Illumination-Robust Variational Optical Flow with Photometric Invariants

Yana Mileva, Andrés Bruhn, and Joachim Weickert

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

Abstract. Since years variational methods belong to the most accurate techniques for computing the optical flow in image sequences. However, if based on the grey value constancy assumption only, such techniques are not robust enough to cope with typical illumination changes in real-world data. In our paper we tackle this problem in two ways: First we discuss different photometric invariants for the design of illumination-robust variational optical flow methods. These invariants are based on colour information and include such concepts as spherical/conical transforms, normalisation strategies and the differentiation of logarithms. Secondly, we embed them into a suitable multichannel generalisation of the highly accurate variational optical flow technique of Brox *et al.* This in turn allows us to access the true potential of such invariants for estimating the optical flow. Experiments with synthetic and real-world data demonstrate the success of combining accuracy and robustness: Even under strongly varying illumination, reliable and precise results are obtained.

1 Introduction

The recovery of the displacement vector field (optical flow) between two consecutive frames of an image sequence is a classical problem in computer vision. In this context, variational methods play an important role, since they allow to incorporate various model assumptions in a transparent way and they yield dense flow fields. Numerous modifications have been introduced since the first variational approaches of Horn and Schunck [8] and Nagel [12]: More recent techniques such as [3, 11, 4] combine discontinuity-preserving regularisers that respect motion boundaries, robust data terms that improve the performance with respect to outliers and noise, and hierarchical optimisation strategies that handle large displacements. This has led to highly accurate methods. Moreover, efficient numerical schemes allow for a real-time computation of the results [5].

However, there is one topic that has hardly been addressed in the literature on variational optical flow methods, but which is of fundamental importance for their applicability in practice: the robustness of the estimation under realistic illumination changes. Such illumination changes include for instance shadow/shading, specular reflections and globally varying illumination [7, 18]. They can provide severe perturbations for important applications such as robot navigation or driver assistance systems.

So far, most of the illumination-robust optical flow techniques in the literature are local methods: They are easy to implement, but they give non-dense flow fields and do not belong to the currently best performing techniques in terms of error measures. In particular, estimation techniques for colour image sequences are very popular [14, 6, 18]. Usually, such methods make use of photometric invariants that are derived from the HSI colour space [6, 18], from normalised RGB channels [6], or from the $r\phi\theta$ representation that is obtained via the spherical coordinate transform (SCT) [18]. These expressions are in general invariant under illumination changes of multiplicative and/or additive type. Alternatively, in the context of grey value image sequences, different methods have been proposed that tackle the illumination problem by an explicit modelling of the underlying physical process [13, 7]. In this case, the optical flow field and the parameters of the illumination model have to be estimated simultaneously. A last class of methods, that are also applicable to grey value image sequences, makes use of image derivatives [17]. However, one should note that derivatives are only invariant under additive illumination changes. Thus, they may not be optimal with respect to realistically varying illumination that always contains a multiplicative part [18].

In face of these strategies and the increasing accuracy of variational methods in the last few years, it becomes evident why recently more efforts have been made to embed such concepts into a suitable variational framework: Prominent examples are the techniques of Brox *et al.* [4] and Papenberg *et al.* [15] that are based on higher order image derivatives, as well as the method of Kim *et al.* [9] that models the illumination changes in an explicitly way. However, with respect to variational techniques that make use of photometric invariants, only one approach is known to us: Barron and Klette [2] incorporate a single invariant expression as part of a multichannel framework into the classical method of Horn and Schunck [8]. Since photometric invariants combine most of the advantages of derivative and model based approaches – they allow for the modelling of multiplicative and additive illumination changes without requiring the estimation of any additional model parameters – it is surprising that there has been no further research done in this direction.

Thus, the goal of the present paper is twofold: First, it shall provide an overview of the most important concepts to design photometric invariants for colour sequences. Secondly, by embedding these invariants into the highly accurate optical flow technique of Brox *et al.* [4], it shall investigate the true potential of recent variational optical flow methods under realistic illumination conditions.

Our paper is organised as follows. In Section 2 we give a short review on the dichromatic reflection model and discuss the basic properties of photometric invariants. This discussion allows us to propose five different invariants in Section 3. How these invariants can be incorporated into a suitable variational framework is then demonstrated in Section 4. Finally, we investigate the performance of the new method in Section 5. The summary in Section 6 concludes this paper.

2   The Dichromatic Reflection Model

In order to understand the basic concept behind photometric invariants, it makes sense to start by giving a short review of the *dichromatic reflection model* [16, 18]. This model describes the observed RGB colour $\mathbf{c}(\mathbf{x}) = (R(\mathbf{x}), G(\mathbf{x}), B(\mathbf{x}))^{\top}$ at a certain location

$\mathbf{x} = (x, y)^\top$ as sum of an interface reflection component $\mathbf{c}_i(\mathbf{x})$ and a body reflection component $\mathbf{c}_b(\mathbf{x})$:

$$\mathbf{c}(\mathbf{x}) = \mathbf{c}_i(\mathbf{x}) + \mathbf{c}_b(\mathbf{x}). \tag{1}$$

While the interface reflection component is caused by specularities or highlights, the body reflection is directly related to the (Lambertian) reflection of the matte body. Physical characteristics of the camera are not modelled explicitly by this equation.

Under *spectrally uniform illumination*, these two terms can be decomposed further. They can be factorized into the overall intensity e, the geometrical reflection factor $m(\mathbf{x})$ and the reflectance colour $\widehat{\mathbf{c}}(\mathbf{x})$. Thus, equation (1) becomes

$$\mathbf{c}(\mathbf{x}) = e\left(m_i(\mathbf{x})\,\widehat{\mathbf{c}}_i(\mathbf{x}) + m_b(\mathbf{x})\,\widehat{\mathbf{c}}_b(\mathbf{x})\right), \tag{2}$$

which actually describes a linear combination of the two reflectance colours $\widehat{\mathbf{c}}_i$ and $\widehat{\mathbf{c}}_b$ with the corresponding geometric reflection factors m_i and m_b as weights. At this point one should note that the interface reflectance colour $\widehat{\mathbf{c}}_i$ cannot be arbitrary: Since we have assumed a spectrally uniform illumination, it is restricted to pure achromatic colours, i.e. grey values of any type. This in turn means that all three channels of $\widehat{\mathbf{c}}_i$ have equal contributions, i.e. $\widehat{R}_i(\mathbf{x}) = \widehat{G}_i(\mathbf{x}) = \widehat{B}_i(\mathbf{x}) =: w_i(\mathbf{x})$.

If we furthermore assume a *neutral interface reflection (NIR)* [10], the value $w_i(\mathbf{x})$ of all three interface channels becomes independent of the location. By defining the vector $\mathbf{1} = (1, 1, 1)^\top$ we can thus rewrite the dichromatic reflection model as

$$\mathbf{c}(\mathbf{x}) = e\left(m_i(\mathbf{x})\,w_i\,\mathbf{1} + m_b(\mathbf{x})\,\widehat{\mathbf{c}}_b(\mathbf{x})\right). \tag{3}$$

One should note that this equation is equivalent to the formulation considered in [18]. However, for a better understanding, we have made all the simplifications explicit.

Now we are in the position to give a concrete definition of photometric invariants: Photometric invariants are those expressions that are constructed from the observed colour $\mathbf{c}$ and that are at least independent of one of the three photometric variables e, m_b or m_i. In general, three different classes of photometric invariants can be distinguished: (i) Invariants with respect to *global multiplicative illumination changes* – these expressions are only independent of the light source intensity e. (ii) Invariants with respect to *shadow* and *shading* – these expressions are independent of the light source intensity e and the geometric body reflection factor m_b, at least for matte surfaces (i.e. $m_i = 0$). (iii) Invariants with respect to *highlights* and *specular reflections* – these expressions are independent of all three photometric variables e, m_b and m_i.

3   Photometric Invariants

After we have discussed the dichromatic reflection model, let us now investigate the main design principles behind photometric invariants. To this end, we consider the colour $\mathbf{c}(\mathbf{x})$ at a certain point $\mathbf{x}$ in terms of its three components $R(\mathbf{x})$, $G(\mathbf{x})$ and $B(\mathbf{x})$, respectively. Then, three main strategies for designing invariants are proposed in the literature: normalisation techniques, the differentiation of logarithmised channels, and the transformation to other colour spaces in terms of spherical/conical coordinates. Let us now discuss all three concepts in detail.

3.1 Normalisation Techniques

The first concept that we consider for designing photometric invariants is the transformation of the RGB colour space by means of normalisation [6]. In general, this transformation can be formulated as

$$(R, G, B)^\top \mapsto \left(\frac{R}{N}, \frac{G}{N}, \frac{B}{N}\right)^\top,$$ (4)

where N is a normalisation factor that depends on R, G and B. Such a proceeding yields a so-called *chromaticity space*. Two popular representatives for chromaticity spaces are the *arithmetic* and the *geometric chromaticity space* that are based on the normalisation by the arithmetic mean $N = (R + G + B)/3$ or the geometric mean $N = \sqrt[3]{RGB}$, respectively. However, with respect to the degree of invariance their behaviour is identical: By plugging the dichromatic reflection model (3) into equation (4), one can see that in both cases the photometric variables e and m_b cancel out (if $m_i = 0$). Thus, both the arithmetic and the geometric normalisation strategy yield expressions that are invariant under *shadow* and *shading*.

3.2 Log-Derivatives Strategies

A second class of strategies for creating photometric invariants is the computation of derivatives of the logarithmised colour channels. In the case of first order differential operators this yields the mapping

$$(R, G, B)^\top \mapsto \left((\ln R)_x, (\ln R)_y, (\ln G)_x, (\ln G)_y, (\ln B)_x, (\ln B)_y\right)^\top,$$ (5)

where subscripts denote partial derivatives, i.e. $G_x = \partial G/\partial x$. However, in contrast to the previous strategy this concept is not invariant with respect to shadow and shading: Only the overall intensity e is eliminated for $m_i = 0$, since the geometric reflection factor m_b depends on the location $\mathbf{x}$ and thus does not vanish. Therefore log-derivatives are only invariant under changes of the *image intensity*.

Nevertheless, this strategy may be an interesting upgrade possibility for techniques that are originally based on image derivatives such as the ones in [17, 4, 15]. By logarithmising the colour channels before the computation, such methods are able to handle global multiplicative illumination changes instead of global additive ones. Moreover, if the spatial variations of the geometric reflection factor m_b are rather small, such a strategy also provides a reasonable degree of invariance with respect to *shadow* and *shading*.

3.3 Spherical and Conical Transforms

The last concept for designing invariants that we discuss in this section is the consideration of other colour spaces that are obtained via spherical or conical transforms. Such colour spaces are e.g. the HSV and the $r\phi\theta$ colour space [6, 18]. Let us start our discussion with the HSV colour space. This colour space represents each colour in terms of hue, saturation and value. While the hue describes the pure colour and the saturation stands for the acromatic/grey component, the value corresponds to the actual brightness.

If we define $M = \max(R, G, B)$ and $m = \min(R, G, B)$, the corresponding transformation is given by:

$$(R, G, B)^\top \mapsto \begin{cases} H = \begin{cases} \frac{G-B}{M-m} \times 60°, & R \geq G, B, \\ (2 + \frac{B-R}{M-m}) \times 60°, & G \geq R, B, \\ (4 + \frac{R-G}{M-m}) \times 60°, & B \geq R, G, \end{cases} \pmod{360°}, \\ S = \frac{M-m}{M} \\ V = M. \end{cases} \qquad (6)$$

Evidently, the hue is invariant under both *shadow* and *shading* as well as *highlights* and *specularities*. However, since it involves the ratio of colour channel differences, it also discards the most information of all invariants. The other two channels are less robust: While the saturation allows to cope at least with *shadow* and *shading*, the value channel is not invariant at all.

In contrast to the HSV colour space that describes the RGB colours in terms of a cone, the $r\phi\theta$ colour space is obtained via a spherical transformation of the RGB coordinates. This transformation is given by

$$(R, G, B)^\top \mapsto \begin{cases} r = \sqrt{R^2 + G^2 + B^2} \\ \theta = \arctan\left(\frac{G}{R}\right) \\ \phi = \arcsin\left(\frac{\sqrt{R^2+G^2}}{\sqrt{R^2+G^2+B^2}}\right) \end{cases}. \qquad (7)$$

Here, r denotes the magnitude of the colour vector and θ and ϕ are the two angles that describe longitude and latitude, respectively. As one can easily verify, both angles θ and ϕ are invariant with respect to *shadow* and *shading*. However, the colour magnitude r is no photometric invariant.

4 Variational Optical Flow Computation

Since we are interested in incorporating the previously discussed photometric invariants into a variational framework, let us briefly recall the basic idea behind variational methods. To this end, let us consider an image sequence $f(\mathbf{x}, t)$, where $\mathbf{x} = (x, y)^\top$ denotes the location within a rectangular image domain Ω and $t \geq 0$ denotes time. Then, variational optical flow methods compute the dense displacement field $\mathbf{u} = (u, v)^\top$ between two consecutive frames $f(\mathbf{x}, t)$ and $f(\mathbf{x}, t+1)$ as minimiser of an energy functional with the general structure

$$E(\mathbf{u}) = E_D(\mathbf{u}) + \alpha \, E_S(\mathbf{u}), \qquad (8)$$

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 (piecewise) smoothness of the result.

4.1 A Multichannel Approach for Photometric Invariants

Having explained the main idea behind variational methods, let us now derive a suitable model for computing the optical flow. Such a model must not only permit the integration of our photometric invariants into the data term, it should also allow the estimation of highly accurate optical flow fields. In order to satisfy both demands, we propose to compute the optical flow as minimiser of an energy functional $E(\mathbf{u})$ with data term

$$E_D(\mathbf{u}) = \int_\Omega \psi_D \Big(\sum_{i=1}^N \gamma_i \, |f_i(\mathbf{x}+\mathbf{u}, t+1) - f_i(\mathbf{x}, t)|^2 \Big) \, dx \, dy \tag{9}$$

and smoothness term

$$E_S(\mathbf{u}) = \int_\Omega \psi_S \left(|\nabla u|^2 + |\nabla v|^2 \right) \, dx \, dy \ . \tag{10}$$

Here, $\nabla u = (u_x, u_y)^\top$ and $\nabla v = (v_x, v_y)^\top$ denotes the spatial gradient of the flow component u and v, respectively.

This energy functional can be considered as a 2-D multichannel extension of the high accuracy technique of Brox *et al.* [4]. However, instead of assuming constancy on the grey value and its spatial derivatives, our method is based on the assumption that for corresponding objects in both frames the N different photometric invariants given by the channels f_i remain constant. Thus, for instance, we consider in the case of the spherical coordinate transform constancy assumptions on both channels of the image sequence $\mathbf{f} = (f_1, f_2) = (\phi, \theta)$ which can be obtained from the original colour image sequence $\mathbf{f} = (R, G, B)$ using equation (7) . In this context, the scalars $\gamma_i > 0$ serve as weights that steer the importance of the different channels. In order to allow for a correct estimation of large displacements, all photometric constancy assumptions are employed in their original nonlinear form. Moreover, both the data and the smoothness term are penalised in a non-quadratic way - to render the approach more robust to outliers and noise in the case of the data term and to preserve 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. It is given by $\psi(s^2) = \sqrt{s^2 + \epsilon^2}$, where ϵ is a small regularisation parameter. In our case ϵ is set to 10^{-3}.

4.2 Minimisation

In order to minimise the previously proposed energy functional, one has to solve its Euler-Lagrange equations. These equations are given by the following coupled pair of nonlinear partial differential equations (PDEs):

$$0 = \psi_D'(...) \Big(\sum_{i=1}^N \gamma_i \, \big(f_i(\mathbf{x}+\mathbf{u}, t+1) - f_i(\mathbf{x}, t) \big) \frac{\partial}{\partial x} f_i(\mathbf{x}+\mathbf{u}, t+1) \Big)$$
$$+ \, \alpha \, \mathrm{div} \left(\psi_S' \left(|\nabla u|^2 + |\nabla v|^2 \right) \nabla u \right) , \tag{11}$$

$$0 = \psi_D'(...) \Big(\sum_{i=1}^N \gamma_i \, \big(f_i(\mathbf{x}+\mathbf{u}, t+1) - f_i(\mathbf{x}, t) \big) \frac{\partial}{\partial y} f_i(\mathbf{x}+\mathbf{u}, t+1) \Big)$$
$$+ \, \alpha \, \mathrm{div} \left(\psi_S' \left(|\nabla u|^2 + |\nabla v|^2 \right) \nabla v \right) , \tag{12}$$

where $\psi'_D(...)$ is an abbreviation for

$$\psi'_D(...) = \psi'_D\Big(\sum_{i=1}^{N}\gamma_i\,|f_i(\mathbf{x}+\mathbf{u},t+1)-f_i(\mathbf{x},t)|^2\Big)\,.$$

After discretising these equations by means of finite difference approximations, the resulting nonlinear system of equations is solved via two nested fixed point iterations and a coarse-to-fine warping strategy as proposed in [4]. Alternatively, also a real-time capable multigrid scheme could be used [5].

5 Experiments

In our first experiment, we investigate the usefulness of different photometric constancy assumptions with respect to spatially varying multiplicative and additive illumination changes (using a Gaussian model). To this end, we consider frame 10 and 11 of the *Street* sequence available at `http://of-eval.sourceforge.net` and create two strongly degraded variants of frame 11 with heavily varying illumination (cf. Figure 1). The different photometric constancy assumptions and the corresponding results in terms of the average angular error [1] are listed in Table 1. As one can see,
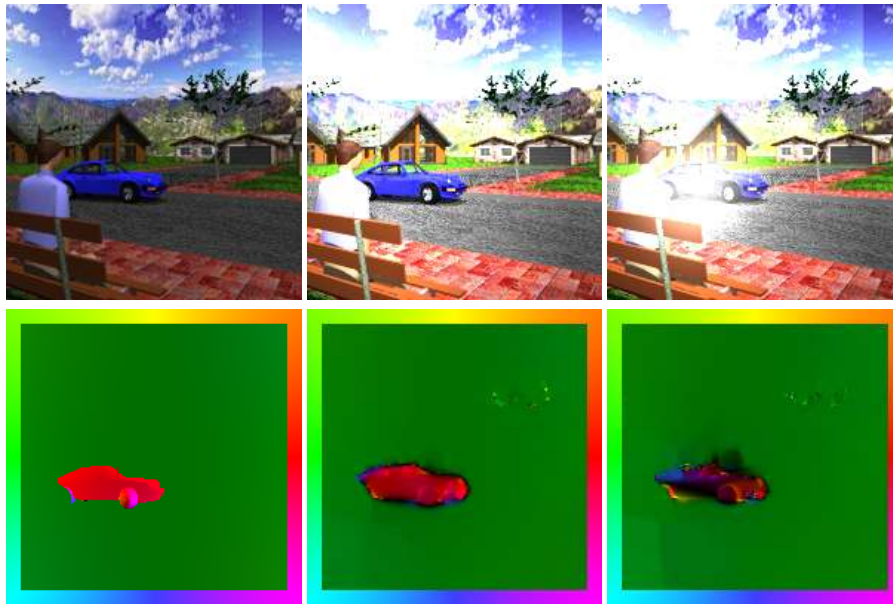
Fig. 1. Robustness of the $\phi\theta$ constancy assumption under varying illumination. *Top Row: (a)* Frame 11 of the *Street* sequence (200×200). *(b)* Frame 11 with spatially varying multiplicative illumination. *(c)* Frame 11 with spatially varying multiplicative and additive illumination. *Bottom Row: (d)* Ground truth. *(e)* Computed result for (a) and (b). *(f)* Computed result for (c). The direction is encoded by the colour (as shown at the boundaries), the magnitude by the brightness.

Table 1. Comparison of the different illumination invariants for the *Street* sequence in its original form (orig.), with locally varying multiplicative illumination (mult.) and with locally varying multiplicative and additive illumination (mult.+add.). All weights γ_i have been set to one. The remaining parameters have been optimised with respect to the average angular error (AAE). #Ch = number of channels.

Concept		#Ch	AAE orig.	AAE mult.	AAE mult. + add.
Standard	RGB	3	$2.65°$	$43.44°$	$43.44°$
Colour Space	HSL (Hue)	1	$4.28°$	$4.28°$	$4.28°$
	spherical (ϕ, θ)	2	$\mathbf{2.07°}$	$\mathbf{2.07°}$	$\mathbf{3.37°}$
Normalisation	RGB (arithm. mean)	3	$2.22°$	$2.22°$	$3.71°$
	RGB (geom. mean)	3	$2.26°$	$2.26°$	$5.64°$
Log-Derivatives	$\nabla \ln(\text{RGB})$	6	$2.89°$	$3.04°$	$4.35°$
Brox et al. (2-D)	RGB + ∇ RGB	9	$2.64°$	$3.89°$	$3.92°$

the standard RGB constancy assumption fails completely under varying illumination. Thereby, the error of $43.44°$ refers to a zero displacement field which means in turn that the underlying method could not make any use of the provided information. In contrast, all techniques based on photometric invariants perform favourably. In particular the constancy assumption on the $\phi\theta$ channels gives excellent results: With average angular errors up to $2.07°$, it does not only outperform the hue channel, that offers a higher degree of invariance at the expense of discarding too much information, it also provides better results for the sequence *with* spatially varying multiplicative illumination changes than the 2-D RGB Brox *et al.* for the sequence *without*. Compared to the best result in the literature that is known to the authors – the result of $4.85°$ by Weickert and Schnörr [19] – this improvement is even more drastical. Thus it is not surprising that the corresponding flow fields of the $\phi\theta$-channels in Figure 1 show a precise estimation of the optical flow: The shape of the car is well preserved and the camera motion is also estimated accurately.

In our second experiment, we analyse the performance of the different photometric invariants with respect to typical illumination changes in real-world data. To this end, we consider the left frames 205 and 207 of the *DIPLODOC Road* stereo sequence available at `http://tev.itc.it/DATABASES/road.html`. As one can see from the computed results in Figure 2, the method based on the standard RGB constancy assumptions has again severe problems. Instead of compensating for the varying illumination between both frames, it interprets this change as a global motion in upward direction (street). Since the ego-motion of the camera system induces a divergent flow field, this estimation is completely wrong. However, once again our techniques based on photometric constancy assumptions give very good results. The $\phi\theta$-channels and the normalised RGB values (using the geometric mean) even allow to detect the pedestrian at the lower left border of the image – in spite of the severely changed illumination conditions. This confirms our findings from the first experiment: If suitable photometric invariants are embedded within an accurate variational framework, they may render the underlying method highly robust with respect to realistic changes of the illumination. However, as the first experiment has also shown, one has to be careful not to discard too much information, since otherwise the quality of the estimation decreases.
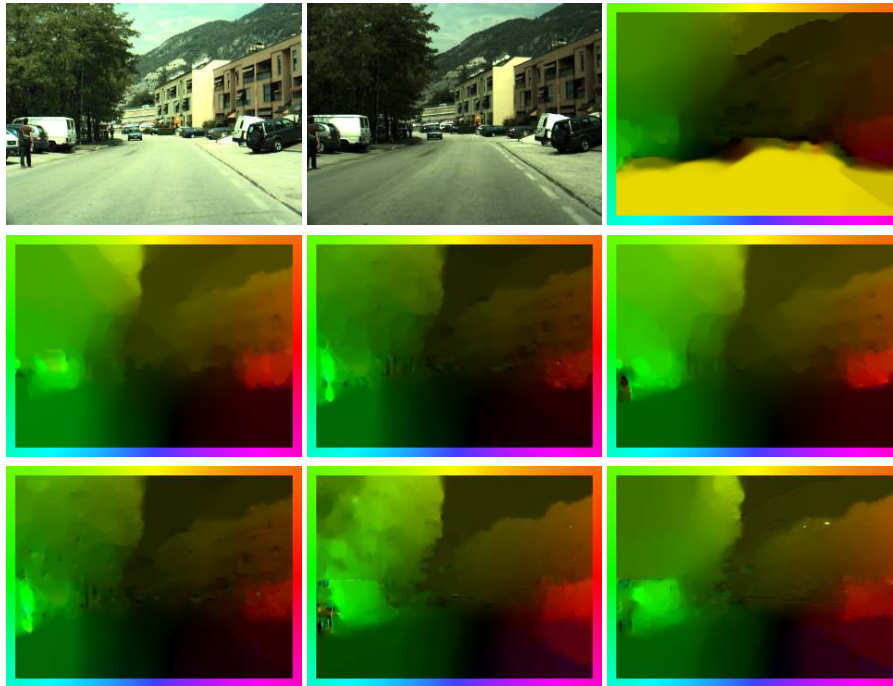
Fig. 2. Results under real illumination conditions *Top Row:* (a) Left frame 205 of the *Road* stereo sequence of the DIPLODOC project (size 320×240). (b) Left frame 207. (c) Flow with RGB constancy assumption. *Middle Row:* (d) Hue constancy assumption. (e) $\phi\theta$ constancy assumption. (f) Normalised RGB constancy assumption (arithm.). *Bottom Row:* (g) Normalised RGB constancy assumption (geom.). (h) Log-derivative constancy assumption. (i) 2-D Brox *et al.* (2-D).

6   Summary and Conclusions

Photometric invariants and variational methods are two successful concepts in image analysis that have emerged without many interactions so far. In our paper we have demonstrated the benefits of combining them in order to solve a challenging computer vision problem: dense and highly accurate motion estimation under realistic changes of the illumination conditions. We have thereby shown that the performance of variational optical flow methods can be significantly improved, if traditional constancy assumptions are replaced by photometric invariants.

It is our hope that this research serves as another step that helps to bridge the gap between mathematically well-founded theories and more robust real-life applications.

Acknowledgements

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. L. Barron and R. Klette. Quantitative colour optical flow. In *Proc. 16th International Conference on Pattern Recognition*, volume 4, pages 251–255, Quebec City, Canada, Aug. 2002. IEEE Computer Society Press.

3. M. J. Black and P. Anandan. The robust estimation of multiple motions: parametric and piecewise smooth flow fields. *Computer Vision and Image Understanding*, 63(1):75–104, Jan. 1996.

4. 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.

5. 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.

6. P. Golland and A. M. Bruckstein. Motion from color. *Computer Vision and Image Understanding*, 68(3):346–362, Dec. 1997.

7. H. Haußecker and D. Fleet. Estimating optical flow with physical models of brightness variation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6):661–673, 2001.

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

9. Y.-H. Kim, A. M. Martínez, and A. C. Kak. Robust motion estimation under varying illumination. *Image and Vision Computing*, 23(1):365–375, 2005.

10. H. C. Lee, E. J. Breneman, and C. P. Schulte. Modeling light reflection for computer vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:402–409, 1990.

11. E. Mémin and P. Pérez. Hierarchical estimation and segmentation of dense motion fields. *International Journal of Computer Vision*, 46(2):129–155, 2002.

12. H.-H. Nagel. Constraints for the estimation of displacement vector fields from image sequences. In *Proc. Eighth International Joint Conference on Artificial Intelligence*, volume 2, pages 945–951, Karlsruhe, West Germany, August 1983.

13. S. Negahdaripour. Revised definition of optical flow: integration of radiometric and geometric clues for dynamic scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(9):961–979, 1998.

14. N. Ohta. Optical flow detection by color images. In *Proc. Tenth International Conference on Pattern Recognition*, pages 801–805, Singapore, Sept. 1989.

15. N. Papenberg, A. Bruhn, T. Brox, S. Didas, and J. Weickert. Highly accurate optic flow computation with theoretically justified warping. *International Journal of Computer Vision*, 67(2):141–158, Apr. 2006.

16. S. A. Shafer. Using color to seperate reflection components. *Color Research and Applications*, 10(4):210–218, 1985.

17. S. Uras, F. Girosi, A. Verri, and V. Torre. A computational approach to motion perception. *Biological Cybernetics*, 60:79–87, 1988.

18. J. van de Weijer and T. Gevers. Robust optical flow from photometric invariants. In *Proc. 2004 IEEE International Conference on Image Processing*, volume 3, pages 1835–1838, Singapore, Oct. 2004.

19. 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.