Illumination-Robust Dense Optical Flow Using Census Signatures

Thomas Müller[1,2], Clemens Rabe[1], Jens Rannacher[1],
Uwe Franke[1], Rudolf Mester[2]

[1]Daimler Research, Sindelfingen
[2]Goethe-Universität Frankfurt am Main

Abstract. Vision-based motion perception builds primarily on the concept of optical flow. Modern optical flow approaches suffer from several shortcomings, especially in real, non-ideal scenarios such as traffic scenes. Non-constant illumination conditions in consecutive frames of the input image sequence are among these shortcomings. We propose and evaluate the application of intrinsically illumination-invariant census transforms within a dense state-of-the-art variational optical flow computation scheme. Our technique improves robustness against illumination changes, caused either by altering physical illumination or camera parameter adjustments. Since census signatures can be implemented quite efficiently, the resulting optical flow fields can be computed in real-time.

1 Introduction

1.1 Dense Optical Flow in Real Scenes

Reliable motion estimation in real-time is a key task for a variety of applications, e.g., in robotics or automotive driver assistance. While many new variants of dense variational optical flow algorithms have been proposed in recent years, they have focused mainly on accuracy under ideal conditions—which is benchmarked on the Middlebury optical flow data set [1]—rather than on robustness in practical applications.

Algorithms for optical flow estimation in real scenes under non-ideal conditions still suffer from the following issues in particular: non-constant illumination conditions in consecutive frames, large displacements (some improvements have been made here in recent years, e.g., [2]), weakly textured areas, and model violations such as transparency or reflections. In this work, we focus on the first topic and provide a new optical flow technique to cope with non-constant illumination conditions, caused either by physical illumination changes (including shadows) or by unanticipated and unknown adjustments of the camera parameters (e.g. the exposure time).

1.2 Related Work

Since most variation-based optical flow algorithms exploit the brightness constancy constraint in consecutive frames, they are not stable under changing

illumination conditions. During the last three decades, since Horn and Schunck [8] introduced their dense optical flow model, several approaches have been proposed to overcome this illumination sensitivity.

One might first think of a simple global mean-and-variance equalization to overcome the illumination sensibility. Although this can help remove global illumination offsets and can compensate the gain of a global transition function, the approach has some drawbacks. First, the device-dependent transition functions which map the incident light energy at the sensor elements to electronic signals are often non-linear. Second, changes of the mean intensity of the image can also result from real changes of the scene rather than the illumination (think of a big dark truck moving into the scene in front of the observer) and a mean-and-variance equalization would lead to incorrect results. Last, even a more generalized histogram adaptation approach could not cope with local changes of the physical illumination of a scene.

In [14], based on the improved model from [15], the authors propose the preprocessing of the original images with an ROF[10] denoising scheme (computed according to [4]) and take the difference to the original images as the new input. This structure/texture decomposition leads to slightly better results than the application of a simple Gauss-based high-pass band filter, which also yields some illumination-change resistance and is computationally more efficient. The shortcoming of an ROF denoising based high-pass filtering of the input images is two-fold: even when a pyramid scheme is used, the structure-texture decomposition leads to problems with larger displacements and an ROF denoising is computationally still quite expensive.

A qualitatively different approach is proposed in [5], where an additional scalar function is estimated together with the optical flow field in a joint optimization process. This function is then expected to cover all illumination inconsistencies. Since this function must be very smooth on the image domain, good results require many iterations. This eliminated the real-time capability in our implementation. An interesting and completely different approach for color images was presented in [9], where the authors use the constancy of a set of photometric invariances from color space in a variational scheme. The use of an advanced data term for illumination robustness in variational optical flow is mentioned in [12], where the normalized cross correlation is used as residual and leads to robustness against multiplicative illumination changes.

2 Dense Motion Estimation

Given two image functions $I_{\{1,2\}} : \Omega \to \mathbb{R}^+$ on the image domain $\Omega \subset \mathbb{R}^2$, the optical flow is defined as the apparent motion of the pixels from I_1 to I_2. Neglecting transparency or reflections, the projected motion of the objects of the real world onto the two-dimensional image plane is an element of the set of possible optical flow fields (ambiguities in the optical flow arise when there are textureless areas in the image sequence).

A reasonable optical flow field $\boldsymbol{u} : \Omega \to \mathbb{R}^2$ is received by solving

$$\arg\min_{\boldsymbol{u}} \left\{ \lambda E_{\mathrm{D}}\left[\boldsymbol{u}\right] + E_{\mathrm{S}}\left[\boldsymbol{u}\right] \right\} , \tag{1}$$

with a m.a.p. expectation maximization in mind. The data term

$$E_{\mathrm{D}}\left[\boldsymbol{u}\right] = \int_{\Omega} \Psi\left(\rho\left(\boldsymbol{x}, \boldsymbol{u}\left(\boldsymbol{x}\right)\right)\right) \mathrm{d}\boldsymbol{x} \tag{2}$$

consists of a norm $\Psi : \mathbb{R} \to \mathbb{R}^+$ and the residual function $\rho : \Omega \times \mathbb{R}^2 \to \mathbb{R}$. The exact form of ρ can vary and determines the behavior under illumination changes as we will show in the next section. The regularizing smoothness term E_{S} helps to provide the most probable solution for Eq. (1), given our model conception in the exact form of E_{S}. Since, per definition, the regularizing smoothness term is independent of the input data $I_{\{1,2\}}$, it is not in the remainder of this paper.

2.1 A General Numerical Solution Scheme

Following the algorithm proposed in [15], we use the coupling term $E_{\mathrm{C}}\left[\boldsymbol{u}, \boldsymbol{v}\right] = (1/2\theta) \cdot \int_{\Omega} \left(\boldsymbol{u}\left(\boldsymbol{x}\right) - \boldsymbol{v}\left(\boldsymbol{x}\right)\right)^2 \mathrm{d}\boldsymbol{x}$ with the coupling constant θ to separate the optical flow functional Eq. (1) into two parts, which are then solved iteratively. Given the result of the previous computation step in $\boldsymbol{v}$, the first part

$$\arg\min_{\boldsymbol{u}} \left\{ \lambda E_{\mathrm{D}}\left[\boldsymbol{u}\right] + E_{\mathrm{C}}\left[\boldsymbol{u}, \boldsymbol{v}\right] \right\} \tag{3}$$

contains the data term and can be solved pointwise. Having the result of (3) in $\boldsymbol{u}$, the second part

$$\arg\min_{\boldsymbol{v}} \left\{ E_{\mathrm{C}}\left[\boldsymbol{u}, \boldsymbol{v}\right] + E_{\mathrm{S}}\left[\boldsymbol{v}\right] \right\} \tag{4}$$

contains the regularization and is solved depending on the exact form of E_{S}. The iterative solution of Eqs. (3) and (4) is performed until the desired accuracy is achieved or a fixed number of iterations have been executed. As proposed in [14], this iterative scheme is combined with median filtering for robustness and a pyramid scheme to cope with larger displacements.

2.1.1 Solving the Data Part for Arbitrary Residuals Taking a closer look at the solution of the data part Eq. (3) allows describing a general straightforward gradient-descent solution scheme. This can be applied later for special forms of ρ and Ψ, especially for illumination-robust ones. Linearizing the residual ρ around the start value $\boldsymbol{v}$ in the second argument yields

$$\rho\left(\boldsymbol{x}, \boldsymbol{u}\right) \approx \tilde{\rho}\left(\boldsymbol{x}, \boldsymbol{u}\right) = \rho\left(\boldsymbol{x}, \boldsymbol{v}\right) + \boldsymbol{\nabla}^{\top} \rho\left(\boldsymbol{x}, \boldsymbol{v}\right) \cdot \left(\boldsymbol{u} - \boldsymbol{v}\right) . \tag{5}$$

With the notation $\left(u_1, u_2\right)^{\top} \equiv \boldsymbol{u}\left(\boldsymbol{x}\right)$ and solving for $i \in \{1, 2\}$

$$0 = \frac{\partial}{\partial u_i} \left\{ \lambda E_{\mathrm{D}}\left[\boldsymbol{u}\right]\left(\boldsymbol{x}\right) + E_{\mathrm{C}}\left[\boldsymbol{u}, \boldsymbol{v}\right]\left(\boldsymbol{x}\right) \right\}$$

establishes the gradient descent step. Noting $\boldsymbol{u} \equiv \boldsymbol{u}(\boldsymbol{x}), \boldsymbol{v} \equiv \boldsymbol{v}(\boldsymbol{x}), (\rho_1, \rho_2)^\top \equiv \boldsymbol{\nabla}\rho(\boldsymbol{x}, \boldsymbol{v})$ leads to

$$0 = \lambda f \cdot \tilde{\rho}(\boldsymbol{x}, \boldsymbol{u}) \cdot \rho_i + \frac{1}{\theta}(u_i - v_i) \tag{6}$$

with the factor f depending on the exact form of the norm Ψ and computed by $f \equiv \Psi'(\rho(\boldsymbol{x}, \boldsymbol{v}))/\rho(\boldsymbol{x}, \boldsymbol{v})$. Since Eq. (6) is linear in u_i, it is possible to formulate a linear equation system $\mathbf{A} \cdot \boldsymbol{u} = \boldsymbol{b}$ with

$$\mathbf{A} = \begin{pmatrix} \frac{1}{\theta} + \lambda f \rho_1^2 & \lambda f \rho_1 \rho_2 \\ \lambda f \rho_1 \rho_2 & \frac{1}{\theta} + \lambda f \rho_2^2 \end{pmatrix}$$

and

$$\boldsymbol{b} = \frac{1}{\theta}\boldsymbol{v} - \lambda f \cdot \boldsymbol{\nabla}\rho(\boldsymbol{x}, \boldsymbol{v}) \cdot R$$

with $R = \rho(\boldsymbol{x}, \boldsymbol{v}) - \boldsymbol{\nabla}^\top \rho(\boldsymbol{x}, \boldsymbol{v}) \cdot \boldsymbol{v}$ and solve it with respect to $\boldsymbol{u}$ by a standard algorithm.

2.1.2 Special Data Terms in the Literature In the classical approach by Horn and Schunck [8], we have $\rho(\boldsymbol{x}, \boldsymbol{u}) = I_2(\boldsymbol{x} + \boldsymbol{u}) - I_1(\boldsymbol{x})$, the well-known brightness constancy constraint $\tilde{\rho}(\boldsymbol{x}, \boldsymbol{u}) = \boldsymbol{\nabla} I(\boldsymbol{x}) \cdot \boldsymbol{u} + I_t(\boldsymbol{x})$ (with $I_t(\boldsymbol{x}) = I_2(\boldsymbol{x}) - I_1(\boldsymbol{x})$), $\Psi(\rho) = \rho^2$ and $f = 2$. Approximating $\Psi(\rho) = |\rho|$ with $\Psi \approx \sqrt{\rho^2 + \epsilon}, \epsilon \ll 1$ in the outlier-robust TV-L^1 model by Zach and Pock [15] leads to $f = 1/\Psi(\rho)$. Many other models known from literature (e.g. the photometric invariant model in [9] or the dense Lucas-Kanade approach in [3]) can also be solved by this scheme. The following, we will use $\Psi = \sqrt{\rho^2 + \epsilon}, \epsilon \ll 1$ for our own approach.

3 Illumination-Invariant Motion Estimation

3.1 Illumination Robustness with Local Compensation of the Mean

The numerical framework sketched in the previous section used the residual function ρ which in [8] or [15] corresponds to the grey-value constancy. Using more sophisticated residuals, such as

$$\rho(\boldsymbol{x}, \boldsymbol{u}) = \sum_{i=1}^{n} \left| I_2(\boldsymbol{x} + \boldsymbol{y}_i + \boldsymbol{u}) - \overline{I}_2(\boldsymbol{x} + \boldsymbol{u}) - I_1(\boldsymbol{x} + \boldsymbol{y}_i) + \overline{I}_1(\boldsymbol{x}) \right|^d \tag{7}$$

with $\overline{I}_{\{1,2\}}(\boldsymbol{x}) = 1/n \cdot \sum_{i=1}^{n} I_{\{1,2\}}(\boldsymbol{x} + \boldsymbol{y}_i)$ as the mean grey value and the fixed list of points $(\boldsymbol{y}_1, \boldsymbol{y}_2, \ldots, \boldsymbol{y}_n)^\top \in \mathbb{R}^{n \times 2}, n \in \mathbb{N}$ around the origin, is a first step towards illumination robustness. The pointwise compensation by the local mean leads to an invariance of the results against local offsets of the illumination. In some cases, this is already sufficient to establish an effective protection against erroneous flow vectors due to the change of illumination conditions. The application of the residual in Eq. (7) is similar to a high pass filtering, which is performed

in [14] by a structure-texture decomposition based on the ROF model. The special cases are the zero mean sum of absolute differences (ZSAD) for $d = 1$ and the zero mean squared sum of differences (ZSSD) for $d = 2$, which can both be easily implemented using the numerical scheme presented in the previous section and which are evaluated in the experimental part (Section 4).

3.2 Illumination Model

Before presenting more illumination-robust variants of the residual ρ, first a closer look at the illumination process. The image intensity field $I : \Omega \to \mathbb{R}^+$ of our input images at a point in time is the result of this process. It can be modeled by

$$I(\boldsymbol{x}) = T(J(c(\boldsymbol{x}), \boldsymbol{x})) \tag{8}$$

with the intrinsic, physical color field $c : \Omega \to \mathbb{R}^+$, the physical illumination (which here includes all atmospheric effects, noise etc.) $J : \mathbb{R}^+ \times \Omega \to \mathbb{R}^+$ and the device-specific generally time-dependent and unknown but always monotonic transition function $T : \mathbb{R}^+ \to \mathbb{R}^+$. The physical illumination J varies slowly in the space domain (apart from shadow and object edges) and it is monotonic in its first argument for a fixed point $\boldsymbol{x}$. This leads to the conclusion that data terms of the optical flow functional not relying on the exact grey-values or ratios, but rather on their ordering, lead to the most robust results under non-constant illumination conditions. In other words, though T and J are widely arbitrary and time-dependent, apart from shadow and object edges, their monotonicity can be relied upon.

3.3 Illumination-Invariant Census Based Residuals

How can we profit from the monotonicity of the functions T and J from Eq. (8) which constitute the illumination process? It is necessary to look at the ordering of the grey-values rather than their exact values, differences or ratios. Exactly this is provided by census transforms. Census signatures for sparse optical flow computation have already been used in [11], where they were used to efficiently compute large displacements.

The census transform maps to every pixel $\boldsymbol{x}$ of the image plane Ω one signature vector, $\boldsymbol{s}_t : \Omega \to \{0,1\}^n, n \in \mathbb{N}, t \in \{1,2\}$, which is defined as

$$s_{t,i}(\boldsymbol{x}) = [I_t(\boldsymbol{x} + \boldsymbol{y}_i) - I_t(\boldsymbol{x}) + \epsilon \geq 0], \tag{9}$$

with $i \in [1,n]$ given the fixed list of points $(\boldsymbol{y}_1, \boldsymbol{y}_2, \ldots, \boldsymbol{y}_n)^\top \in \mathbb{R}^{n \times 2}$ near the origin and a small $\epsilon \in \mathbb{R}$. The brackets $[X]$ in Eq. (9) indicate whether the statement X is true ($[X] = 1$, the pixel at $\boldsymbol{x} + \boldsymbol{y}_i$ is clearly brighter than the pixel at $\boldsymbol{x}$) or false ($[X] = 0$, the pixel at $\boldsymbol{x} + \boldsymbol{y}_i$ is darker than or similar to the pixel at $\boldsymbol{x}$). Our implementation uses points in the direct neighborhood of the origin. Depending on the chosen patch size r, the first r rectangles of pixels around the origin are considered, so that $n = (2r+1)^2 - 1$. Our evaluation used a 3×3 ($n = 8$) and a 5×5 ($n = 24$) patch size variant of the census transform.

In the classical case, the residual ρ is based on the grey value difference between two corresponding points of the consecutive images. We will now introduce a residual function which is based on the similarity of the census signatures of the two corresponding points. We propose to use the Hamming distance

$$\rho(\boldsymbol{x}, \boldsymbol{u}) = \sum_{i=1}^{n} \left[s_{2,i}(\boldsymbol{x} + \boldsymbol{u}) \neq s_{1,i}(\boldsymbol{x}) \right] \tag{10}$$

which is zero if and only if $\boldsymbol{s}_1(\boldsymbol{x}) = \boldsymbol{s}_2(\boldsymbol{x} + \boldsymbol{u})$ and which is n if every component of $\boldsymbol{s}_1$ is different to corresponding component of $\boldsymbol{s}_2$. The Hamming distance based on the signature vectors has the nice property that it is inherently invariant under arbitrary changes of the device specific transition function T (since T is monotonic), and it is also invariant under most changes of the physical illumination J (since J is piecewise slowly varying on the space domain), if the point does not lie on an illumination edge; the illumination process does not affect the value of the Hamming distance $\rho(\boldsymbol{x}, \boldsymbol{u})$.

The best choice for the parameter ϵ in Eq. (9) depends on the noise level of the image. Note that the residual function ρ has to be smooth in the second argument (the displacement $\boldsymbol{u}$) to be used in the iterative gradient descent scheme described in Sec. 2.1. While this turns out to be also the case for the Hamming distance in Eq. (10) (we use linear interpolation of the grey values at non-integer image positions for s_2), the smoothness of ρ is further promoted by the small offset ϵ when dealing with image noise. A ternary variant, where equality between the values $I_t(\boldsymbol{x} + \boldsymbol{y}_i)$ and $I_t(\boldsymbol{x})$ is treated as a third case, provides more information and can also be applied.

4 Results

4.1 Evaluation of a Synthetic Scene from the Middlebury Data Set

Testing the illumination-robustness on the `Grove2` sequence from the Middlebury optical flow data set [1], for which the ground-truth optical flow $\boldsymbol{u}_{\mathrm{GT}} : \Omega \to \mathbb{R}^2$ is available, provides a reproducible quantitative evaluation of the method proposed here. Considering the optical flow field between the frames 10 and 11 with the corresponding images I_{10} and I_{11} and varying the γ value of I_{11} from $\gamma = 1$ to $\gamma = 4$ results in the modified images $I_{11,\gamma}$ with $I_{11,\gamma}(\boldsymbol{x}) = 255 \cdot (I_{11}(\boldsymbol{x})/255)^{\gamma}$. For a variety of optical flow methods, the average endpoint error of the flow field, $1/|\Omega| \cdot \int_{\Omega} |\boldsymbol{u}(\boldsymbol{x}) - \boldsymbol{u}_{\mathrm{GT}}(\boldsymbol{x})| \, \mathrm{d}\boldsymbol{x}$, is computed for every γ in the range $[1, 4]$.

Though it can be seen in the table of Fig. 1 that the census-based approaches are slightly outperformed by the grey-value-based TV-L^1 approach from [13] for $\gamma = 1$ (no illumination change), the graph on the left shows that when increasing the γ value, one approach after another—apart from the census based methods—becomes unstable and fails to provide correct flow results. The illumination robustness of the census based methods is only affected by discretization effects of the γ correction.

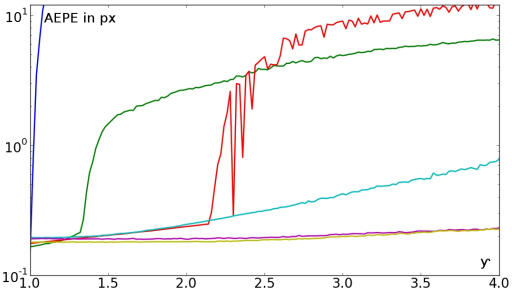

AEPE in px	$\gamma = 1$	$\gamma = 2$	$\gamma = 4$
TV-L^1	0.150	89.6	106
TV-L^1+S/T	0.165	2.69	6.46
Illumination term	0.175	0.236	13.4
ZSAD	0.193	0.246	0.780
3×3 Census	0.191	0.189	0.231
5×5 Census	0.179	0.180	0.227

Fig. 1. Evaluation of the `Grove2` sequence from the Middlebury optical flow dataset from frame 10 to 11. Frame 10 is γ-corrected with $1 \leq \gamma \leq 4$. Left: Average endpoint error, blue: TV-L^1 optical flow without structure/texture decomposition; green: TV-L^1 with structure/texture decomposition, both from [14]; red: optical flow with joint illumination term estimation [5]; cyan: ZSAD based optical flow; magenta: 3×3 census based optical flow; yellow: 5×5 census based optical flow. Right: Exemplary quantitative results.

4.2 Real Scenes

4.2.1 Example: Buggy Scene Two different real image sequences are qualitatively explored using the proposed methods. In the first example, we again perform an artificial gamma adjustment on two input images of a city scene (first image $\gamma = 0.9$, second $\gamma = 0.8$) and compare the flow results of the several algorithms. The results in Fig. 2 are now descussed in detail.

In the second row of Fig. 2, the result of the TV-L^1 approach from [14], applied on the two differently γ-adjusted input images (top row), is shown. Clearly, the illumination difference between the two input images makes it impossible for the grey-value based algorithm to provide correct flow fields. In contrast, the structure/texture decomposition on the right hand side allows for illumination robustness in this case. The joint computation of an additional illumination term, already mentioned in Sec. 1.2 from [5], also solves the problem (third row, right) but leads to a loss of details in the flow field. Note for example the fast moving foot of the woman in the right part of the image which is merged with the environment, or consider the absence of the umbrella.

There are two main aspects when reviewing the census result: First, even without structure/texture precomputation or computation of an additional illumination term, both variants of census based optical flow are indeed inherently illumination-invariant. Second, the resulting flow fields are more detailed than the ones received from approaches known from literature: the fast moving foot can be clearly observed when using census signatures. In addition, the underestimation of the flow vectors on the turning car known from the approaches above vanishes almost completely. Note that census signatures are patch-based and are therefore suffering from one major drawback: the loss of resolution in the result. As seen in Fig. 2, a patch structure is clearly visible in the resulting optical flow field.

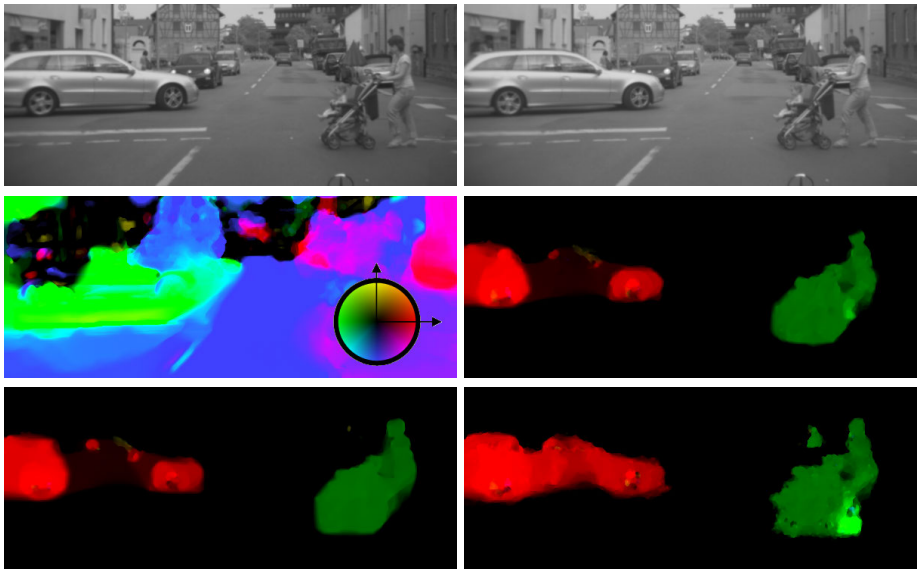

Fig. 2. Comparison of flow results of a typical traffic scene. Top row: The two differently gamma adjusted input images (note that the right one has a slightly lower gamma value). Second row: TV-L^1 optical flow from [14] without (left, color encoding: color indicates direction, intensity indicates magnitude) and with structure/texture decomposition (right). Third row, left: image intensity based optical flow with joint illumination term computation from [5]. Right: binary 3×3 census based optical flow (left) without illumination term or structure/texture decomposition.

4.2.2 Example: Highway Scene We also explore the results on a second traffic scene taken on a highway [6] which is very challenging for flow estimation. In addition to illumination changes, large flow vectors, large weakly textured regions, and aliasing effects occur. As a consequence, even with constant illumination between two frames, standard optical flow methods (see Fig. 3, third row, left) only yield partially reasonable flow fields (note that we would expect a rainbow-like structure on the highway surface). During the sequence, the camera system is exposed to very different illumination conditions which change abruptly. This causes the device-internal transfer function to change as well in a wide range. Results of several optical flow methods are presented in Fig. 3. No artificial gamma adjustment is performed for this sequence.

We were not able to receive reasonable flow results with standard approaches like the TV-L^1 from [14] with (third row, left) or without structure/texture decomposition (second row, right). However, the joint estimation of an illumination field [5] (third row, right) was able to detect the moving car on the left and the car in front of the camera which is correctly colored black, since it is nearly constantly moving with the same velocity in the focus of expansion. Both census

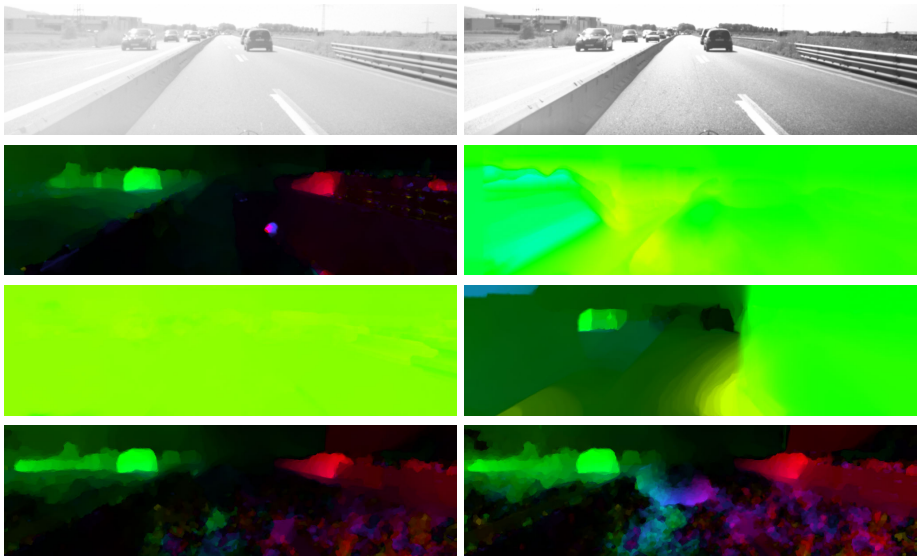

Fig. 3. Comparison of flow results from a very challenging highway scene [6], frame 108 to 109. Top row: The two differently exposed input images (note the different illumination). Second row, left: TV-L^1 flow from [14] without structure/texture decomposition of the previous frame (107 to 108), without illumination changes for reference; right: TV-L^1 optical flow from [14] without structure/texture decomposition. Third row: TV-L^1 optical flow with structure/texture decomposition from [14] (left), image intensity based flow with joint illumination term computation from [5] (right). Bottom row: binary census based optical flow (left), ternary census optical flow (right), both from Sec. 3. Color encoding as in Fig. 2.

approaches (bottom row) are apparently able to detect the correct motion, if the corresponding image region is sufficiently textured.

5 Conclusion

This work proposed a new technique for optical flow computation coping with illumination changes often occurring in real scenes such as traffic scenarios. A variety of different residual functions can be used in a general numerical solution scheme for optical flow. Profiting from the monotonicity of the illumination process, illumination-invariant residuals can be constructed when considering the grey value ordering rather than their exact differences. Hamming distances, based on several different census signatures, replace the simple grey value constancy constraint from classical optical flow approaches. In the second part, results from the proposed novel technique are compared to those based on work known from literature. The application of census signatures in dense variational optical flow leads to the best results both on synthetic and real traffic scenes during illumination changes. While census signatures may have their shortcom-

ings when used as single correspondence measures, they generally seem to provide very robust results when applied in a regularized global optimization scheme (see also the application of census transforms in stereo vision [7]). Future work will combine the census based optical flow computation method with large displacement support using feature correspondences and extended regularizing models to make it possible to apply dense motion estimation in real scenarios, e.g. in traffic scenes for driver assistance systems.

References

1. Baker, S., Scharstein, D., Lewis, J., Roth, S., Black, M., Szeliski, R.: A database and evaluation methodology for optical flow. International Journal of Computer Vision 92(1), 1–31 (2011)
2. Brox, T., Bregler, C., Malik, J.: Large displacement optical flow. In: IEEE International Conference on Computer Vision and Pattern Recognition. Miami Beach, Florida, USA (2009)
3. Bruhn, A., Weickert, J., Schnörr, C.: Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. International Journal of Computer Vision 61(3), 211–231 (2005)
4. Chambolle, A.: An algorithm for total variation minimization and applications. Journal of Mathematical Imaging and Vision 20(1–2), 89–97 (2004)
5. Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. http://hal.archives-ouvertes.fr/hal-00490826/en/ (2010)
6. Friedrich, H., Rabe, C., Mester, R.: DAGM 2011 AVCC website. http://www.dagm2011.org/adverse-vision-conditions-challenge.html (2011)
7. Hirschmüller, H., Gehrig, S.: Stereo matching in the presence of sub-pixel calibration errors. In: IEEE Conference on Computer Vision and Pattern Recognition (2009)
8. Horn, B.K.P., Schunck, B.G.: Determining optical flow. In: Artificial Intelligence. vol. 17, pp. 185–203 (1981)
9. Mileva, Y., Bruhn, A., Weickert, J.: Illumination-robust variational optical flow with photometric invariants. In: DAGM-Symposium. pp. 152–162 (2007)
10. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. Physica D 60, 259–268 (1992)
11. Stein, F.: Efficient computation of optical flow using the census transform. In: DAGM-Symposium (2004)
12. Steinbrücker, F., Pock, T., Cremers, D.: Advanced data terms for variational optic flow estimation. In: Vision, Modelling, and Visualization Workshop. Braunschweig, Germany (2009)
13. Wedel, A., Meissner, A., Rabe, C., Franke, U., Cremers, D.: Detection and segmentation of independently moving objects from dense scene flow. In: Energy Minimization Methods in Computer Vision and Pattern Recognition (2009)
14. Wedel, A., Pock, T., Zach, C., Bischof, H., Cremers, D.: Statistical and Geometrical Approaches to Visual Motion Analysis, chap. An Improved Algorithm for TV-L1 Optical Flow, pp. 23–45. Springer-Verlag Berlin (2009)
15. Zach, C., Pock, T., Bischof, H.: A duality based approach for realtime TV-L1 optical flow. In: DAGM-Symposium (2007)