Fast Algebraic Multigrid For Discontinuous Optical Flow Estimation

Sugata Ghosal and Petr Č. Vaněk

Abstract

Multiple moving objects, occluded objects, or even a moving object against the static background give rise to discontinuities in the optical flow field in corresponding image sequences. Uniform global smoothing-based regularization cannot provide accurate estimates if the optical flow is discontinuous. A 'weighted anisotropic smoothness'-based regularization technique is proposed for accurately estimating dense optical flow field. Weighted sum of the magnitude and first-order spatial derivatives of the optic flow field is used for regularization. Less regularization is performed at points where strong intensity gradient information is available. The solution at any point is interpolated more from those at neighboring points along the weaker intensity gradient-component. A new scalable algebraic multigrid algorithm is used to efficiently solve resulting second-order elliptic partial differential equations with discontinuities and anisotropies in coefficients. This multigrid technique utilizes information about coefficients of the stiffness matrix in generating coarse spaces. The proposed optical flow computation technique outperforms Horn–Schunck's and Nagel–Enkelmann's techniques in terms of accuracy as well as speed. Experimental results and execution times are reported for a variety of synthetic and real sequences to demonstrate the effectiveness of the weighted anisotropic smoothness-based regularization and the algebraic multigrid solver.

KeyWords— **Motion analysis, discontinuous optical flow estimation, weighted anisotropic smoothness, partial differential equation (PDE), scalable algorithm, algebraic multigrid.**

I. Introduction

Distribution of the temporal displacement of intensity patterns between two frames in an image sequence is known as the *optical flow*. Computation of optical flow has drawn considerable attention in computer vision. Popular optical flow techniques can be broadly classified into four categories: (1) *Derivative*-based, (2) *Region matching*-based, (3) *Energy*-based, and (4) *Phase*-based methods [2]. Derivative-based techniques are most popular to date because of their relative speed and their ability to produce good qualitative solution. These techniques compute optical flow field from spatial and temporal derivatives of intensity values (or other gray-level properties) in an image sequence assuming the conservation of intensity values (or properties) in time. Single *data constraint equation* at any image point cannot provide unique solution to the optical flow problem. Thus, additional constraints are required to solve the problem. These constraints can be either *global* or *local* in nature. Most popular global constraints require spatial smoothness of the solution. Horn and Schunck employed the first-order spatial smoothness of optical flow components, and the solution is obtained by solving the corresponding Euler-Lagrangian equations by iterative Gauss-Seidel relaxation [10]. Nagel suggested an *oriented smoothness* constraint in which smoothness requirement is only imposed orthogonal to the intensity gradient (or principal curvature of local intensity profile) in an attempt to handle occlusion and motion discontinuity [15],[16]. Second-order spatial derivatives of intensity values were also used for accurately solving optical flow problem at gray-level corner points and the solution is propagated to other regions of the image [15]. On the other hand, local constraints assume some parametric model

The authors are with the Center for Computational Mathematics, University of Colorado at Denver, Denver, CO 80217.

This research is supported by the National Science Foundation grants ASC–9121431, ASC–9217394, and ASC–9404734.

of the underlying optical flow field in the neighborhood of an image point. The flow field is typically obtained by solving an overdetermined system of intensity (or other gray-level properties) constraint equations [5],[8],[11],[21],[24].

The optical flow field and the actual projected 2-D velocity match closely at image points with high intensity gradient values [23], [9], and high gradient points frequently correspond to locations of 3-D motion discontinuities. Thus, it is essential to calculate the optical flow as accurately as possible at these points. At the same time, high density of the flow vectors is desirable for segmenting 3-D scenes into moving and stationary regions with applications to IVHS and mobile robotics, and for data compression of video sequences for multimedia applications etc. In spite of wide-spread popularity, regularization-based techniques suffer from inaccuracies near motion discontinuities that result from multiple objects in the scene moving with different relative velocities w.r.t. the sensor, or in regions of occlusion between different moving objects. The usual smoothness constraints have adverse effect on the estimated flow field because they tend to blur motion discontinuities. For example, Horn and Schunck's smoothness constraint is globally uniform in nature, and penalizes the deviation from first-order spatial variations of flow components equally at every image point. Nagel and Enkelmann, on the other hand, employed an oriented smoothness term which essentially smoothes the optical flow field in the direction perpendicular to the gradient and principal curvature of local intensity profiles. The smoothness term involving intensity curvatures may result in numerical instability (Schnörr in [20] advocated multiple data constraints to provide stable solution). Also, oriented smoothness-based techniques are highly sensitive to the gradient direction and prone to serious problems, if the local intensity gradient is not perpendicular to the local edge contour (edge contours with varying intensity gradient). On the other hand, local methods are usually faster and may provide more accurate solution at image points where strong intensity gradient information is available. However, these are less stable than global regularization-based techniques. Also, these techniques do not allow for propagating the optical flow field to uniformly bright (zero or low intensity gradient) image regions from nearby points. Thus, the detected flow field can be sparse. Markov random field (MRF) models are often used to alleviate the shortcomings of derivative-based techniques [9] (and references therein). These techniques, however, often suffer from the presence of multiple local minima and consequently very high solution time.

In this paper, we propose a *weighted anisotropic* smoothness-based regularization technique for discontinuous and dense optical flow estimation, that attempts to combine the strengths of both global and local derivative-based methods. This smoothness constraint regularizes the solution more at locations of low 'confidence', where little information is available in terms of spatial intensity variation. The amount of regularization is small at locations of high 'confidence' with strong intensity gradient, e.g., an edge, to allow for potential motion discontinuities. The relative strength of regularization along x- and y-axis is dependent on the relative 'confidence' of available spatial intensity gradients along these axes. If the intensity gradient-component is stronger along one axis, it is reasonable to allow the data constraint term to play a greater role along that axis, and thus, the effect of smoothness term is kept small. Also, a zero-order smoothness (i.e., small magnitude) of the solution is imposed to attenuate the flow field in static regions of the image.

The weighted anisotropic smoothness-based regularization gives rise to second-order elliptic partial differential equations (PDEs) with discontinuities and anisotropies in coefficients. The rate of convergence of popular Gauss-Seidel iteration depends strongly (two-fold) on the size of discretization, h (approximately, the square root of number of image points) and also on the ratio of maximum to minimum values of coefficients, that can be extremely large for real images.

Gauss-Seidel relaxation methods has a time complexity of $O(n^2)$, which is not at all conducive for real-time processing. Consequently, multigrid methods have been used for solving computer vision problems with $O(n)$ time complexity [7], [17], [18], [19]. The rate of convergence of standard multigrid algorithms is h-independent for second-order problems; however, it depends adversely on the coefficients. We employ an algebraic multigrid algorithm, reported in [13] that can handle second-order elliptic problems with general coefficients, including *jumps* (discontinuities) and anisotropies, without deterioration in the rate of convergence. This multigrid technique is based on the concept of *smoothed aggregation*, introduced recently by Vaněk [12] and employs coarse spaces by explicitly utilizing the information about the coefficients of the *stiffness* matrix, arising from discretization of the second-order elliptic PDEs. This multigrid solver requires information about matrix entries only; thus it is fully black-box.

The rest of this paper is organized as follows. Section II deals with the proposed optical flow technique that employs a weighted anisotropic smoothness constraint for regularization. The algebraic multigrid algorithm for solving the system of linear algebraic equations, arising from finite-difference discretization of the Euler-Lagrangian PDEs is presented in Section III. Experimental results are presented in Section IV with several synthetic and real images to demonstrate the effectiveness of the proposed smoothness constraint. Speed and scalability of the algebraic multigrid algorithm are experimentally demonstrated on a serial machine. Conclusions and future research issues are discussed in Section V.

II. Regularization for Discontinuous Optical Flow Estimation

Optical flow computation is a well-known ill-posed problem where available spatiotemporal information do not constrain the solution sufficiently. Therefore, smoothness constraint-based regularization is performed to compute the optical flow. In this section, a weighted and anisotropic smoothness-based regularization technique is proposed that provides relatively accurate estimates of discontinuous optical flow field from a two-frame sequence of images. Such regularization-based optical flow computation technique results in a strictly convex variational problem, that allows application of multilevel preconditioner-based iterative methods for very fast numerical solution.

A. Weighted anisotropic smoothness term

Global regularization-based techniques compute the optical flow field by minimizing a quadratic functional. This functional, typically, is a sum of data constraint term and a smoothness term. The data constraint term arises from the conservation of intensity $E(x,y)$ of a moving point $(x(t),y(t))$ over two frames, i.e.,

$$\begin{aligned} & \frac{d}{dt}E(x(t),y(t),t)=0 \\ \text{or,}\quad & \frac{\partial E}{\partial x}\frac{dx(t)}{dt}+\frac{\partial E}{\partial y}\frac{dy(t)}{dt}+\frac{\partial E}{\partial t}=0 \\ \text{or,}\quad & E_x u+E_y v+E_t=0 \end{aligned} \tag{1}$$

where, (u,v) is the so-called optical flow field at point (x,y). This data constraint term has been used extensively in optical flow literature. Smoothness-based global constraints are used in practice. Horn and Schunck proposed a first-order smoothness term for regularization,

$$\sum_{sm}=\lambda[u_x^2+u_y^2+v_x^2+v_y^2] \tag{2}$$

Thus, the optical flow is computed by minimizing the functional,

$$F = \int_{\Omega} (E_x u + E_y v + E_t)^2 d\Omega + \int_{\Omega} \sum_{sm} d\Omega \tag{3}$$

The first-order spatial variations of optical flow components are penalized uniformly by a constant, λ at every image point. Such uniform smoothness term penalizes the formation of sharp motion discontinuities which may frequently occur in a scene with multiple moving objects or in occluded areas. In an attempt to alleviate this problem, Nagel and Enkelmann proposed an oriented smoothness constraint of the form [16],

$$\sum_{sm} = \frac{\lambda}{E_x^2 + E_y^2 + 2\delta} \left[(E_y u_x - E_x u_y)^2 + (E_y v_x - E_x v_y)^2 + \delta(u_x^2 + u_y^2 + v_x^2 + v_y^2) \right] \tag{4}$$

If the image gradient, $E_x^2 + E_y^2$ is small, the smoothness term is dominated by an isotropic component (same as Horn and Schunck's), and the solution (u, v) at that point is obtained by interpolating equally (which we can call *Laplacian smoothing*) from all neighboring image points. If the gradient is large, the isotropic component negligibly contributes to the regularization process, and smoothing is determined by

$$\begin{aligned} \sum_{sm} &\approx \frac{\lambda}{E_x^2 + E_y^2} [(E_y u_x - E_x u_y)^2 + (E_y v_x - E_x v_y)^2] \\ &= \frac{\lambda}{k^2 + 1} [(k u_x - u_y)^2 + (k v_x - v_y)^2] \end{aligned} \tag{5}$$

where, $E_y = k E_x$. In principle, such a smoothness term is very attractive, and ensures that "less smoothing is done in the direction of the gradient ... of intensity than in the orthogonal direction", and that "we smooth more where the gradient is small" [1].

A closer look at this oriented smoothness term reveals some unsuitable restrictions. It is very reasonable to enforce smoothness along the edge contour only. Oriented smoothness constraint enforces smoothness of the solution in the direction orthogonal to that of the intensity gradient. However, the topological gradient may not be parallel to the intensity gradient. In such case, oriented smoothness-based regularization will enforce smoothness of the solution in inappropriate directions, and consequently the resulting solution may not be smooth along the topological edge contour (as demonstrated in Fig. 1). Consider the occluded patterns, shown in Fig. 1(a). One of the patterns is moving to the upper left corner with velocity $(-1, 1)$, and the other one to the lower right with velocity $(1, -1)$. Intensity is varying sinusoidally along x-axis inside both the patterns, and is *exactly* conserved in the entire sequence. The optical flow field detected using (4) is shown in Fig. 1(d), with constants λ and δ so chosen that the effect of isotropic component of smoothness is small. The technique suffers seriously from the so-called *aperture* problem, and fails to recover correctly the optical flow component normal to the intensity gradient. Let us describe the behavior of the oriented smoothness term near the horizontal edge contours. The topological gradient is oriented along the y-axis at any of these edge points. However, the intensity gradients at these edge points are varying sinusoidally. The oriented smoothness-based regularization imposes smoothness perpendicular to the gradient direction at every edge point, and incorrectly computes the y-component of the velocity field and propagates the components inside the entire pattern. Since, the direction of topological gradient matches that of intensity gradient at vertical edge points, x-component of the optical field is reasonably recovered. Moreover, computation of the direction of gradient is prone to errors. This sensitivity can be suppressed by high amount

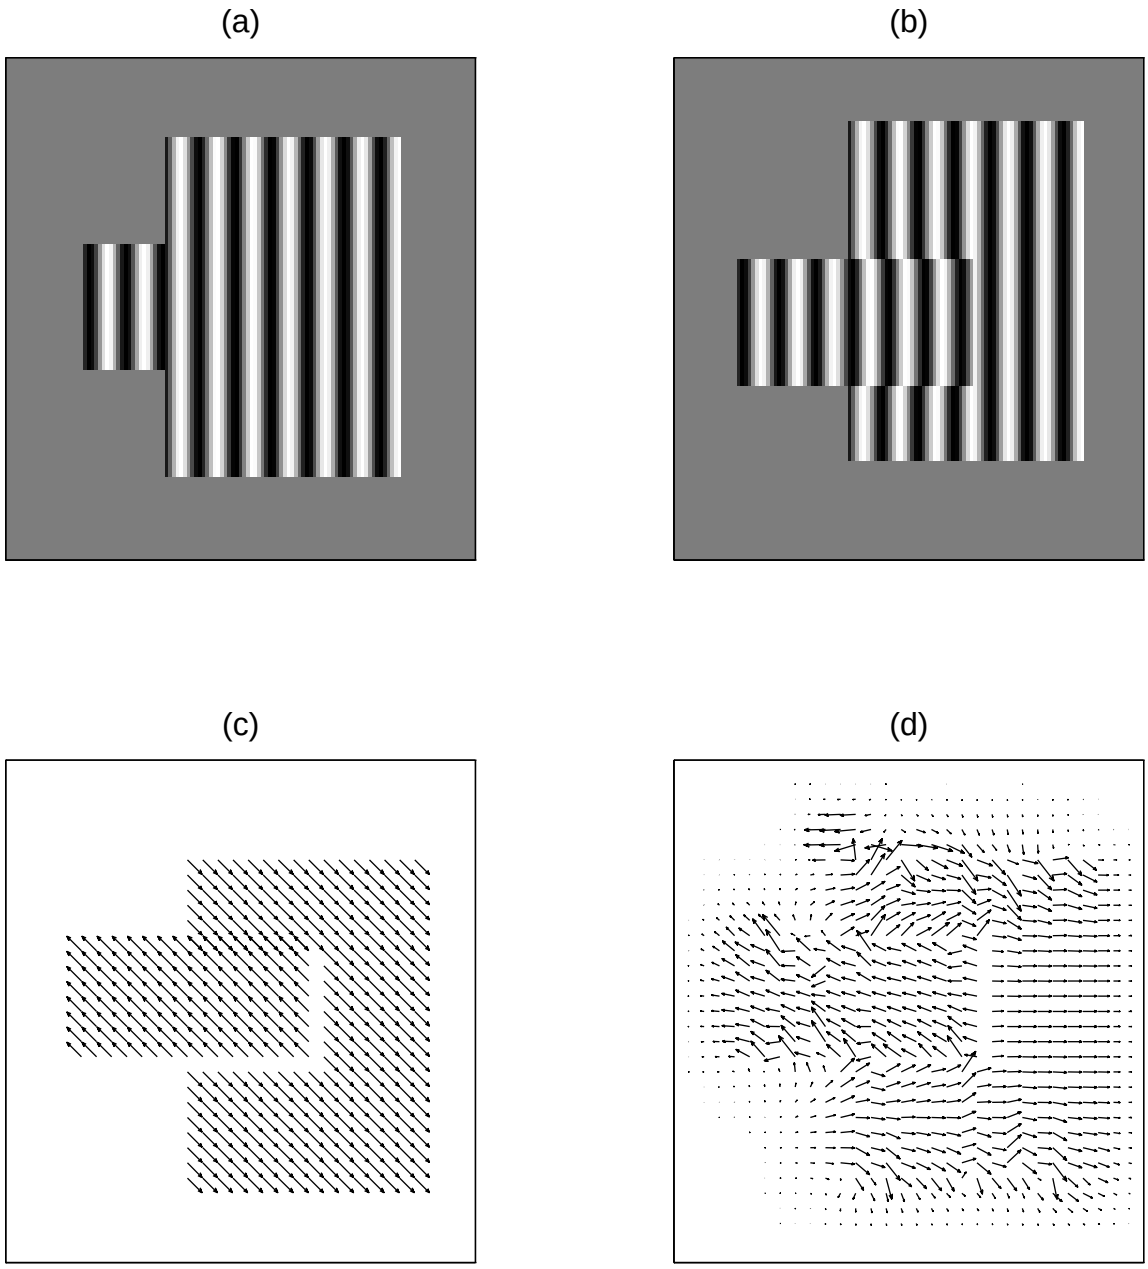

Figure 1: Behavior of oriented smoothness-based regularization. (a) First frame of occluded patterns, one translating to the upper left and the other to the lower right corner. (b) Last frame of the sequence. (c) Actual optical flow field. (d) Recovered optical flow by oriented smoothness constraints.

of presmoothing, as mentioned in [2]. Large presmoothing may in turn reduce the resolution of the entire problem, and the accuracy in the recovered optical flow field.

Let us now formulate a weighted anisotropic smoothness constraint, which we believe is more appropriate for dense but discontinuous optical flow estimation. Uniform smoothing-based regularization mechanism is able to propagate the optical flow field to regions of uniform brightness from all directions; but in this process it tends to blur actual motion discontinuities present in the image sequence. Also, the accuracy of optical flow components suffers due to uniform smoothing at points where the intensity variation is primarily in one spatial direction only. Local techniques, on the other hand, determine optical flow components locally by assuming some kind of parametric model of the underlying flow field. These cannot propagate the flow field to regions with little spatial information. In image regions where intensity is varying only in one direction, local techniques, in general, compute only the so-called *normal* component of the flow. However, these do provide with the so-called 'confidence' measure at image points to denote the availability of spatial information, and the accuracy of computation. Consider the optical flow functional without the regularization term,

$$\mathcal{F}_{data} = \int_{\Omega_i} (E_x u + E_y v + E_t)^2 d\Omega_i$$

in a local image neighborhood Ω_i. Minimization of this functional w.r.t. (u, v) yields,

$$\begin{bmatrix} \sum_{\Omega_i} E_x^2 & \sum E_x E_y \\ \sum E_x E_y & \sum E_y^2 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = - \begin{bmatrix} \sum E_x E_t \\ \sum E_y E_t \end{bmatrix} \tag{6}$$

If there is only one image point $(x, y) \in \Omega_i$, the eigenvalues can be easily computed to be

$$\lambda_1 = 0 \qquad \text{and} \qquad \lambda_2 = E_x^2 + E_y^2.$$

If $|E_x| \gg 0$ and $E_y = 0$, u can be calculated as $-E_t/E_x$, and v cannot be recovered uniquely. If $|E_y| \gg 0$ and $E_x = 0$, v can be obtained uniquely as $-E_t/E_y$, but u cannot be recovered uniquely. If both E_x and E_y are zeros, infinitely many solutions exist for (u, v). If both are nonzeros, only normal component of the optical field can be obtained. Thus, the optical flow field (u, v) can be partly recovered from the data constraint term alone. In order to obtain optical flow field with a higher accuracy, it is necessary to apply different amounts of smoothness-based regularization at different points along different axes. Firstly, $\lambda_1 + \lambda_2$ provides a direct measure of *goodness* of the data constraint term. If the sum is small at some image point, high amount of regularization should be performed, and *vice versa*. Similar measures, computed in a local image neighborhood have been used in [21], [22] in the context of local optimization-based and region matching-based optical flow estimation. Relative strength of smoothness along x- and y-axis should be such that less regularization is performed if the intensity is varying strongly along that axis, i.e., the confidence of data along the axis is high. For example, if $|E_x| \gg 0$ and $E_y = 0$ at any point, regularization should be performed at that point along y-axis to ensure that the optical flow component v is strongly propagated from neighboring points, while the flow component u is essentially determined by the data constraint term (oriented smoothness-based regularization in [16] allows for equal smoothing along both axes). Such a regularization term can be written as,

$$\sum_{psm} = \frac{\lambda}{E_x^2 + E_y^2} \left[\, |E_y| u_x^2 + |E_x| u_y^2 + |E_y| v_x^2 + |E_x| v_y^2 \, \right] \tag{7}$$

Note that, this smoothness term is anisotropic in nature coupled with jumps. Obviously, if $|E_y| \gg |E_x|$, then the variation of the solution is penalized less in y- than in x-direction, and *vice versa*. Thus, the data constraint term plays a stronger role along the y-axis in obtaining the solution. In addition, deviation from smoothness is penalized less at locations of strong image gradient to allow for optical flow discontinuities at these locations. High intensity gradients correspond to depth gradients or 3-D edges that potentially are locations of 3-D motion discontinuities. If both E_x and E_y are zero, high amount of isotropic smoothness is imposed to propagate solutions from nearby points. If $|E_x| = |E_y|$ at every image point, the proposed smoothness term becomes isotropic, and equivalent to that used in Horn and Schunck's technique.

This smoothness constraint allows for motion discontinuities to be formed at points with strong gradient information, e.g., edges, by penalizing less spatial variations of solutions. Blurring of motion discontinuities can be further reduced by enforcing a *zero-order* smoothness term (L_2 penalization). The weight associated with this zero-order term should be only a small fraction of that associated with the first-order smoothness term. Such a weighting ensures that the effect of this term is negligible at regions where some spatiotemporal gradient information is available. This term dominates at regions where little information is available and helps to attenuate the propagation of the optical flow field. This is particularly desirable for suppressing the solution in the background points that are nearer to image boundaries. Thus, the proposed functional for estimating the discontinuity preserving optical flow field can be written as,

$$\mathcal{F}_{pr} = \int (E_x u + E_y v + E_t)^2 + \frac{\gamma}{\sqrt{E_x^2 + E_y^2}}(u^2 + v^2) + \frac{\lambda}{E_x^2 + E_y^2}[\ |E_y|(u_x^2 + v_x^2) + |E_x|(u_y^2 + v_y^2)\ ] d\Omega \quad (8)$$

where the weight γ, associated with the zero-order smoothness term is approximately 0.1% of λ. If γ is chosen large, the entire flow field would be reduced to zero. 0.1% weighting of γ seems to be a good one, based on our experience of a diverse set of real and synthetic images. A similar penalization is mentioned in [3] in the context of *shape from shading* problems.

B. Minimization of the optical flow functional

The optical flow field on the image plane can be obtained by minimizing the proposed functional $\mathcal{F}_{pr}$ w.r.t. u and v. $\mathcal{F}_{pr}(u, v)$ is strictly convex on $[H_0^1(\Omega)]^2$, corresponding to the system of second-order partial differential equations given by,

$$(E_x^2 + \frac{\gamma}{\sqrt{E_x^2 + E_y^2}})u + E_x E_y v + \frac{\lambda}{E_x^2 + E_y^2}[\ |E_y| u_{xx} + |E_x| u_{yy}\ ] = -E_x E_t \quad (9)$$

$$(E_y^2 + \frac{\gamma}{\sqrt{E_x^2 + E_y^2}})v + E_x E_y u + \frac{\lambda}{E_x^2 + E_y^2}[\ |E_y| v_{xx} + |E_x| v_{yy}\ ] = -E_y E_t \quad (10)$$

with boundary conditions,

$$u\ |_{\delta\Omega} = \ v\ |_{\delta\Omega} = \ 0.$$

Multigrid algorithm is used to iteratively solve (12) over the entire image domain Ω. A five point stencil is used for finite-difference discretization of the problem, e.g., (9) at point (i, j) is discretized as,

$$(E_x^2(i,j) + \frac{\gamma}{\sqrt{E_x^2 + E_y^2}})u(i,j) + \frac{\lambda}{E_x^2 + E_y^2}[\ |E_y|(u(i,j-1) + u(i,j+1) - 2u(i,j)) + |E_x|(u(i-1,j) + u(i+1,j) - 2u(i,j))\ ] + E_x(i,j)E_y(i,j)v(i,j) \quad (11)$$

where, unknown $(u(i,j), v(i,j))$ is the optical flow field at point (i,j). Thus, finite-difference discretization leads to two linear equations in unknowns $(u(i,j), v(i,j))$ at every point (i,j). Arranging (u,v) at all image points in the lexicographic form, the system of linear equations can be expressed as,

$$Ax = b \tag{12}$$

where, A is a $2n \times 2n$ symmetric positive-definite matrix for image size of n pixels, and x is a vector consisting of optical flow components at all image points. Due to finite-difference discretization, A is highly sparse in nature, having at most five nonzeros per row, corresponding to five-point stencil. A is the so-called *stiffness* matrix.

Any optical flow functional, $\mathcal{F}$ should be strictly convex with a unique minimum such that a unique solution for the optical flow problem exists and multilevel preconditioner-based fast iterative methods can be applied for solving the resulting second-order elliptic partial differential equations. In other word, the finite-difference discretization of $\mathcal{F}$ should yield a stiffness matrix that is positive definite. Now, consider the stiffness matrix A, arising from discretization of Euler-Lagrangian equations corresponding to $\mathcal{F}_{pr}$. It can be expressed as a sum of two matrices A^{data} and A^{sm}, arising from data constraint and smoothness constraint, respectively. A^{data} is a 2×2 block diagonal matrix with entries,

$$A_{ii}^{data} = \begin{pmatrix} E_x^2 + \frac{\gamma}{\sqrt{E_x^2+E_y^2}} & E_x E_y \\ E_x E_y & E_y^2 + \frac{\gamma}{\sqrt{E_x^2+E_y^2}} \end{pmatrix}$$

Clearly, A^{data} is positive semi-definite (all eigenvalues ≥ 0). A^{sm} is the discretization of the second-order differential operator with discontinuous and anisotropic coefficients, which is positive definite. So, A is a sum of semi-definite and positive definite matrices; thus, A is positive definite and a unique discrete solution exists for optical flow field.

III. Black-Box Multigrid for Second-Order Elliptic PDEs

A robust and scalable multigrid algorithm is described in this section, that can solve a large class of systems of linear equations with positive definite matrices, arising from discretization of second-order elliptic PDEs. This technique offers very fast convergence for the weighted and anisotropic smoothness-based regularization term, proposed in section II. The proposed regularization term results in second-order PDEs with anisotropies coupled with jumps in coefficients, that in turn give rise to *singularly perturbated equations* (where ratio of the constant of boundedness of the bilinear form in H^1, and the constant of V-ellipticity in H^1 can be very large). Standard multigrid used in computer vision problems employing uniform averaging-based interpolation fails to converge fast in this case. In standard multigrid techniques, the rate of convergence strongly depends on coefficients of the second-order PDEs, and in turn on coefficients of stiffness matrix A in the corresponding discrete problem. A new automatic coarsening strategy is used that builds interpolation operators based on the information about coefficients of the stiffness matrix A in such a way that this undesirable dependence is eliminated. The multigrid algorithm is automatically adapted according to properties of the problem, and is truly a black-box solver.

A. Multigrid Algorithm

Multigrid methods are fast iterative methods for the solution of large system of linear equations of the form $Ax = b$. These are of particular importance for solving sparse systems arising

from discretization of PDEs. The unknowns and equations typically correspond to points in a physical space (in our case, image) and only unknowns, associated with neighboring points (i.e., they occur together in at least one equation) are coupled. Multigrid methods overcome the problem of local adjustment, found in pointwise iterative methods by utilizing several related systems (levels) with a smaller number of variables which interact with the original system to propagate the information. Information is exchanged between different levels by so-called *transfer* operators. The algebraic multigrid technique presented in this section is based on the so-called *variational* approach, where the hierarchy of coarse-level problems is generated successively using Ritz-Galerkin projection.

An algebraic multigrid algorithm consists of two main stages: (a) *coarsening :* that creates a hierarchy of coarse-level matrices and corresponding family of prolongators $\{P_l\}$; and (b) *iterative adjustment* of approximate solution by high-energy smoothing and transfering of information between adjacent levels. The typical implementation can be described as follows.

Coarsening:

- input: $A_1 = A$ - the stiffness matrix.
- output:
  - $\{A_l\}_{l=2}^{L}$, the hierarchy of coarse-level matrices of size $n_l \times n_l$, where L is the number of levels.
  - $\{P_l\}_{l=1}^{L-1}$, the family of prolongator matrices. P_l is a $n_{l+1} \times n_l$ matrix.

1. $l \leftarrow 1$
2. construct P_l.
3. $A_{l+1} = P_l^T A P_l$
4. $l \leftarrow l + 1$
5. if n_l is large for direct solution of the corresponding system of equations, go to 1.
6. $L \leftarrow l$.

Iteration: Iteratively adjust the current solution x.

- input:
  - $\{A_l\}_{l=1}^{L}$, $\{P_l\}_{l=1}^{L-1}$
  - right hand side of the finest-level problem $b^1 = b$; number of smoothing steps ν_1, ν_2, and initial guess x^1.

One iteration proceeds as,
$l \leftarrow 1$

1. Pre-smoothing: do ν_1 times $x^l \leftarrow \mathcal{S}^l(x^l, b^l)$, where $\mathcal{S}^l(\cdot,\cdot)$ is some pointwise iterative method, e.g., Gauss-Seidel or damped Jacobi.
2. Coarse-grid correction:
   (a) $b^{l+1} \leftarrow P_l^T(b^l - A_l x^l)$
   (b) solve $A_{l+1}x^{l+1} = b^{l+1}$ directly or by a recursive application of this algorithm. The exact way is defined by various multigrid cycles.
   (c) $x^l \leftarrow x^l + P_l x^{l+1}$
3. Post-smoothing: do ν_2 times $x^l \leftarrow \mathcal{S}^l(x^l, b^l)$.

For more details about basic algebraic multigrid components, see [14]. Generation of the family of $\{P_l\}$ is key to the performance of the algebraic multigrid algorithm, since, coarse-level matrices are obtained using $\{P_l\}$. Standard multigrid algorithms used in [17], [18], [19] build the

prolongators based on uniform averaging without obeying the coefficients of the stiffness matrix, which is appropriate only for special cases, e.g., Laplacian regularization. In more general cases, as for our proposed technique (also for oriented smoothness [15],[16], and Medioni's anisotropic smoothing [17]), it is necessary to construct coarse spaces based on the coefficients of the stiffness matrix to maintain a high rate of convergence. Simplified version of this technique is analyzed in [13].

A.1 Requirements for building prolongators

According to the multigrid abstract convergence theory [4], in order to construct efficient multigrid solver for solving second-order elliptic PDEs with general coefficients, the following conditions should be satisfied while building the prolongators such that the algorithm is able to achieve high rate of convergence.

Decomposition of unity: Let P_l denote the prolongator matrix between two adjacent levels l and $l+1$. The decomposition of unity requires,

$$\sum_{j=1}^{m} P_{ij} = 1, \qquad i = 1, .., n_l.$$

Let, $x_c \in \Re^{n_{l+1}}$ be the coarse-level vector with all entries equal to one. The decomposition of unity enforces that $P_l x_c$ is the vector with all entries equal to one again. In other word, the prolongator should be built in such a way that fine-level constant vectors can be approximated using coarse-level constant vectors exactly. This property can be violated for the fine-level nodes that are close to the boundary where Dirichlet boundary condition is imposed. For more details, see [13].

Small Fill–in means small number of nonzero entries in coarse-level matrices. This is important from both practical as well as theoretical points of view. Large number of nonzeros reduces the performance of the iterative process (speed and memory), since in sparse matrix algebra, number of operations is of the order of number of nonzeros in the matrix. From the theoretical point of view, the bounded number of nonzeros per column in the coarse-level matrices is closely related to the bounded number of intersections of supports of basis functions. This so-called 'bounded intersection property' is important for the localization of approximation estimates [13]. This condition is one of the basic conditions of the finite element theory [6].

Small Energy of the coarse-level vectors means that the "energy" of prolongated vectors, $< A_l P_l x,\ P_l x >$ must be small. If the hierarchy of coarse-level matrices is scaled properly, then this condition can be expressed as the reduction of spectral bounds of $\{A_l\}$, or, more generally as the inverse inequality on coarse spaces [13]. Reduction in spectral bounds must be in agreement with the reduction of coarse-level problems. During successive coarsening, we lose the ability to approximate precisely the finest-level vectors. This must be compensated by smoothness (low energy) of coarse-level vectors. For details, see Theorem 1 in [4].

Shape of supports of basis functions must be such that the support of any basis function (one column of P_l) contains strongly coupled nodes only. For example, strongly anisotropic equations can be solved in such a way that coarsening will be done in the direction of anisotropy only (so-called *semi-coarsening*). Algebraically, the anisotropy is reflected in the strong coupling of coefficients of stiffness matrix, describing the communication in the direction of strong anisotropy.

A.2 Automatic coarsening and generation of prolongators

In this subsection, we present the algorithm for automatic coarsening strategy that utilizes

the information about coefficients of stiffness matrix A in generation of prologators, $\{P_l\}$. The proposed coarsening algorithm is based on combination of two techniques: (a) *unknown aggregation* to generate auxiliary prolongators, followed by (b) *smoothing* of auxiliary prolongators. Unknown aggregation generates prolongators that satisfy the requirements of decomposition of unity, small size of coarse-level matrices, and the shape of supports. However, these auxiliary prolongators do not satisfy the requirement of small energy of coarse-level vectors. Smoothing is performed to ensure small energy of coarse-level vectors. The unknown aggregation technique works as follows. First, the disjoint covering of the set of nodes $\{C_i\}_{i=1}^m$ is generated, where each set C_i contains strongly coupled nodes only. The information about the coupling of nodes can be obtained from node entries of the stiffness matrix. The strongly coupled neighborhood, N_i^l of node i is defined as,

$$N_i^l = \{\ j :\ \ |a_{ij}| \geq \epsilon \sqrt{a_{ii}a_{jj}}\ \}.$$

If the stiffness matrix is diagonally dominant, it is more appropriate to use,

$$N_i^l = \{\ j :\ \ |a_{ij}| \geq \epsilon \sqrt{q_i q_j}\ \}.$$

where

$$q_i = \max_{k \neq i} |a_{ik}|, \quad q_j = \max_{k \neq j} |a_{jk}|,$$

and $\epsilon \in (0,1)$. These sets are small and contain only a few nodes. To each set C_i, we assign one of its nodes (the so-called C-point) that plays the role of coarse-level node. The prolongator is built by simply copying the degree of freedom associated with the C-point to all nodes in the corresponding $\{C_i\}$. Thus, prolongators can be expressed as,

$$P_{ij} = \begin{cases} 1 & \text{if } i \in C_j \\ 0 & \text{otherwise} \end{cases} \tag{13}$$

P is very sparse and contains one nonzero entry per row. Detailed computational steps are summarized below:

Auxiliary prolongators: Generate the auxiliary prolongation operators, $\{\hat{P}_l\}$.

- input
  - A_l, the coarse-level matrix of order n_l on level l
- output
  - $\hat{P}_l$, the auxiliary prolongator matrix of size $n_{l+1} \times n_l$

$R \leftarrow \{1, .., n_l\}, \quad j \leftarrow 0.$
for $i := 1$ to n_l do
  if $N_i^l \subset R$ then
    $j \leftarrow j + 1,$
    $C_j \leftarrow N_i^l,$
    $R \leftarrow R \setminus C_j$
  end if
end for

for $i := 1$ to n_l do
  if $i \in R$ then
    find $C_k :\ N_i^l \cap C_k \neq \emptyset,$

$C_k \leftarrow C_k \cup \{i\}$
end if
end for
set $n_{l+1} \leftarrow j$

Generate $\hat{P}_l : \Re^{n_{l+1}} \rightarrow \Re^{n_l}$ by

$$(\hat{P}_l)_{ij} = \begin{cases} 1 & \text{if } i \in C_j \\ 0 & \text{otherwise} \end{cases}$$

Improvement of the auxiliary prolongator by smoother

- input
  - A_l, $\hat{P}_l$
- output
  - P_l, the improved prolongator matrix
- Generate the filtered matrix $A_l^f = \{b_{ij}\}_{i,j=1}^{n_l}$ from the matrix $A_l = \{a_{ij}\}_{i,j=1}^{n_l}$:
  - offdiagonal entries, i.e., $i \neq j$:

$$b_{ij} = \begin{cases} a_{ij} & j \in N_i^l \\ 0 & \text{otherwise} \end{cases}$$

  - diagonal entries, i.e., $i = j$:

$$b_{ii} = a_{ii} - \sum_{j=1, j\neq i}^{n_l} (a_{ij} - b_{ij})$$

- Smooth the auxiliary prolongators:
  - Define the smoother

$$S_l = I - \omega D^{-1} A_l^f$$

where $D = \text{diag}(A_l^f)$, and compute the new prolongator P_l by

$$P_l = S_l \hat{P}_l$$

For a large class of problems, we have experimented with, reasonable choices of parameters are

$$\epsilon = 0.2, \quad \omega = \frac{4}{3}.$$

Note that the smoothing step does not violate the decomposition of unity at nodes where the constant is a local kernel of A_l.

The coarsening process described above is appropriate for the case of one degree of freedom (number of unknowns) per node. However, in many applications, e.g., optical flow estimation problem, we need to solve problems with more than one degree of freedom on each node (pixel). If the number of degrees of freedom on each node in the entire domain is constant, the scalar algorithm can be easily generalized using so-called *block* approach. The idea is to compute with matrices of order equal to the number of degrees of freedom per node instead of computing with a single number at each node, which describes the communication between neighboring nodal points. Let $df(k)$ denote the list of degrees of freedom associated with node k. The matrix

describing the communication between nodes k and l is the selection $A_{kl} = A(df(k), df(l))$. In order to generalize the algorithm, two minor changes are required. Let n_d be the number of degrees of freedom per node. The definition of the set of strongly coupled nodes can be generalized as

$$N_i^l = \{\ j :\ \ \rho(A_{ij}^l) \geq \epsilon \sqrt{\rho(A_{ii}^l)\ \rho(A_{jj}^l)}\ \};$$

and for diagonally dominant matrices as,

$$N_i^l = \{\ j :\ \ \rho(A_{ij}^l) \geq \epsilon \sqrt{q_i q_j}\ \}$$

where

$$q_i = \max_{k \neq i} \rho(A_{ik}).$$

$\rho(.)$ denotes the spectral bound of square matrix. Also, the 0 and 1's are replaced by zero and identity matrices of order n_d, respectively.

IV. Experimental Results

Numerous experiments have been conducted with synthetic and real image sequences to demonstrate the efficacy of the proposed weighted anisotropic smoothness-based regularization and the multigrid solution technique. Some of these are reported here. All experiments are conducted on an RS6000 workstation. Detected flows are pictorially presented for the proposed, Horn and Schunck's isotropic regularization-based, and Nagel and Enkelmann's techniques. For synthetic image sequences, where optical flow fields are known *a priori*, errors are calculated for moving regions and background regions separately and then averaged. These errors are calculated in degree, as suggested in [2]. Partial derivatives are computed using 4-point central difference with $\frac{1}{12}[-1, 8, 0, -8, 1]$. Real sequences are presmoothed with a spatiotemporal Gaussian to reduce temporal aliasing. No smoothing is done for the proposed or Horn and Schunck's techniques in the case of synthetic sequences. Presmoothing of synthetic images is done for Nagel and Enkelmann's technique only when the iterative algorithm did not converge. [1] Horn and Schunck's and the proposed techniques were terminated when the relative residual (measured in L_2-norm) drops below 10^{-5}. All results are reported with suboptimal choices of parameters.

Figs. 2(a)–(e) show five pairs of synthetic image-frames used in our numerical experiments. Subsampled actual flow fields are shown in Figs. 3 – 7(a), respectively. Subsampled detected flow fields are shown in Figs. 3 – 7(b), (c) and (d) for Horn and Schunck's, Nagel and Enkelmann's, and the weighted anisotropic smoothness-based optical flow techniques. Numerical errors are shown in Tables 1–5, respectively.

The plaid patterns, shown in Fig. 2(a) and Fig. 2(d) are created by superimposing two sinusoids in two spatial directions, whereas Fig. 2(b), Fig. 2(c), and Fig. 2(e) contain plaid pattern(s) with intensity variation in one direction only. Actual optical flow field is rotational in Fig. 2(a). The vertical component of the flow field is perpendicular to the intensity variation in Fig. 2(b). In Fig. 2(c) the correct flow field is either in x- or in y-direction, perpendicular to intensity variations in the corresponding plaid. Fig. 2(d) consists of two occluded plaid patterns with similar intensity structure, moving in perpendicular directions. In the occluded patterns shown in Fig. 2(e), no intensity gradient information is present in the y-direction inside plaid patterns, which are moving in diagonally opposite directions.

Intensity gradient information is available in both spatial directions in Fig. 2(a) at all moving points. The proposed, Horn and Schunck's, and Nagel and Enkelmann's techniques all perform

[1] This implementation has been obtained from anonymous ftp site ftp.csd.uwo.ca:pub/vision/NAGEL; courtesy: John Barron.

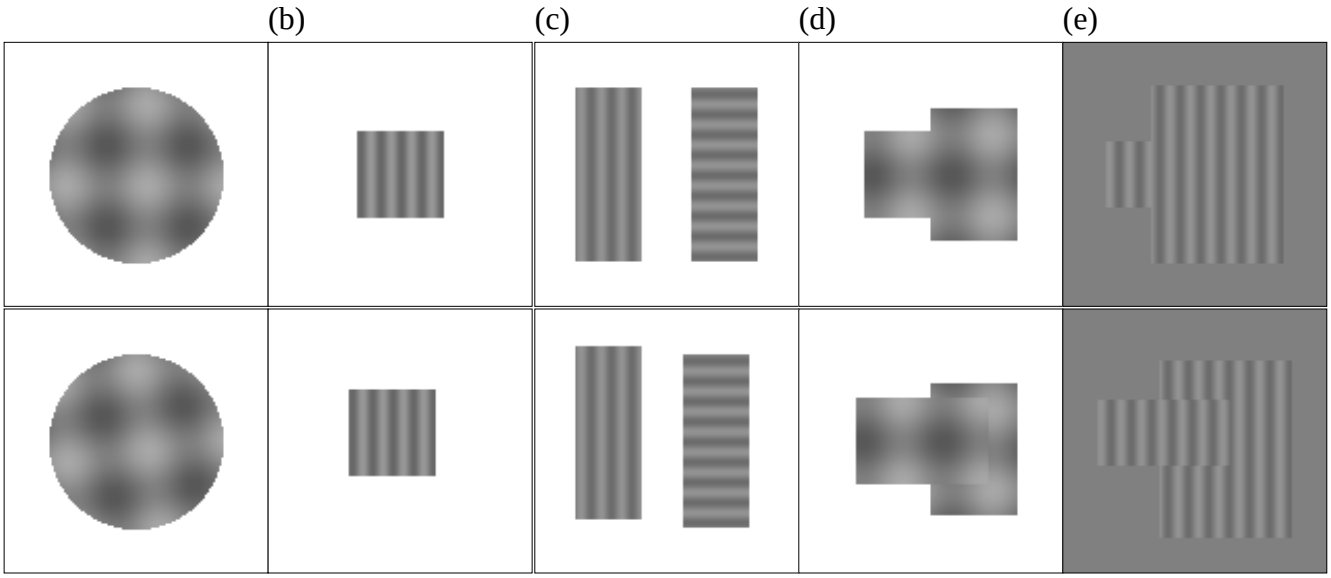

Figure 2: First and fifth-frame of synthetic image sequences. Optical flow is computed for the third frame. (a) Rotating plaid pattern. (b) Translating plaid pattern with intensity variation in x-direction only. (c) Two plaid patterns moving horizontally and vertically. (d) Partially occluded plaid patterns moving horizontally and vertically. (e) Two partially occluded blocks moving towards the upper left and lower right corners.

	Horn–Schunck	Nagel–Enkelmann	Proposed		
			A	AZ	DAZ
Object	9.02	12.16	8.20	8.31	6.35
Background	5.10	2.38	2.59	2.49	0.76
Average	7.060	7.270	5.394	5.400	3.555

Table 1: Errors in optical flow for the sequence shown in Fig. 2(a). `A` denotes anisotropic smoothness-based regularization. `AZ` denotes anisotropic regularization with zero-order term. `DAZ` denotes discontinuous anisotropic smoothness-based regularization along with zero-order term.

reasonably well. The proposed technique is able to suppress the optical flow field more in the background compared to Nagel and Enkelmann's and specially to Horn and Schunck's techniques. Fig. 2(b) and Fig. 2(c) are challenging sequences because the vertical component of actual optical flow fields is perpendicular to the direction of intensity variation. Since the oriented smoothness-based regularization strongly couples the spatial variation of flow field with that of intensity, Nagel and Enkelmann's optical flow estimates are inaccurate inside the object region compared to Horn and Schunck's and the anisotropic smoothness-based proposed technique. Oriented smoothness-based regularization is however able to suppress the flow field in the background region quite effectively. The proposed technique performs the best in both Fig. 2(b) and Fig. 2(c). Flow estimates are inaccurate for Fig. 2(d) and Fig. 2(e) due to strongly occluded regions. Intensity gradient information is available in both spatial directions inside moving regions in Fig. 2(d); thus the accuracy of flow estimates is more than that in Fig. 2(e), where gradient information is only available in x-direction. The proposed technique outperforms Horn and Schunck's and Nagel and Enkelmann's techniques in these sequences too.

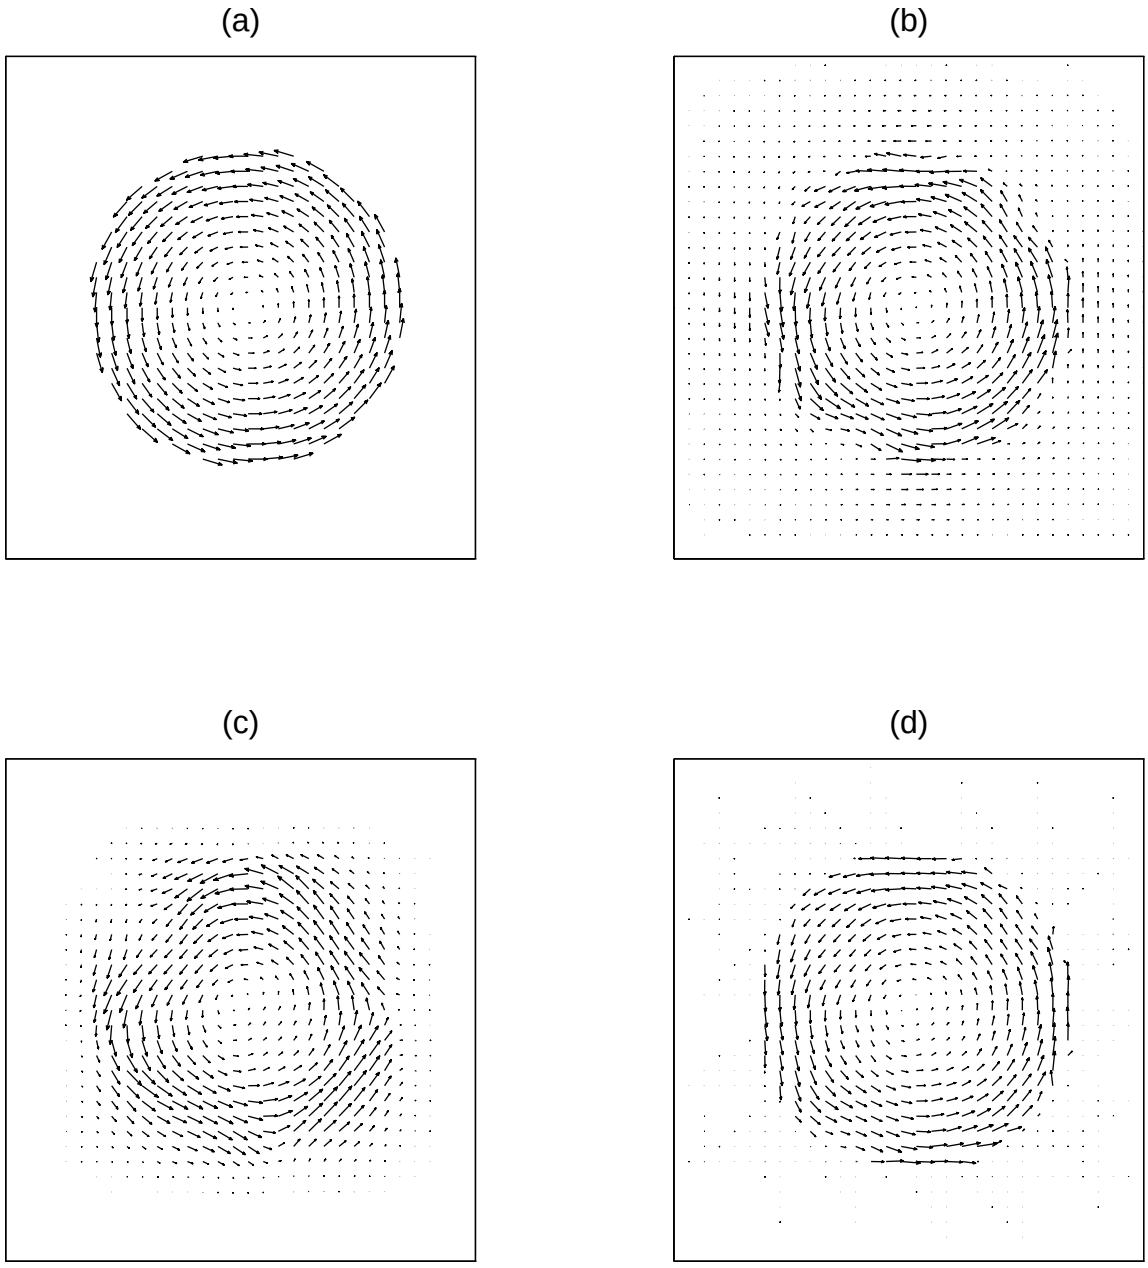

Figure 3: Optical flow computation for the sequence shown in Fig. 2(a). (a) Correct flow. (b) Flow computed using Horn–Schunck's technique. (c) Using Nagel–Enkelmann's technique. (d) Using weighted anisotropic smoothness with zero-order term.

	Horn–Schunck	Nagel–Enkelmann	Proposed		
			A	AZ	DAZ
Object	4.92	21.11	2.53	4.79	5.66
Background	20.72	3.99	14.44	10.30	8.07
Average	12.820	12.55	8.485	7.545	6.865

Table 2: Errors in optical flow for the sequence shown in Fig. 2(b).

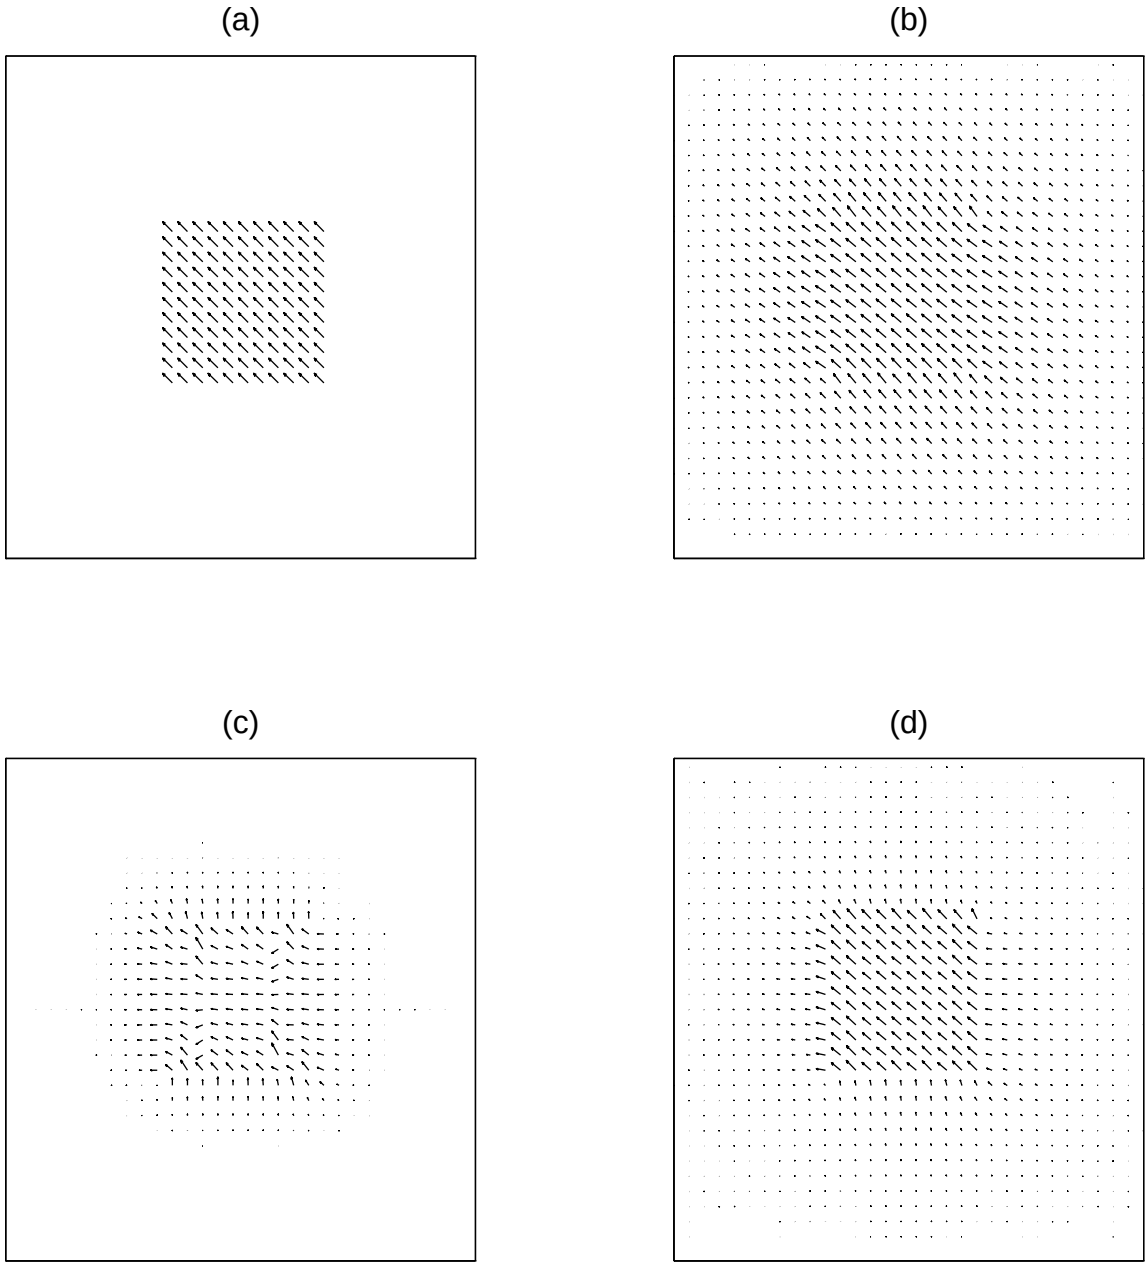

Figure 4: Optical flow computation for the sequence shown in Fig. 2(b). (a) Correct flow. (b) Flow computed using Horn–Schunck's technique. (c) Using Nagel–Enkelmann's technique. (d) Using weighted anisotropic smoothness with zero-order term.

	Horn–Schunck	Nagel–Enkelmann	Proposed
Object	9.56	17.17	2.69
Background	21.63	12.73	12.92
Average	15.595	14.950	7.805

Table 3: Errors in optical flow for the sequence shown in Fig. 2(c).

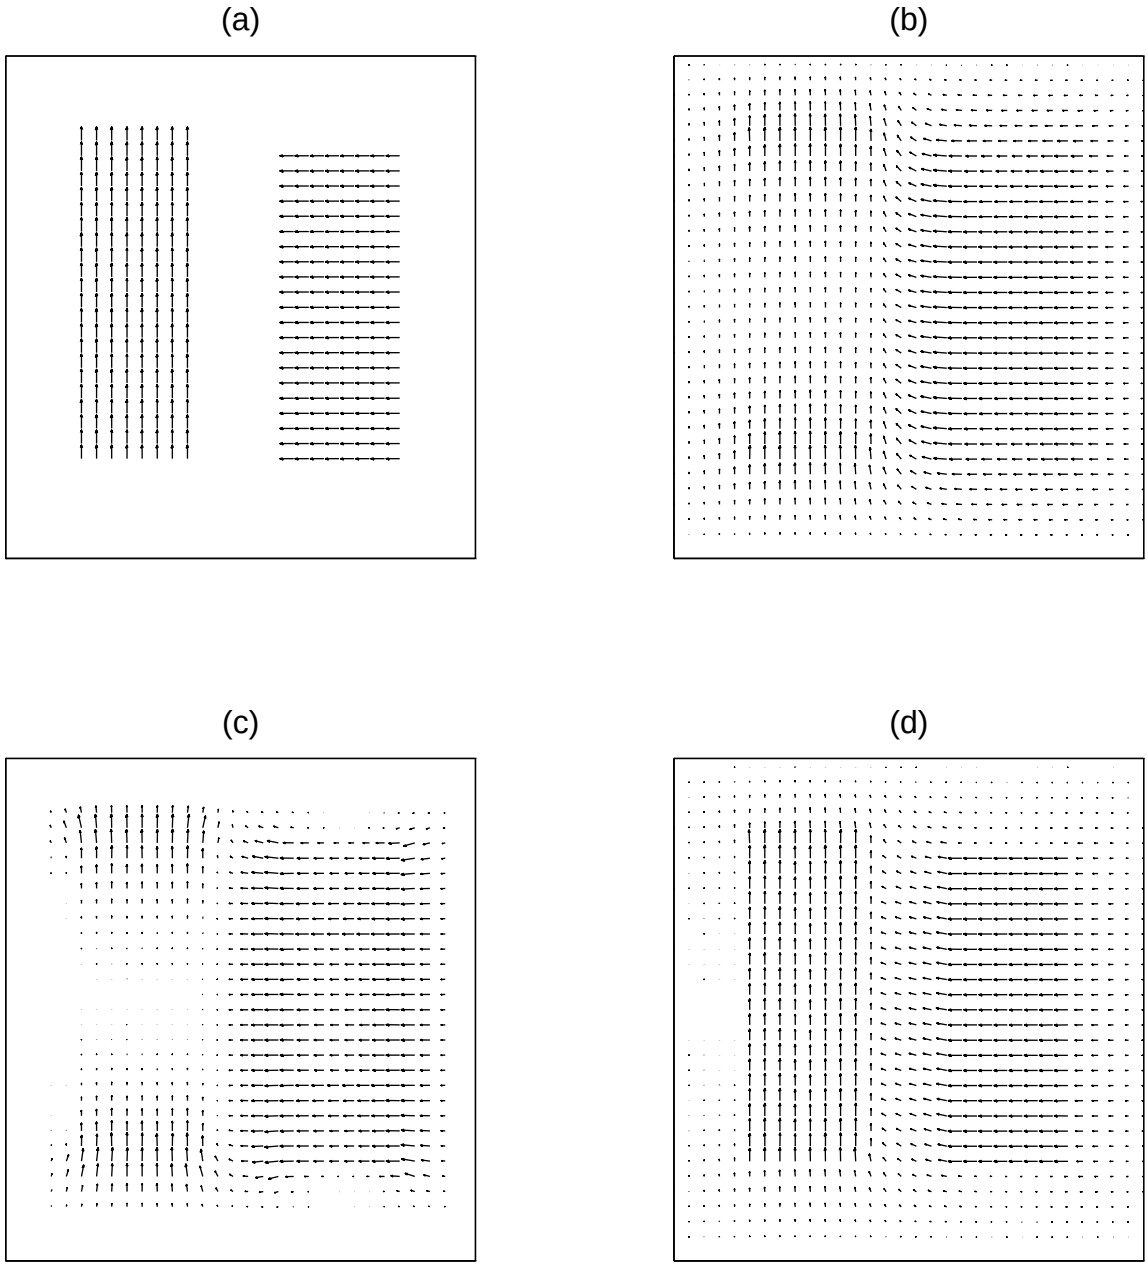

Figure 5: Optical flow computation for the sequence shown in Fig. 2(c). (a) Correct flow. (b) Flow computed using Horn–Schunck's technique. (c) Using Nagel–Enkelmann's technique. (d) Using weighted anisotropic smoothness with zero-order term.

	Horn–Schunck	Nagel–Enkelmann	Proposed
Object	6.61	11.76	4.76
Background	19.35	9.60	13.37
Average	12.980	10.680	9.065

Table 4: Errors in optical flow for the sequence shown in Fig. 2(d).

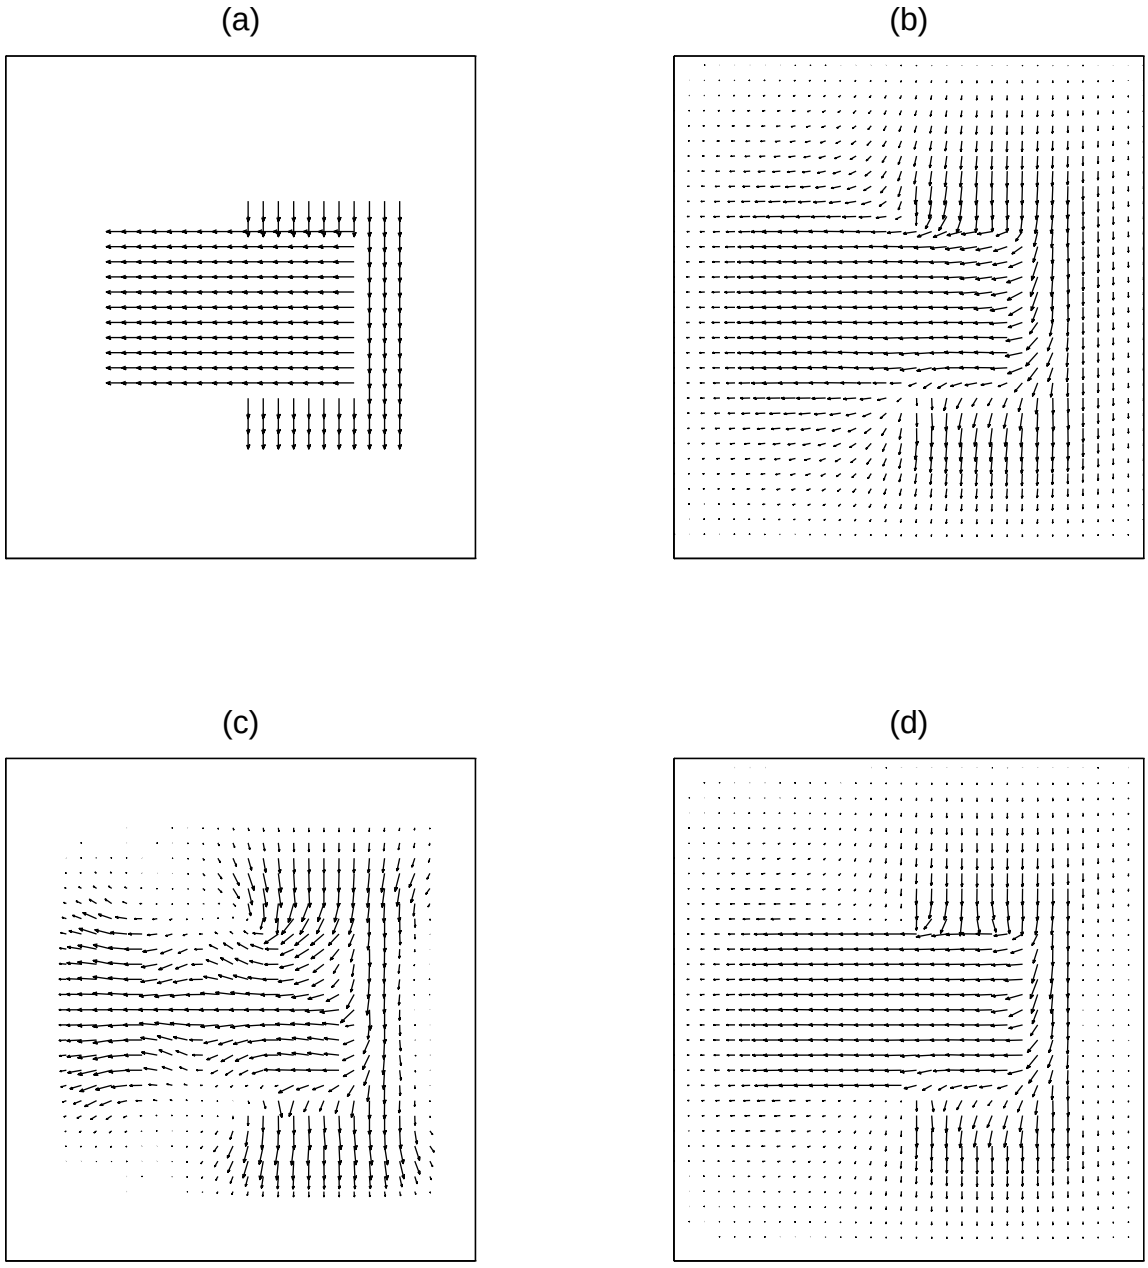

Figure 6: Optical flow computation for the sequence shown in Fig. 2(d). (a) Correct flow. (b) Flow computed using Horn–Schunck's technique. (c) Using Nagel–Enkelmann's technique. (d) Using weighted anisotropic smoothness with zero-order term.

	Horn–Schunck	Nagel–Enkelmann	Proposed
Object	25.54	28.93	17.99
Background	21.42	11.65	3.20
Average	23.48	20.29	10.595

Table 5: Errors in optical flow for the sequence shown in Fig. 2(e).

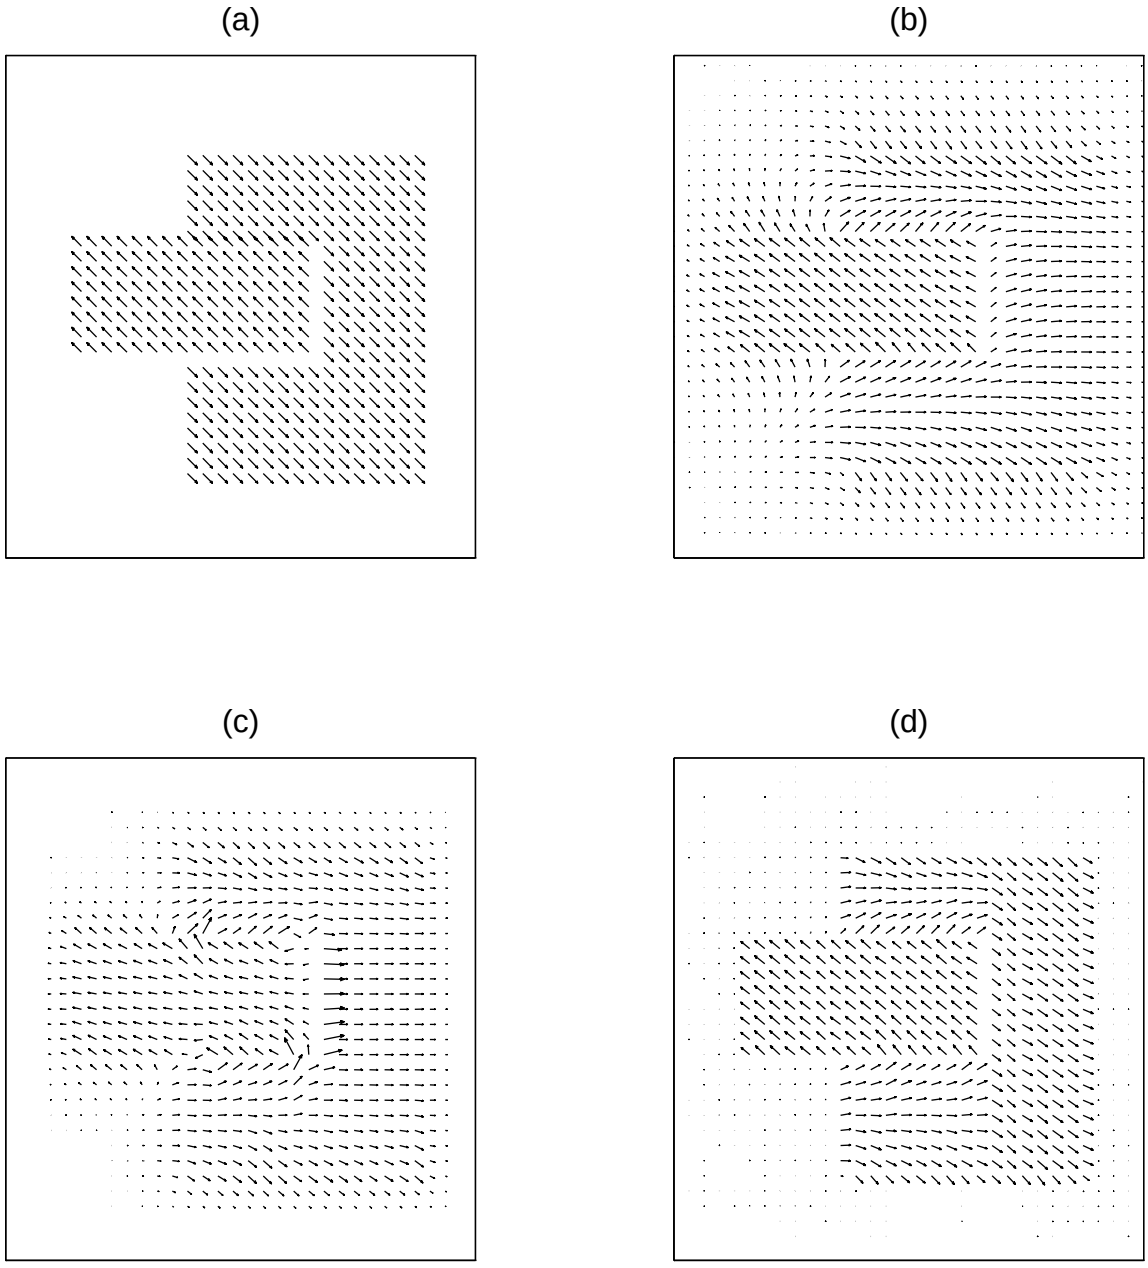

Figure 7: Optical flow computation for the sequence shown in Fig. 2(e). (a) Correct flow. (b) Flow computed using Horn–Schunck's technique. (c) Using Nagel–Enkelmann's technique. (d) Using weighted anisotropic smoothness with zero-order term.

Fig. 8(a)–(d) show first and last frames of four real sequences. [2] The optical flow is estimated in the middle frame. In the Hamburg Taxi sequence, a car in the lower left is moving to right, a van in the lower right backing up, and a taxi turning the corner. In the Nasa sequence, the camera is moving along its line of sight toward the Coke Can near the center of the image. Thus, the motion is primarily dilational. The turntable is rotating counter-clockwise in the Rotating Cube sequence, whereas the camera is moving perpendicular to its line of sight in the Pepsi sequence. The turntable and and the Pepsi Can are displaced more on the image plane than others due to perspective. Results obtained using the proposed weighted anisotropic smoothness and magnitude penalization are shown in Fig. 9(a)–(d), respectively. For results obtained using Horn and Schunck's and Nagel and Enkelmann's techniques, see [2].

All these results clearly demonstrate the effectiveness of the proposed smoothness-based regularization for discontinuous optical flow estimation.

A. Speed and Scalability

The speed of the entire optical flow algorithm is of immense importance in real-world applications. Numerical solution of the resulting Euler-Lagrangian equations is the most time consuming part, and time complexity of the solution grows as high as the square of problem size, e.g., in popular Gauss-Seidel iterative method. As a result, in spite of a large number of available processors in the optimistic case, solution times will be larger for larger images. Fast and scalable algorithms are required to ensure execution times to be only linearly dependent on image size in a serial computing environment. Scalable algorithms involving matrix-vector operations are perfectly parallelizable on vector processors, and thus image size-independent solution times can be obtained. Uniform averaging-based multigrid techniques are scalable for isotropic smoothness-based regularization, e.g., Horn and Schunck's, which gives rise to second-order elliptic PDEs with constant coefficients. Second-order elliptic PDEs with discontinuities and anisotropies in coefficients cannot be solved efficiently using uniform averaging-based multigrid techniques. The multigrid solution technique employed here generates coarse spaces utilizing the information about entries of the stiffness matrix, and thus works efficiently for elliptic PDEs with anisotropies coupled with jumps in coefficients. Fig. 10 shows first and last frames of a 256×512 synthetic image sequence. The image frames are subsampled (uniformly in both directions) to create smaller-size images. Optical flows are computed for images of different size and user times are plotted in Fig. 11. The linear nature of the plot confirms the scalability of the algebraic multigrid solution technique. [3] Execution times are plotted in Fig. 12 for some of the real sequences shown in Fig. 8. Corresponding times for a block Gauss-Seidel method are shown in Fig. 12. For maximum allowable tolerance of 10^{-5}, the multigrid algorithm (with $V(1,1)$ cycle) took at most 5 iterations to converge for all images reported here, whereas the block Gauss-Seidel (more efficient than usual) method took as high as 1800 iterations to converge. Cost of each multigrid iteration is approximately 5 times that of the Gauss-Seidel's. All these results clearly establish the very fast rate of convergence and scalability of the algebraic multigrid algorithm for solving second-order elliptic PDEs with discontinuities and anisotropies in coefficients.

V. Conclusions

A weighted anisotropic smoothness-based regularization technique is proposed in this paper for estimating discontinuous optical flows from two-frame sequences. Penalization of the mag-

[2] Obtained from anonymous ftp site ftp.csd.uwo.ca.

[3] Implementation can be obtained by contacting the authors via sghosal or pvanek@tiger.denver.colorado.edu.

Seq

(b) Nasa Sequence

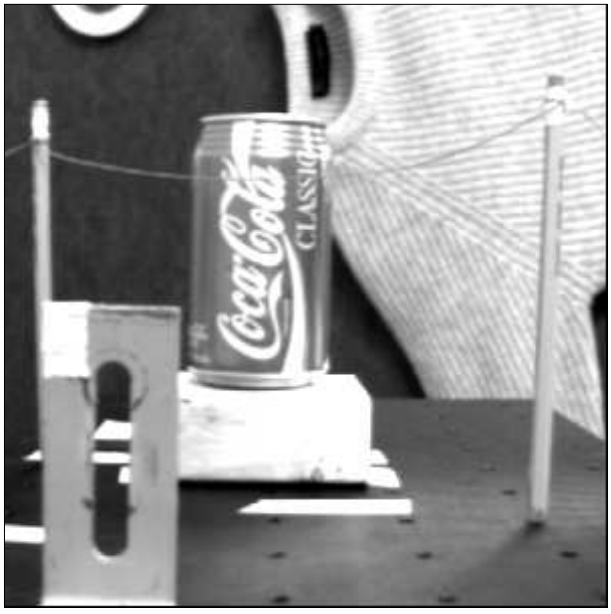

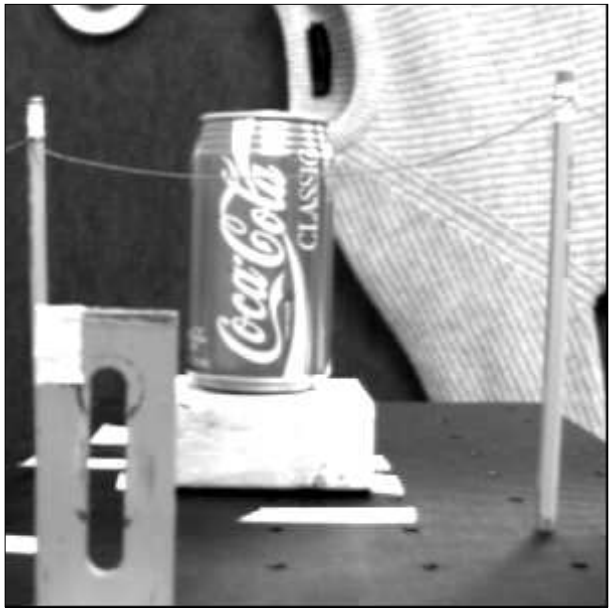

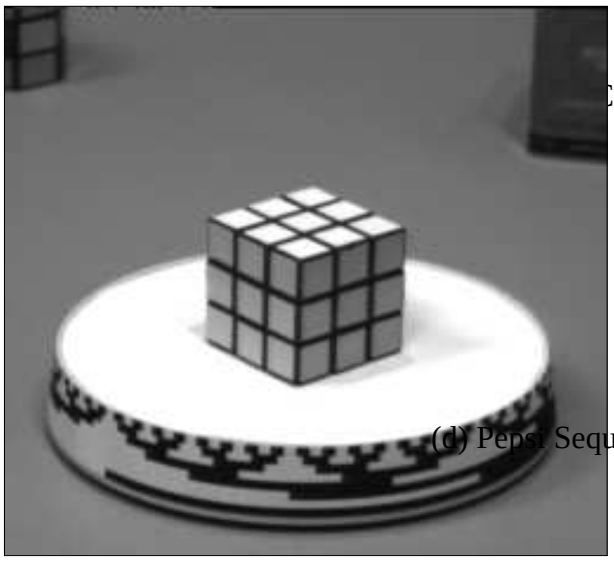

Cube Se

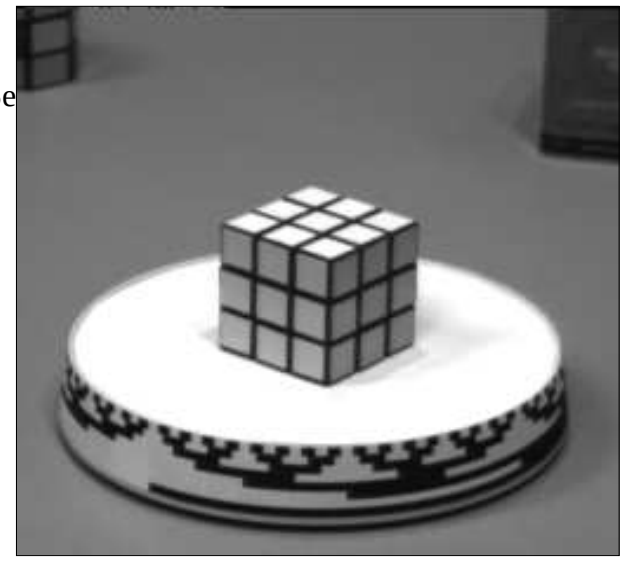

(d) Pepsi Sequence

Figure 8: Real image sequences. (a) Hamburg Taxi sequence. (b) Nasa sequence. (c) Rotating cube sequence. (d) Pepsi sequence.

(a)

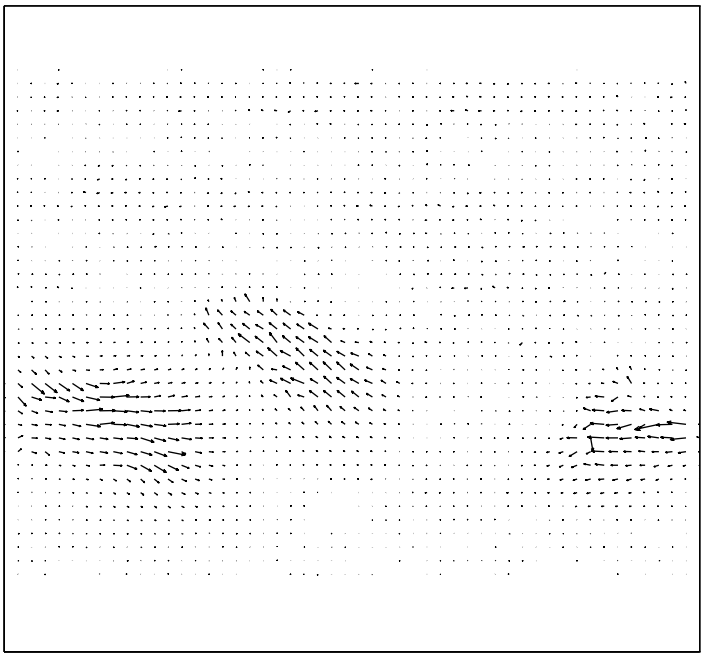

(b)

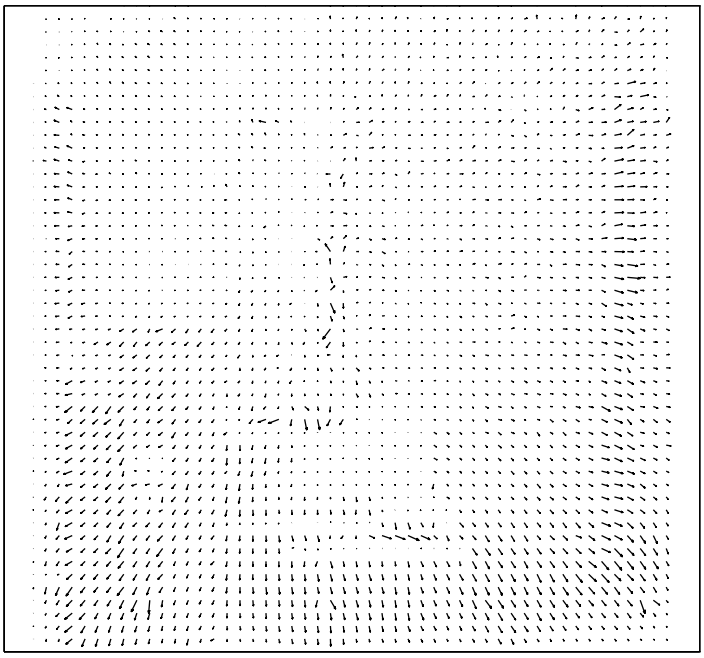

Optical flow computation for real sequences using the proposed technique.

(c)

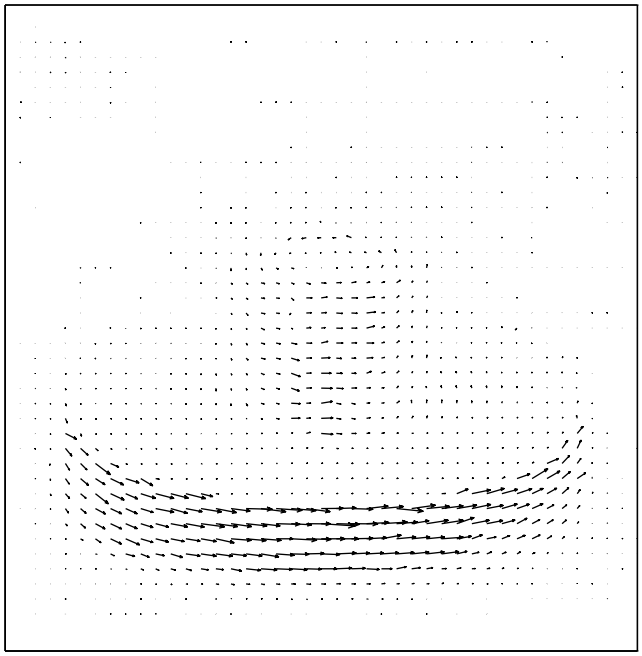

(d)

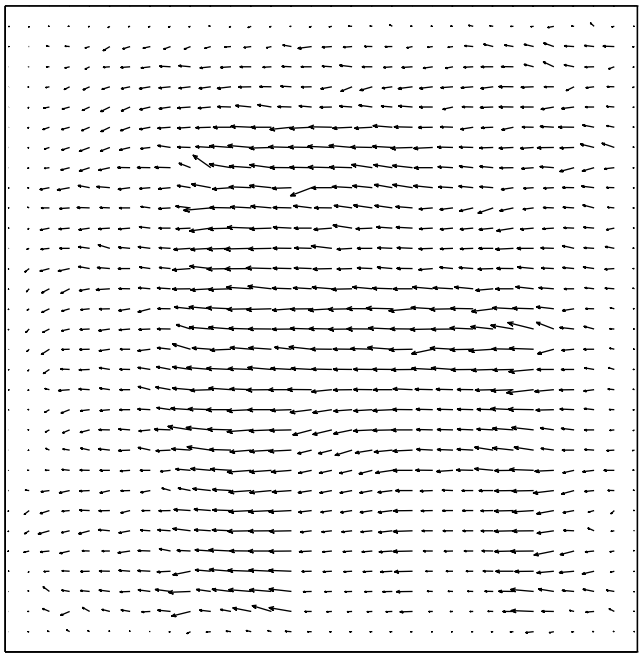

Figure 9: Optical flow computation for real sequences using the proposed technique. Flow field in (a) Hamburg Taxi sequence. (b) Nasa sequence. (c) Rotating cube sequence. (d) Pepsi sequence.

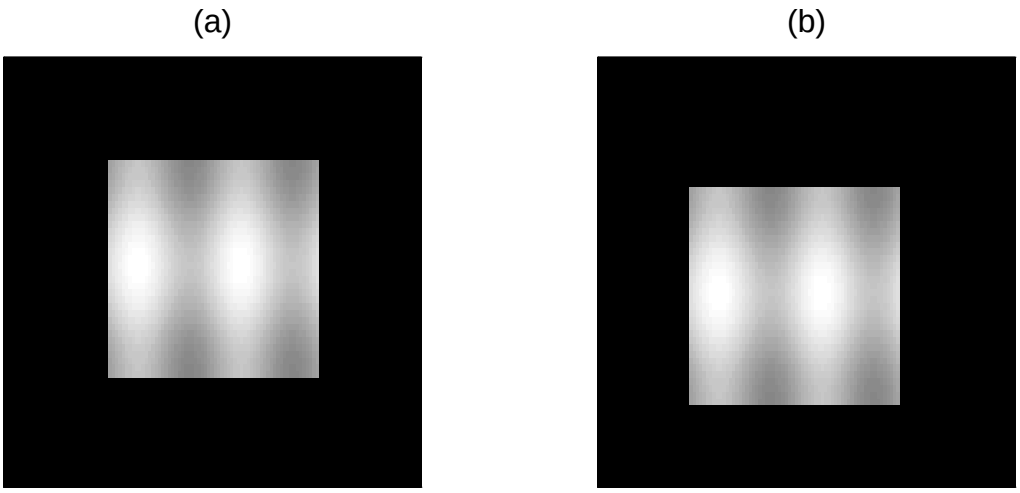

Figure 10: Synthetic image sequence used to test the scalability of the proposed multigrid technique.

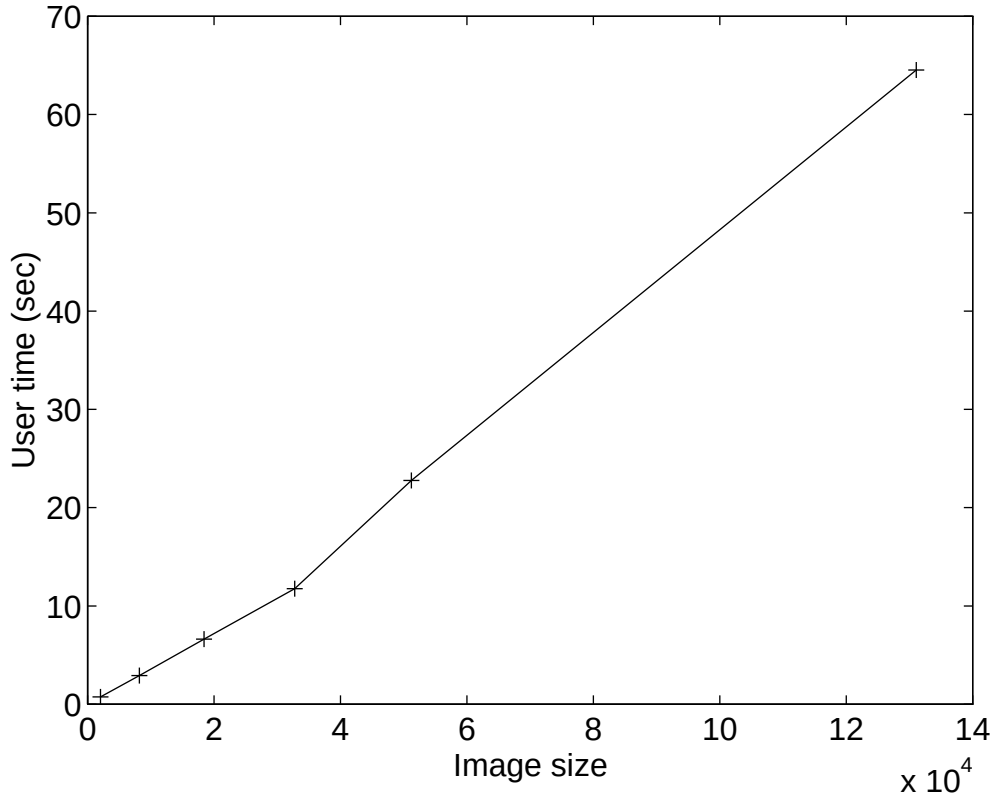

Figure 11: Scalability of the proposed multigrid technique for optical flow computation.

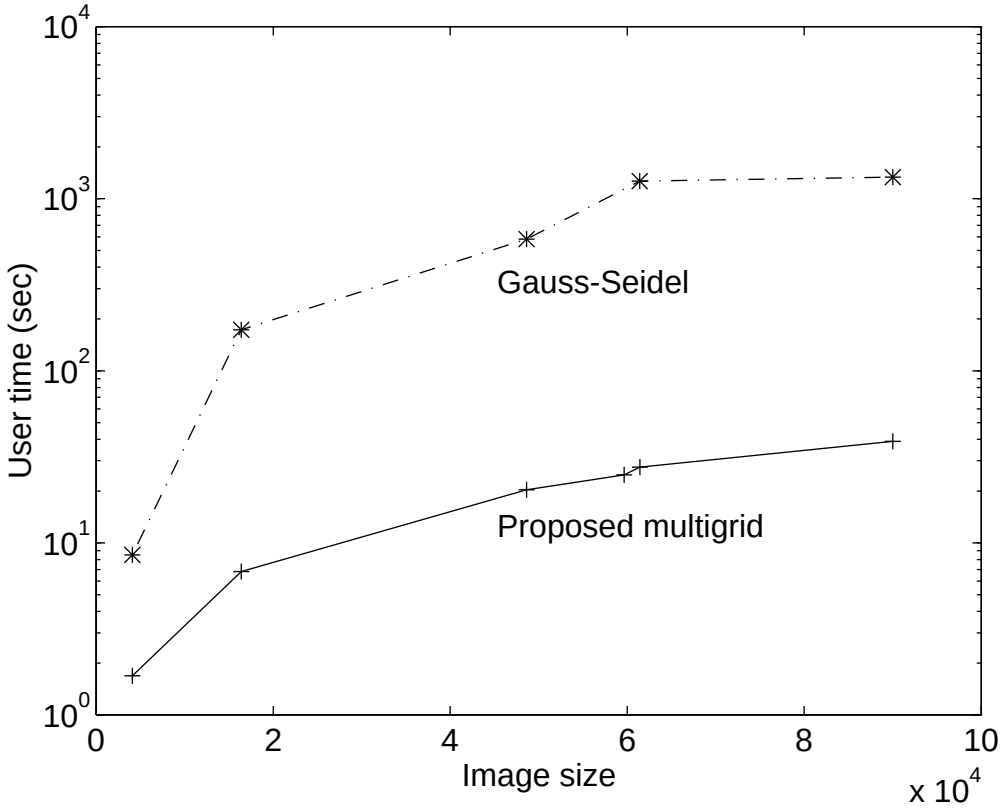

Figure 12: Comparison of execution speeds between Gauss-Seidel and the proposed multigrid for real sequences.

nitude of the solution is found to be quite effective in suppressing flow field in static image regions. A new algebraic multigrid method that employs coarse spaces based on coefficients of the stiffness matrix is used for efficiently solving second-order elliptic PDEs with discontinuities and anisotropies in coefficients, resulting from corresponding Euler-Lagrangian equations. Scalability and speed of the multigrid technique are demonstrated experimentally for both synthetic and real image sequences. Future research will include investigation of higher-order smoothness terms, development of robust multigrids for solving higher-order elliptic PDEs; and intelligent schemes for adapting the constant of smoothness [25] based on the present solution, to account for the inherent nonlinearity in the optical flow problem in the framework of linear PDEs.

Acknowledgment

The authors like to thank Jan Mandel of University of Colorado at Denver for his comments on an earlier draft of this manuscript, and John Barron of University of Western Ontario, Canada for providing ftp access to real image sequences and implementation of Nagel and Enkelmann's optical flow technique.

References

[1] Y. Aloimonos, D. Shulman, "Learning early-vision computations," *J. Opt. Soc. America A*, vol. 6, no. 6, pp. 908–919, 1989.

[2] J.L. Barron, D.J. Fleet, S.S. Beauchemin, "Systems and experiments: Performance of optical flow techniques," *Int. J. Comput. Vision*, vol. 12, no. 1, pp. 43–77, 1994.

[3] M. Bertero, T.A. Poggio, V. Torre, "Ill-posed problems in early vision," *Proc. IEEE*, vol. 76, no. 8, pp. 869–889, 1988.

[4] J.H. Bramble, J.E. Pasciak, J. Wang, J. Xu, "Convergence estimates for multigrid algorithm without regularity assumptions," *Math. Comput.*, vol. 57, 1991.

[5] M. Campani, A. Verri, "Computing optic flow from an overconstrained system of linear algebraic equations," *Proc. Intl. Conf. Comput. Vision*, pp. 22–26, 1987.

[6] P.G. Ciarlet, *The Finite-Element Method for Elliptic Problems*, North-Holland, Amsterdam, 1978.

[7] W. Enkelmann, "Investigations of multigrid algorithms for the estimation of optical flow fields in image sequences," *Comput. Vision, Graphics, Image Processing*, vol. 43, pp. 150–177, 1988.

[8] S. Ghosal, R. Mehrotra, "Robust optical flow estimation," *Proc. IEEE Intl. Conf. Image Processing*, Austin, 1994, to appear.

[9] F. Heitz, P. Bouthemy, "Multimodal estimation of discontinuous optical flow using Markov random fields," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 15, no. 12, pp. 1217–1232, 1993.

[10] B.K.P. Horn, B.G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, pp. 185–204, 1981.

[11] B.D. Lucas, *Generalized image matching by the method of differences*, Ph.D. Dissertation, Carnegie-Mellon Univ., 1984.

[12] P. Vaněk, "Acceleration of convergence of a two-level algorithm by smoothing transfer operator," *Applications of Mathematics*, vol. 37, 1992.

[13] P. Vaněk, J. Mandel, M. Brezina, "Algebraic Multigrid on Unstructured Meshes," in preparation.

[14] S.F. McCormick (ed.), *Multigrid Methods*, SIAM, Philadelphia, 1987.

[15] H.-H. Nagel, W. Enkelmann, "An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 8, pp. 565–593, 1986.

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

[17] P. Saint-Marc, J.-S. Chen, G. Medioni, "Adaptive smoothing: A general tool for early vision," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, no. 6, pp. 514–529, 1991.

[18] D. Terzopoulos, "Image analysis using multigrid relaxation methods," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 8, no. 2, pp. 129–138, 1986.

[19] D. Terzopoulos, "The computation of visible-surface representations," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 10, no. 4, pp. 417–438.

[20] C. Schnörr, "On functionals with greyvalue-controlled smoothness terms for determining optical flow," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 15, no. 10, pp. 1074–1079, 1993.

[21] E. Simoncelli, E. Adelson, D. Heeger, "Probability distributions of optical flow," *Proc. IEEE Conf. Comput. Vision Pattern Recog.*, pp. 310–315, 1991.

[22] A. Singh, "An estimation-theoretic framework for image-flow computation," *Proc. Intl. Conf. Comput. Vision*, Osaka, pp. 168–177, 1990.

[23] A. Verri, T. Poggio, "Motion field and optical flow: Qualitative properties," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 11, no. 5, pp. 490–498, 1989.

[24] A.M. Waxman, K. Wohn, "Contour evolution, neighborhood deformation and global image flow: Planar surface in motion," *Intl J. Rob. Research*, vol. 4, pp. 95–108, 1985.

[25] H. Zheng, S.D. Blostein, "An error-weighted regularization algorithm for image motion-field estimation," *IEEE Trans. Image Processing*, vol. 2, no. 2, pp. 246–252, 1993.