EVALUATION OF COMPUTER VISION ALGORITHMS

The Role of Theory

S. S. BEAUCHEMIN AND R. BAJCSY
GRASP Laboratory
Department of Engineering and Information Sciences
University of Pennsylvania
Philadelphia PA 19104-6228
USA

Abstract. Undeniably, the numerical evaluation of Computer Vision algorithms is of utmost importance. However, often neglected is the role of theoretical knowledge to interpret the numerical performance of those algorithms. In addition, the lack of theoretical research in Computer Vision has long been recognized. In this contribution, we demonstrate that extended theoretical knowledge of a phenomenon enables one to design algorithms that are better suited for the task at hand and to evaluate the theoretical assumptions of other, similar algorithms. For instance, the problem posed by multiple image motions was poorly understood in the frequency domain yet frequency-based multiple motions algorithms were developed. We present algorithms for computing multiple image motions arising from occlusion and translucency which are capable of extracting the information-content of occlusion boundaries and distinguish between those and additive translucency phenomena. These algorithms are based on recent theoretical results on occlusion and translucency in the frequency domain and demonstrate that a complete theoretical understanding of a phenomenon is required in order to design adequate algorithms. We conclude by proposing an evaluation protocol which includes theoretical considerations and their influence on the numerical evaluation of algorithms.

1. Introduction

The importance of motion in image processing cannot be understated: in particular, approximations to image motion may be used to estimate 3D scene properties and motion parameters from a moving visual sensor

(Longuet-Higgins, 81), to perform motion segmentation (Murray and Buxton, 87), to compute the focus of expansion and time-to-collision (Overington, 87), to perform motion-compensated image encoding (Musmann *et al.*, 85), to compute stereo disparity (Jenkin *et al.*, 91) and to measure biological parameters in medical imagery (Prince and McVeigh, 92).

Based on recent theoretical developments in discontinuous motion, we devise multiple motion algorithms. We consider 1D and 2D signals, adopt a constant model of velocity and use a robust statistical procedure to extract multiple motions from local frequency spectra. The motion information provided by the algorithms includes single velocity, multiple (2) velocities, assessment of transparency versus occlusion, and upon occlusion events, the orientation of the occlusion boundary and the identification of the occluding signal.

1.1. LITERATURE SURVEY

Computing multiple motions is a complex and rarely undertaken task. Indeed, most of the existing optical flow methods that have appeared in the literature make an explicit use of the single motion hypothesis. However, at motion discontinuities, where the information content of a signal mostly resides, the hypothesis is violated. Area-based and feature-based correlation techniques are equally sensitive to occlusion as local image structures and features appear and disappear from one image to the next. To further complicate matters, regularization techniques which impose a degree of continuity to optical flow are also clearly inadequate over occlusion boundaries. However, in the more recent research in optical flow, the non-linear, discontinuous and multiple-valued nature of image motion in the coordinates of the image plane has been recognized (Barron and Beauchemin, 95).

In order to allow multiple motion events in optical flow estimation processes, a number of strategies have been devised, such as strong intensity gradients acting as inhibitors of flow coherence (Nagel, 87) and robust estimators designed to capture dominant motions (Black, 97). Other techniques such as clustering (Schunck, 89), superposed motion layers and distributions (Shizawa and Mase, 91; Wang and Adelson, 93), parametric models of motion with discontinuous functions (Negahdaripour and Lee, 92; Black and Jepson, 94) and mixtures of probability densities (Jepson and Black, 91) have appeared.

Our approach emanates from recent theoretical results (Beauchemin and Barron, 97; Beauchemin and Barron, 97; Beauchemin *et al.*, 97) describing the Fourier structure of occlusion and translucency phenomena for constant and linear models of optical flow and points to an evaluation protocol capable of assessing the impact of theoretical considerations onto the

numerical evaluation of algorithms.

2. Structure of Occlusion

We proceed to describe the structure of occlusion events in the frequency domain for 1D and 2D signals composed of an arbitrary number of distinct frequencies.

Let $\mathbf{I}_1(x)$ and $\mathbf{I}_2(x)$ be 1D functions satisfying Dirichlet conditions such that they may be expressed as complex exponential series expansions:

$$\mathbf{I}_1(x) = \sum_{n=-\infty}^{\infty} c_{1n} e^{ink_1 x} \quad \text{and} \quad \mathbf{I}_2(x) = \sum_{n=-\infty}^{\infty} c_{2n} e^{ink_2 x}, \tag{1}$$

where n is integer, c_{1n} and c_{2n} are complex coefficients and k_1 and k_2 are the fundamental frequencies of both signals.

Let $\mathbf{I}_1(x,t) = \mathbf{I}_1(v_1^{(0)}(x,t))$ and $\mathbf{I}_2(x,t) = \mathbf{I}_2(v_2^{(0)}(x,t))$, where $v_i^{(0)} = x - a_i t$. The frequency spectrum of the occlusion is:

$$\begin{aligned}
\hat{\mathbf{I}}(k,\omega) &= \pi \sum_{n=i-\infty}^{\infty} c_{1n}\delta(k - nk_1, \omega + nk_1 a_{16}) \\
&+ (1-\pi) \sum_{n=-\infty}^{\infty} c_{2n}\delta(k - nk_2, \omega + nk_2 a_{26}) \\
&+ i \sum_{n=-\infty}^{\infty} \left(\frac{c_{2n}\delta(ka_{12} + \omega - nk_2\Delta a_6)}{(k - nk_2)} - \frac{c_{1n}\delta(ka_{12} + \omega)}{(k - nk_1)} \right),
\end{aligned} \tag{2}$$

where $\Delta a_6 = a_{16} - a_{26}$.

In the 1D case, equation (2) reveals that the frequency spectra of both signals are preserved to within scaling factors. In addition, the Dirac delta functions $\delta(ka_{12} + \omega)$ and $\delta(ka_{12} + \omega \pm k_2\Delta a_{22})$ constitute linear spectra, intersecting the frequencies of both the occluding and occluded signals, and are oriented in the direction of the constraint line pertaining to the occluding signal. Figure 1 shows a typical example with 1D translating sinusoids in an occlusion scene.

Similarly, The frequency spectra for 2D signals are planar and preserved to within scaling factors under occlusion. In addition, the distortion cast by the occlusion boundary fits oriented planes parallel to the planar spectrum of the occluding signal.

Let $\mathbf{I}_1(\mathbf{x})$ and $\mathbf{I}_2(\mathbf{x})$ be 2D functions satisfying Dirichlet conditions such that they may be expressed as complex exponential series expansions:

$$\mathbf{I}_1(\mathbf{x}) = \sum_{\mathbf{n}=-\vec{\infty}}^{\vec{\infty}} c_{1\mathbf{n}} e^{i\mathbf{x}^T N\mathbf{k}_1} \quad \text{and} \quad \mathbf{I}_2(\mathbf{x}) = \sum_{\mathbf{n}=-\vec{\infty}}^{\vec{\infty}} c_{2\mathbf{n}} e^{i\mathbf{x}^T N\mathbf{k}_2}, \tag{3}$$

where $\mathbf{n} = (n_x, n_y)^T$ and $N = \mathbf{n}^T I$ are integers, $\mathbf{x}$ are spatial coordinates, $\mathbf{k}_1 = (k_{1x}, k_{1y})^T$ and $\mathbf{k}_2 = (k_{2x}, k_{2y})^T$ are fundamental frequencies and $c_{1\mathbf{n}}$ and $c_{2\mathbf{n}}$ are complex coefficients. Let $\mathbf{I}_1(\mathbf{x}, t) = \mathbf{I}_1(\mathbf{v}_1^{(0)}(\mathbf{x}, t))$, $\mathbf{I}_2(\mathbf{x}, t) = \mathbf{I}_2(\mathbf{v}_2^{(0)}(\mathbf{x}, t))$, where $\mathbf{v}_i^{(0)} = \mathbf{x} - \mathbf{a}_i t$, and the occluding boundary be locally represented by:

$$\mathbf{U}(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x}^T\mathbf{n}_1 \geq 0 \\ 0 & \text{otherwise,} \end{cases} \tag{4}$$

where $\mathbf{n}_1$ is a vector normal to the occluding boundary at $\mathbf{x}$. The frequency spectrum of the occlusion is:

$$\begin{aligned} \hat{\mathbf{I}}(\mathbf{k}, \omega) &= \pi \sum_{\mathbf{n}=-\vec{\infty}}^{\vec{\infty}} c_{1\mathbf{n}} \delta(\mathbf{k} - N\mathbf{k}_1, \omega + \mathbf{a}_1^T N\mathbf{k}_1) \\ &+ (1-\pi) \sum_{\mathbf{n}=-\vec{\infty}}^{\vec{\infty}} c_{2\mathbf{n}} \delta(\mathbf{k} - N\mathbf{k}_2, \omega + \mathbf{a}_2^T N\mathbf{k}_2) \\ &+ i \sum_{\mathbf{n}=-\vec{\infty}}^{\vec{\infty},} \left(\frac{c_{2\mathbf{n}} \delta((\mathbf{k} - N\mathbf{k}_2)^T \mathbf{n}_1^-, \mathbf{k}^T\mathbf{a}_1 + \omega - \Delta\mathbf{a}_6^T N\mathbf{k}_2)}{(\mathbf{k} - N\mathbf{k}_2)^T \mathbf{n}_1} \right. \\ &- \left. \frac{c_{1\mathbf{n}} \delta((\mathbf{k} - N\mathbf{k}_1)^T \mathbf{n}_1^-, \mathbf{k}^T\mathbf{a}_1 + \omega)}{(\mathbf{k} - N\mathbf{k}_1)^T \mathbf{n}_1} \right) \end{aligned} \tag{5}$$

where $\Delta\mathbf{a}_6 = \mathbf{a}_1 - \mathbf{a}_2$.

Equation (5) is a generalization of equation (2) from 1D to 2D signals and its geometric interpretation is similar. For instance, frequencies $(N\mathbf{k}_1, -\mathbf{a}_1^T N\mathbf{k}_1)$ and $(N\mathbf{k}_2, -\mathbf{a}_2^T N\mathbf{k}_2)$ fit the constraint planes of the occluding and occluded signals, defined as $\mathbf{k}_1^T\mathbf{a}_1 + \omega = 0$ and $\mathbf{k}^T\mathbf{a}_2 + \omega = 0$. In the distortion term, the Dirac δ function with arguments $(\mathbf{k} - N\mathbf{k}_2)^T\mathbf{n}_1^-$ and $\mathbf{k}^T\mathbf{a}_1 + \omega - \Delta\mathbf{a}_6^T N\mathbf{k}_2$ represent a set of lines parallel to the constraint plane of the occluding signal $\mathbf{k}^T\mathbf{a}_1 + \omega = 0$ and, for every discrete frequency $N\mathbf{k}_1$ and $N\mathbf{k}_2$ exhibited by both signals, there is a frequency spectrum fitting the lines given by the intersection of planes $\mathbf{k}^T\mathbf{a}_1 + \omega - \Delta\mathbf{a}_6^T N\mathbf{k}_2 = 0$ and $(\mathbf{k} - N\mathbf{k}_1)^T\mathbf{n}_1^- = 0$. The magnitudes of these spectra are determined by their corresponding scaling functions $c_{1\mathbf{n}}[(\mathbf{k} - N\mathbf{k}_1)^T\mathbf{n}_1\ ]^{-1}$ and $c_{2\mathbf{n}}[(\mathbf{k} - N\mathbf{k}_2)^T\mathbf{n}_1\ ]^{-1}$.

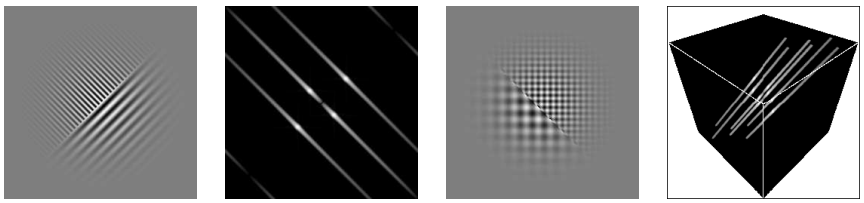

Figure 1. **(from left to right): a)** *Gaussian-windowed 1D signal with sinusoidals acting as occluding and occluded surfaces. The occluding signal has spatial frequency* $k_1 = \frac{2\pi}{16}$ *and velocity* $\mathbf{v}_1 = (1, 1)$. *The occluded signal has frequency* $k_2 = \frac{2\pi}{8}$ *and velocity* $\mathbf{v}_2 = (-1, 1)$. **b)** *Fourier spectrum of* **a)**. **c)** *Gaussian-windowed 2D signal with sinusoidals acting as occluding and occluded surfaces. The occluding signal has spatial frequency* $\mathbf{k}_1 = (\frac{2\pi}{16}, \frac{2\pi}{16})$ *and velocity* $\mathbf{v}_1 = (1, 1, 1)$. *The occluded signal has frequency* $\mathbf{k}_2 = (\frac{2\pi}{8}, \frac{2\pi}{8})$ *and velocity* $\mathbf{v}_2 = (-1, -1, 1)$. **d)** *Fourier spectrum of* **c)**.

3. Estimation of Multiple Image Motion

Equations (2) and (5) provide a model of the Fourier spectrum at an occlusion boundary. We devise, for both 1D and 2D image signals, algorithms capable of extracting multiple velocity measurements along with the information-content of occlusion boundaries.

3.1. 1D ALGORITHM

Given a frequency measurement $\hat{\mathbf{m}}_j = (\hat{k}_j, \hat{\omega}_j)^T$, its corresponding velocity estimate is given by $\hat{\mathbf{v}}_i = (-\hat{\omega}_j/\hat{k}_j, 1)^T$ and an error metric corresponding to the angular deviation between a measurement $\hat{\mathbf{m}}_j$ and an estimate of the i^{th} velocity $\hat{\mathbf{v}}_i$, under the assumption that $\theta \approx \sin\theta$, may be defined as (Jepson and Black, 91):

$$\xi(\hat{\mathbf{m}}_j, \hat{\mathbf{v}}_i) = \frac{\hat{\mathbf{m}}_j^T \hat{\mathbf{v}}_i}{\|\hat{\mathbf{m}}_j\|_2 \|\hat{\mathbf{v}}_i\|_2}. \tag{6}$$

Under the assumption that angular errors are normally distributed, we define a mixture of normal distributions to account for multiple motions as:

$$G(\hat{\mathbf{m}}_j) = \sum_{i=0}^{2} \pi_i f_i(\hat{\mathbf{m}}_j, \hat{\mathbf{v}}_i), \tag{7}$$

which is the PDF for measurement $\hat{\mathbf{m}}_j$ and where π_i is a mixture probability. Aside from posing homoscedasticity, we use a uniform outlier distribution. Measurements at a predetermined distance from other distributions should be considered as outliers and not enter the estimation process. The

constant probability of observing a noisy measurement is expressed as:

$$\frac{1}{\sqrt{2\pi}\sigma_v} e^{\frac{-\lambda^2}{2\sigma_n^2}}, \tag{8}$$

from which it is noted that measurements at λ standard deviations from the means of the normal distributions are considered to be corrupted by noise.

With the hypothesis of homoscedasticity, constant standard deviation and uniform distribution of noisy measurements, we establish the iterative equations for the Expectation-Maximization algorithm. The expectation step is the computation of posterior probabilities, which we write as:

$$\hat{\tau}_{ij}^{(k)} = \frac{\hat{\pi}_i^{(k)} e^{\frac{-1}{2\sigma_v^2}\xi^2(\hat{\mathbf{m}}_j, \hat{\mathbf{v}}_i)}}{\sum_{t=1}^{2} \hat{\pi}_t^{(k)} e^{\frac{-1}{2\sigma_v^2}\xi^2(\hat{\mathbf{m}}_j, \hat{\mathbf{v}}_t)} + \hat{\pi}_0^{(k)} e^{\frac{-\lambda^2}{2\sigma_n^2}}} \tag{9}$$

for $i = 1, 2$, the number of normal distributions and $j = 1, \ldots, n$ the number of measurements. For the uniform distribution of noisy measurements we write

$$\hat{\tau}_{0j}^{(k)} = \frac{\hat{\pi}_0^{(k)} e^{\frac{-\lambda^2}{2\sigma_v^2}}}{\sum_{t=1}^{2} \hat{\pi}_t^{(k)} e^{\frac{-1}{2\sigma_v^2}\xi^2(\hat{\mathbf{m}}_j, \hat{\mathbf{v}}_t)} + \hat{\pi}_0^{(k)} e^{\frac{-\lambda^2}{2\sigma_n^2}}} \tag{10}$$

for $j = 1, \ldots, n$. The equations for the maximization step, in which the velocities and mixture probabilities are updated, we write:

$$\hat{\mathbf{v}}_i^{(k+1)} = \frac{\sum_{j=1}^{n} \hat{\tau}_{ij}^{(k)} \kappa(\hat{\mathbf{m}}_j) \hat{\mathbf{m}}_j^{-}}{\hat{\pi}_i^{(k)} \sum_{j=1}^{n} \kappa(\hat{\mathbf{m}}_j)} \qquad \hat{\pi}_i^{(k+1)} = \frac{\sum_{j=1}^{n} \kappa(\hat{\mathbf{m}}_j) \hat{\tau}_{ij}^{(k)}}{\sum_{j=1}^{n} \kappa(\hat{\mathbf{m}}_j)}, \tag{11}$$

where $\kappa(\hat{\mathbf{m}}_j)$ is the magnitude of $\hat{\mathbf{m}}_j$ and $\hat{\mathbf{m}}_j^{-}$ is its negative reciprocal.

In order to identify the spectra associated with occluding boundaries, we first find peak frequency measurements for both signals. That is to say, we find for signal t, the frequency $\hat{\mathbf{m}}_t$ such that $\tau_{tk} > \tau_{ik}$ for $t \neq i$ and $\kappa(\hat{\mathbf{m}}_t)$ is maximal and determine the strength of measurements $\hat{\mathbf{m}}_j$ along the direction perpendicular to the hypothesized occluding velocity at the peak frequency of the hypothesized occluded signal.

To test for the signal corresponding to velocity $\mathbf{v}_i$ as occluding, the procedure is to first consider only those measurements $\hat{\mathbf{m}}_j$ belonging to the uniform distribution of the mixture: $\tau_{0j} > \tau_{ij}$, for $i = 1, 2$ and $j = 1 \ldots n$, as determined by the EM algorithm and the peak frequency of the signal corresponding to velocity $\mathbf{v}_t$, where $t \neq i$. We then proceed with the computation of the strengths of measurements confirming this hypothesis, once again using mixture probabilities.

Among measurements belonging to the uniform noise distribution, we compute their posterior probability of being part of the distortion spectra cast by the hypothesized occlusion and we also determine the posterior probabilities of the measurements to be from the uniform noise distribution to the exclusion of the spectra of the occlusion. Mixture proportions may be obtained from these posterior probabilities that assess the hypothesis under test. Thus, if velocity $\hat{\mathbf{v}}_i$ is occluding, then the strengths of measurements confirming this hypothesis outnumber those pertaining to its contrary. This hypothesis-testing method is applied to determine the image events giving rise to multiple velocities.

3.2. 2D ALGORITHM

The algorithm for 2D signals is essentially similar to the 1D algorithm we described. The measurements $\hat{\mathbf{m}}_j = (k_{xj}, k_{yj}, \omega_j)^T$ and velocity estimates $\hat{\mathbf{v}}_i = (v_x, v_y, v_t)^T$ are used in the error metric (6) to deterimine the posterior probabilities τ_{ij}, as is the case with the 1D algorithm. However, the choice of velocity estimates differs substantially. In the case of 2D signals, the velocity estimates at each EM iteration must maximize the numerator exponential of (9). In this case, we follow the approach adopted by Jepson and Black (Jepson and Black, 91), and consider the square of the error metric (6) as the equation for which the solutions yield velocity estimates. We observe that $\xi^2(\hat{\mathbf{m}}_j, \hat{\mathbf{v}}_i)$ may be written in matrix form as

$$(\mathbf{m}_j^T \mathbf{v}_i)^2 = \mathbf{v}_i^T M_j \mathbf{v}_i \tag{12}$$

where $M_j = \hat{\mathbf{m}}_j \hat{\mathbf{m}}_j^T$. By selecting the eigenvector corresponding to the minimum eigenvalue of M_j for $\mathbf{v}_i$, we minimize (12). Since M_j is real and symmetric, its eigenvalues are real and non-degenerate and the eigenvectors form an orthogonal basis in the space of measurements. In light of these observations, we define

$$\Sigma_i^{(k+1)} = \frac{\sum_{j=1}^n \tau_{ij}^{(k)} \kappa(\hat{\mathbf{m}}_j) M_j}{\sum_{j=1}^n \kappa(\hat{\mathbf{m}}_j)} \tag{13}$$

as the matrix from which the velocity estimate $\mathbf{v}_i^{(k+1)}$ is to be obtained in the form of the eigenvector $\mathbf{e}_i^{(k+1)}$ corresponding to the minimum eigenvalue $e_i^{(k+1)}$ of Σ_i. The minimum eigenvalue holds information about the velocity estimate obtained from its corresponding eigenvector. A zero value for e_i indicates that the velocity measurement is normal, whereas a non zero value indicates a full velocity measurement (Jahne, 90). To see this, consider a set of observations consisting of collinear measurements, consistent with

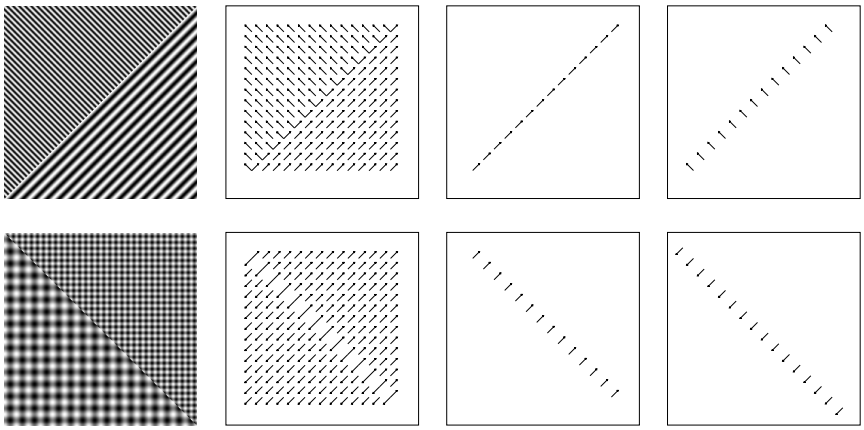

Figure 2. Experimental results. Occluding frequencies and velocities are labeled as k_1 *or* $\mathbf{k}_1$ *and* $\mathbf{v}_1$. **(top, left to right): a)** *1D occlusion imagery with spatial frequencies* $k_1 = \frac{2\pi}{16}$ *and* $k_2 = \frac{2\pi}{8}$ *and velcities* $\mathbf{v}_1 = (1, 1)$ *and* $\mathbf{v}_2 = (-1, 1)$. **b)** *Full optical flow.* **c)** *Occluding velocities.* **d)** *Occluded velocities.* **(bottom, left to right): e)** *2D occlusion imagery with spatial frequencies* $\mathbf{k}_1 = (\frac{2\pi}{16}, \frac{2\pi}{16})$ *and* $\mathbf{k}_2 = (\frac{2\pi}{8}, \frac{2\pi}{8})$ *and velcities* $\mathbf{v}_1 = (1, 1, 1)$ *and* $\mathbf{v}_2 = (-1, -1, 1)$. **f)** *Full optical flow.* **g)** *Occluding velocities.* **h)** *Occluded velocities.*

a normal velocity. It is observed that in such circumstances, the lines of matrix Σ_i are linearly dependent, leading to a minimum eigenvalue of value zero. The final eigenvalues e_i contain information on the nature of the measured velocities that is very relevant in most uses of image velocity.

Under the hypothesis of a straight-edged occlusion boundary, its normal may be estimated from the frequency structure of the occlusion (Beauchemin and Barron, 97). To perform this estimation, the algorithm recovers the orientation of the spectrum cast by the occlusion about the maximum frequency of the occluded signal, within a plane parallel to that of the occluding signal. To perform this estimation, it is necessary to include an EM iteration which converges to this linear orientation within the specified constraint plane.

4. Experiments

We report two numerical experiments on synthetic sinusoidal imagery composed of a 1D and a 2D occlusion scene as described in Figure 2. The images used for these experiments are noise-free and so are the computed optical flow fields. It has been clear for some time that a number of vision algorithms still fail to meet this fundamental criterion (Barron *et al.*, 94). Local frequency measurements are obtained for an image location by computing

a local Fast Fourier Transform within regions of side size 32. We observed that 30 iterations were sufficient for the EM algorithm to converge. The initial estimates for velocities and mixture proportions may be chosen randomly, but we prefer to have initial velocity estimates as apart as possible in order to avoid convergence of both estimates to a single peak.

4.1. DISCUSSION

The nature of image motion, most particularly discontinuous motion in frequency space, has long been unclear. The algorithms proposed herein are based on a firm theoretical framework which describes the coherent behavior of occlusion events in frequency space. Indeed, we strongly believe that further developments in the field of optical flow and motion analysis ought to be based on firmly established theoretical backgrounds rather than incidental evidence, as is sometimes the case (Jain and Binford, 90).

In light of this, we propose a novel evaluation protocol which includes theoretical considerations and their impact on numerical evaluations. By considering the body of theoretical knowledge concerning a phonomenon, such as image motion for instance, one is able to assess the numerical impact that each part provided by a theoretical model has. For example, a formulation of the algorithms in this contribution that would use the single motion hypothesis may be numerically compared with the complete algorithms as described on a relevant set of test images, thus numerically assessing the gain provided by a more complete theoretical model. The cycle in which this protocol operates may be defined by the following steps:

- Incrementally incorporate theoretical knowledge of the phenomenon under study into a descriptive, computational model.
- Derive the corresponding algorithms based on the computational model.
- Numerically evaluate these algorithms on an appropriate set of images and assess their accuracy.
- Repeat the steps until the computational model is entirely compliant with the current body of applicable theoretical knowledge about the phenomenon under study.

This iterative protocol allows to assess the impact that theoretical knowledge has on numerical evaluations of algorithms and we believe it to be of importance because of its ability to provide the necessary insights pertaining to observed accuracy differentials of various algortihms.

However, by proposing such a protocol, we do not imply to shift the emphasis away from pure numerical evaluations. On the contrary, we merely suggest to include such comparative numerical evaluations within this framework in order to understand the numerical impact of computational models derived from various bodies of theoretical knowledge.

References

J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *IJCV*, 12(1):43–77, 1994.

S. S. Beauchemin and J. L. Barron. The computation of optical flow. *ACM Computing Surveys*, 27(3):433–467, 1995.

S. S. Beauchemin and J. L. Barron. The local frequency structure of 1d occluding image signals. *IEEE PAMI*, submitted, 1997.

S. S. Beauchemin and J. L. Barron. A theoretical framework for discontinuous optical flow. submitted, 1997.

S. S. Beauchemin, A. Chalifour, and J. L. Barron. Discontinuous optical flow: Recent theoretical results. In *Vision Interface*, pages 57–64, Kelowna, Canada, May 1997.

M. J. Black. A robust gradient-method for determining optical flow. Technical Report YALEU/DCS/RR-891, Yale University, New-Haven, CT, 1991.

M. J. Black and A. Jepson. Estimating optical flow in segmented images using variable-order parametric models with local deformations. Technical Report SPL-94-053, Xerox Systems and Practices Laboratory, Palo Alto, California, 1994.

B. Jahne. Motion determination in space-time images. In *Proceedings of ECCV*, pages 161–173, Antibes, France, April 1990.

R. C. Jain and T. O. Binford. Ignorance, myopia and naivete in computer vision systems. *CVGIP:IU*, 53:112–117, 1991.

M. R. M. Jenkin, A. D. Jepson, and J. K. Tsotsos. Techniques for disparity measurement. *CVGIP*, 53(1):14–30, 1991.

A. D. Jepson and M. Black. Mixture models for optical flow computation. In *IEEE Proceedings of CVPR*, pages 760–761, New York, New York, June 1993.

H. C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 223:133–135, 1981.

D. W. Murray and B. F. Buxton. Scene segmentation from visual motion using global optimization. *IEEE PAMI*, 9(2):220–228, 1987.

H. G. Musmann, P. Pirsch, and H. J. Grallert. Advances in picture coding. *Proc of IEEE*, 73(4):523–548, 1985.

H.-H. Nagel. On the estimation of optical flow: Relations between different approaches and some new results. *Artificial Intelligence*, 33:299–324, 1987.

S. Negahdaripour and S. Lee. Motion recovery from image sequences using only first order optical flow information. *IJCV*, 9(3):163–184, 1992.

I. Overington. Gradient-based flow segmentation and location of the focus of expansion. In *Alvey Vision Conference*, pages 860–870, University of Cambridge, England, September 1987.

J. L. Prince and E. R. McVeigh. Motion estimation from tagged mr image sequences. *IEEE Trans. on Medical Images*, 11(2):238–249, 1992.

B. G. Schunck. Image flow segmentation and estimation by constraint line clustering. *IEEE PAMI*, 11(10):1010–1027, 1989.

M. Shizawa and K. Mase. Principle of superposition: A common computational framework for analysis of multiple motion. In *IEEE Proceedings of Workshop on Visual Motion*, pages 164–172, Princeton, New Jersey, October 1991.

J. Y. A. Wang and E. H. Adelson. Layered representation for motion analysis. In *Proceedings of CVPR*, 1993.