

OPTICAL FLOW ESTIMATION WITH CONFIDENCE MEASURES FOR SUPER-RESOLUTION BASED ON RECURSIVE ROBUST TOTAL LEAST SQUARES

Tobias Schuchert and Fabian Oser
Fraunhofer IOSB, Fraunhoferstr. 1, Karlsruhe, Germany

Keywords: Optical flow, Motion estimation, Super resolution, Confidence measures.

Abstract: In this paper we propose a novel optical flow estimation method accompanied by confidence measures. Our main goal is fast and highly accurate motion estimation in regions where information is available and a confidence measure which identifies these regions. Therefore we extend the structure tensor method to robust recursive total least squares (RRTLS) and run it on a GPU for real-time processing. Based on a coarse-to-fine framework we propagate not only the motion estimates to finer scales but also the covariance matrices, which may be used as confidence measures. Experiments on synthetic data show the benefits of our approach. We applied the RRTLS framework to a real-time super-resolution method for deforming objects which incorporates the confidence measures and demonstrates that propagating the covariances through the pyramid improves super-resolution results.

1 INTRODUCTION

Many methods for motion estimation between two or more images have been proposed in the last years. In 2007 the Middlebury benchmark (Baker et al., 2011) was proposed, where to date about 50 motion estimation methods have been compared (including implementations of the well known optical flow methods of (Horn and Schunck, 1981), and (Lucas and Kanade, 1981)). The benchmark consists of several synthetic and real world sequences and handles many different use cases. Also different masks for discontinuities and untextured regions are available and used for comparing the algorithms for these special problems. For more details on optical flow methods we refer to the surveys in (Barron et al., 1994) and (Weickert et al., 2006).

Nevertheless there exist applications where it is important to know, where an estimation is reliable and where not. In other applications a confidence measure may help to improve further processing steps. For example (1) a precise estimation of unknown motions, e.g. cell growth, may be heavily affected by a smoothness constraint, as small important motions may be over-smoothed or (2) estimation of super-resolution images from image sequences based on motion estimates may include the reliability of the motion estimates to weight the different image and motion pa-

rameters.

In this paper we present a total least squares (TLS) framework for optical flow estimation. In order to cope with discontinuities and large motions we developed a robust recursive TLS in a coarse-to-fine framework. We estimate motion and the reliability of these motions on each pyramid step and propagate the estimation results and the confidence measures through the different scales.

The Middlebury benchmark (Baker et al., 2011) gives an overview of current state-of-the art motion estimation techniques. Most of these techniques are based on variational approaches which yield full flow fields and high accuracy. Moreover there exist efficient parallel implementation frameworks for variational methods making them real-time efficient on current GPUs¹ (Werlberger et al., 2009).

Local methods consider only neighbouring image regions and are therefore even more suited for parallelization. However, only few local methods are present on the Middlebury benchmark website. Mostly as reference methods and with unsatisfying results, although they proved to yield high accuracy (cmp. (Haussecker and Spies, 1999)). Recently (Senst et al., 2011) proposed a local method for optical flow estimation and feature tracking implemented on a

¹Graphics Processing Unit

GPU. This method is based on ordinary least squares and does not provide confidence measures.

To the best of our knowledge none of the methods on the Middlebury website yields special reliability measures for motion estimates. The publicly available version of (Werlberger et al., 2009) offers the possibility to calculate a geometric reliability measurement for each estimate. There the flow field is compared in both directions for inconsistency and the result is mapped to a probability distribution. Another method is to compute the inverse of the variational energy in a local region, to identify regions, where the energy is still large (Bruhn and Weickert, 2006). There exist a number of confidence measures for local methods, e.g. spatial coherence or corner measure (cmp. (Haussecker and Spies, 1999)). (Kondermann, 2009) compared different confidence measures for local and global methods, but did not find a satisfying one. Therefore Kondermann proposed two confidence measures based on motion statistics from sample data and a hypothesis test, which yield superior results compared to the two methods mentioned before. These measures are applicable to all optical flow fields (local/global) afterwards. However, these measures must be trained before and are therefore no alternatives in our case as learning motion statistics is expensive and training data must be available.

Based on the structure tensor approach (Foerster and Guelch, 1987) we developed a novel robust recursive total least squares framework, which allows accurate motion estimation in structured regions and computing the covariance matrix of each estimate. This framework is related to the robust total least squares method presented by (Bab-Hadiashar and Suter, 1998), but compared to their method we (1) combine the structure tensor approach with a robust function and do not use LMSOD or LMedS for outlier detection, which speeds up computation without loss of accuracy, (2) use a coarse-to-fine framework to handle large motions, and (3) use the covariance matrix based on the approximation of (Nestares et al., 2000) as a reliable confidence measure. Therefore we propagate not only the motion estimates but also the covariance matrices through the different pyramid steps, similar to a Kalman filter approach (cmp. (Simoncelli, 1999)). In experiments we demonstrate the applicability of our approach.

In the first section of this paper the optical flow estimation framework is presented. This includes how the structure tensor approach is combined with a robust implementation and how it is integrated in a coarse-to-fine framework based on recursive total least squares (TLS). Then in Sec. 3 computation and distribution of covariance matrices on different scales

of the pyramid are shown. Experiments in Sec. 4 show the accuracy of the optical flow estimator on synthetic sinusoidal sequences and the Middlebury benchmark and is compared to the method of (Werlberger et al., 2009), one of the fastest global estimator, which is still ranked high in the Middlebury benchmark. Section 5 shows application of the new framework for super-resolution. We show how the covariance is incorporated into the super-resolution framework improving the results on real images. A summary and an outlook on future work follows in Sec. 6.

2 ESTIMATION FRAMEWORK

We present an extended version of the well known structure tensor method (Foerster and Guelch, 1987) for optical flow estimation. On the one hand we integrate a robust function to handle outliers following (Black and Anandan, 1996) and (Schuchert, 2010) and on the other hand we include the renormalization technique of (Kanatani, 1996) to avoid systematic bias following (Chojnacki et al., 2001). This modified structure tensor method is then embedded in a coarse-to-fine strategy based on recursive total least squares (cmp. (Boley et al., 1996)). In the following we explain these steps in more detail.

The structure tensor approach is based on the brightness constancy constraint equation (BCCE)

$$g_x u_x + g_y u_y + g_t = \mathbf{g}^T \mathbf{p} = 0 \quad (1)$$

with image gradients $\mathbf{g} = (g_x, g_y, g_t)^T$, optical flow $\mathbf{u} = (u_x, u_y)^T$ and parameter vector $\mathbf{p} = (p_1, p_2, p_3)^T$. The parameter vector and the optical flow vector are related by $\mathbf{u} = 1/p_3(p_1, p_2)^T$. It is assumed that solution vector $\tilde{\mathbf{p}}$ approximately solves all constraint equations in the local neighbourhood Λ with size N and therefore $\mathbf{g}_i^T \mathbf{p}$ with $i \in [1, N]$ only approximately equals 0. We get

$$\mathbf{g}_i^T \mathbf{p} = e_i \quad \forall i \in \{1, \dots, N\} \quad (2)$$

with errors e_i which have to be minimized by the sought solution $\tilde{\mathbf{p}}$. We minimize \mathbf{e} in a weighted 2-norm

$$\begin{aligned} \mathbf{p} &= \arg \min \|\mathbf{e}\|_2^2, \text{ with} \\ \|\mathbf{e}\|_2^2 &= \sum_{i=1}^N \left(\frac{\mathbf{p}^T \mathbf{g}_i w_i \mathbf{g}_i^T \mathbf{p}}{\mathbf{p}^T \mathbf{p}} \right) =: \mathbf{p}^T \mathbf{J} \mathbf{p} \end{aligned} \quad (3)$$

where w_i is the weight for the i -th constraint and where matrix \mathbf{J} is called structure tensor. In general Gaussian weights are used to reduce the influence of constraints far away from the center pixel. The additional constraint $|\mathbf{p}| = 1$ is introduced to avoid the trivial solution $\mathbf{p} = \mathbf{0}$. An eigenvalue analysis yields the

minimum solution, i.e. the eigenvector to the smallest eigenvalue of \mathbf{J} .

In order to reduce the influence of outliers in the solution we introduce a robust function ρ into the estimation

$$\begin{aligned} \mathbf{p} &= \arg \min \|\mathbf{e}\|_\rho, \text{ with} \\ \|\mathbf{e}\|_\rho &= \sum_{i=1}^N \rho(\ell_i, \sigma) = \sum_{i=1}^N \rho(\mathbf{g}_i^T \mathbf{p}, \sigma) \end{aligned} \quad (4)$$

with a variable parameter σ . As robust function we choose the Lorentzian

$$\rho_{Lor}(x, \sigma) = \log \left(1 + \frac{1}{2} \left(\frac{x}{\sigma} \right)^2 \right). \quad (5)$$

where σ is adapted following (Black and Anandan, 1996). We iteratively solve the structure tensor method and update the weights according to (4).

(Chojnacki et al., 2001) presented a new derivation of the renormalization method of Kanatani. We use this method in our robust total least squares framework as follows. Starting from the standard TLS estimation $\mathbf{p}_{TLS} = \mathbf{p}_0$ we iteratively solve the generalized eigenvalue problem

$$\mathbf{M}_{\mathbf{p}_k} \xi = \lambda \mathbf{N}_{\mathbf{p}_k} \xi \quad (6)$$

with $\mathbf{p}_{k+1} = \xi_{\min}$ the smallest eigenvector to the smallest eigenvalue λ_{\min} . The matrices are given by

$$\mathbf{M}_{\mathbf{p}_k} = \sum_{i=1}^N \frac{\mathbf{A}_i}{\mathbf{p}_k^T \mathbf{C}_i \mathbf{p}_k} \quad (7)$$

and

$$\mathbf{N}_{\mathbf{p}_k} = \sum_{i=1}^N \frac{\mathbf{C}_i}{\mathbf{p}_k^T \mathbf{C}_i \mathbf{p}_k}, \quad (8)$$

where \mathbf{A} is the unsmoothed structure tensor, i.e. \mathbf{J} with $w_i = 1 \forall i \in \{1, \dots, N\}$, and \mathbf{C} is the covariance matrix of the image gradients. The iteration is stopped, when $|\mathbf{p}_k - \mathbf{p}_{k-1}| < \varepsilon$ where ε is a small fixed threshold. In case of a translational motion model the covariance matrix \mathbf{C}_i is constant in a neighbourhood, so the denominator can be excluded from the sum, which simplifies computation. For further details we refer to (Chojnacki et al., 2001).

For estimation of large distances a coarse-to-fine framework based on Gaussian pyramids is used. In order to propagate the result through the scales we use recursive TLS following (Boley et al., 1996). Then the structure tensor on the current scale l is achieved by

$$\mathbf{J}_l = \hat{\mathbf{J}}_l + \beta \check{\mathbf{J}}_{l-1} \quad (9)$$

with recursive weighting factor β . $\hat{\mathbf{J}}_l$ is the standard structure tensor of the current scale, but with warped temporal gradients g_t . These temporal gradients are computed following (Simoncelli, 1999), i.e. the filters

are sheared in the direction of the flow. The temporal derivatives g_t then include the motion estimate from the previous scale and therefore $\check{\mathbf{J}}_{l-1}$ contains only the 2×2 spatial structure tensor of the previous scale. The last row and column of $\check{\mathbf{J}}_{l-1}$ are set to zeros.

3 COVARIANCE ESTIMATION

Following (Nestares et al., 2000) and (Schuchert et al., 2010) we approximate the error covariance matrices of the optical flow estimates by the inverse of the Hessian, given by

$$\mathbf{H} = \frac{\gamma}{\sigma_n^2 \|\mathbf{p}\|^2} \left(\mathbf{S} - \frac{1}{\|\mathbf{p}\|^2} (\mathbf{p}^T \mathbf{J} \mathbf{p}) \mathbf{I} \right) \quad (10)$$

where \mathbf{S} is given by the 2×2 spatial structure tensor, \mathbf{I} is an identity matrix of the same size as \mathbf{S} . The parameter γ is proportional to the signal to noise ratio. We set this parameter for all our experiments to $\gamma = 1$.

The Covariance matrix is propagated through the different scales analogue to the structure tensor using the recursive technique proposed in (9).

4 EXPERIMENTS

We evaluated our algorithm (*RRTLS*) on synthetic sinusoidal sequences and on the Middlebury benchmark (Baker et al., 2011). We used 3×3 Scharr filters (Scharr, 2005) for spatial gradient estimation. The image pyramid starts by sizes of 20 pixel on the smaller side and we used a downsampling factor of 0.85 for all sequences. The recursive weighting factor is set to $\beta = 0.4$ for both, the structure tensor and the covariance matrix.

We compare our results with the Anisotropic-Huber-L1 method proposed in (Werlberger et al., 2009). This algorithm is among the best and fastest algorithms on the Middlebury benchmark and is named *AHLI* in the remainder of this paper. *AHLI* is based on a variational framework using a smoothness term in order to estimate motion also in regions where no or not much data is available. The method is publicly available and offers the possibility to calculate the geometric confidence measure (cmp. Sec. 1). We use the proposed high accuracy parameters from (Werlberger et al., 2009) and leave all other parameters in the standard configuration for all sequences. We use two frames for experiments with *AHLI* and three images for *RRTLS* wherever possible. (Werlberger et al., 2009) stated that the two frames version of their estimator yields better results, on the Middlebury benchmark. However *RRTLS* produced better results, es-

pecially more reliable covariance estimations, if three frames are used.

4.1 Sinusoidal Sequences

In order to evaluate the performance of our *RRTLS* and its covariance estimation, we used two sinusoidal patterns which move contrary to each other with a non structured background and a non structured rectangle in the middle (cmp. Fig. 1a). The colour coded ground truth motion is displayed in Fig. 1b and the corresponding colour code is shown in Fig. 1g. The magnitude of the motion is below one pixel. Nevertheless both algorithms are using the coarse-to-fine framework. The estimates of *AHLI* and of *RRTLS* are shown in Fig. 1c and e, respectively.

The effect of the smoothness constraint in *AHLI* can be clearly seen as motion is also estimated in regions, where neither structure nor motion is present. *RRTLS* estimates no or only small motion in regions where no motion is present, however at borders the estimates are heavily degraded. Figure 1h shows the estimated trace of the covariances from *RRTLS* coded with a grey scale. White means low and black high confidence.

The regions where estimates are highly unreliable are clearly visible, i.e. white and grey. Also the *AHLI* yields a geometric confidence measure. Figures 1d and f show 50 percent of the best motion estimates masked by the corresponding confidence measures of *AHLI* and *RRTLS*. The geometric confidence measure of *AHLI* masked some regions, where no structure is available, but not all. Moreover It removes regions on the patterns, where estimates should be reliable. Our proposed method masks unreliable estimates much better and keeps most of the regions on the patterns.

We also tested other confidence measures for structure tensor methods, proposed in (Haussecker and Spies, 1999), where the corner measures c_c yields best results. The corner measure c_c is the difference between the total coherence measure and the spatial coherence measure and is computed from the eigenvalues of the structure tensor method.

Table 1 summarizes average endpoint errors (cmp. (Baker et al., 2011)) for 50% masks by the trace of the propagated covariance, the trace of the covariance only on the finest grid, i.e. not propagated, and the corner measure. Errors are calculated for different noise standard deviations σ_n . The propagated covariance shows to be clearly the most reliable confidence measure.

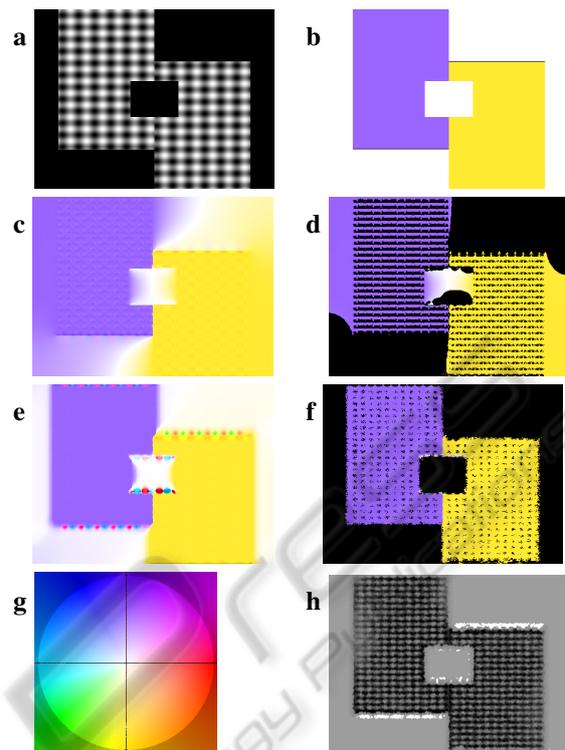


Figure 1: Sinusoidal pattern data with noise and estimated flow and confidence measure. **a**: one image of the input sequence, **b**: is the true flow (colour coded), **c**: flow estimates of *AHLI* and **d**: masked flow estimates of *AHLI* using their confidence mask. **e** and **f** are flow estimates of *RRTLS* and the masked results with the *RRTLS* confidence measure. **g** and **h** show the colour code and the estimated covariance grey scale coded of our algorithm, respectively.

Table 1: Comparison of different confidence measures. Average endpoint error for the 50% most reliable estimates of *RRTLS* on the sinusoidal pattern sequence.

Confidence measure	$\sigma_n = 0$	$\sigma_n = 1$	$\sigma_n = 5$
propagated covariance	0.008	0.024	0.075
covariance	0.012	0.179	0.212
corner measure (c_c)	0.032	0.161	0.208

4.2 Middlebury Benchmark

We evaluated our algorithm on the well known Middlebury benchmark (Baker et al., 2011). We are not interested in achieving overall best results on this benchmark, but to have best results in regions where the confidence of our measurement is high, i.e. where the trace of the estimated covariance matrix is low.

Figure 2 shows the input image of the Urban3 sequence (Fig. 2a), the colour coded ground truth flow (Fig. 2d) and the estimated flows of *AHLI* (Fig. 2c) and *RRTLS* (Fig. 2e). In Fig. 2d and f only the best 50% of the estimates according to their confidence

Table 2: Average Endpoint Error on the Middlebury data set using standard masks (disc and untext) and masks based on confidence measures. Results in brackets are for the 2 frame *RRTLS*, because only 2 frames are available.

Algorithm	Mask	Dim	Grove2	Grove3	Hyd	RW	Urban2	Urban3	Venus
<i>AHLI</i>	all	0.175	0.161	0.641	0.202	0.155	0.389	0.721	0.332
<i>RRTLS</i> (sym)	all	(0.375)	0.207	0.751	0.449	0.271	1.159	0.707	(0.931)
<i>AHLI</i>	disc	0.269	0.473	1.222	0.502	0.528	1.734	2.298	1.098
<i>RRTLS</i> (sym)	disc	(0.537)	0.562	1.494	1.053	0.758	2.358	2.797	(2.132)
<i>AHLI</i>	untext	0.200	0.138	0.631	0.129	0.137	0.402	0.778	0.350
<i>RRTLS</i> (sym)	untext	(0.438)	0.201	0.637	0.293	0.257	1.210	0.704	(1.215)
<i>AHLI</i>	95%	0.174	0.123	0.537	0.176	0.126	0.246	0.612	0.278
<i>RRTLS</i> (sym)	95%	(0.352)	0.168	0.647	0.405	0.223	0.977	0.660	(0.855)
<i>AHLI</i>	75%	0.176	0.088	0.316	0.115	0.093	0.185	0.403	0.243
<i>RRTLS</i> (sym)	75%	(0.262)	0.124	0.337	0.268	0.157	0.643	0.360	(0.599)
<i>AHLI</i>	50%	0.175	0.068	0.171	0.066	0.084	0.138	0.244	0.237
<i>RRTLS</i> (sym)	50%	(0.209)	0.105	0.193	0.147	0.125	0.594	0.168	(0.339)

measures are displayed. The same effects as on the sinusoidal sequence can be seen here. *AHLI* over-smooths motion borders, whereas *RRTLS* yields highly erroneous estimates in some parts of the image, e.g. lower left border. Nevertheless the covariance of these estimates is high and can be masked out, whereas wrong estimates of *AHLI* are not recognized, e.g. the over-smoothing in the lower bottom left part of the image.

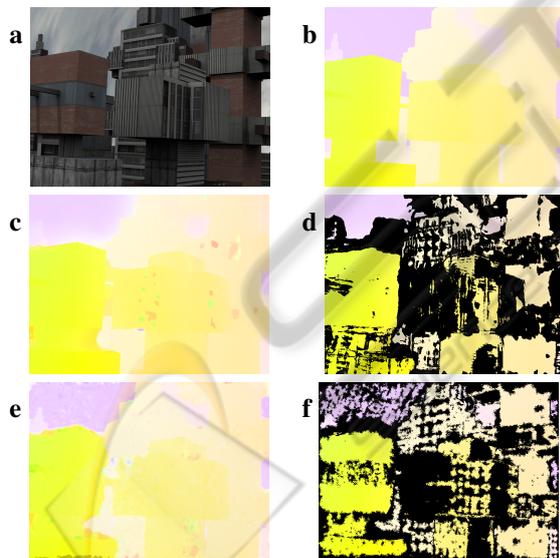


Figure 2: One sequence of the Middlebury benchmark. **a**: one input image, **b**: true flow (colour coded), **c** and **e**: flow estimates *AHLI* and *RRTLS* respectively. **d** and **f**: masked estimates using the corresponding confidence masks.

Table 2 shows results on the other Middlebury benchmarks for both estimators. Additional to the masks from Middlebury, i.e. disc (discontinuities) and untext (untextured regions), the masks filter 90%, 75% and 50% of the best estimates according to the

confidence measure. It can be clearly seen, that *AHLI* outperforms our method in almost all cases. Using confidence masks, our method achieves on almost all sequences considerable better results and the errors get near or even below the results of *AHLI*. This demonstrates the reliability of our covariance measure and that our estimator can compete with *AHLI* in accuracy, if only 50% of the best estimates are considered. Moreover our method is more than 3 times faster than *AHLI* on images of size 640×480 pixel.

5 APPLICATION

A reliable confidence measure of the motion estimation may be used to improve further processing steps. Super-resolution from multiple images needs reliable motion estimation. Most approaches assume constant or affine motions over the whole image. These assumptions do not hold for deformable surfaces, i.e. faces. Moreover we have the problem, that not all facial images in a sequences are achieved from front view. Some super-resolution techniques therefore filter incoming images before computation in order to reject images with large motions or achieved from different view points.

In order to cope with these problems, we integrated the *RRTLS*-algorithm into a super-resolution framework, based on the multi-frame super-resolution technique proposed by (Farsiu et al., 2003). There atmospheric blur is neglected and a blur-warp-model² is

²Farsiu starts with a warp-blur-model but translates it into a blur-warp-model because of a translational motion assumption. This is not the case for deformable surfaces like faces. However (Wang and Qi, 2004) showed, that when motion is estimated on low-resolution images, the warp-blur-model has a systematic error and the blur-warp-model may be used.

used

$$\mathbf{y}_k = \mathbf{D}\mathbf{F}_k\mathbf{H}\mathbf{x} + \mathbf{v}_k \quad k \in [1, N], \quad (11)$$

where \mathbf{x} is the high resolution image of size $[rM_1 \times rM_2]$, \mathbf{y}_k are the k low resolution images of size $[M_1 \times M_2]$ and \mathbf{v}_k is the additive systematic noise. The Blurring operator \mathbf{H} is the Point-Spread-Function of the camera (here a 5×5 Gaussian filter), the warping-operator \mathbf{F}_k is modelled by a $[r^2M_1M_2 \times r^2M_1M_2]$ matrix and the downsampling-factor \mathbf{D} by a $[M_1M_2 \times r^2M_1M_2]$ matrix. Using a L^p -norm we have to find the minimum of

$$\hat{\mathbf{x}} = \arg \min_x \left[\sum_{k=1}^N \|\mathbf{D}\mathbf{F}_k\mathbf{H}\mathbf{x} - \mathbf{y}_k\|_p^p \right]. \quad (12)$$

Farsiu estimates the minimum solution in two steps. First the low-pass filtered high resolution image $\mathbf{z} = \mathbf{H}\mathbf{x}$ is determined by data fusion of the low-resolution images based on motion estimates. The low-resolution images are registered on the high resolution grid and then filtered by a median for each high resolution pixel. In the second step the image \mathbf{z} is deblurred. We use the estimated covariance of *RRTLS* in both steps. In step one the data is fused based on the estimated motion. The median filter is realized as a weighted median based on the trace of the covariance estimates. In the second step the trace of the covariance also influences the deblurring of the high resolution image when the single high resolution pixels are weighted. For further details on the super-resolution algorithm we refer to (Farsiu et al., 2003).

Figure 3 shows estimation results for a sequence of facial images. The low resolution images (e.g. Fig. 3a) are computed from the high resolution images (Fig. 3b) by downsampling, smoothing and addition of noise. The sequence consists of 9 facial images with small head movements and motion of the lips (speaking) and 58 images showing parts of the head or background only. All low-resolution images are referenced on the low-resolution image of Fig. 3e. The estimated flow for one low resolution image with *RRTLS* is depicted in Fig. 3c and the trace of the corresponding covariance estimates in Fig. 3d. Again white defines high and black low covariances. Figures 3e and f show the estimated high resolution images with and without using the covariance information. In Fig. 3g and h the difference error between the estimated super-resolution images and the original high resolution image is shown for both algorithms. The difference error is scaled by a factor of 6. Differences appear mostly in regions where larger motions are present, i.e. at the sides of the head and in the region around the mouth. Using covariance information for super-resolution highly improves the estimation (cmp. Fig. 3g).

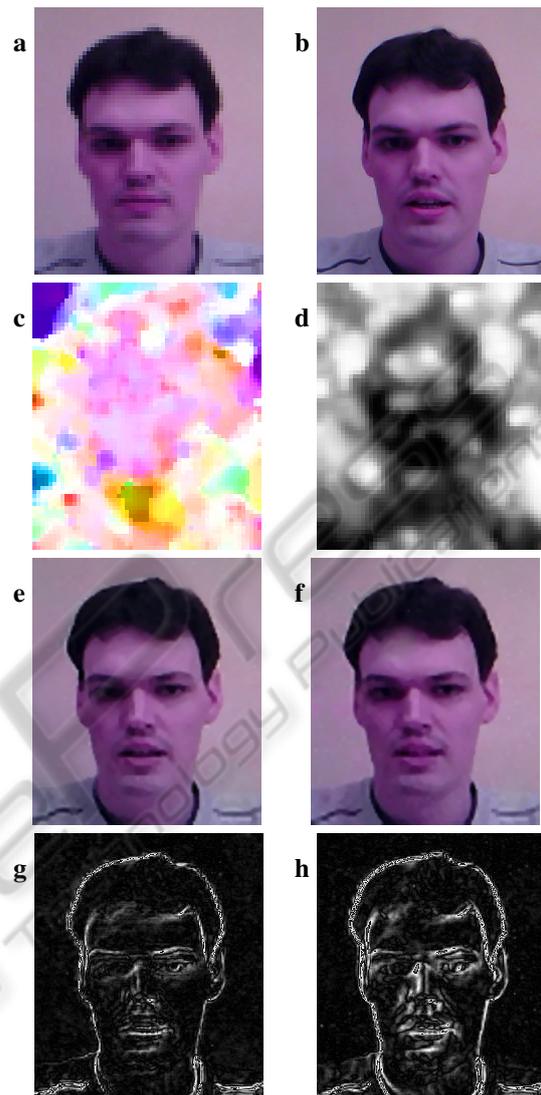


Figure 3: Results for super-resolution. **a**: One low resolution input image, **b**: high resolution image, **c**: estimated flow with *RRTLS*, **d**: estimated covariance grey value coded, **e** and **f**: high resolution images estimated by our algorithm with and without using the covariance, respectively, and **g** and **h**: difference images between computed super-resolution and high resolution image.

6 SUMMARY

In this paper we presented a novel method for motion estimation with confidence measures implemented on GPU. The method is based on robust total-least squares, embedded in a coarse-to-fine framework using a recursive algorithm. We have shown that the estimated covariance matrices which are propagated through the whole pyramid are well suited as confi-

dence measure (e.g. the trace of the covariance matrix). Experiments on synthetic sequences showed the good performance of this approach.

We used the estimated motion and confidence measures to improve a super-resolution method. However, motion estimates are still not as accurate as best estimators in the Middlebury benchmark. In future we plan for improvements for large motions, more complex motion models, e.g. affine motion models, and motions under changing illumination, e.g. by considering gradient information or estimation of illumination changes.

REFERENCES

- Bab-Hadiashar, A. and Suter, D. (1998). Robust optic flow computation. *International Journal of Computer Vision*, 29(1):59–77.
- Baker, S., Scharstein, D., Lewis, J., Roth, S., Black, M., and Szeliski, R. (2011). A database and evaluation methodology for optical flow. *International Journal of Computer Vision*, pages 1–31. 10.1007/s11263-010-0390-2.
- Barron, J. L., Fleet, D. J., and Beauchemin, S. S. (1994). Performance of optical flow techniques. *Int. Journal of Computer Vision*, 12:43–77.
- Black, M. and Anandan, P. (1996). The robust estimation of multiple motions: parametric and piecewise-smooth flow fields. In *Computer Vision and Image Understanding*, volume 63(1), pages 75–104.
- Boley, D., Steinmetz, E., and Sutherland, K. (1996). Recursive total least squares: An alternative to using the discrete kalman filter in robot navigation. In *Reasoning with Uncertainty in Robotics*, volume 1093 of *Lecture Notes in Computer Science*, pages 221–234. Springer Berlin / Heidelberg. 10.1007/BFb0013963.
- Bruhn, A. and Weickert, J. (2006). A confidence measure for variational optic flow methods. In *Properties for Incomplete data*, pages 283–298. Springer.
- Chojnacki, W., Brooks, M. J., and Hengel, A. V. D. (2001). Rationalising the renormalisation method of kanatani. *Journal of Mathematical Imaging and Vision*, 14:21–38.
- Farsiu, S., Robinson, D., Elad, M., and Milanfar, P. (2003). Fast and robust multi-frame super-resolution. *IEEE Transactions on Image Processing*, 13:1327–1344.
- Foerstner, W. and Guelch, E. (1987). A fast operator for detection and precise location of distinct points, corners and centers of circular features. In *Proceedings of the ISPRS Intercommission Workshop on Fast Processing of Photogrammetric Data*, pages 281–305.
- Haussecker, H. and Spies, H. (1999). Motion. In *Handbook of Computer Vision and Applications*, volume 2, chapter 13, pages 309–396. Academic Press, 1 edition.
- Horn, B. and Schunck, B. (1981). Determining optical flow. *Artificial Intelligence*, 17:185–204.
- Kanatani, K. (1996). *Statistical Optimization for Geometric Computation: Theory and Practice*. Elsevier.
- Kondermann, C. (2009). *Postprocessing and Restoration of Optical Flows*. PhD thesis, Heidelberg University.
- Lucas, B. D. and Kanade, T. (1981). An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th international joint conference on Artificial intelligence - Volume 2*, pages 674–679, San Francisco, CA, USA.
- Nestares, O., Fleet, D. J., and Heeger, D. J. (2000). Likelihood functions and confidence bounds for Total Least Squares estimation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'2000)*.
- Scharr, H. (2005). Optimal filters for extended optical flow. In *Complex Motion, 1. Int. Workshop, Günzburg, Oct. 2004*, volume 3417 of *Lecture Notes in Computer Science*, Berlin. Springer Verlag.
- Schuchert, T. (2010). *Plant leaf motion estimation using a 5D affine optical flow model*. PhD thesis, RWTH Aachen University.
- Schuchert, T., Aach, T., and Scharr, H. (2010). Range flow in varying illumination: Algorithms and comparisons. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32:1646–1658.
- Senst, T., Eiselein, V., Evangelio, R. H., and Sikora, T. (2011). Robust modified l2 local optical flow estimation and feature tracking. In *IEEE Workshop on Motion and Video Computing*, pages 685–690, Kona, USA.
- Simoncelli, E. P. (1999). Bayesian multi-scale differential optical flow. In *Handbook of Computer Vision and Applications*, volume 2, chapter 14, pages 397–422. Academic Press, San Diego.
- Wang, Z. and Qi, F. (2004). On ambiguities in super-resolution modeling. *Signal Processing Letters, IEEE*, 11(8):678–681.
- Weickert, J., Bruhn, A., Brox, T., and Papenberg, N. (2006). A survey on variational optic flow methods for small displacements. In *Mathematical Models for Registration and Applications to Medical Imaging*, volume 10 of *Mathematics in Industry*, pages 103–136. Springer Berlin Heidelberg.
- Werlberger, M., Trobin, W., Pock, T., Wedel, A., Cremers, D., and Bischof, H. (2009). Anisotropic huber-l1 optical flow. In *Proceedings of the British Machine Vision Conference (BMVC)*, London, UK.