nabla_3 v|^2 + \epsilon^2} \Big) \; dx \, dy \, dt,
\end{aligned}
\tag{8}
$$

where ϵ serves as small regularisation parameter. Related spatial energy functionals have been proposed by [6, 7, 12], and similar spatiotemporal functionals have been investigated in [22].

2.3 The Filling-In Effect

The strategy of regularising the solution by a smoothness assumption has a useful side-effect: Variational methods always yield 100 % dense flow fields. At locations with $\nabla_3 f \approx 0$, the data term does not allow a reliable computation of a local flow estimate. However, the smoothness term fills in information from the neighbourhood. This can be explained as follows: Since the contribution of the data term to the energy functional is very small at these locations, the smoothness term becomes relatively more important. As a consequence, the local flow estimate is adjusted to its neighbourhood in accordance with the smoothness constraint. This propagation of neighbourhood information is the so called *filling-in* effect.

3 The Gradient-Based Confidence Measure

As we have seen in the previous section, the aperture problem does only allow the direct computation of the normal flow. Since this requires the gradient at the corresponding pixel to be different from zero, [1] proposed to connect the reliability of a flow estimate to the magnitude of the underlying image gradient. Thus, the following confidence measure is obtained:

$$c_{\text{grad}} = |\nabla f|. \tag{9}$$

However, this ad-hoc criterion suffers from two drawbacks. Large gradients often result from noise and occlusions. Therefore, evaluating the magnitude of the gradient rewards exactly those locations, where the estimation of the optic flow is particularly problematic. It is not surprising, that the application of such a measure can only be of limited success. Moreover, it is clear that "a-priori" measures that only judge the initial situation *before* the computation, are not in the best position to decide on the reliability of a local flow estimate. They are simply not capable of considering the propagation of neighbourhood information for solving the aperture problem. In particular with regard to variational optic flow methods that rely on the global filling-in effect of the regulariser. This constitutes another drawback.

4 A Novel Energy-Based Confidence Measure

In order to capture the filling-in effect of the regulariser in a better way, one should think of involving the computed flow in the decision process. Let us now demonstrate how this can be accomplished in a natural way. As we know from Section 2, variational methods are based on the minimisation of an energy functional. This energy functional penalises deviations from model assumptions by summing up the local deviations from the image domain. At locations where this deviation is small, the computed flow respects the underlying model. At locations where the deviation is large, on the other hand, the model assumptions are violated severely. In this context it appears very natural to use this indicator for assessing the local reliability of the computation. Thus, we propose a confidence measure where the reliability is inversely proportional to the local energy contribution:

$$c_{\text{ener}} = \frac{1}{D\left(u, v, \nabla_3 f\right) + \alpha\, S\left(\nabla_3 u, \nabla_3 v, \nabla_3 f\right) + \epsilon^2}, \tag{10}$$

where ϵ serves as small regularisation parameter that prevents the denominator from becoming singular. Its actual value is not important since we are only interested in a ranking of the confidence at different locations.

Apart from its simplicity, this confidence measure has several additional advantages. Firstly, it is a consequent continuation of the concept of variational methods: *The confidence measure is based on exactly the same assumptions as the underlying energy functional.* There is no reason, why other constraints should be used for evaluating the reliability of the estimated flow field: If other constraints are considered important, they should have been taken into account earlier by incorporating them in the variational model for computing the flow field. Secondly, the energy-based confidence measure allows to consider the filling-in effect of the regulariser: In contrast to the image gradient it is based on the evaluation of the flow field. This is the only data where the filling-in effect is present. Thirdly, it allows to detect noise and occlusions to a certain degree. At those locations contradictory information does either allow to fulfill the smoothness or the data term. As a consequence, these locations have a relatively high local energy contribution and thus can be identified easily. Fourthly, the proposed confidence measure can be derived in a straightforward way from any energy functional. This makes it applicable to the entire class of energy minimising optic flow techniques. And finally, it is parameter-free. Since the parameter have already been set before the computation of the flow field there is no need to readjust them afterwards.

5 Results

In order to be able to quantify the reliability of confidence measures in the experimental section we restrict ourselves to image sequences for which the ground truth is available. In particular, this allows us to compute error measures such as the frequently used average angular error. It is defined as the arithmetic mean of

$$\arccos\left(\frac{u_c u_e + v_c v_e + 1}{\sqrt{(u_c^2 + v_c^2 + 1)(u_e^2 + v_e^2 + 1)}}\right) \tag{11}$$

where (u_c, v_c) denotes the correct flow, and (u_e, v_e) is the estimated flow (cf. also [1]).
In our first experiment we compare the performance of the gradient and the energy based confidence measures. To this end we use the spatiotemporal approach with locally integrated data term and homogeneous regularisation (CLG) and compute the flow field between frame 8 and 9 of the famous *Yosemite* sequence with clouds (`ftp://csd.uwo.ca` under `/pub/vision/`). Then, we successively sparsify the estimated flow field by applying the confidence measures independently, and calculate the corresponding average angular errors within a density range from 100 % to 1 %.

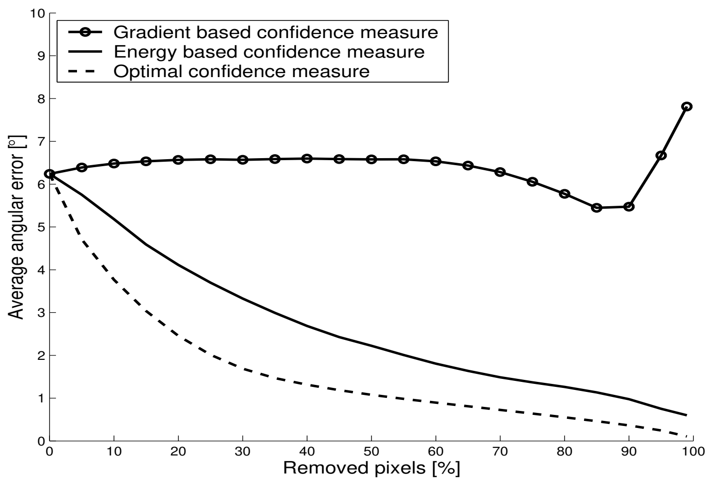

Figure 1: Comparison of the gradient, eigenvalue and energy-based confidence measures for the Yosemite sequence with clouds using the 3-D approach with locally integrated data term and homogeneous regularisation (CLG).

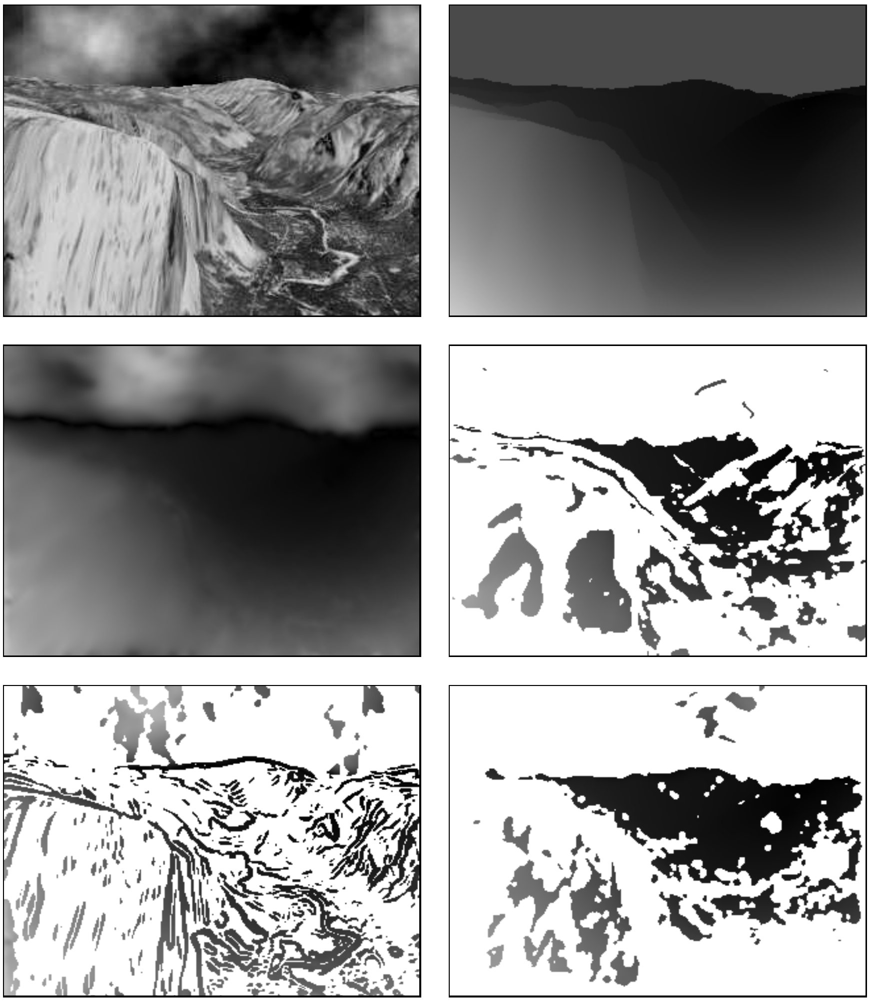

Figure 2: *From left to right, and from top to bottom: (a)* Frame 8 of the Yosemite sequence with clouds (316 $\times$ 256 pixels). *(b)* Magnitude of the ground truth. Brighter structures indicate larger values. *(c)* Magnitude of the computed field for a spatiotemporal approach with locally integrated data term and homogeneous regularisation (CLG). *(d)* 25 % quantile sparsified using the optimal confidence measure. *(e)* Ditto for the gradient-based confidence measure. *(f)* Ditto for the energy-based confidence measure.

The resulting graphs for both confidence measures are depicted in Figure 1. Moreover, a third graph is shown that illustrates the optimal sparsification performance with respect to the average angular error. It serves as theoretical bound for all other confidence measures. As one can see, the proposed energy based criterion performs very favourably. It outperforms the gradient based confidence measure by far. One can also observe that the angular error decreases monotonically under sparsification over the entire range from 100 % down to 1 %. This is a clear indication for an interesting finding that may seem counterintuitive at first glance: *Regions in which the filling-in effect dominates give particularly small angular errors.* At such regions the data term vanishes and only the smoothness term contributes to the local energy. However, this contribution is often very small, since the regulariser allows a smooth extension of the flow field in most cases.
The results also confirm our expectation that $|\nabla f|$ is not necessarily a good confidence measure: Areas with large gradients may represent noise or occlusions, where reliable flow information is difficult to obtain. The filling-in effect, however, may create more reliable information in flat regions by averaging less reliable information from all the surrounding high-gradient regions. The corresponding flow fields with a density of 25 % shown in Figure 2 confirm these considerations. Obviously, only the energy-based criterion allows a realistic sparsification of the computed flow field. The result of the gradient based confidence measure, however, does not coincide very well with the flow field obtained from the optimal sparsification criterion.
Our second experiment investigates the performance of the energy-based confidence measure for a variety of image sequences. To this end we used the spatiotemporal approach with isotropic flow-driven regularisation (TV) and computed flow fields for the following image sequences:

- The *Marble* sequence by Otte and Nagel shown in Figure 3(a)-(b) (`http://i21www.ira.uka.de/image_sequences/`)
- The *Office* sequence by [9] shown in Figure 3(c)-(d) (`http://www.cs.otago.ac/nz/research/vision/`)
- The *Diverging Trees* sequence by Fleet shown in Figure 3(e)-(f)

Figure 4 shows that the application of the energy-based confidence measure improves the estimation significantly in all three cases. In particular at the beginning of the sparsification process a fast decay of the average angular error can be observed. The reason for this behaviour lies in the removal of wrong flow estimates caused by areas with high noise or occlusions. *Due to*

Figure 3: *From left to right, and from top to bottom: (a)* Frame 16 of the 512×512 Marble sequence. *(b)* Magnitude of the ground truth. *(c)* Frame 10 of the 200×200 Office sequence. *(d)* Magnitude of the ground truth. *(e)* Frame 20 of the 150×150 Translating Trees sequence. *(f)* Magnitude of the ground truth.

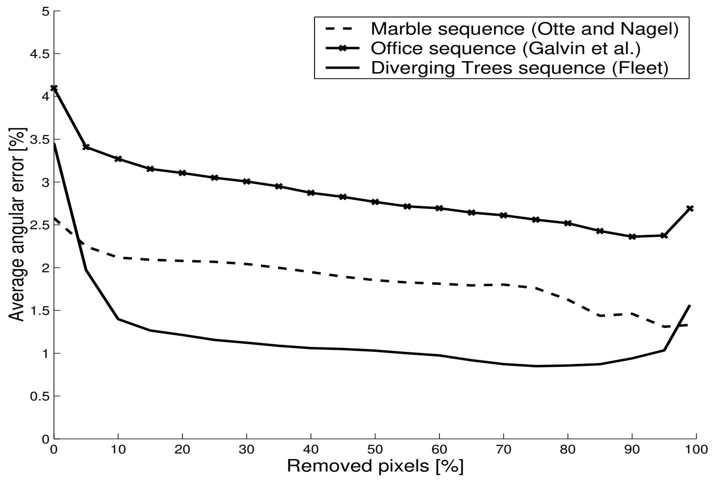

Figure 4: Performance of the energy-based confidence measures for different sequences using the 3-D approach with isotropic flow-driven regularisation (TV).

the massive occurrence of contradictory information in these areas, either the smoothness term or the data term are large. As a consequence, these locations are considered very unreliable by the confidence measure and are already removed at an early stage of the sparsification.

We have seen that the proposed confidence measure based on the evaluation of the local energy contribution performs well for a variety of sequences. Since it is applicable to the entire class of energy minimising optical flow methods, let us now investigate its performance for different variational techniques. To this end we considered all three global approaches introduced in Subsection 2.2 and used our energy-based confidence measure to sparsify the computed flow fields for the *Yosemite* sequence with clouds. The corresponding graphs are presented in Figure 4. As one can see, they show once more an almost monotonic decay of the average angular error under sparsification. In particular, the observed behaviour is independent of the underlying variational approach. This is another confirmation of our findings that the evaluation of the local energy contribution is a simple confidence indicator that is efficient and widely applicable at the same time.

In our final experiment we compare our sparsified flow fields to the best non-dense results from the literature (see also [5]). To this end we use the spatiotemporal approach with locally integrated data term and homogeneous regularisation (CLG) and compute the flow field for the *Yosemite* sequence

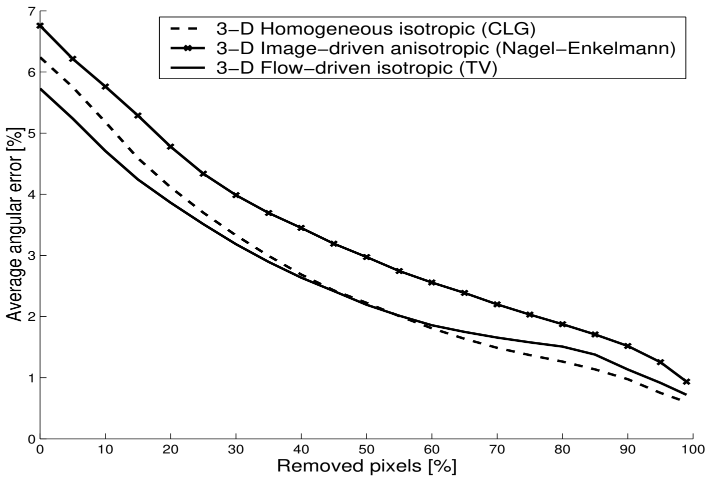

Figure 5: Performance of the energy-based confidence measure for the Yosemite sequence with clouds using different variational approaches.

with clouds. Using the energy-based confidence measure, the obtained flow field is then sparsified in such a way that the reduced densities coincide with the densities of other optic flow methods from the literature. The corresponding average angular errors are presented in Table 1. As one can see, our sparsified flow fields have a significantly lower angular error than all other methods with the same density. In this case errors down to 0.76 ° are reached for a flow density of 2.4 %. To our knowledge, these are the best values that have been obtained by non-dense methods for this sequence in the entire literature.

6 Summary and Conclusion

The absence of good confidence measures is regarded as one of the main drawbacks of variational optic flow methods. The goal of the present paper was to address this problem.

We have seen why the popular gradient-based measure fails: Its rewards high-gradient regions where noise and occlusions dominate, and it ignores the filling-in effect of the regulariser, since it is an "a-priori" measure that does not take into account the estimated flow field.

As a remedy we have proposed a novel energy-based alternative that is both natural and simple: The confidence is chosen to be inversely proportional to

Table 1: Comparison between the "non-dense" results from [1], [20], [16] and our results for the *Yosemite* sequence with cloudy sky (adapted from [5]). AAE = average angular error. CLG = average angular error of the spatiotemporal approach with locally integrated data term an homogeneous regularisation (CLG) with the same density. The sparse flow field has been created using our energy-based confidence criterion. The table shows that using this criterion clearly outperforms all results of non-dense methods.

Technique	Density	AAE	CLG
Singh, step 2, $\lambda_1 \leq 0.1$	97.7 %	10.03°	6.04°
Ong/Spann	89.9 %	5.76°	5.26°
Heeger, level 0	64.2 %	22.82°	3.00°
Weber/Malik	64.2 %	4.31°	3.00°
Horn/Schunck, original, $\lvert\nabla f\rvert \geq 5$	59.6 %	25.33°	2.72°
Ong/Spann, tresholded	58.4 %	4.16°	2.66°
Heeger, combined	44.8 %	15.93°	2.07°
Lucas/Kanade, $\lambda_2 \geq 1.0$	35.1 %	4.28°	1.71°
Fleet/Jepson, $\tau = 2.5$	34.1 %	4.63°	1.67°
Horn/Schunck, modified, $\lvert\nabla f\rvert \geq 5$	32.9 %	5.59°	1.63°
Nagel, $\lvert\nabla f\rvert \geq 5$	32.9 %	6.06°	1.63°
Fleet/Jepson, $\tau = 1.25$	30.6 %	5.28°	1.55°
Heeger, level 1	15.2 %	9.87°	1.15°
Uras *et al.*, $\det(H) \geq 1$	14.7 %	7.55°	1.14°
Singh, step 1, $\lambda_1 \leq 6.5$	11.3 %	12.01°	1.07°
Waxman *et al.*, $\sigma_f = 2.0$	7.4 %	20.05°	0.95°
Heeger, level 2	2.4 %	12.93°	0.76°

the local energy contribution. This measure is applicable to the entire class of energy minimising optic flow techniques and it does not require additional parameters. It puts highest confidence to those locations where the model assumptions are satisfied most.

Our experiments have shown that the energy-based confidence measure performs significantly better than the gradient-based one. It may lead to excellent sparsification results, independently of the image sequence or the underlying variational approach. This is also confirmed by a final comparison to results from the literature, in which our sparsified flow fields proved to be more accurate than those of all other non-dense methods.

Acknowledgements

Our optic flow research is partly funded by the *Deutsche Forschungsgemeinschaft (DFG)* under the project WE 2602/3-1. This is gratefully acknowledged.

References

[1] J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, Feb. 1994.

[2] J. Bigün, G. H. Granlund, and J. Wiklund. Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(8):775–790, Aug. 1991.

[3] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optic flow estimation based on a theory for warping. In T. Pajdla and J. Matas, editors, *Computer Vision – ECCV 2004*, volume 3024 of *Lecture Notes in Computer Science*, pages 25–36. Springer, Berlin, 2004.

[4] A. Bruhn, J. Weickert, and C. Schnörr. Combining the advantages of local and global optic flow methods. In L. Van Gool, editor, *Pattern Recognition*, volume 2449 of *Lecture Notes in Computer Science*, pages 454–462. Springer, Berlin, 2002.

[5] A. Bruhn, J. Weickert, and C. Schnörr. Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. *International Journal of Computer Vision*, 2004. In press.

[6] I. Cohen. Nonlinear variational method for optical flow computation. In *Proc. Eighth Scandinavian Conference on Image Analysis*, volume 1, pages 523–530, Tromsø, Norway, May 1993.

[7] R. Deriche, P. Kornprobst, and G. Aubert. Optical-flow estimation while preserving its discontinuities: a variational approach. In *Proc. Second Asian Conference on Computer Vision*, volume 2, pages 290–295, Singapore, Dec. 1995.

[8] W. Förstner and E. Gülch. A fast operator for detection and precise location of distinct points, corners and centres of circular features. In

Proc. ISPRS Intercommission Conference on Fast Processing of Photogrammetric Data, pages 281–305, Interlaken, Switzerland, June 1987.

[9] B. Galvin, B. McCane, K. Novins, D. Mason, and S. Mills. Recovering motion fields: an analysis of eight optical flow algorithms. In *Proc. 1998 British Machine Vision Conference*, Southampton, England, Sept. 1998.

[10] H. Haußecker and H. Spies. Motion. In B. Jähne, H. Haußecker, and P. Geißler, editors, *Handbook on Computer Vision and Applications. Vol. 2: Signal Processing and Pattern Recognition*, pages 309–396. Academic Press, San Diego, 1999.

[11] B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.

[12] A. Kumar, A. R. Tannenbaum, and G. J. Balas. Optic flow: a curve evolution approach. *IEEE Transactions on Image Processing*, 5(4):598–610, Apr. 1996.

[13] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proc. Seventh International Joint Conference on Artificial Intelligence*, pages 674–679, Vancouver, Canada, Aug. 1981.

[14] H.-H. Nagel. Extending the 'oriented smoothness constraint' into the temporal domain and the estimation of derivatives of optical flow. In O. Faugeras, editor, *Computer Vision – ECCV '90*, volume 427 of *Lecture Notes in Computer Science*, pages 139–148. Springer, Berlin, 1990.

[15] 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.

[16] E. P. Ong and M. Spann. Robust optical flow computation based on least-median-of-squares regression. *International Journal of Computer Vision*, 31(1):51–82, 1999.

[17] A. R. Rao and B. G. Schunck. Computing oriented texture fields. *CVGIP: Graphical Models and Image Processing*, 53:157–185, 1991.

[18] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.

[19] E. P. Simoncelli, E. H. Adelson, and D. J. Heeger. Probability distributions of optical flow. In *Proc. 1991 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 310–315, Maui, HI, June 1991. IEEE Computer Society Press.

[20] J. Weber and J. Malik. Robust computation of optical flow in a multiscale differential framework. *International Journal of Computer Vision*, 14:67–81, 1995.

[21] J. Weickert and C. Schnörr. A theoretical framework for convex regularizers in PDE-based computation of image motion. *International Journal of Computer Vision*, 45(3):245–264, Dec. 2001.

[22] J. Weickert and C. Schnörr. Variational optic flow computation with a spatio-temporal smoothness constraint. *Journal of Mathematical Imaging and Vision*, 14(3):245–255, May 2001.