Is Dense Optic Flow Useful to Compute the Fundamental Matrix?

Markus Mainberger, Andrés Bruhn, and Joachim Weickert

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

Abstract. Estimating the fundamental matrix from a pair of stereo images is one of the central problems in stereo vision. Typically, this estimation is based on a sparse set of point correspondences that has been obtained by a matching of characteristic image features. In this paper, however, we propose a completely different strategy: Motivated by the high precision of recent variational methods for computing the optic flow, we investigate the usefulness of their dense flow fields for recovering the fundamental matrix. To this end, we consider the state-of-the-art optic flow method of Brox et al. (ECCV 2004). Using non-robust and robust estimation techniques for determining the fundamental matrix, we compare the results computed from its dense flow fields to the ones estimated from a RANSAC method that is based on a sparse set of SIFT-matches. Scenarios for both converging and ortho-parallel camera settings are considered. In all cases, the computed results are significantly better than the ones obtained by the RANSAC method – even without the explicit removal of outliers.

1 Introduction

The recovery of depth information from two different views of the same scene is an essential task in computer vision. However, this task cannot be solved without additional information about the position, orientation and internal parameters of the camera system. In this context, the estimation of the fundamental matrix plays a very important role, since it describes the projective-geometric relation between both views in terms of the so-called epipolar constraint [15, 5]. This information not only allows to establish more reliable correspondences between both views, it is also essential to finally set up a projective reconstruction of the original 3-D scene.

First approaches to estimate the fundamental matrix or its Euclidean equivalent – the essential matrix – go back to Longuet-Higgins and his eight point algorithm [13]. This linear method and its variants compute the fundamental matrix from a set of given correspondences via least squares or total least squares fits subject to the epipolar constraint [5]. Some years later, nonlinear methods followed that allowed to interpret deviations from the epipolar constraint in a geometrical sense [15] - e.g. as distances to associated epipolar lines. However, apart from a few exceptions that suggest to estimate the fundamental matrix directly without the explicit use of any point matches [12] all these methods rely exclusively on the use of a *sparse* set of point correspondences

that has been obtained by a previous matching of characteristic image features. Typical features in this context are corners found by the Harris corner detector [8], or, more recently, SIFT features [14] that offer certain invariances with respect to scale, viewpoint and illumination. One main drawback of such feature-based matching approaches is the relatively large number of false matches that are estimated - mainly due to the lack of global contextual information. Thus it is not surprising that this well-known sensitivity of the previously mentioned linear and nonlinear techniques with respect to outliers and noise [21] has triggered an extensive research in robust methods. They include M-estimators [11], case deletion algorithms [4] or even more robust, random sampling techniques such as RANSAC [6, 17, 22] and its improved variant for quasi-degenerated cases QDEGSAC [7].

In view of this development, it is astonishing that mainly the issue of estimating the fundamental matrix from feature-based point correspondences is addressed in the literature [20, 5], but not the question which types of methods are actually most suitable to provide these correspondences. In fact, recent progress in optic flow estimation due to highly accurate variational methods [2, 16] has shown that there is a variety of precise alternatives available for reliably computing correspondences from two images without any prior knowledge - such as the epipolar geometry of the scene. Apart from their high accuracy, variational methods offer at least two additional theoretical advantages compared to traditional feature-based matching techniques: (i) They provide *dense* flow fields with hundreds of thousands of correspondences. In particular with respect to the robustness of the estimation this should be a very important issue. (ii) They do not create strong outliers, since they combine statistically robust data terms with global smoothness assumptions on the solution [1, 2, 16]. This in turn may make sophisticated robust algorithms for estimating the fundamental matrix such as RANSAC obsolete.

The goal of the present paper is to analyse the potential of such dense optical flow fields for estimating the fundamental matrix. To this end, we consider displacement fields that are computed by the recent variational technique of Brox *et al.* [2]. Experiments with respect to both quality and robustness of the estimation show the advantages of such dense optic flow fields: Even if only simple linear methods are used to estimate the fundamental matrix, the results are significantly better than the ones of SIFT-matches using a robust RANSAC approach.

Our paper is organised as follows. In Section 2 we shortly review the required basics of the epipolar geometry and the fundamental matrix. The highly accurate variational method that provides us with a set of correspondences is explained in Section 3. while non-robust and robust techniques to estimate the fundamental matrix from them are discussed in Section 4. Finally, experiments with different camera settings and a comparison to results based on SIFT-matches are presented in Section 5. The paper is concluded with a summary in Section 6.

2 Problem Statement

Let us start by recalling the basic concept of the epipolar constraint [5]. To this end, we consider a pair of stereo images $f_l(\mathbf{x})$ and $f_r(\mathbf{x})$, where subscripts stand for the left and the right image, respectively, and $\mathbf{x} = (x, y)^\top$ denotes the location within a rectangular image domain Ω. Given a point $\mathbf{x}$ in the left image, the epipolar geometry tells us that its

corresponding point $\mathbf{x}'$ in the right image is constrained to lie on the associated epipolar line. This relation between both views is known as the epipolar constraint [5]. It can be written as

$$\hat{\mathbf{x}}'^{\top} F \hat{\mathbf{x}} = 0\,, \tag{1}$$

where $\hat{\mathbf{x}} = (x, y, 1)^{\top}$ and $\hat{\mathbf{x}}' = (x', y', 1)^{\top}$ are projective coordinates of corresponding points, and F is the fundamental matrix [15] - a 3×3 matrix of rank two that is only defined up to a scale. Due to the two additional constraints on F, the epipolar geometry of the scene is described by 7 independent parameters. Therefore, excluding degenerated configurations due to coplanar points, a minimal set of 7 correspondences is required to estimate the fundamental matrix from the epipolar constraint.

In the following we are interested in investigating the usefulness of dense optic flow for estimating these parameters from a given pair of stereo images. To this end, we propose a simple two step strategy: In a first step, we use a recent variational optic flow method to compute a dense displacement field between both views. Since such methods provide exactly one match per pixel, this step will give us a huge set of point correspondences. In a second step, we then make use of these displacements and estimate the actual fundamental matrix from them. Due to the intrinsic robustness of the large amount of correspondences, we only consider standard linear approaches for this task. However, we will also discuss two alternative approaches based on SIFT-matches, RANSAC and LMedS (least median of squares). These feature-based techniques shall serve as references with respect to methods that are frequently used in practice.

3 Dense Variational Optic Flow

First approaches for computing the optic flow with variational methods go back to Horn and Schunck [10] two decades ago. Since then the estimation quality of such techniques improved significantly. Recent variational methods even belong to the most accurate techniques in terms of error measure in the entire literature [2, 16].Such methods compute the dense displacement field $\mathbf{u} = (u, v)^{\top}$ between two images f_l and f_r as minimiser of a suitable energy functional. In general, this energy functional has the form

$$E(\mathbf{u}) = E_D(\mathbf{u}) + \alpha\, E_S(\mathbf{u})\,, \tag{2}$$

where $E_D(\mathbf{u})$ and $E_S(\mathbf{u})$ denote the data and the smoothness term, respectively, and $\alpha > 0$ is a scalar weight that steers the degree of smoothness. While the data term penalises deviations from constancy assumptions – e.g the constancy of the grey value of objects – the smoothness term regularises the often non-unique local solution of the data term by assuming a (piecewise) smoothness of the result.

3.1 The Method of Brox et al.

As our representative for the class of variational methods that provide accurate and dense optic flow fields, let us consider the recent technique of Brox, Bruhn, Papenberg and Weickert (BPPW) [2]. Formulated separately as data and smoothness term, the

energy functional corresponding to the 2-D variant of this technique is given by

$$E_D(\mathbf{u}) = \int_\Omega \psi_D\Big(\left|f_r(\mathbf{x}+\mathbf{u}) - f_l(\mathbf{x})\right|^2 + \gamma \left|\nabla f_r(\mathbf{x}+\mathbf{u}) - \nabla f_l(\mathbf{x})\right|^2 \Big)\, d\mathbf{x} \qquad (3)$$

and

$$E_S(\mathbf{u}) = \int_\Omega \psi_S\left(|\nabla u|^2 + |\nabla v|^2\right)\, d\mathbf{x}\,. \qquad (4)$$

While the first expression in the data term models the assumption that the grey value of objects is constant over time, the second one renders the approach more robust against varying illumination. This is achieved by assuming constancy of the spatial image gradient given by $\nabla f = (f_x, f_y)^\top$. The weighting between the two assumptions is realised with a positive scalar γ. In order to allow for a correct estimation of large displacements, both assumptions are used in their original nonlinear form – all linearisations are postponed to the numerical scheme where they do not compromise the performance [2]. Finally, deviations from both the data and the smoothness term are penalised in a non-quadratic way via a robust function ψ. This improves the performance of the approach with respect to outliers and noise in the case of the data term and preserves motion boundaries by modelling a piecewise smooth flow field in the case of the smoothness term. For both purposes the regularised version of the L_1-norm is used, which, for the smoothness term, comes down to the total variation (TV) regulariser [18] $\psi(s^2) = \sqrt{s^2 + \epsilon^2}$. The regularisation parameter ϵ is set to 10^{-3}.

3.2 Optimisation

The energy functional of the BBPW method is minimised via a multiscale warping strategy as described in [2]. This allows for an accurate estimation of large displacements that are usually present in wide baseline stereo. Moreover, we followed the multigrid framework suggested in [3] in order to speed up the computation of the resulting nonlinear systems of equations. Thus, typical runtimes are in the order of two to three seconds for images of size 640×480 which allow for a usage in practical applications.

4 Fundamental Matrix Estimation

4.1 Least Squares Fit

After we have explained how our sets of correspondences are obtained, let us now discuss the methods we are using to determine the fundamental matrix. Due to the expected robustness of our optic flow results, we restrict ourselves to a simple least squares fit [5].This requires to rewrite the epipolar constraint as

$$0 = \hat{\mathbf{x}}'^\top F \hat{\mathbf{x}} = \begin{pmatrix} x', y', 1 \end{pmatrix} \begin{pmatrix} f_{11}\, f_{12}\, f_{13} \\ f_{21}\, f_{22}\, f_{23} \\ f_{31}\, f_{32}\, f_{33} \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} =: \mathbf{s}^\top \mathbf{f} \qquad (5)$$

where the vector $\mathbf{s} := (xx', yx', x', xy', yy', y', x, y, 1)^\top$ contains all information on the point pair $\hat{\mathbf{x}}'$ and $\hat{\mathbf{x}}$, and the vector $\mathbf{f} := (f_{11}, f_{12}, f_{13}, f_{21}, f_{22}, f_{23}, f_{31}, f_{32}, f_{33})^\top$

represents the fundamental matrix F. Then, using $N \geq 8$ correspondences, the fundamental matrix can be computed by minimizing the least squares fit

$$E(\mathbf{f}) = \sum_{i=1}^{N} \left(\mathbf{s}_i^\top \mathbf{f}\right)^2 , \tag{6}$$

where one of the entries of $\mathbf{f}$ is set to 1 in order to exclude the trivial solution $\mathbf{f} = \mathbf{0}$. Please note that this parameterisation of one of the entries also excludes all fundamental matrices, where the selected entry is 0. Thus all nine possible least squares fits must be considered – each one with a different entry of $\mathbf{f}$ set to 1. Finally, from the results, the normalised fundamental matrix with the smallest residual is selected. In this context, one should note that this approach was chosen, since it gave slightly better results in our experiments than the widely used total least square fit with the $|\mathbf{f}| = 1$ constraint [13]. Furthermore, experiments with convex M-estimators such as the regularized L_1 norm [11] indicated that hardly improvements in terms of accuracy are possible by such types of statistically robust strategies.

4.2 Data Normalisation and Rank-2 Constraint

Since we are interested in using the complete set of correspondences provided by the dense optic flow techniques, we cannot ensure the rank-2 constraint in terms of $\det(F) = 0$ during the minimisation (as is done by many minimal-set-approaches). Instead it has to be enforced afterwards by means of projection, i.e. by zeroing the smallest eigenvalue via a singular value decomposition. The resulting matrix is then given by the closest solution to the estimated matrix with respect to the Frobenius norm. Finally, we also normalised the input data as suggested in [9]. It is well-known that such a step is essential, since it may improve the overall performance of linear methods drastically.

4.3 The SIFT Reference Techniques

In order to compare our method to traditional approaches that are based on a sparse set of point correspondences, we have chosen two techniques that make use of matches from the scale invariant feature transform (SIFT) [14]. The first method that we consider applies the previously discussed least squares fit directly to SIFT-matches and thus allows for a direct comparison to the BBPW method w.r.t. the quality of the underlying correspondences. In the experiments this method will be denoted by SIFT. The second approach additionally usees the robust RANSAC framework [6] in order to remove outliers that occur during the estimation. For this purpose we have chosen a RANSAC variant based on the least median of squares (LMedS) [17]. This technique relates the quality of the fundamental matrix estimated for each drawn subset with the median of its residuals for the remaining samples. It can be summarised in four steps:

1. Draw n subsets of 8 samples.
2. Compute for each subset the fundamental matrix as previously described.
3. Rate each matrix by the median of its residuals.
4. Chose the matrix with the smallest median.

As proposed in [5] the obtained result was only used to identify outliers in the SIFT-correspondences. After these outliers had been removed, the fundamental matrix was then recomputed according to our previous algorithm. In the experiments this method will be referred to as RANSAC-SIFT.

5 Results

Let us now compare the results from our dense optic flow technique to the ones of the two SIFT-reference-methods. To this end, we have considered both an ortho-parallel and a converging camera setting. In order to assess the quality of the different methods, we have computed in all cases the distance $\Delta F_{\mathrm{truth}}$ to the ground truth fundamental matrix according to distance measure presented in [5]. This measure uses one matrix to create a large number of correspondences and the other matrix to establish the corresponding epipolar lines. After the distances between points and lines have been computed, the roles of the two matrices are reversed so that a symmetric measure is obtained that describes an average deviation between the epipolar geometry of both matrices in terms of pixels. Additionally, we evaluated for each method the distance $\Delta F_{\mathrm{intra}}$ of its correspondences to its *own* epipolar lines. Please note that this error measure – in contrast to the first one – does not give information on the quality of the estimated fundamental matrix. However, it gives information on the homogeneity of the point matches, since large values reflect here correspondences that contradict themselves w.r.t. the epipolar geometry of the scene. With respect to the RANSAC approach, we demanded a success probability of 99% and assumed an outlier rate of 40% resulting in about 300 random subsets of samples. The SIFT matches itself have been obtained by using the implementation by David Lowe that is publicly available at the internet address `http://www.cs.ubc.ca/~lowe/keypoints/`.

5.1 Ortho-Parallel Camera Setup

In our fist experiment we have evaluated the quality of our approach for the *Teddy* image pair of the Middlebury stereo benchmark [19]. In Figure 1 the epipolar lines of our approach and the two reference techniques are compared to the ones of the ground truth for both the left and the right frame. As one can see, the lines of our approach and the RANSAC-SIFT method are perfectly in accordance with the ground truth. The pure SIFT method, however, performs quite badly due to outliers in the estimation. These impressions are validated by the corresponding qualitative evaluation of the estimated fundamental matrices that is presented in in Table 1(a). Apart from the SIFT

Table 1. Comparison to the ground truth fundamental matrix for the *Teddy* and the *Javier* stereo pair. Deviations are given in pixels. Parameters have been optimised for $\Delta F_{\mathrm{truth}}$.

(a) Ortho-Parallel Setup: **Teddy**

	SIFT	RANSAC-SIFT	BBPW
$\Delta F_{\mathrm{truth}}$	95.014	0.804	**0.109**
$\Delta F_{\mathrm{intra}}$	3.957	0.123	0.171

(b) Converging Setup: **Javier**

	SIFT	RANSAC-SIFT	BBPW
$\Delta F_{\mathrm{truth}}$	151.915	10.202	**1.441**
$\Delta F_{\mathrm{intra}}$	24.106	0.717	0.890

variant without RANSAC, all methods give reasonable results. However, the fundamental matrices obtained from dense optic flow *without* any robust strategy are much more accurate than the one by the SIFT approach *with* RANSAC. This confirms our considerations that estimation techniques based on dense flow fields may benefit strongly from the large amount of correspondences. Additionally, the small values for $\Delta F_{\mathrm{intra}}$ show that the estimated fundamental matrices from optic flow and RANSAC-SIFT hardly suffer from outliers and thus that in their case a higher accuracy of the estimation is directly related to a better quality of the matches.

5.2 Converging Camera Setup

Let us in our second experiment analyse the quality of the estimation for the case of a converging camera setup. To this end, we have used the *Javier* stereo image pair

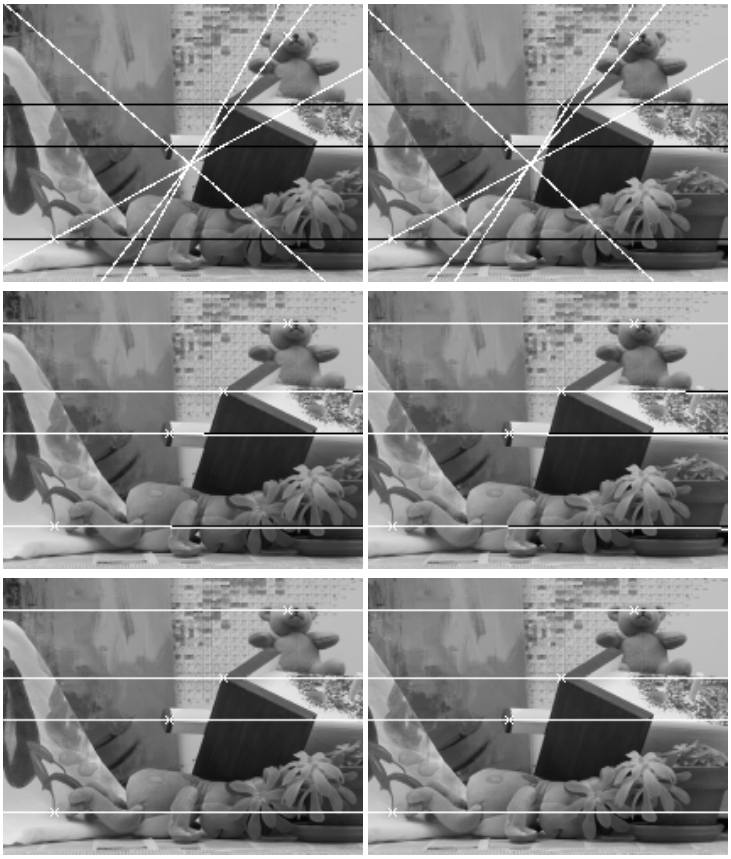

Fig. 1. *Top Row:* Left and right frame of the *Teddy* stereo pair (size 450×375) with epipolar lines shown for the SIFT method (white) and the ground truth (black). *Center Row:* Ditto for the RANSAC-SIFT method (white) and the ground truth (black). *Bottom Row:* Ditto for the BBPW method (white) and the ground truth (black). Parameters have been optimised for $\Delta F_{\mathrm{truth}}$.

that is available together with the ground truth fundamental matrix at the internet address `http://serdis.dis.ulpgc.es/~jsanchez/research/`. The computed results are shown in Figure 2. As in the previous case, the left and the right frame of the stereo pair are presented and epipolar lines are shown that allow for a visual comparison of the estimated fundamental matrices to the ground truth. This time, one can see, that only the epipolar geometry based on the BBPW method is very accurate. This is also validated by the corresponding qualitative comparison in Table 1(b): While the results from SIFT and RANSAC-SIFT are not convincing at all, the estimation from the BBPW correspondences is still very precise. Also the small value for $\Delta F_{\mathrm{intra}}$ confirms that the precision is directly related to the quality of the correspondences and not obtained by a lucky averaging of false matches. Once again, the result for the dense optic flow method has been obtained *without* any robust on-top-strategy.

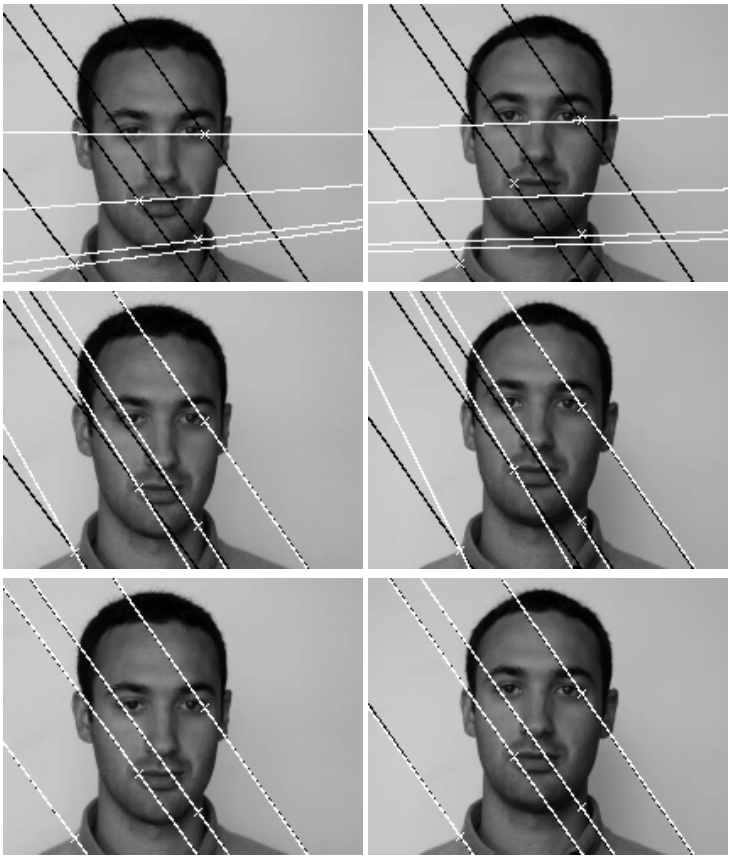

Fig. 2. *Top Row:* Left and right frame of the *Javier* stereo pair (size 640×480) with epipolar lines shown for the SIFT method (white) and the ground truth (black). *Center Row:* Ditto for the RANSAC-SIFT method (white) and the ground truth (black). *Bottom Row:* Ditto for the BBPW method (white) and the ground truth (black). Parameters have been optimised for $\Delta F_{\mathrm{truth}}$.

Table 2. Comparison to the ground truth fundamental matrix for the *Javier* stereo pair for Gaussian noise of standard deviations $\sigma_{\mathrm{n}} = 0, 10, 20$. Deviations w.r.t. $\Delta F_{\mathrm{truth}}$ is given in pixels.

Sensitivity to Noise: **Javier**

Noise	SIFT	RANSAC-SIFT	BBPW
$\sigma_{\mathrm{n}} = 0$	151.915	9.948	**1.441**
$\sigma_{\mathrm{n}} = 10$	92.389	6.913	**2.540**
$\sigma_{\mathrm{n}} = 20$	140.366	7.125	**4.322**

5.3 Sensitivity to Noise

In our last experiment we have investigated the robustness of our approach with respect to noise. To this end, we have added Gaussian noise of zero mean and standard deviation $\sigma_{\mathrm{n}} = 0, 10$ and 20 to the original *Javier* stereo pair with grey values in the range $[0, 255]$ and recomputed the fundamental matrices. The corresponding results are compared in Table 2. Evidently, the optic flow technique is much more robust against noise than the SIFT approach. In particular, the result *with* noise of $\sigma_{\mathrm{n}} = 20$ based on the BBPW flow field is still four times more accurate then the result of the RANSAC-SIFT approach *without* noise. This, however, is not surprising: On one hand, the estimation based on global optic flow techniques benefits twice from an averaging process – (i) via the smoothness term and (ii) via the least squares fit using a huge set of correspondences. On the other hand, in the presence of noise the SIFT approach produces rather delocalisation errors with respect to all matches than real outliers. Thus in most cases, robust extensions such as RANSAC cannot overcome the poorer quality of the underlying correspondences. Even a presmoothing of the noisy images by convolution with a Gaussian kernel did not give better results: This, in turn, can be explained by the fact that SIFT-features are based on a scale-space approach anyway and thus cannot benefit from additional smoothing procedures.

6 Summary and Conclusions

Traditional approaches for estimating the fundamental matrix use a sparse set of point correspondences, and many efforts are spent to find clever strategies for selecting the most reliable matches. Our paper shows that there is an interesting alternative: Modern dense optic flow methods have reached such a high degree of accuracy and robustness that one can incorporate "unscrupulously" the matches for *all* pixels. In this way, even a simple linear estimation strategy outperforms fairly sophisticated methods such as a robust LMedS-RANSAC method based on SIFT correspondences. Moreover, the large number of correspondences renders the dense approach intrinsically insensitive to noise. It is evident that the idea of using dense flow fields for computing the fundamental matrix is not limited to linear estimation techniques only: It can be extended in a straightforward way to any other type of method that is originally based on a sparse set of correspondences. Investigating this potential is a topic of our ongoing research.

It is our hope that this paper helps to reconsider some traditional preferences of sparse over dense approaches: Incorporating dense variational methods into geometric computer vision strategies is an area that is still widely unexplored. However, since it combines two very advanced methodologies, great progress can be expected.

References

1. M. J. Black and P. Anandan. Robust dynamic motion estimation over time. In *Proc. 1991 IEEE Conference on Computer Vision and Pattern Recognition*, pages 292–302, Maui, HI, June 1991. IEEE Computer Society Press.
2. 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.
3. A. Bruhn, J. Weickert, T. Kohlberger, and C. Schnörr. A multigrid platform for real-time motion computation with discontinuity-preserving variational methods. *International Journal of Computer Vision*, 70(3):257–277, Dec. 2006.
4. S. Chaterjee and A. Hadi. *Sensitivity Analysis in Linear Regression*. Wiley, New York, 1988.
5. O. Faugeras, Q.-T. Luong, and T. Papadopoulo. *The Geometry of Multiple Images*. MIT Press, Cambridge, MA, 2001.
6. M. Fischler and R. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24:381–385, 1981.
7. J.-M. Frahm and M. Pollefeys. RANSAC for (quasi-) degenerate data (QDEGSAC). In *Proc. 2006 IEEE Conference on Computer Vision and Pattern Recognition*, pages 453–460, New York, NY, June 2006. IEEE Computer Society Press.
8. C. G. Harris and M. Stephens. A combined corner and edge detector. In *Proc. Fourth Alvey Vision Conference*, pages 147–152, Manchester, England, Aug. 1988.
9. R. Hartley. In defense of the eight-point algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19:580–593, 1997.
10. B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
11. P. J. Huber. *Robust Statistics*. Wiley, New York, 1981.
12. S. Lehmann, A. P. Bradley, V. L. Clarkson, J. Williams, and P. J. Kootsookos. Correspondence free determination of the affine fundamental matrix. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(1):82–97, 2007.
13. H. C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293:133–135, Sept. 1981.
14. D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
15. Q.-T. Luong and O. D. Faugeras. The fundamental matrix: theory, algorithms, and stability analysis. *International Journal of Computer Vision*, 17(1):43–75, Jan. 1996.
16. T. Nir, R. Kimmel, and A. Bruckstein. Over-parameterized variational optical flow. *International Journal of Computer Vision*, 76(2):205–216, 2008.
17. P. Rousseeuw and A. Leroy. *Robust Regression and Outlier Detection*. Wiley, 1987.
18. L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
19. D. Scharstein and R. Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision*, 47(1-3):7–42, 2002.
20. P. Torr and D. Murray. The development and comparison of robust methods for estimating the fundamental matrix. *International Journal of Computer Vision*, 24(3):271–300, 1997.
21. J. Weng, T. Huang, and N. Ahuja. Motion and structure from two perspecitve views: algorithms, error analysis and error estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(5):451–476, 1989.
22. Z. Zhang, R. Deriche, O. Faugeras, and Q.-T. Luong. A robust technique for matching two uncalibrated images through the recovery of the unknown epipolar geometry. *Artificial Intelligence*, 78:87–119, 1995.

Errata for the paper

Is Dense Optical Flow Useful to Compute The Fundamental Matrix?

On page 9, Table 2 The entries of the RANSAC-SIFT reference technique in Table 2 are not correct. The updated table with the correct error values is presented below.

Table 2. Comparison to the ground truth fundamental matrix for the *Javier* stereo pair for Gaussian noise of standard deviations $\sigma_{\mathrm{n}} = 0, 10, 20$. Deviations w.r.t. $\Delta F_{\mathrm{truth}}$ is given in pixels.

Sensitivity to Noise: **Javier**

Noise	SIFT	RANSAC-SIFT	BBPW
$\sigma_{\mathrm{n}} = 0$	151.915	10.202	**1.441**
$\sigma_{\mathrm{n}} = 10$	92.389	12.095	**2.540**
$\sigma_{\mathrm{n}} = 20$	140.366	16.012	**4.322**