Proceedings Pattern Recognition 2003 (Annual Meeting of the Deutsche Arbeitsgemeinschaft für Mustererkennung, Magdeburg, September 2003, Springer Verlag)

A new view at differential and tensor-based motion estimation schemes

Rudolf Mester

Institute for Applied Physics, University of Frankfurt, Robert-Mayer-Str. 2–4, D-60054 Frankfurt, Germany
mester@iap.uni-frankfurt.de, www.uni-frankfurt.de/fb13/iap/cvg

Abstract. This paper contributes to the recent attempts to generalize classical differential and tensor-based motion estimation methods; it provides means for a better adaptation to the statistical structure of signal and noise. We show that conventional differential constraint equations often used for motion analysis capture the essence of translational motion only partially, and we propose more expressive formulations using higher order derivatives, finally leading to steerable nulling filters. Such filters consist of a 'static' prefilter and a steerable filter; we show how the prefilter can be optimized for given autocovariance structure of signal and noise, and how the steerable filter can be related to classical tensor-based approaches, leading to a constrained eigensystem problem.

1 Introduction

After an initial phase during which many different proposals for motion estimation have been made, a certain tendency to join different classes of motion algorithms in a common framework can be observed in the recent years. It is a sign of growing maturity of the field that modern approaches to motion analysis more and more deal with the theoretical background of the motion analysis task, and increasingly exploit the perfected methods of numerical analysis, signal theory and statistics, which are standard tools in modern signal processing.

This paper shall contribute to the theoretical foundations of motion analysis. We start in section 2 with a short review on the differential approach to motion analysis, which can be understood as a special case of tensor-based methods [1]. We model signal and noise probabilistically and formulate the motion estimation problem as a statistical estimation problem. If, furthermore, prior knowledge on the average distribution of the entity to be estimated (i.e. the motion vector) is available, our problem can be tackled using Bayesian theory. From a statistical point of view, there is no other way to use the information contained in the image data in a better way, but of course, there are quite many design parameters in the process that make it difficult to devise *the* optimum motion estimator. Therefore, this paper is more the exposition of a theoretical framework rather than the description of a specific algorithm.

2 Differential approaches to motion analysis

The general principle behind all differential approaches to motion estimation is that the conservation of some local image characteristic throughout its temporal evolution is reflected in terms of differential-geometric descriptors. In its simplest form, the assumed conservation of brightness along the motion trajectory through space–time leads to the well-known *brightness constancy constraint equation* (BCCE), where $\mathbf{g}(\mathbf{x})$ is the gradient of the gray value signal $s(\mathbf{x})$:

$$\left(\frac{\partial s}{\partial x_1}, \frac{\partial s}{\partial x_2}, \frac{\partial s}{\partial x_3}\right) \cdot \mathbf{r} = 0 \quad \Leftrightarrow \quad \mathbf{g}^T(\mathbf{x}) \cdot \mathbf{r} = 0. \tag{1}$$

Since $\mathbf{g}^T(\mathbf{x}) \cdot \mathbf{r}$ is proportional to the directional derivative of s in direction $\mathbf{r}$, the BCCE states that this derivative vanishes in the direction of motion.

In order to cope with the *aperture problem* , and in order to decrease the variance of the motion vector estimate, usually some kind of weighted averaging is performed in a neighborhood V, using a weight function $w(\mathbf{x})$.

$$\int_V w(\mathbf{x}) \cdot \left|\mathbf{g}^T(\mathbf{x}) \cdot \mathbf{r}\right|^2 \, d\mathbf{x} \longrightarrow \min$$

$$\Rightarrow \quad \mathbf{r}^T \cdot \mathbf{C}_g \cdot \mathbf{r} \longrightarrow \min \qquad \text{with} \quad \mathbf{C}_g := \int_V \mathbf{g}(\mathbf{x}) \cdot w(\mathbf{x}) \cdot \mathbf{g}^T(\mathbf{x}) \, d\mathbf{x}$$

The solution vector $\hat{\mathbf{r}}$ is the eigenvector corresponding to the minimum eigenvalue of the *structure tensor* [1] $\mathbf{C}_g$ (cf. [4], [5], p.366; [6]). The minimization criterion according to eq.2 can be replaced by modified criteria which yield exactly the same solution $\mathbf{r}$ in case of ideal rank-2 signals, but different solutions for the realistic case of perturbations in matrix $\mathbf{C}_g$. One motivation for doing so is for compensating the typically complicated error statistics in the vector $\mathbf{g}$ which are due to anisotropic image statistics, colored noise, and overlapping filter support. Furthermore, since the matrix $\mathbf{C}_g$ is built from outer products of the vectors $\mathbf{g}$, the covariances between elements of $\mathbf{C}_g$ are hard to analyze [7].

In the classical theory, the (discrete!) implementation of the derivative operators itself is a formidable problem, even though some early authors apparently overlooked the crucial importance of this point. Both [8–10] as well as [11] treat this explicitly as an optimization problem; Elad et al. [11] point out that the closeness of the derivatives to the 'ideal' derivatives is not the correct goal, but rather the quality of the direction estimate obtained by using them. Furthermore, they show that prior knowledge on the distribution of motion vectors significantly changes the shape of the computed filter masks.

In the present contribution, the same problem is approached with a different strategy: we stay with the conventional scheme of building the gradient operators with a prefilter and a subsequent derivative operator, but we let the prefilter be completely determined by the image statistics (see section 5).

---

[1] The tensor approach to motion analysis and representation has been extensively investigated by Granlund, Knutsson, Bigün and many others at Linköping university, Sweden. For two rather recent accounts on that subject, see [2] and [3].

2.1 How conventional differential approaches can be generalized

The differential formulation of brightness constancy along the motion trajectory is not the unique and presumably not the most expressive way of specifying a relation between the entity that is sought (the motion vector $\mathbf{r}$) and the data that can be observed. What is definitely to be taken into account are the *spectral characteristics of the image signal* and the *spectral characteristics of the noise*. Following the principles introduced into signal processing by Nobert Wiener already in the 1940ies, those spectral components where the signal-to-noise ratio is high should be paid more attention than those with a low S/N ratio. As this is shown in section 5, this can be understood also as consequently applying weighted-least-squares theory to the given estimation problem, using the autocovariance functions of signal and noise.

Quite obviously, since any pure rank-2 signal s may be prefiltered by (almost) any filter with arbitrary transfer function $P(\mathbf{f})$ (except complete nullification of the signal) without changing the *direction* of the eigenvectors of $\mathbf{C_g}$, the interpretation of $\mathbf{g}$ as being the local gradient is much too narrow. There is also more to these prefilters than just a 'regularization' of gradient computation: they can (and should) be designed according to Wiener's criteria, i.e. minimizing the S/N ratio of the gradient estimates.

The same applies also to the classical (and mostly undisputed) usage of 1st order derivatives in optic flow algorithms: Assuming brightness constancy along the motion trajectory, all higher order directional derivatives vanish in the motion direction:

$$\frac{\partial s}{\partial \mathbf{r}} \stackrel{!}{=} 0 \quad \bigcap \quad \frac{\partial^2 s}{\partial \mathbf{r}^2} \stackrel{!}{=} 0 \quad \bigcap \quad \ldots \tag{2}$$

For an ideally oriented signal, anything that could be said is already expressed by the 1st order condition, but for real signals *none* of the requirements in eq. 2 is met. A condition which is less stringent than eq. 2, but nevertheless comprises as much as possible from these multitude of conditions in a single linear equation can be obtained by summing up the constraints:

$$\alpha_1 \frac{\partial s}{\partial \mathbf{r}} + \alpha_2 \frac{\partial^2 s}{\partial \mathbf{r}^2} + \alpha_3 \frac{\partial^3 s}{\partial \mathbf{r}^3} + \ldots \quad \stackrel{!}{=} 0 \tag{3}$$

The left hand side of this equation is nothing else than a generator for a very rich class of filter operators, parameterized by direction vector $\mathbf{r}$:

$$h(\mathbf{x} \mid \mathbf{r}) * s(\mathbf{x}) \stackrel{!}{=} 0$$

This means that all linear operators that do not let an ideal oriented signal $s(\mathbf{x})$ pass have the structure of eq. 3. Since all oriented signals have power spectra that are concentrated on lines or planes in the Fourier domain [12], we can denote these filters as *oriented nulling filters*, and the Fourier transform of eq.3 describes the Taylor series expansion of the transfer function of such an oriented nulling filter.

If a signal $s(\mathbf{x})$ has a single dominant orientation (i.e. if it is a rank-2 signal), it may be prefiltered by any filter $p(\mathbf{x})$ with radially symmetric transfer function $P(\mathbf{f})$ without changing the *direction* of the eigenvectors of $\mathbf{C}_g$ (the non-vanishing eigen*values* will in general change). Even the application of an (almost) arbitrary *anisotropic* linear filter $p(\mathbf{x})$ on $s(\mathbf{x})$ will not change the orientation (except that $p(\mathbf{x})$ nullifies $s(\mathbf{x})$ totally). This means that the interpretation of $\mathbf{g}$ in eq.2 as being the local gradient is much too narrow and we can use the pre-filter $p(\mathbf{x})$ for tuning the approach. Coming back to eq.3, we can replace $s(\mathbf{x})$ by $s(\mathbf{x}) * p(\mathbf{x})$, and the orientation remains invariant (for rank 2 signals!) Thus we may formulate a new type of constraint equation:

$$\left(\alpha_1 \frac{\partial}{\partial \mathbf{r}} + \alpha_2 \frac{\partial^2}{\partial \mathbf{r}^2} + \ldots\right) * (s(\mathbf{x}) * p(\mathbf{x})) \stackrel{!}{=} 0 \qquad \Rightarrow \qquad \tilde{h}(\mathbf{x} \mid \mathbf{r}) * (s(\mathbf{x}) * p(\mathbf{x})) \stackrel{!}{=} 0$$

Of course, the more degrees of freedom we obtain in formulating constraint equations from which (hopefully) the motion direction can be deduced, the more parameter values have to be selected. For real image data neither the traditional BCCE nor any generalization will hold; instead of insisting in the left hand side to *vanish*, the requirement is relaxed to *minimizing it on an average*. This is exactly the point where the precise design of the various imaginable constraint equations makes a difference: it implicitly defines a metric on the space of direction vectors $\mathbf{r}$. We obtain now: which, subsuming all the reasoning made before, can be generalized to

$$\int_{\mathbf{x}} w(\mathbf{x}) \cdot |h(\mathbf{x} \mid \mathbf{r}) * s(\mathbf{x})|^2 \; d\mathbf{x} \quad \longrightarrow \quad \min \tag{4}$$

where $h(\mathbf{x} \mid \mathbf{r})$ comprises the combination of directional derivatives of different order, and an optional *pre-filter* $p(\mathbf{x})$. This means: *The frequency-weighted directional variation of the signal is minimized in the direction of motion.* Since we are aiming at estimates for $\mathbf{r}$ which are optimum with respect to e.g. minimizing the variance of the residual error in the estimate $\hat{\mathbf{r}}$, the parameterized filter $h(\mathbf{x} \mid \mathbf{r})$ should be selected in a way that considers the known statistics of signal and noise. We have recently proposed [13] a motion analysis scheme that incorporates a steerable directional approximation of the given signal and can be extended in the spirit of Wiener filtering. The approach in the present paper is to split $h(\mathbf{x} \mid \mathbf{r})$ into a *pre-filter* which is optimally designed estimating the 'true' image signal by minimizing the mean squared error, and a *directional post-filter*. This leaves the possibility to design the (steerable) post-filter according to signal-independent optimality requirements discussed extensively elsewhere (e.g. in [10]).

3 Generalizing structure tensors using steerable filters

We proceed by restating the relation between directional derivatives and steerable filters, which have been explored e.g. in [8, 14, 9]. The partial derivative in a direction specified by a unit vector $\mathbf{e}_r \in \mathbb{R}^3$ parameterized via spherical angles $\theta = (\theta_1, \theta_2)$ as $\mathbf{e}_r = (a_1(\theta),\ a_2(\theta),\ a_3(\theta))$ is given by $\frac{\partial}{\partial \mathbf{e}_r} s(\mathbf{x}) \;=\; \mathbf{e}_r^T \cdot \mathbf{g}(x) = \mathbf{e}_r^T \cdot \nabla s(\mathbf{x}) \;=\; \sum_{i=1}^{3} a_i(\theta) \cdot \frac{\partial s(\mathbf{x})}{\partial x_i}$

3.1 Generalized directional derivatives

Following the reasoning on pre-filters presented in section 2.1, we may insert a prefilter $p(\mathbf{x})$ (see e.g. [8, 10])

$$\frac{\partial}{\partial \mathbf{e}_r}(s(\mathbf{x}) * p(\mathbf{x})) \;=\; \left(\sum_{i=1}^{3} a_i(\theta_1, \theta_2) \cdot \left(\frac{\partial}{\partial x_i} p(\mathbf{x}) \right) \right) * s(\mathbf{x}) \qquad (5)$$

For $p(\mathbf{x})$ there are therefore many more functions under consideration than only a simple Gaussian kernel[2]. We can design $p(\mathbf{x})$ in a way that optimizes the signal/noise ratio at the output of the prefilter; this is the *Wiener-type prefilter approach* [15]. On the other hand, we may generalize the structure of the analysis scheme described by eq.5 and arrive at a generalized class of structure tensors, as will be shown in the following.

3.2 Steerable oriented signal energy determination

We abstract now from derivative filters and regard a family of steerable filter operators which can be written in the form [14] $h(\mathbf{x} \mid \theta) = \sum_{i=1}^{N} a_i(\theta) \cdot b_i(\mathbf{x})$ Since the original signal is *sheared* instead of being rotated by motion, it is appropriate to design $h(\mathbf{x} \mid \theta)$ accordingly; however, we will not deal here with details of such *shearable filters*. The symbol θ stands for a general parameter (or parameter vector) that controls the direction in which the filter operator is being steered. The $b_i(\mathbf{x})$ are basis functions; $a_i(\theta)$ and $b_i(\mathbf{x})$ are subject to certain conditions discussed in [14]. This operator will now be applied to a signal $s(\mathbf{x})$:

$$h(\mathbf{x} \mid \theta) * s(\mathbf{x}) = \sum_{i=1}^{N} a_i(\theta) \cdot (b_i(\mathbf{x}) * s(\mathbf{x}))$$

As before, the local energy of the resulting signal will be computed. The localization of the computation is again ensured by the weight function $w(\mathbf{x})$:

$$Q(\theta) = \int_{\mathbf{x}} w(\mathbf{x}) \cdot (h(\mathbf{x} \mid \theta) * s(\mathbf{x}))^2 \; d\mathbf{x}$$

A closer look reveals (using $g_i(\mathbf{x}) \equiv s(\mathbf{x}) * b_i(\mathbf{x})$):

$$(h(\mathbf{x} \mid \theta) * s(\mathbf{x}))^2 \;=\; \left(\sum_{i=1}^{N} a_i(\theta) \cdot (b_i(\mathbf{x}) * s(\mathbf{x})) \right)^2$$

$$= \left(\sum_{i=1}^{N} a_i(\theta) \cdot g_i(\mathbf{x}) \right)^2 \;=\; \sum_{i=1}^{N} \sum_{k=1}^{N} a_i(\theta) \cdot a_k(\theta) \cdot g_i(\mathbf{x}) \cdot g_k(\mathbf{x})$$

[2] In general, a binomial filter does much better than a sampled (and truncated) Gaussian.

If now a local integration is performed across this squared signal we obtain the quadratic form:

$$Q(\theta) = \sum_{i=1}^{N} \sum_{k=1}^{N} a_i(\theta) \cdot a_k(\theta) \int_{\mathbf{x}} w(\mathbf{x}) \cdot g_i(\mathbf{x}) \cdot g_k(\mathbf{x}) \, d\mathbf{x} \quad = \quad \mathbf{a}^T(\theta) \cdot \mathbf{J} \cdot \mathbf{a}(\theta) \quad (6)$$

In the standard structure tensor approach, $N = 3$, and $h(\mathbf{x} \mid \theta)$ is the first order directional derivative which can be represented by a steerable set of $N = 3$ filters (each of them representing the directional derivative in one of the principal directions of space-time). It is not very surprising that in this case $\mathbf{a}(\theta)$ is a unit vector in $\mathrm{I\!R}^3$, and the determination of $Q(\theta) \longrightarrow \min$ boils down into a simple eigensystem problem, as given already in eq.2. For synthesizing and steering a more general filter operator $h(\mathbf{x} \mid \theta)$, we know that the basis functions $b_i(\mathbf{x})$ should be polar-separable harmonic functions.The coefficient functions $a_i(\theta)$ will then be trigonometric functions of different (harmonic) frequencies [16], and the optimization problem eq.(3.2) will not be so simple to solve, though well-behaved. The design of the localization function $w(\mathbf{x})$ and the generalization of the directional derivative can be adapted to the signal and noise power spectra, respectively. Within this framework, a wide class of orientation selective steerable filters can be used to find principal orientations; if necessary they can equipped with a much more pronounced selectivity, offering the potential for higher accuracy.

4 Covariance structure of video signals

Since the autocovariance structure of the image and the noise are needed for designing the Wiener prefilter, we have to obtain that information for the class of image sequences under consideration. In order to obtain the required model for the (auto)covariance function *(acf)* of video signals, it turns out to be useful to analyze the video acquisition process in the continuous domain.

Let us assume that the image signals we are dealing with are generated by shifting a given twodimensional image with constant speed (v_x, v_y) in an certain direction. We assume that the twodimensional image can be characterized by a *twodimensional* autocovariance function $\tilde{\varphi}_{ss}(x, y)$. Obviously, the resulting three-dimensional autocovariance function $\varphi_{ss}(x, y, t)$ is then

$$\varphi_{ss}(x, y, t) = \tilde{\varphi}_{ss}(x - v_x t, \;\; y - v_y t)$$

If the motion vector (v_x, v_y) itself is a generated by a random process (as it is reasonable to assume for real video signals), there is a distribution for the new position $(x(t), y(t))$ at time instant t of the point which was located at $x = 0, y = 0$ at time instant 0. This distribution $\zeta(x, y, t)$ depends on t, and under normal circumstances its variance in x, y-direction will increase with increasing time t. The overall 3D autocovariance results from a convolution of the *purely spatial* autocovariance function $\tilde{\varphi}_{ss}(x, y)$ with the *position distribution function* $\zeta(x, y, t)$.

5 The prefilter as a Least Squares (Wiener) restoration operator for image blocks

Our model states that the observed image signal z in a spatio-temporal block of dimension $N \times N \times N$ is given by

$$z(i,j,k) = s(i,j,k) + v(i,j,k) \; .$$

Here, $v(i,j,k)$ denotes the observation (measurement) noise and we assume that this is discrete white noise, also denoted as *independent identically distributed (i.i.d.)* noise. For the subsequent steps, it is convenient to arrange the elements of the blocks s, v, and z in vectors $\mathbf{s}, \mathbf{v}$, and $\mathbf{z}$.

5.1 The canonical basis

The canonical basis is the coordinate frame of vectors $\mathbf{y}$ which are obtained from vectors $\mathbf{s}$ by a rotation $\mathbf{A}$ according to

$$\mathbf{y} = \mathbf{A} \cdot \mathbf{s}, \quad |\mathbf{y}| = |\mathbf{s}|$$

such that the covariance matrix $\mathbf{C}_y \equiv \mathsf{Cov}[\mathbf{y}]$ is diagonal. It is well known that this rotation is performed by the *principal component analysis (PCA)* or *Karhunen-Loève transform (KLT)*. In this new coordinate frame, we have

$$\mathbf{C}_y \stackrel{def}{=} \mathsf{Cov}[\mathbf{y}] = \mathsf{diag}\left\{\sigma_{yi}^2\right\} \tag{7}$$

The row vectors of the orthonormal (i.e. rotation) matrix $\mathbf{A}$ are given by the unit norm eigenvectors of the covariance matrix $\mathbf{C}_s \equiv \mathsf{Cov}[\mathbf{s}]$, i.e. $\mathbf{A}$ is the solution to the eigensystem problem

$$\mathbf{A} \cdot \mathsf{Cov}[\mathbf{s}] \cdot \mathbf{A}^T = \mathsf{diag}\{\lambda_i\} \quad \mid \; \mathbf{A} \cdot \mathbf{A}^T = \mathbf{I}$$

In the new coordinate frame, we have

$$\text{signal vector (unobservable):} \quad \mathbf{y} = \mathbf{A}\mathbf{s} \tag{8}$$

$$\text{noise vector (unobservable):} \quad \mathbf{u} = \mathbf{A}\mathbf{v} \tag{9}$$

$$\text{observed vector:} \quad \mathbf{w} = \mathbf{y} + \mathbf{u} = \mathbf{A}\mathbf{z} \tag{10}$$

Since the original noise covariance matrix was a scaled unit matrix, and since $\mathbf{A}$ is a rotation, the covariance matrix of the noise vector $\mathbf{u}$ contained in $\mathbf{w}$ remains simple:

$$\mathbf{C}_u \stackrel{def}{=} \mathsf{Cov}[\mathbf{u}] = \sigma_v^2 \cdot \mathbf{I} \tag{11}$$

Estimating the value of $\mathbf{y}$ means now to minimize the loss function

$$J(\hat{\mathbf{y}}) = (\hat{\mathbf{y}} - \mathsf{E}[\mathbf{y}])^T \mathbf{C}_w^{-1} (\hat{\mathbf{y}} - \mathsf{E}[\mathbf{y}]) + (\hat{\mathbf{y}} - \mathbf{w})^T \cdot \mathbf{C}_v^{-1} \cdot (\hat{\mathbf{y}} - \mathbf{w})$$

The *minimum mean squared error (MMSE)* estimate of $\mathbf{y}$ is then given by

$$\hat{\mathbf{y}} = \left(\mathbf{C}_y^{-1} + \mathbf{C}_u^{-1}\right)^{-1} \cdot \left(\mathbf{C}_y^{-1} \cdot \mathsf{E}\left[\mathbf{y}\right] + \mathbf{C}_u^{-1} \cdot \mathbf{w}\right) \tag{12}$$

It is relatively simple to compute (and to interpret!) this estimate in the canonical coordinate frame. For the white noise case, the covariance matrix of the noise term remains proportional to a unity matrix under any arbitrary rotation of the coordinate frame. With the specific covariance matrices which we have here (see eq.(7) and eq.(11)), we obtain:

$$\hat{\mathbf{y}} = \left(\mathsf{diag}\left\{\frac{1}{\sigma_{yi}^2}\right\} + \frac{1}{\sigma_v^2} \cdot \mathbf{I}\right)^{-1} \cdot \left(\mathbf{C}_y^{-1} \cdot \mathsf{E}\left[\mathbf{y}\right] + \mathbf{C}_u^{-1} \cdot \mathbf{w}\right) \tag{13}$$

Assuming $\mathsf{E}\left[\mathbf{y}\right] = \mathbf{0}$, this leads to

$$\hat{\mathbf{y}} \;=\; \left(\mathsf{diag}\left\{\frac{1}{\sigma_{yi}^2} + \frac{1}{\sigma_v^2}\right\}\right)^{-1} \cdot \frac{1}{\sigma_v^2}\mathbf{w} \qquad \Rightarrow \quad \hat{\mathbf{y}} \;=\; \mathsf{diag}\left\{\frac{\sigma_{yi}^2}{\sigma_v^2 + \sigma_{yi}^2}\right\} \cdot \mathbf{w} \tag{14}$$

Since σ_{yi}^2 is the power of the signal and σ_v^2 the power of the noise in the regarded 'spectral component', this result is also according to our intuitive expectation. Using eq.(10) we obtain:

$$\hat{\mathbf{s}} \;=\; \mathbf{A}^{-1}\hat{\mathbf{y}} = \mathbf{A}^{-1}\mathsf{diag}\left\{\frac{\sigma_{yi}^2}{\sigma_v^2 + \sigma_{yi}^2}\right\} \cdot \mathbf{w} \qquad = \; \mathbf{A}^T\mathsf{diag}\left\{\frac{\sigma_{yi}^2}{\sigma_v^2 + \sigma_{yi}^2}\right\} \cdot \mathbf{A} \cdot \mathbf{z} \tag{15}$$

This means that for obtaining a MMSE estimate of the signal, the observed signal vector $\mathbf{z}$ has to be transformed (rotated) into the canonical coordinate frame, the canonical coordinates have to be attenuated according to the fraction $\sigma_{yi}^2/(\sigma_v^2 + \sigma_{yi}^2)$ and finally rotated back into the original coordinate frame.

6 Conclusions

I have presented two main contributions to motion analysis in this paper. The first one is the generalization of the differential constraint equation, the exposition of its relation to directional variation measures, and its usage in terms of generalized tensor-based methods.

The second contribution is the prefilter design approach based on classical Wiener theory. It provides the possibility to adapt differential or tensor-based motion estimation schemes to the specific statistics of the given signal and the noise contained in it. This explicitly includes the design of different filters for the case of prior knowledge on the typical range of motion vectors; the spatio-temporal distribution of the displacement is directly reflected in the autocovariance function, and thus also considered in the optimum pre-filter derived from the autocovariance function.

References

1. Johansson, B., Farnebäck, G.: A theoretical comparison of different orientation tensors. In: Proceedings SSAB02 Symposium on Image Analysis, Lund (2002) 69–73
2. Knutsson, H.: What is so special about quadrature filters? In: International Conference on Image Processing, ICIP, Barcelona, IEEE Computer Society Press (2003)
3. Nordberg, K., Farnebäck, G.: An overview of methods for estimation and representation of local orientation. In: International Conference on Image Processing, ICIP, Barcelona, IEEE Computer Society Press (2003)
4. Bigün, J., Granlund, G.H.: Optimal orientation detection of linear symmetry. In: First International Conference on Computer Vision, ICCV, Washington, DC., IEEE Computer Society Press (1987) 433–438
5. Haussecker, H., Spies, H.: Motion. In: Handbook of Computer Vision and Applications, Academic Press (1999)
6. Jähne, B.: Digital Image Processing. 4th edn. Springer Verlag (1998)
7. Mühlich, M., Mester, R.: Subspace methods and equilibration in computer vision. In: Proc. Scandinavian Conference on Image Analysis. (2001)
8. Simoncelli, E.P.: Design of multi-dimensional derivatives filters. In: Intern. Conf. on Image Proc., Austin TX (1994)
9. Farid, H., Simoncelli, E.P.: Optimally rotation-equivariant directional derivative kernels. In: 7th Intern. Conf. on Computer Analysis of Images and Patterns, Kiel, Germany (1997)
10. Scharr, H.: Optimal Operators in Digital Image Processing. PhD thesis, Interdisciplinary Center for Scientific Computing, Univ. of Heidelberg (2000)
11. Elad, M., Teo, P., Hel-Or, Y.: Optimal filters for gradient-based motion estimation. In: Proc. Intern. Conf. on Computer Vision (ICCV'99). (1999)
12. Granlund, G.H., Knutsson, H.: Signal processing for computer vision. Kluwer (1995)
13. Mester, R.: A system-theoretic view on local motion estimation. In: Proc. IEEE SouthWest Symposium on Image Analysis and Interpretation, Santa Fé (NM), IEEE Computer Society (2002)
14. Freeman, W.T., Adelson, E.H.: The design and use of steerable filters. IEEE Transactions on Pattern Analysis and Machine Intelligence **13** (1991)
15. Mester, R.: A generalization of differential optical flow estimation using 3D covariance functions of signal and noise. In: Submitted in Nov. 2002 to IEEE Conf. on Computer Vision and Pattern Recognition. (2003)
16. Yu, W., Daniilidis, K., Sommer, G.: A new 3D orientation steerable filter. In: DAGM 2000, Springer (2000)
17. Dong, D.D., Atick, J.J.: Statistics of natural time-varying images. Computation in Neural Systems **6(3)** (1995) 345–358
18. Therrien, C.W.: Discrete random signals and statistical signal processing. Prentice Hall (1992)