

Line Search Multilevel Optimization as Computational Methods for Dense Optical Flow*

El Mostafa Kalmoun[†], Luis Garrido[‡], and Vicent Caselles[§]

Abstract. We evaluate the performance of different optimization techniques developed in the context of optical flow computation with different variational models. In particular, based on truncated Newton (TN) methods that have been an effective approach for large-scale unconstrained optimization, we develop the use of efficient multilevel schemes for computing the optical flow. More precisely, we compare the performance of a standard unidirectional multilevel algorithm—called multiresolution optimization (MR/Opt)—with that of a bidirectional multilevel algorithm—called full multigrid optimization (FMG/Opt). The FMG/Opt algorithm treats the coarse grid correction as an optimization search direction and eventually scales it using a line search. Experimental results on three image sequences using four models of optical flow with different computational efforts show that the FMG/Opt algorithm outperforms both the TN and MR/Opt algorithms in terms of the computational work and the quality of the optical flow estimation.

Key words. optical flow, line search multigrid optimization, multiresolution, truncated Newton

AMS subject classifications. 68U10, 65N55, 65K10, 49M15

DOI. 10.1137/100807405

1. Introduction. The problem of optical flow computation consists of finding the two-dimensional displacement field that represents apparent motion of objects in a sequence of images. Many efforts have been devoted to this in computer vision and applied mathematics [27, 37, 20, 3, 4, 13, 42, 45, 14, 12, 33, 29, 1, 2, 8, 34, 35, 38]. The computation of optical flow is usually based on the conservation of some property during motion, either the objects' gray level or their shape properties. In a variational setting, the problem is usually formulated as a minimization of an energy function, which is a weighted sum of two terms: a data term coming from the motion modeling and a smoothing term as a result of the regularization process. For standard optical flow algorithms, the data term is usually based on a brightness constancy assumption, which assumes that the object illumination does not change along its motion trajectory. The regularization process ensures that the optical flow estimation

*Received by the editors September 2, 2010; accepted for publication (in revised form) April 5, 2011; published electronically June 23, 2011. This research was initiated while the first author was visiting Centre de Recerca Matemàtica and Universitat Pompeu Fabra in Barcelona during 2008. The first author was partially supported by FRGS grant UUM/RIMC/P-30 (S/O code 11872) offered by Ministry of Higher Education Malaysia. The second and third authors were partially supported by MICINN project, reference MTM2009-08171, and by GRC, reference 2009 SGR 773. The third author also was partially supported by "ICREA Acadèmia" prize for excellence in research funded by the Generalitat de Catalunya.

<http://www.siam.org/journals/siims/4-2/80740.html>

[†]Faculty of Quantitative Sciences, College of Arts and Sciences, Universiti Utara Malaysia, 06010 UUM Sintok, Malaysia (elmostafa@uum.edu.my).

[‡]Departament de Matemàtica Aplicada i Anàlisi, Universitat de Barcelona, 08007 Barcelona, Spain (luis.garrido@ub.edu).

[§]Departament de Tecnologia, Universitat Pompeu Fabra, 08003 Barcelona, Spain (vicent.caselles@upf.edu).

problem is well-posed. Many regularization terms have been investigated in the last two decades ranging from isotropic to anisotropic. Isotropic smoothers tend to blur motion edges, while the anisotropic ones require additional computational resources.

Two strategies might be conducted for the numerical minimization of an energy functional arising in a variational framework like the one above. The first, which is called optimize-discretize, is achieved by discretizing and solving the corresponding Euler–Lagrange equations. In the second approach, called discretize-optimize, one can use directly numerical optimization methods for solving a discrete version of the problem. While the first approach has been commonly used for variational optical flow computation, to the best of our knowledge, the second computational strategy has been less investigated in this context.

Adopting the first strategy, efficient algorithms have been designed to approximate the optical flow. In particular, in [1, 2] the authors used a multiscale approximation to solve the corresponding Euler–Lagrange equations using the Nagel–Enkelmann regularization term [38]. Thus, they computed a series of approximations where each one solves a regularized version of the Euler–Lagrange equations starting from the previous one while keeping the original grid fixed. A different approach is proposed in [11, 43], where the authors use fixed point iterations to solve the corresponding Euler–Lagrange equations, fully implicit in the smoothness term and semi-implicit in the data term. Still, this fixed point iteration leads to implicit equations, and they linearize some of their terms using a warping technique. The equations obtained are fully linearized using a lagged diffusivity method [11, 43, 12]. The final linear system is solved using a linear solver such as a Gauss–Seidel-type method, or a successive overrelaxation (SOR) method. The connections with warping are detailed in [11, 43].

On the other hand, in order to develop efficient and accurate algorithms working in real time, there have been recent efforts to improve performance of optical flow algorithms using multilevel techniques. We distinguish at this stage between two classes of multilevel algorithms. The first, known as a coarse-to-fine multiresolution, uses a sequence of coarse grid subproblems to find a good initialization for the finest grid problem that avoids possible local minima. In this paper, we shall refer to this strategy as multiresolution. The second strategy alternates between solution relaxations using the underlying algorithm and solution corrections obtained from a sequence of subproblems defined on coarse grids. This leads to recursive algorithms such as the so-called V- or W-cycle, which traverse between fine and coarse grids in the mesh hierarchy. We will reserve the term multigrid for this class of two-way multilevel methods. In the case of elliptic PDEs—and for a wide class of problems—multigrid methods are known to outperform multiresolution methods [7, 25].

However, being straightforward to implement, multiresolution methods were more widely used in computer vision and particularly for motion estimation; see, e.g., [20, 3, 5, 42, 45, 17]. For instance, Enkelmann [20] developed a coarse-to-fine algorithm for oriented smoothed optical flow. Also, Cohen and Herlin [17] recently used a multiresolution technique with nonuniform sampling for total variation (TV) regularization.

For the sake of optimal performance, multigrid schemes were developed to solve the resulting Euler–Lagrange equations for both isotropic and anisotropic regularizations; see [23, 46, 22, 29, 28, 14, 12] and the references therein. The first attempts are due to Glazer [23] and Terzopoulos [46] using standard multigrid components. Improvements in performance were reported on simple synthetic images, and later [5], standard multigrids were stated to

be not appropriate due to a possible information conflict between different scales. However, the method was recently better adapted for optical flow computation by tuning the multigrid components. In [22], an entirely algebraic multigrid approach was developed for a weighted anisotropic regularization-based model. A geometrical multigrid based on Galerkin coarse grid discretization approximation was developed for the classical Horn–Schunck method in [29]. This algorithm was extended and parallelized in [28] for real-time computation of vector motion fields for three-dimensional images. Another recent geometric multigrid investigation was presented in [14, 12] with anisotropic regularization (a full account of it can be found in [12]).

All these works have considered data terms based on the brightness constancy assumption, which lead to less accurate optical flow fields when the image sequence contains illumination variations in the temporal domain, which may be often found in real images. In [16], a model is proposed for illumination-invariant optical flow computation, previously introduced in [19] in the context of image registration. The brightness constancy assumption is replaced by the assumption that the shapes of the image move along the sequence. In this context, the terms of the Euler–Lagrange equation corresponding to the data attachment term which contains derivatives of the unit normal vector fields are highly nonlinear. They will not produce systems of equations with a symmetric and positive semidefinite matrix (usually after linearization), the basic systems to which the previous multigrid methods for optical flow have been applied.

For that reason, in this paper we follow the second strategy of discretize-optimize that allows us to handle such variational problems. Instead of computing the Euler–Lagrange equations of the energy model and discretizing and solving them, our approach is based on the use of numerical optimization methods to solve the discrete version of the energy, based on either gray level constancy or shape properties. This leads to the need for developing efficient algorithms for the numerical resolution of a large-scale optimization problem. Therefore only large-scale unconstrained optimization methods are relevant to solving this variational problem. One of these is the truncated Newton (TN) method [18]. It requires only the computation of function and gradient values and then has suitable storage requirements for large-scale problems. However, due to the intensive computations that are required to solve the energy minimization problem, only multilevel versions of the method (multiresolution and multigrid) are expected to provide suitable performance.

In our context, the bidirectional multigrid scheme should be adapted directly to an optimization problem. Current research is still ongoing in this direction for solving some variational problems lacking a governing PDE in different fields. Recently [40, 31], the multigrid strategy has been extended to optimization problems for TN methods. Motivated by variational problems lacking a governing PDE, the multigrid optimization was derived from the full approximation storage (FAS) [9] for nonlinear PDEs but applied directly in an optimization setting. With regard to nonlinear multigrids, the multigrid optimization (MG/Opt) algorithm includes two safeguards that guarantee convergence: bounds on the coarse grid correction that introduce a constrained optimization subproblem, and a line search that eventually scales the coarse grid correction treated as a search direction. In [31], it has been shown that the coarse grid subproblem is a first-order approximation of the fine grid problem. This somehow justifies the introduction of the two safeguards. The first-order approximation suggests that the correction will be reliable only near the restricted approximation, and at the same time it relates

the MG/Opt to the steepest descent method. The latter connection indicates that the coarse grid correction typically will not be a well-scaled descent direction, which, in turn, implies that a line search should be performed to adjust the scale of the multigrid search direction. This may not be necessary when the MG/Opt is near to convergence since in that case it will provide a Newton-like search direction for which a search step equal to 1 will likely be accepted. These connections to both steepest descent and the Newton method suggest that the MG/Opt will perform well far from and near the solution.

Our aim is to develop the MG/Opt method for the computation of optical flow, a problem which has considerable computational resource requirements. To the best of our knowledge, the MG/Opt method is studied here for the first time for optical flow computation. We note here that four different optical flow models are used for testing the performance of the proposed algorithm combining linear and nonlinear data terms with quadratic (or Tikhonov) and TV regularization. Several components of the MG/Opt technique have been tuned for high efficiency, and the algorithm is fully evaluated with respect to one-way multiresolution optimization (MR/Opt). The proposed numerical strategy can be adapted to the minimization of other nonlinear energy functionals such as the illumination invariant model proposed in [16] or depth estimation in stereo problems. Although they are an important motivation for the development of the present techniques, they will not be considered in this paper.

The outline of our paper is as follows. In section 2, we start off with a review of the variational formulation of the optical flow problem. In section 3, we recall the basics of the TN method. In section 4, we present multilevel algorithms applied to optimization problems. First, we discuss the coarse-to-fine multiresolution strategy. Then, after recalling the idea of multigrid for linear systems, we describe its application to optimization-based problems. In section 5, we outline some of the implementation details, namely, the calculation of the gradient of the objective functional for each of the four models considered in this paper, the computation of the Hessian, and the image derivatives. In section 6, we report our experimental results by considering three classical sequences of synthetic images: the translating tree, the diverging tree, and the Yosemite sequences. In section 7, we conclude the paper and indicate future research directions.

2. Variational models for optical flow. Let us consider a sequence of gray level images $I(t, x, y)$, $t \in [0, T]$, $(x, y) \in Q$, where Q denotes the image domain, which we assume to be a rectangle in \mathbb{R}^2 . We shall consider either the case where $t \in [0, T]$ or the case where the sequence is sampled at the times $t_j = j\Delta t$, $j = 0, \dots, K$.

Assuming that the gray level of a point does not change over time, we may write the constraint

$$I(t, x(t), y(t)) = I(0, x, y),$$

where $(x(t), y(t))$ is the apparent trajectory of the point $(x(0), y(0)) = (x, y)$. Differentiating with respect to t and denoting $(u(t, x, y), v(t, x, y)) = (x'(t), y'(t))$, we obtain the optical flow constraint

$$(2.1) \quad I^t + uI^x + vI^y = 0.$$

The vector field $w(t, x, y) := (u(t, x, y), v(t, x, y))$ is called optic flow, and I^t, I^x, I^y denote the partial derivatives of I with respect to t, x, y , respectively. Clearly, the single constraint

(2.1) is not sufficient to uniquely compute the two components (u, v) of the optic flow (this is called the aperture problem) and only gives the component of the flow normal to the image gradient, i.e., to the level lines of the image. As usual, in order to recover a unique flow field a regularization constraint is added. For that, we assume that the optic flow varies smoothly in space, or better, that it is piecewise smooth in Q . This can be achieved by including a smoothness term of the form

$$(2.2) \quad R(w) := \int_Q G(\nabla I, \nabla u, \nabla v) \, dx dy,$$

where $G : \mathbb{R}^2 \times \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}$ is a suitable function. The case $G = \|\nabla u\|^2 + \|\nabla v\|^2$ corresponds to the Horn–Schunk model [27], the case $G = \text{trace}((\nabla w)^T D(\nabla I) \nabla w)$ corresponds to the Nagel–Enkelmann model [38], and the case $G = \sqrt{\|\nabla u\|^2 + \|\nabla v\|^2}$ or $G = \|\nabla u\| + \|\nabla v\|$ corresponds to TV regularization models. For a full account of this and the associated taxonomy, we refer the reader to [47, 26].

Both data attachment and regularization terms can be combined into a single energy functional

$$(2.3) \quad \int_Q (I^t + uI^x + vI^y)^2 \, dx dy + \alpha \int_Q G(\nabla I, \nabla u, \nabla v) \, dx dy,$$

where $\alpha > 0$ is the regularization parameter weighting the relative importance of both terms.

In case of illumination changes, the gray level constancy assumption is violated and may be replaced by the constancy of the gradient [6, 43], which can be expressed in differential form as

$$(2.4) \quad \begin{aligned} \langle w, \nabla I^x \rangle &= 0, \\ \langle w, \nabla I^y \rangle &= 0. \end{aligned}$$

Other cases include the constancy of the gradient direction [15] or the assumption that the shapes of the image (identified as the level lines) move along the sequence [16] (this assumption has been used in [19] in the context of image registration). Higher derivative models have been studied in [43].

The models (2.3), (2.4) do not take into account that video sequences are sampled in time so that our data has the form $I_j(x, y) := I(t_j, x, y)$, $t_j = j\Delta t$, $j = 0, \dots, K$. Without loss of generality, let us assume that $\Delta t = 1$. In that case, the gray level constancy may be expressed as

$$(2.5) \quad I_{j-1}(x, y) = I_j(x + u_j(x, y), y + v_j(x, y)),$$

where $(u_j(x, y), v_j(x, y))$ is the optical flow between images I_{j-1} and I_j . As argued in [1, 2], the linearized gray level constancy constraint (2.1) may not be a good approximation in case of large displacements and the form (2.5) may be more appropriate [12].

A corresponding energy functional can be obtained by combining the nonlinearized form of the brightness constancy assumption and a regularization term. For convenience, we assume

that we want to compute the optical flow between two images $I_1(x, y)$ and $I_2(x, y)$. We may write the energy

$$(2.6) \quad \int_Q \Psi \left((I_1(x, y) - I_2(x + u(x, y), y + v(x, y)))^2 \right) dx dy + \alpha \int_Q G(\nabla I, \nabla u, \nabla v) dx dy,$$

where $\Psi : \mathbb{R} \rightarrow \mathbb{R}$ is an increasing smooth function and $\alpha > 0$. Examples of function G have been given above (see [47, 26, 12] for an account of the many different possibilities).

Observe that the energy (2.6) is nonlinear and nonconvex. In order to develop the basic numerical optimization methods, the variational optical flow problem is set into the simplified minimization form

$$(2.7) \quad \min_w f(w),$$

where $f(w) := D(w) + \alpha R(w)$. Here D denotes a given data term based on either (2.1) or (2.6), and R denotes a regularization term that is either quadratic or the TV. More precisely, we consider in this paper the case of $G = \|\nabla u\|^2 + \|\nabla v\|^2$ or $G = \|\nabla u\| + \|\nabla v\|$. We also consider $\Psi(s^2) = s^2$ for $|s| \leq \gamma$ and $\Psi(s^2) = \gamma^2$ for $|s| > \gamma$.

As discussed in the introduction, we adopt the strategy discretize-optimize. This means that we shall first discretize the objective functional and then solve a finite-dimensional but large-scale optimization problem. Therefore only large-scale unconstrained optimization methods are relevant to solving this variational problem. One of these methods that requires only the computation of function and gradient values and has suitable storage requirements for large-scale problems is the TN method [18]. We note that other choices of the underlying optimization procedure—such as the BFGS quasi-Newton method—are possible. The TN algorithm is embedded in a multilevel strategy, both multiresolution and multigrid.

3. Truncated Newton optimization. As is well known, Newton methods are based on the second-order Taylor approximation of the objective function to build an iterative process for approaching a local minimum. The step s_k to move from the current point w_k to a new iterate w_{k+1} is chosen to be a minimum of the quadratic model of f given by

$$(3.1) \quad q_k(s) = f_k + g_k^T s + \frac{1}{2} s^T H_k s,$$

where $f_k = f(w_k)$, $g_k = \nabla f(w_k)$, and $H_k = \nabla^2 f(w_k)$. The Newton step s_k is then obtained by solving the linear system

$$(3.2) \quad H_k s = -g_k.$$

For large-scale problems, solving this linear system exactly will be very expensive. Instead, TN methods use an iterative method to find an approximate solution to (3.2) and truncate the iterates as soon as a required accuracy is reached or whenever (in the case when the Hessian matrix H_k is not positive definite) a negative curvature is detected. One of the most well known iterative methods among TN methods is the preconditioned conjugate gradient (PCG) algorithm (see Algorithm 2), due to its efficiency and modest memory requirements. In this context, we will refer to the process of finding the step s_k as inner iterations, while the process

of updating w_k using the computed s_k will be called outer iterations. Our discussion in what follows depends on the use of the PCG method as an inner solver.

Depending on how the step solution s_k of the quadratic model (3.1) is exploited, two broad classes of algorithms are distinguished: line search methods and trust region methods. Line search methods scale the step s_k by a factor α_k that approximately minimizes f along the line that passes through w_k in the direction s_k . On the other hand, trust region methods restrict the search for s_k to some region \mathcal{B}_k around w_k in which the algorithm “trusts” that the model function q_k behaves like the objective function f . This paper focuses on the former, the line search method, whereas the comparison between both methods (line search and trust region) will be the object of future work.

There are three components in the TN method related to speeding up the convergence of inner iterations that might have a great impact on its overall efficiency: (a) the truncation criterion, (b) the preconditioning strategy, and (c) how to handle the case when the Hessian matrix is not positive definite. Indeed, for the latter, one of the advantages of trust region over line search is that negative curvature directions can be properly exploited. In the PCG algorithm (Algorithm 2), the inner iterations are truncated as soon as a negative curvature direction is detected, that is, when $p_j^T H_k p_j$ is negative. In our case, we replace the negative curvature test with the equivalent descent direction test [49]; see lines 8–11 of Algorithm 2. From our practical experience, the descent direction test has a better numerical behavior than the negative curvature test.

For the other two components (a) and (b), we use a scaled two-step limited memory BFGS [39] with a diagonal scaling for preconditioning the CG method, and we truncate the inner iterations when the following criterion is satisfied:

$$\|r_j\|_{M_k^{-1}} \leq \zeta_k \|r_0\|_{M_k^{-1}},$$

where r_j is the PCG residual at inner iteration j and $\|r_j\|_{M_k^{-1}} = \sqrt{r_j^T M_k^{-1} r_j}$. M_k is the preconditioning matrix and

$$(3.3) \quad \zeta_k = \max\left(0.5/(k + 1), \|r_0\|_{M_k^{-1}}\right)$$

are both being provided at outer iteration k . The matrix M_k^{-1} may be computed easily if M_k is updated using the BFGS method [41]. In this case, however, M_k^{-1} does not need to be computed since $M_k^{-1} r_j = v_j$, and thus $r_j^T M_k^{-1} r_j = r_j^T v_j$.

Line search TN methods ensure that the new TN step s_k provides a good descent by performing a line minimization along this direction, and then the new outer iterate becomes

$$w_{k+1} = w_k + \alpha_k s_k.$$

Exact line search is avoided due to expensive function evaluations, and normally a sufficient function decrease is obtained by imposing the Wolfe conditions

$$(3.4) \quad f_{k+1} \leq f_k + c_1 \alpha_k g_k^T s_k,$$

$$(3.5) \quad g_{k+1}^T s_k \geq c_2 g_k^T s_k,$$

where $0 < c_1 < c_2 < 1$. In order to obtain a maximum decrease with a minimum number of function evaluations, interpolating polynomials is usually employed. Here in TN a cubic interpolation was used [36].

The algorithm associated to the outer iterations of the line search TN is shown in Algorithm 1. The algorithm iteratively updates w_k by computing a search direction and performing a line search until a given tolerance is reached. In our case we set tolerances on the gradient, ϵ_g , to detect local minima, and on the function values and iterate values, ϵ_f and ϵ_x , to detect convergence. Note also that the preconditioning matrix M_k is updated here, which can be easily done from previous values of w_k and $g(w_k)$; see [41].

Algorithm 1. Line search truncated Newton (outer iterations of TN).

```

1: Initialization:  $w_0, M_0 = Id, \epsilon_g, \epsilon_f, \epsilon_x$ 
2: for  $k = 0$  to  $max\_outer$  do
3:    $g_k = \nabla f(x_k)$ 
4:   if  $\|g_k\| < \epsilon_g$  then
5:     exit with solution  $w_k$ 
6:   end if
7:   Compute  $s_k$  by calling Algorithm 2.
8:   Perform a line search to scale the step  $s_k$  by  $\alpha_k$ .
9:    $w_{k+1} = w_k + \alpha_k s_k$ .
10:  Update  $M_{k+1}$  by the BFGS formula.
11:  if  $|f_{k+1} - f_k| < \epsilon_f$  or  $\|w_{k+1} - w_k\| < \epsilon_x$  then
12:    exit with solution  $w_k$ 
13:  end if
14: end for

```

4. Multilevel methods. Multilevel methods use a set of subproblems defined on coarser levels of resolution in order to improve the convergence to the global optimum of (2.7) when compared to using only one level of resolution. Using multiple levels of resolution allows us to deal, in the case of optical flow computation, with large displacements and enables the convergence to the global minimum since some local minima may disappear at sufficient coarse resolutions.

Let us denote by Ω_i the image domain at level i , where $i = 0, \dots, r$, where $i = 0$ corresponds to the finest level of resolution and $i = r$ to the coarsest one. Grid spacing at the coarser grid Ω_{i+1} is usually twice the spacing at the grid Ω_i . Two approaches are currently widely used in the computer vision field, namely, the multiresolution and the multigrid methods. Both are explained below.

4.1. Multiresolution methods. Multiresolution methods have been applied successfully in many computer vision and image analysis problems where the problem can be expressed as a global optimization problem [20, 3, 5, 42, 45, 17]. Multiresolution methods use a series of coarse-to-fine resolution levels to obtain an estimate of the solution to the problem. An initial estimate is obtained at the coarsest level. In our case, this estimate may be obtained by applying the TN algorithm (see Algorithm 1). The estimate is then extended to the next level

Algorithm 2. Preconditioned conjugate gradient (inner iterations of TN).

```

1: Initialization:  $z_0 = 0$ ,  $r_0 = -g_k$ ,  $v_0 = M_k^{-1}r_0$ ,  $p_0 = v_0$ ,  $\epsilon = 10^{-10}$ ,  $\zeta_k$  (see (3.3))
2: for  $j = 0$  to  $max\_inner$  do
3:   // Singularity test
4:   if ( $|r_j^T v_j| < \epsilon$  or  $|p_j^T H_k p_j| < \epsilon$ ) then
5:     exit with  $s_k = z_j$  (for  $j = 0$  take  $s_k = -g_k$ )
6:   end if
7:    $\alpha_j = r_j^T v_j / p_j^T H_k p_j$ ,  $z_{j+1} = z_j + \alpha_j p_j$ 
8:   // descent direction test replaces negative curvature test
9:   if ( $g_k^T z_{j+1} \geq g_k^T z_j - \epsilon$ ) then
10:    exit with  $s_k = z_j$  (for  $j = 0$  take  $s_k = -g_k$ )
11:   end if
12:    $r_{j+1} = r_j - \alpha_j H_k p_j$ ,  $v_{j+1} = M_k^{-1} r_{j+1}$ 
13:   // truncation test (note that  $\|r_{j+1}\|_{M_k^{-1}} = r_{j+1}^T v_{j+1}$ )
14:   if ( $r_{j+1}^T v_{j+1} \leq \zeta_k g_0^T v_0$ ) then
15:     exit with  $s_k = z_{j+1}$ 
16:   end if
17:    $\beta_j = r_{j+1}^T (v_{j+1} - v_j) / r_j^T v_j$ ,  $p_{j+1} = v_{j+1} + \beta_j p_j$ .
18: end for
19: exit with  $s_k = z_j$ 

```

of resolution, where it is refined. This process is repeated until the finest level is reached, where the final estimate is obtained (see Algorithm 3), where the MR (multiresolution) method is shown. This function is called $MR(L-1, x_{L-1,0})$, where L is the number of resolution levels and $x_{L-1,0}$ is the initial estimate on the coarsest level.

The advantage of coarse-to-fine multiresolution is that a good initial guess may be obtained for the finest grid problem by estimating a solution to the problem using the coarser grids. However, for linear systems it is currently known that one-way multilevel methods do not reach the optimal efficiency of standard bidirectional multigrid methods, which are detailed in the next section.

4.2. Multigrid methods. Multigrid methods were originally developed to solve elliptic PDEs and at present are known to be among the most powerful numerical methods for improving computational efficiency of a wide class of equations. For a classical reference on multigrids, we refer to [9, 10]. The main characteristic of multigrid algorithms is based on the observation that different frequencies are present in the error of the solution of the finest grid problem. Some algorithms, called *smoothers* (such as Gauss–Seidel), are known to efficiently reduce the high frequency components of the error on a grid (or, in other words, the components whose “wavelength” is comparable to the grid’s mesh size). However, these algorithms have a small effect on the low frequency error components. This is the reason why the application of schemes such as Gauss–Seidel to solve a problem, for a given grid, effectively reduces the error in the first iterations of the procedure (due to smoothing of the high frequency errors) but then converges slowly to the solution (due to smoothing of the low

Algorithm 3. Multiresolution method.

```

1: function MR( $i, x_{i,0}$ )
2:    $i = L - 1$ : coarsest level of resolution
3:    $x_{i,0}$ : initial estimate
4:   repeat
5:     Apply TN with initial estimate  $x_{i,0}$ .
6:     Let  $x_{i,1}$  be the result.
7:     if  $i > 0$  then
8:       // Prolongate current estimate to next finer level
9:        $x_{i-1,0} = Px_{i,1}$ 
10:    end if
11:     $i = i - 1$ 
12:  until  $i < 0$ 
13:  return  $x_{0,1}$ 

```

frequency errors).

One may observe that low frequency errors appear as higher frequencies on coarser grids. The latter may be effectively reduced using a smoother in the coarse grid. Moreover, smoothing at a coarser level typically has a much lower cost than on finer grids. The core idea of multigrid is to use a sequence of subproblems defined on coarser grids as a means to accelerate the solution process (by relaxation such as Gauss–Seidel) on the finest grid. This leads to recursive algorithms such as the so-called V- or W-cycle, which traverse between fine and coarse grids in the mesh hierarchy. Since ultimately only a small number of relaxation steps must be performed on each level, multigrid provides an asymptotically optimal method whose complexity is only $O(N)$, where N is the number of mesh points.

4.2.1. Multigrid for linear systems. We recall first the basics of multigrid techniques. We consider a sparse linear system which typically results from the discretization of a PDE on a fine grid Ω_i with a given grid spacing h :

$$(4.1) \quad A_i x_i = b_i.$$

Let $x_{i,0}$ be an initial approximation of (4.1) and x_i be the exact solution. The first step of the algorithm consists of smoothing the error $x_i - x_{i,0}$ by applying N_0 iterations of a relaxation scheme S to (4.1) that has the smoothing property [25]. Examples of the smoother S are Richardson, Jacobi, or Gauss–Seidel. The obtained smooth approximation $\bar{x}_{i,1}$ satisfies (4.1) up to a residual r_i :

$$A_i x_{i,1} = b_i - r_i.$$

The corresponding error equation is therefore given by

$$(4.2) \quad A_i e_i = r_i,$$

where $e_i = x_{i,1} - x_i$ is the unknown smooth error. Equation (4.2) is then solved on a coarser grid Ω_{i+1} with a grid spacing that has to be larger than h (typical choice is $2h$, but other

choices are possible). For this, the fine grid error equation must be approximated by a coarse grid equation:

$$(4.3) \quad A_{i+1}e_{i+1} = r_{i+1}.$$

We therefore need to transfer the residual to the coarse grid and construct a coarse version of the fine matrix. Let R be a *restriction* operator mapping functions on Ω_i to functions on Ω_{i+1} (common examples are injection and full weighting; see [10]). The coarse residual is then given by

$$r_{i+1} = Rr_i.$$

The coarse grid matrix may be obtained by discretization on Ω_{i+1} or by using the Galerkin coarse grid approximation $A_{i+1} = RA_iP$. Here P is a *prolongation (or interpolation)* operator mapping functions from Ω_{i+1} to Ω_i (standard examples are linear and bilinear interpolation).

The result $e_{i+1,*}$ due to solving (4.3) is transferred back to the fine grid to obtain

$$e_{i,*} = Pe_{i+1,*}.$$

With this approximation of the fine grid error, the approximate solution is corrected to obtain

$$x_{i,2} = x_{i,1} + e_{i,*}.$$

In order to damp high frequency error components that might arise due to the interpolation, N_1 iterations of the smoother S are applied on the fine grid to get the new iterate $x_{i,3}$.

The coarse grid problem (4.3) has a lower dimension than the original problem (4.1), but it must be solved accurately for each iteration, which can be very costly. In a typical multigrid with three levels or more, this problem is solved by recursively calling the algorithm γ times until reaching the coarsest grid, where we solve (4.3) exactly with a negligible cost. The steps of this multigrid algorithm are summarized in Algorithm 4.

The resulting $x_{i,3}$ may be injected iteratively as initialization in Algorithm 4 until the residual on the finest grid,

$$r_{i,3} = b_i - A_i x_{i,3},$$

is smaller than a given tolerance.

Note that for the smoothing steps N_0 and N_1 , only the sum is important in a convergence analysis. Typical values for the tuple (N_0, N_1) are $(1, 1)$, $(2, 1)$, and $(2, 2)$. For the cycling parameter γ , only the values 1 and 2 are commonly used. The so-called V-cycle corresponds to $\gamma = 1$, and the W-cycle corresponds to $\gamma = 2$. This is illustrated in Figure 1 for three levels.

An important variation of multigrids, which is known as the *full multigrid method* (FMG) [9] or *nested iteration technique* [25], combines a multiresolution approach with a standard multigrid cycle. The FMG starts at the coarsest grid level, solves a very low-dimensional problem, extends the solution to a finer space, performs a multigrid cycle, and repeats the process until a multigrid cycle is performed on the finest grid level. In this way, a good initialization is obtained to start the multigrid cycle on the finest level, which usually reduces the total number of iterations required. The method is illustrated in Figure 1 for three levels using a V-cycle.

Algorithm 4. Linear multigrid V-cycle.

```

1: function MG-cycle( $i, A_i, b_i, x_{i,0}$ )
2:    $i$ : level
3:    $A_i, b_i$ : defines linear system to be solved
4:    $x_{i,0}$ : initialization
5:   if  $i$  is the coarsest level then
6:     Solve  $x_{i,1} = A_i^{-1}b_i$  // Solve exactly
7:   else
8:      $x_{i,1} := S(x_{i,0}, A_i, b_i, N_0)$  // Smooth with  $N_0$  iterations
9:      $x_{i+1,0} = 0$ 
10:     $r_i := b_i - A_i x_{i,1}$  // Residual
11:     $r_{i+1} := R r_i$  // Restrict
12:    Call MG( $i + 1, A_{i+1}, r_{i+1}, x_{i+1,0}$ )  $\gamma$  times // Recursive call
13:    Let  $e_{i+1,*}$  be the returned value
14:     $e_{i,*} := P e_{i+1,*}$  // Extend
15:     $x_{i,2} := x_{i,1} + e_{i,*}$  // Apply correction step
16:     $x_{i,3} := S(x_{i,2}, A_i, b_i, N_1)$  // Postsmooth with  $N_1$  iterations
17:    Return  $x_{i,3}$ 
18:  end if

```

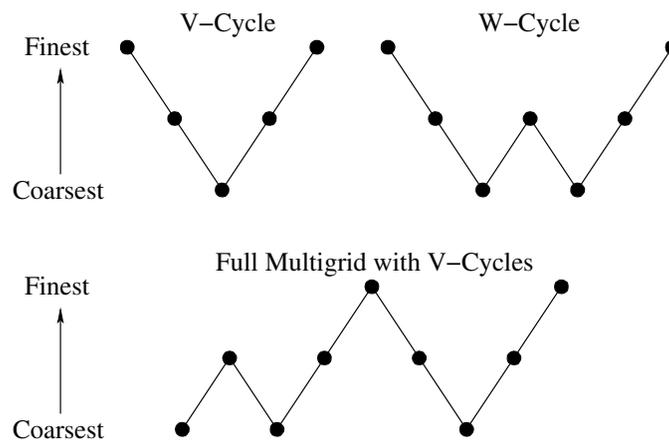


Figure 1. *V-cycle, W-cycle, and FMG with V-cycle.*

4.2.2. Multigrid for optimization problems. As commented previously, the multigrid strategy has been recently extended to optimization problems for both line search methods and trust region problems. Previous works on nonlinear multigrid methods applied the techniques directly to systems of nonlinear equations obtained by solving the first-order optimality conditions. This approach is not suitable for problems that are not easily transformed into systems of nonlinear equations. Another possible way to use multigrid for optimization problems is to apply the linear multigrid as inner iterations for solving the linear system (3.2) for a given iterate w_k . This idea implicitly assumes that the Newton method is the underlying

optimization algorithm and that the Hessian matrices can be explicitly computed. Since many large-scale optimization algorithms require only the computation of function and gradient values, it is less obvious how such multigrid algorithms can be applied. Moreover, solving the Newton equations using multigrid may lead only to a better unilevel optimization algorithm as long as the multigrid technique is not applied in the outer iterations.

We deal here with a multigrid algorithm that works directly with the optimization problem, enabling us to work with problems that are not easily transformed into systems of nonlinear equations. The line search MG/Opt strategy is described in the following.

As with multigrid for linear systems, in order to solve the optimization problem (2.7) over an original (finest) grid level $i = 0$, a sequence of optimization subproblems are considered on nested coarser grids. Given a fine grid level $i \geq 0$, let f_i denote a representation of the objective function f on this level. Let $w_{i,0}$ be an initial fine approximation to the optimization problem at level i . For the finest level $i = 0$, the optimization problem corresponds to the minimization of f_0 , the finest representation of the objective function f . However, for a coarser level i , the optimization problem corresponds to the minimization of a function h_i that shall be specified later. The first step in the multigrid procedure is called a preoptimization phase (analogous to presmoothing in linear multigrid) and consists of applying N_0 iterations of an optimization procedure such as TN (in our case, line search TN) to h_i to obtain w_{i,N_0} . As for nonlinear multigrid, this w_{i,N_0} is transferred to a coarser grid to obtain $w_{i+1,0} := R w_{i,N_0}$. The residual at this level is given by

$$r_{i+1} := \nabla f_{i+1}(w_{i+1,0}) - R \nabla h_i(w_{i,N_0}).$$

The function

$$(4.4) \quad h_{i+1}(w_{i+1}) = f_{i+1}(w_{i+1}) - r_{i+1}^T w_{i+1}$$

is the function to be minimized on the coarse grid level $i + 1$. We take $h_0 := f_0 = f$ on the finest level.

Assume that $w_{i+1,*}$ is a solution to the optimization of (4.4). The error between $w_{i+1,*}$ and the initial approximation $w_{i+1,0}$ is called coarse grid correction. This correction is extended back to level i :

$$s_{i,N_0} = P(w_{i+1,*} - w_{i+1,0}).$$

In an optimization context, this correction step used to update the current solution w_{i,N_0} to w_{i,N_0+1} is considered as a search direction called *recursive* to distinguish it from the *direct* search direction step that is computed by a given optimization procedure on the same grid level.

Finally, in order to remove the oscillatory components that may have been introduced by the correction step, one may finish with a postoptimization phase by applying N_1 iterations of the optimization procedure (in our work, line search TN) to h_i with initial guess w_{i,N_0+1} , obtaining w_{i,N_0+N_1+1} .

The coarse optimization subproblem for h_{i+1} can be seen as a first-order approximation to the fine grid problem h_i since their gradients coincide on $w_{i+1,0}$:

$$\nabla h_{i+1}(w_{i+1,0}) = \nabla f_{i+1}(w_{i+1,0}) - r_{i+1} = R \nabla h_i(w_{i,N_0}).$$

The first-order approximation suggests that the correction will be reliable only near the restricted approximation, and it relates at the same time the MG/Opt algorithm to the steepest descent method. The latter connection indicates that the coarse grid correction will not typically be a well-scaled descent direction, which, in turn, implies that a line search should be performed to adjust the scale of the recursive search direction. However, it can be demonstrated that the multigrid algorithm is related to the Newton method, in the sense that the recursive search direction s_i is an approximate Newton direction [31]. Accordingly, in order to improve computational efficiency, the line search for a recursive direction step s_i is performed only if $w_{i,k} + \alpha_k s_i$ with $\alpha_k = 1$ does not reduce the value of h_i . That is, if $h_i(w_{i,k} + s_i) < h_i(w_{i,k})$, we update with $w_{i,k+1} = w_{i,k} + s_i$. Otherwise a line search is performed.

In [31] bound constraints are proposed for the optimization subproblem. In the context of our work, we have seen that bound constraints may improve robustness of the MG/Opt algorithm. These bound constraints may be implemented by means of active sets [41]. However, in our case we have seen that we do not need to set up such bounds since the line search TN algorithm used to optimize the subproblem h_{i+1} already provides similar constraints. Line search algorithms restrict the search at each iteration $w_{i+1,k}$ to an upper bound which depends on the gradient norm values, and thus ensure that the update $w_{i+1,k+1}$ is not far away from $w_{i+1,k}$.

We have implemented the MG/Opt algorithm using a full multigrid method (FMG) to solve the problem, and the resulting algorithm will be denoted by FMG/Opt. As for the linear case, the FMG/Opt starts at the coarsest grid level where enough TN iterations are performed and prolongates the solution to the next finer level i , where V_i iterations of the MG/Opt cycle are performed.

Algorithm 5 shows the V-cycle algorithm used for FMG/Opt. At each iteration the algorithm computes a step s_i either directly using the inner iteration of the TN method (Algorithm 2) on the current level, or recursively by means of the multigrid strategy. However, as noted in [48, 24], the recursive call is useful only if $\|Rg_{i,k}\|$ is large enough compared to $\|g_{i,k}\|$, where $g_{i,k} = h_i(w_{i,k})$ and $k \leq N_0$. Thus we restrict the use of a coarser level $i + 1$ to the case where

$$(4.5) \quad \|Rg_{i,k}\| > \kappa_g \|g_{i,k}\| \text{ and } \|Rg_{i,k}\| > \epsilon_g$$

for some constant $\kappa_g \in (0, \min \|R\|)$ and where $\epsilon_g \in (0, 1)$ is a measure of first-order criticality for h_{i+1} . The latter condition is easy to check before trying to compute a step at level $i + 1$.

5. Implementation issues. The numerical algorithms TN, MR/Opt, and FMG/Opt have been implemented in C using the MegaWave2 library. In this section, we provide details about derivative computation for both the objective function and the image.

5.1. Functional gradient calculation. The gradient of the objective function in (2.7) is calculated analytically and given by

$$g = \nabla f = \begin{pmatrix} f^u \\ f^v \end{pmatrix} = \begin{pmatrix} D^u + \alpha R^u \\ D^v + \alpha R^v \end{pmatrix}.$$

Algorithm 5. The V-cycle for the FMG/Opt algorithm.

```

1: function MG/Opt-cycle( $i, h_i, w_{i,0}$ )
2:    $i$ : level
3:    $h_i$ : coarse objective function at level  $i$  //  $h_0 = f$ 
4:    $w_{i,0}$ : initial approximation to  $h_i$ 
5:   if  $i$  is coarsest level then
6:     task = optimize,  $N = \text{max\_outer}$ 
7:   else
8:     task = preoptimize,  $N = N_0$ 
9:   end if
10:   $k_0 = k = 0$ 
11:  loop
12:    if task is optimize, preoptimize, or postoptimize then
13:      Compute  $s_{i,k}$  by calling Algorithm 2
14:    else
15:       $w_{i+1,0} = R w_{i,k}$  // Restrict
16:       $r_{i+1} = \nabla f_{i+1}(w_{i+1,0}) - R \nabla h_i(w_{i,k})$  // Residual
17:       $h_{i+1}(w_{i+1}) = f_{i+1}(w_{i+1}) - r_{i+1}^T w_{i+1}$  // Next coarse objective function
18:      Call MG/Opt-cycle( $i+1, h_{i+1}, w_{i+1,0}$ )
19:      Let  $w_{i+1,*}$  be the returned solution
20:       $s_{i,k} = P(w_{i+1,*} - w_{i+1,0})$  // Correction step
21:    end if
22:    Perform a line search to scale the step  $s_{i,k}$  by  $\alpha_k$ 
23:     $w_{i,k+1} = w_{i,k} + \alpha_{i,k} s_{i,k}$ 
24:    Update  $M_{i,k+1}$  by the BFGS formula
25:    if  $|f_{i,k+1} - f_{i,k}| < \epsilon_f$  or  $\|w_{i,k+1} - w_{i,k}\| < \epsilon_x$  then
26:      return  $w_{i,k+1}$ 
27:    end if
28:     $k := k + 1$ 
29:    // Select next task to do
30:    if task is preoptimize and (4.5) is satisfied then
31:      task = recursive-call
32:    else if task is recursive-call then
33:      task = postoptimize,  $N = N_1, k_0 = k$ 
34:    end if
35:    // Check if maximum number of outer iterations has been reached
36:    if (task is optimize, preoptimize, or postoptimize) and  $(k - k_0 = N)$  then
37:      return  $w_{i,k}$ 
38:    end if
39:  end loop

```

5.1.1. Horn–Schunck data term. For the linear data term we have

$$D(u, v) = \frac{1}{2} \sum_{i,j} \psi \left([I^x u_{i,j} + I^y v_{i,j} + I^t]^2 \right),$$

where i (respectively, j) corresponds to the discrete column (respectively, row) of the image, with the coordinate origin located in the top-left corner of the image. The function ψ is used to enhance robustness with respect to outliers. In our work we have used

$$\psi(x^2) = \begin{cases} x^2 & \text{if } |x| \leq \gamma, \\ \gamma^2 & \text{otherwise,} \end{cases}$$

where γ is a given threshold. The gradient D for $|x| \leq \gamma$ is therefore given by

$$\begin{pmatrix} D_{i,j}^u \\ D_{i,j}^v \end{pmatrix} = \begin{pmatrix} I^x [I^x u_{i,j} + I^y v_{i,j} + I^t] \\ I^y [I^x u_{i,j} + I^y v_{i,j} + I^t] \end{pmatrix},$$

where $D_{i,j}^u$ and $D_{i,j}^v$ refer to the partial derivative of $D(u, v)$ with respect to variables $u_{i,j}$ and $v_{i,j}$, respectively. Here I^x, I^y, I^t are the spatial and temporal image derivatives for which the computation is explained in section 5.3. Note that for $|x| > \gamma$ the gradient D is $(D_{i,j}^u, D_{i,j}^v)^T = (0, 0)^T$.

5.1.2. Intensity constancy-based data term. The nonlinear data term based on the constancy assumption is as follows:

$$D(u, v) = \frac{1}{2} \sum_{i,j} \psi \left([I_2(i + u_{i,j}, j + v_{i,j}) - I_1(i, j)]^2 \right).$$

The gradient of this functional for $|x| \leq \gamma$ is given by

$$\begin{pmatrix} D_{i,j}^u \\ D_{i,j}^v \end{pmatrix} = \begin{pmatrix} I_2^x(i + u_{i,j}, j + v_{i,j}) [I_2(i + u_{i,j}, j + v_{i,j}) - I_1(i, j)] \\ I_2^y(i + u_{i,j}, j + v_{i,j}) [I_2(i + u_{i,j}, j + v_{i,j}) - I_1(i, j)] \end{pmatrix}.$$

5.1.3. Quadratic regularization term. Now we consider the regularization term and start with the quadratic functional. We have

$$R(u, v) = \frac{1}{2} \sum_{i,j} \|\nabla_{i,j}(u, v)\|^2,$$

where $\|\nabla_{i,j}(u, v)\| = \sqrt{(u_{i,j}^x)^2 + (u_{i,j}^y)^2 + (v_{i,j}^x)^2 + (v_{i,j}^y)^2}$. The partial derivatives of u, v are computed by forward finite differences with a discretization step h , that is, $u_{i,j}^x = h^{-1}(u_{i+1,j} - u_{i,j})$ and $u_{i,j}^y = h^{-1}(u_{i,j+1} - u_{i,j})$ (derivatives v^x and v^y are computed similarly). The gradient of this functional is obtained as

$$\begin{pmatrix} R_{i,j}^u \\ R_{i,j}^v \end{pmatrix} = \frac{1}{h^2} \begin{pmatrix} 4u_{i,j} - u_{i-1,j} - u_{i,j-1} - u_{i+1,j} - u_{i,j+1} \\ 4v_{i,j} - v_{i-1,j} - v_{i,j-1} - v_{i+1,j} - v_{i,j+1} \end{pmatrix}.$$

5.1.4. Total variation (TV) term. In this case we suppose that

$$R(u, v) = \sum_{i,j} \|\nabla_{i,j}(u, v)\|.$$

To overcome the problem of nondifferentiability of the TV, a widely used technique consists of approximating R by a differentiable function:

$$R(u, v) \approx \sum_{i,j} \psi_{i,j},$$

where $\psi_{i,j} = \sqrt{(u_{i,j}^x)^2 + (u_{i,j}^y)^2 + (v_{i,j}^x)^2 + (v_{i,j}^y)^2 + \mu}$ and μ is a small positive parameter. Using again forward finite differences, the gradient of the last approximation is given by

$$\begin{pmatrix} R_{i,j}^u \\ R_{i,j}^v \end{pmatrix} = \frac{1}{h} \begin{pmatrix} \frac{u_{i-1,j}^x}{\psi_{i-1,j}} + \frac{u_{i,j-1}^y}{\psi_{i,j-1}} - \frac{u_{i,j}^x + u_{i,j}^y}{\psi_{i,j}} \\ \frac{v_{i-1,j}^x}{\psi_{i-1,j}} + \frac{v_{i,j-1}^y}{\psi_{i,j-1}} - \frac{v_{i,j}^x + v_{i,j}^y}{\psi_{i,j}} \end{pmatrix}.$$

Another approximation of the TV is given by

$$R(u, v) \approx \sum_{i,j} \varphi_\mu(\|\nabla_{i,j}(u, v)\|), \quad \text{where } \varphi_\mu(x) = \begin{cases} |x| & \text{if } |x| \geq \mu, \\ x^2/2\mu + \mu/2 & \text{otherwise.} \end{cases}$$

Theoretically, this approximation is two times better than the first standard approximation. However, numerically in our implementation, both approximations lead to almost the same results.

5.2. Hessian calculation. For computing the Newton direction in the TN method, the linear conjugate gradient is a Hessian-free procedure (see Algorithm 2) and needs to supply only a routine for computing the product of the Hessian with Newton direction p . This matrix-vector product is computed via forward finite differences:

$$H(w) p = \frac{g(w + \epsilon p) - g(w)}{\epsilon},$$

where ϵ is chosen to be the square root of the machine precision divided by the norm of w .

5.3. Image gradient calculation. Differentiation is an ill-posed problem [6], and regularization may be used to obtain good numerical derivatives. Such regularization may be accomplished with a low-pass filter such as the Gaussian and is essential for motion estimation [4, 44]. More recently, [21] proposed using a matched pair of low-pass and differentiation filters as a gradient operator, which are the ones used in this work.

Usually, the derivative at an integer $I_{i,j}^x$ is computed by applying a separable filter composed of a Gaussian filter (alternatively, a matched low-pass filter) in the y direction and the derivative of the Gaussian (alternatively, a matched derivative) in the x direction. Conversely, the computation of the derivative in the y direction at an integer point is performed by applying a separable filter composed of a Gaussian (or matched low-pass filter) in the x direction and the derivative of the Gaussian (or a matched derivative) in the y direction.

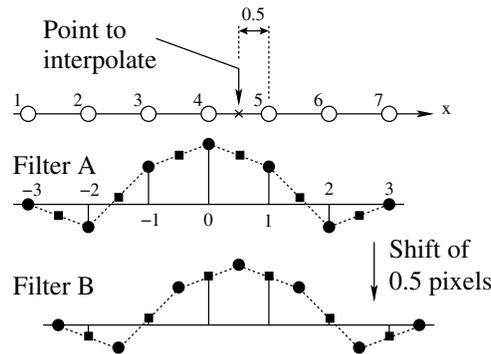


Figure 2. Filter tap interpolation by means of a shift in the Fourier domain.

For motion estimation applications it may be necessary to compute the gradient at non-integer points, since noninteger displacements are allowed. This is the case, for instance, of the nonlinear data term of section 5.1.2. In such cases, a simple way to proceed is a two-step process: in a first step, the original image is interpolated at the required points using a bilinear interpolation or an interpolation kernel such as [30], and then the derivative is computed using the latter interpolated points. Another way to proceed is to first compute the gradient at integer points and then apply an interpolation kernel over the gradient values using these latter values. Both procedures are theoretically equivalent since differentiation and interpolation are based on linear operators (and thus, they are interchangeable). Let us call this approach linear gradient interpolation.

In this work the computation of the gradient at noninteger points is done by means of a shift in the Fourier domain. Moreover, rather than shifting the image (or gradient) to obtain the interpolated values, the derivative filter taps are shifted so as to obtain the filter coefficients that have to be applied on the integer image values in order to obtain the corresponding interpolated gradient value.

Assume that the gradient and interpolation kernels are linear, shift-invariant, and separable. Such kernels may be found in [21]. Without loss of generality, the interpolation problem in the Fourier domain can thus be restricted to one dimension. Let us consider Figure 2, which shows the proposed technique. On top several samples of a one-dimensional signal are shown. Filter taps A are used to obtain the gradient at integer points (in the example the filter A is centered on $x = 4$), whereas filter taps B may be used to obtain the gradient at every noninteger point halfway between two integer points (in the example the filter B is centered on $x = 4.5$). The filter taps B are obtained from filter taps A by means of a shift of 0.5 in the Fourier domain. They can thus be applied directly on the original data. We will call this procedure Fourier gradient interpolation.

Figure 3 shows an example in which a set of matched filters [21] of size 9 are interpolated at noninteger points by a shift in the Fourier domain and by a linear interpolation. Performing a linear interpolation on the filter taps in order to apply them to the original data values is equivalent to the gradient linear interpolation approach described above. Note that, as expected, the obtained filter taps are different for both methods. The experimental section will show that Fourier-based interpolation leads to a better performance than linear interpolation,

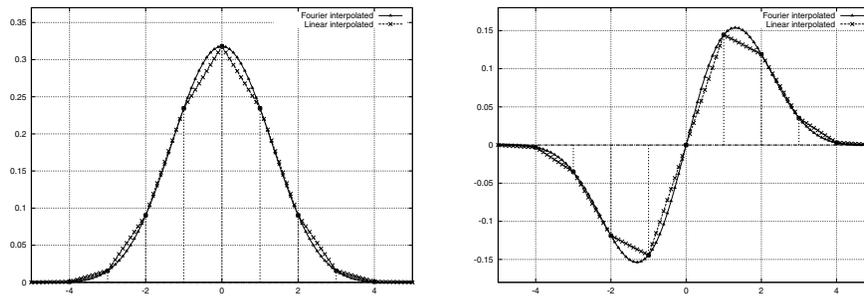


Figure 3. A set of matched filters of size 9 are interpolated at noninteger points by a shift in the Fourier domain and by a linear interpolation.

especially in the multigrid approach.

In the two-dimensional case, if I^x has to be computed at point $(i + \Delta i, j + \Delta j)$, where (i, j) is an integer discrete image position and $\Delta i < 1$ and $\Delta j < 1$, the matched low-pass filter in the y direction (respectively, the matched derivative in the x direction) is obtained by shifting the corresponding taps Δi (respectively, Δj) in the Fourier domain. A similar procedure is used to compute I^y at a noninteger point.

The previous scheme has a high computational load if the Fourier shift has to be applied to each noninteger position where one needs to interpolate the gradient. The computational load can be reduced significantly if, at the initialization of the optimization algorithm, the gradient is computed with the Fourier gradient interpolation at all points of a grid of size h/D . Then, each time one has to compute the gradient at a noninteger point, a bilinear interpolation with the neighboring pixels is used. Figure 3 shows the case where $D = 10$.

6. Experimental results. To assess the performance of the proposed algorithms, we use three classical sequences of synthetic images that consist of scenes of various complexity, namely, the translating tree, the diverging tree, and the Yosemite sequences. The reference frame and the corresponding ground truth of the synthetic sequences are shown in Figures 4 and 5. The image size of the tree sequences is 150×150 , while the Yosemite sequence is of size 316×252 .

Since we are more interested in the computational complexity of the proposed algorithms, we first compare the CPU time needed by each numerical algorithm to reach a similar accuracy using four optical flow models. All tests are performed on a PC with a 2.0 GHz Mobile Intel DualCore processor and 1 GB RAM. We also compute the number of function and gradient evaluations that were performed by each algorithm to reach the estimated optical flow. We measure the overall number of function evaluations as

$$N_f = \sum_{i=0}^{L-1} \frac{N_{f,i}}{F_i},$$

where $N_{f,i}$ is the number of function evaluations performed by the optimization algorithm at resolution level i . For the TN and our MR/Opt and FMG/Opt algorithms, the function is evaluated during the line search procedure. F_i is the mesh resolution ratio of a given level i with respect to the finest level 0. In this work $F_i = 2^{2i}$.

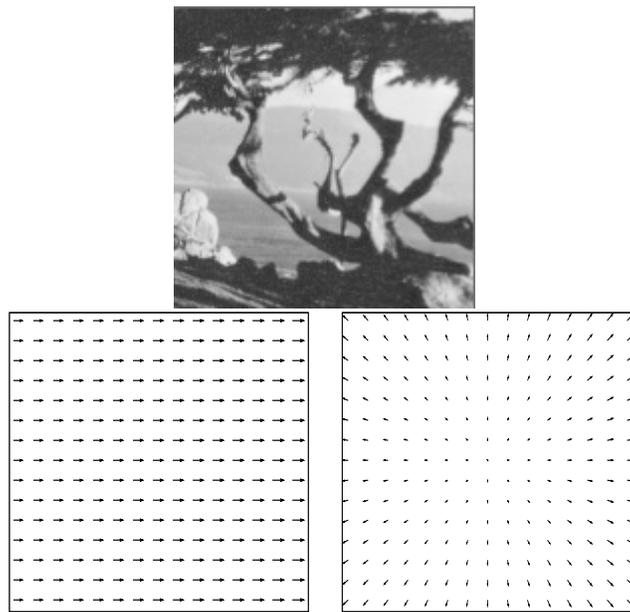


Figure 4. At the top, one frame of the original sequence is shown. At the bottom, the ground truth for the corresponding translating (left) and diverging (right) is shown. Motion vectors have been scaled by a factor of 2.5 for better visibility.

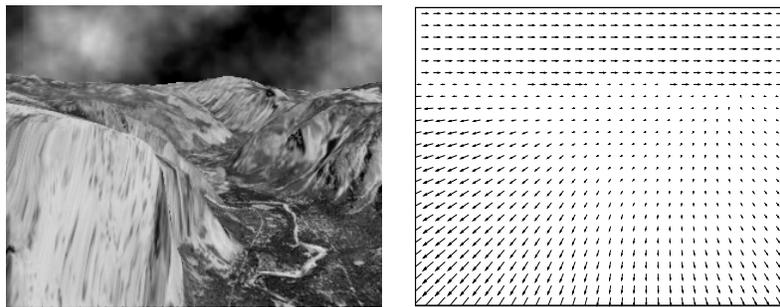


Figure 5. At left, one frame of the Yosemite sequence is shown. At right, the corresponding ground truth is depicted, where motion vectors have been scaled by a factor of 2.5 for better visibility.

The number of gradient evaluations is defined in a similar manner. While the function is evaluated only during the line search procedure, the gradient is additionally evaluated during the inner iterations of the TN algorithm (i.e., the Hessian computation). Thus, we expect the number of function evaluations to always be lower than the gradient evaluations. Note also that the number of gradient evaluations is more crucial to the overall computational work than the number of objective function evaluations. Here one gradient evaluation is approximately equivalent to two function evaluations when using quadratic regularization and three function evaluations in the case of TV regularization.

Moreover, we compare the quality of the optical flow estimations of the numerical algorithms. For this, we measure the average angular error (AAE) and the standard deviation

(STD) of the estimated flow w^e with respect to the ground truth w^c . For a given pixel (i, j) , the angular error (AE) between the ground truth motion vector, $w_{i,j}^c$, and the estimated flow, $w_{i,j}^e$, is computed as

$$\text{AE}(w_{i,j}^c, w_{i,j}^e) = \cos^{-1} \left(\frac{u_{i,j}^c u_{i,j}^e + v_{i,j}^c v_{i,j}^e + 1}{\sqrt{(u_{i,j}^c)^2 + (v_{i,j}^c)^2 + 1} \sqrt{(u_{i,j}^e)^2 + (v_{i,j}^e)^2 + 1}} \right).$$

The AAE is the mean of the AE over all pixels N_{np} of the image:

$$\text{AAE}(w_c, w_e) = \frac{1}{N_{np}} \sum_{i,j} \text{AE}(w_{i,j}^c, w_{i,j}^e).$$

The STD is computed as

$$\text{STD}(w_c, w_e) = \sqrt{\frac{1}{N_{np}} \sum_{i,j} \left(\text{AE}(w_{i,j}^c, w_{i,j}^e) - \text{AAE}(w_c, w_e) \right)^2}.$$

Before going through the performance evaluation of multilevel optimization algorithms, we first demonstrate the competitiveness of the adopted discretize-optimize approach (numerical optimization) versus the standard optimize-discretize approach (Gauss–Seidel) and also justify the selection of the chosen numerical optimization algorithm. To this end, we compare the CPU time needed to reach a similar AAE by the following algorithms: (1) the proposed line search two-step preconditioned truncated Newton algorithm approach of section 3 (TN1), and (2) the line search quasi-Newton limited memory BFGS (L-BFGS) approach of [32] (QN). A positive-definite approximation to the Hessian is obtained at each iteration w_k by storing the previous steps w_k and the BFGS approach. Instead of solving at each outer iteration the Newton equation, the L-BFGS method takes advantage of the fact that the obtained BFGS preconditioning matrix is easily invertible. Thus, in the L-BFGS approach line 7 of Algorithm 1 is replaced by $s_k = M_k^{-1} g_k$. The proposed algorithm TN1 is also compared to 3) a line search L-BFGS preconditioned truncated Newton method (TN2). This corresponds to the same algorithm as TN1, but the L-BFGS preconditioning approach [32] is used instead of the two-step BFGS approach [39].

It should be noted that in these three numerical optimization algorithms (TN1, TN2, and QN) the same line search approach is used [36]. The previous algorithms were compared using the linear data term and quadratic regularization, which corresponds to the classical Horn–Schunck model. Thus we also implemented a Gauss–Seidel Horn–Schunck smoother [27]. We call the latter algorithm GS.

The experimental results for the four algorithms (GS, TN1, TN2, and QN) are shown in Table 1. Best results for QN were obtained using four steps, and TN2 was set up to the same number of steps for the preconditioner. As expected from a linear relaxation scheme, GS has problems near the solution, and indeed in the tests it can reach an AAE below 8.00 in less than 1 second but will get stuck and does not reach the optimal solution. Overall, TN algorithms perform better, justifying the choice of the TN as the underlying optimization smoother within multilevel algorithms to determine direct search directions. TN1 and TN2

Table 1

CPU time needed by the approaches GS, TN1, TN2, and QN to reach similar accuracy for the Yosemite sequence using the Horn–Schunck model.

Algorithm		GS	TN1	TN2	QN
CPU time (s)	unilevel	45.22	8.90	9.36	17.20
	multiresolution	10.74	5.21	5.65	9.50
AAE		7.40	6.81	6.82	6.85

perform almost the same with a slightly better performance by TN1. This algorithm will be used in what follows as the smoother and will be denoted by Opt.

Now we report the performance evaluation of the unilevel, the multiresolution, and the multigrid optimization algorithms applied to estimate the optical flow between two successive frames of the above three sequences—translating tree, diverging tree, and Yosemite. We considered four optical flow models as described above in section 2 and subsection 5.1. We note that we used thresholding in all the data terms to remove outliers. In all experiments we considered $L = 6$ levels of resolution for multilevel algorithms (multiresolution and multigrid), and the initial guess for optical flow was zero. The stopping criteria for all algorithms were set on the relative error of the objective function, the gradient norm, or the solution norm with a tolerance of 10^{-5} . For the MR/Opt algorithm, the maximum number of outer iterations was set to 10 iterations per resolution for all the optical flow models except when using the nonlinear data term and the TV regularization, for which we used a maximum of 15 iterations. For the FMG/Opt, we performed two or three V-cycles when using the nonlinear data term, while only 1 V-cycle was sufficient in the case of the linear data term. We note here that for the purpose of a fair computational work comparison, we stopped the outer iterations once a similar accuracy on the optical flow estimation was reached. For all the algorithms, the maximum number of inner iterations within the Opt method was set to 20.

In Tables 2–4, we summarize the quality of the solution and the computational costs of the three numerical algorithms for four optical flow models. Model 1 refers to the linear data term plus the quadratic regularization; Model 2 refers to the nonlinear data term plus the quadratic regularization; Model 3 refers to the linear data term plus the TV regularization; and, finally, Model 4 refers to the nonlinear data term plus the TV regularization. The latter is the most accurate for optical flow estimation but computationally the most expensive of the four models. Indeed, Model 4 contains the nonlinear data term as well as the TV regularization, which provides accurate optical flow estimations but is numerically more challenging than the quadratic regularization.

In terms of the quality of the solution, by comparing the unilevel algorithm to the multilevel algorithms, we note that the optical flow estimation of the four models is more accurate when computed using the latter algorithms for all the tested images. In this regard, MG/Opt has been shown to provide a more accurate estimation than MR/Opt if we take the AAE as an accuracy measure. In particular, the most accurate optical flow estimations for the three sequences are obtained with Model 4 using the FMG/Opt algorithm. To illustrate the results obtained by using MG/Opt, we show first in Figure 6 the estimated motion fields of Models 2 and 4 for the translating and diverging tree sequences. In Figure 7, we plot the estimated

Table 2

Comparison of computational work and optical flow estimation for two frames of the translating tree sequence using optimization algorithms *Opt*, *MR/Opt*, and *FMG/Opt*. Time is CPU time in seconds.

		Opt	MR/Opt	FMG/Opt
Model 1	Time	11.91	2.73	0.98
	Nf	69	15	12
	Ng	572	133	37
	AAE	1.83	0.90	0.89
	STD	1.23	0.55	0.51
Model 2	Time	14.22	5.40	2.07
	Nf	89	19	25
	Ng	586	221	72
	AAE	0.27	0.23	0.23
	STD	0.29	0.22	0.22
Model 3	Time	27.78	4.90	2.40
	Nf	95	15	21
	Ng	1042	189	87
	AAE	1.54	0.77	0.75
	STD	0.86	0.47	0.42
Model 4	Time	35.32	10.25	4.08
	Nf	125	35	46
	Ng	1130	334	126
	AAE	0.20	0.20	0.20
	STD	0.18	0.18	0.17

Table 3

Comparison of computational work and optical flow estimation for two frames of the diverging tree sequence using optimization algorithms *Opt*, *MR/Opt*, and *FMG/Opt*. Time is CPU time in seconds.

		Opt	MR/Opt	FMG/Opt
Model 1	Time	14.56	4.10	0.72
	Nf	64	16	10
	Ng	695	209	27
	AAE	2.06	1.99	1.82
	STD	1.60	1.57	1.30
Model 2	Time	16.81	5.28	1.70
	Nf	82	20	16
	Ng	664	221	63
	AAE	2.08	2.07	1.90
	STD	1.95	2.09	1.85
Model 3	Time	26.59	6.61	1.06
	Nf	128	20	11
	Ng	1016	253	35
	AAE	2.49	2.46	1.59
	STD	1.86	1.95	1.17
Model 4	Time	31.60	14.73	3.26
	Nf	133	96	30
	Ng	997	492	97
	AAE	2.09	2.07	1.86
	STD	1.84	1.89	1.92

Table 4

Comparison of computational work and optical flow estimation for two frames of the Yosemite sequence using line search optimization algorithms (Opt, MR/Opt, FMG/Opt). Time is CPU time in seconds.

		Opt	MR/Opt	FMG/Opt
Model 1	Time	12.07	5.21	2.97
	Nf	13	8	11
	Ng	161	60	26
	AAE	6.79	6.78	6.77
	STD	9.25	9.31	9.49
Model 2	Time	62.20	15.59	7.06
	Nf	108	27	19
	Ng	680	167	67
	AAE	6.47	6.22	6.17
	STD	9.28	9.15	9.18
Model 3	Time	82.27	17.38	7.01
	Nf	83	19	15
	Ng	856	188	81
	AAE	6.33	6.07	6.08
	STD	9.10	8.62	8.40
Model 4	Time	205.56	25.51	14.46
	Nf	218	24	21
	Ng	1821	223	121
	AAE	5.75	5.49	5.45
	STD	8.92	8.57	7.90

motion fields of the four models for the Yosemite sequence. The results obtained with the other methods are visually similar, and for reasons of space we do not show them. We refer to Tables 2–4 for the quantitative comparison.

Next we discuss the computational expense of the FMG/Opt algorithm compared to the Opt and MR/Opt algorithms. From Tables 2–4, we remark that our algorithm significantly outperforms both optimization algorithms in terms of the CPU running time and the number of function and gradient evaluations for the three image sequences and the four optical flow models. For Model 1, while linear multigrid with standard components does not work [5, 29, 28], our FMG/Opt algorithm is around 8 times faster to run than Opt and 2.5 times faster than MR/Opt. The same almost applies for Model 2 which contains a nonlinear data term. As expected, for Model 3 which incorporates a TV regularization term, multilevel optimization is much more effective than unilevel optimization. In fact, our FMG/Opt algorithm takes 13 times less CPU time to run than the Opt algorithm. Finally, when used for Model 4, the FMG/Opt algorithm shows a similar significant improvement of a factor of 12 over Opt and more than 2 over MR/Opt.

Overall, the FMG/Opt algorithm performs at least 2 times better than the MR/Opt algorithm and 10 times better than the Opt algorithm; see Table 5. We also note that the FMG/Opt algorithm is less independent of the image size because it often takes similar numbers of function and gradient evaluations when compared across the same optical flow model.

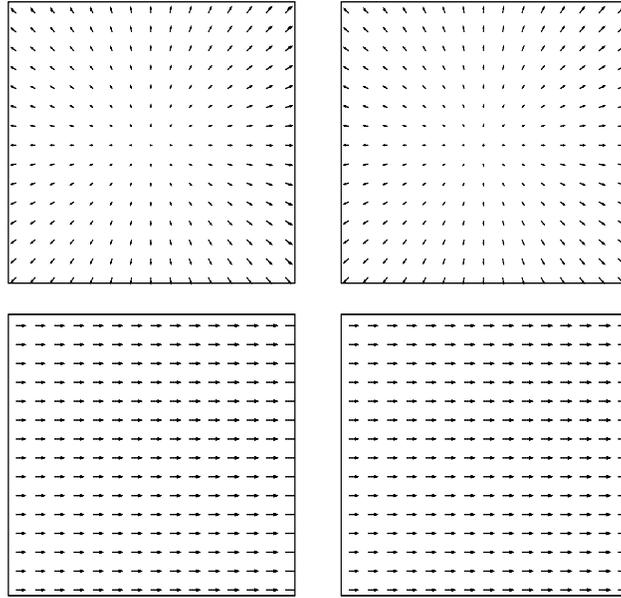


Figure 6. These plots show the motion fields obtained using the FMG/OPT algorithm for the diverging (top) and translating (bottom) tree. On the left (respectively, right) the result for Model 2 (respectively, Model 4) is shown superimposed with the ground truth. Motion vectors have been scaled by a factor of 2.5 for better visibility.

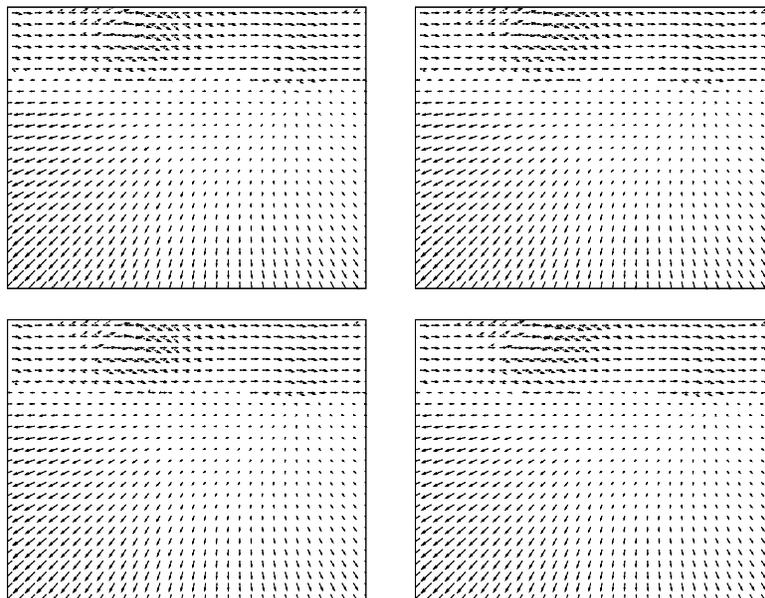


Figure 7. These plots show the motion fields obtained using the FMG/OPT algorithm for the Yosemite sequence. From left to right, and from top to bottom, Models 1–4 are shown superimposed with the ground truth. Motion vectors have been scaled by a factor of 2.5 for better visibility.

Table 5

Global characteristics of *Opt*, *MR/Opt*, and *FMG/Opt* for optical flow models on all the three images. *Nfg* is the total number of function and gradient evaluations.

		Opt	MR/Opt	FMG/Opt
Model 1	Total time (s)	38.54	12.04	4.76
	Total Nfg	1501	421	106
Model 2	Total time (s)	93.63	26.27	10.83
	Total Nfg	2069	642	232
Model 3	Total time (s)	136.64	28.89	10.47
	Total Nfg	3016	648	219
Model 4	Total time (s)	272.48	50.49	21.80
	Total Nfg	3117	1101	376
All models	Total time (s)	541.29	117.69	47.86
	Total Nfg	9703	2812	933

7. Conclusion. Based on the discretize-optimize approach, we have applied different numerical optimization techniques to variational models for optical flow computation. First, we have shown the competitiveness of this strategy compared to the classical optimize-discretize approach. Three Newton-based optimization algorithms were superior to the Gauss–Seidel method when applied to the classical Horn–Schunck model. In particular, TN was shown to be a suitable unilevel optimization algorithm and was chosen as the smoother for optimization-based multilevel methods. We then implemented the FMG/Opt algorithm based on a line search strategy to scale the (direct) Newton or the (recursive) multigrid search direction. Several components of the MG/Opt technique have been tuned for high efficiency, and the algorithm has been fully evaluated with respect to unilevel and (one-way) MR/OPT. Our experimental results have demonstrated that the FMG/Opt algorithm can be effectively used for optical flow computation. Using different models and images, we have observed that the FMG/Opt algorithm was faster and more accurate than both unilevel and multiresolution TN. Further research will investigate the use of line search multigrid versus trust region multigrid in the context of dense optical flow computation. The proposed numerical strategy can be adapted to the minimization of other nonlinear energy functionals such as the illumination-invariant model proposed in [16] or depth estimation in stereo problems.

REFERENCES

- [1] L. ÁLVAREZ, J. WEICKERT, AND J. SÁNCHEZ, *A scale-space approach to nonlocal optical flow calculations*, in *Scale-Space Theories in Computer Vision*, Lecture Notes in Comput. Sci. 1682, Springer-Verlag, Berlin, Heidelberg, 1999, pp. 235–246.
- [2] L. ÁLVAREZ, J. WEICKERT, AND J. SÁNCHEZ, *Reliable estimation of dense optical flow fields with large displacements*, *Int. J. Comput. Vision*, 39 (2000), pp. 41–56.
- [3] P. ANANDAN, *A computational framework and an algorithm for the measurement of visual motion*, *Int. J. Comput. Vision*, 2 (1989), pp. 283–310.
- [4] J. BARRON, D. FLEET, AND S. BEAUCHEMIN, *Performance of optical flow techniques*, *Int. J. Comput. Vision*, 12 (1994), pp. 43–77.
- [5] R. BATTITI, E. AMALDI, AND C. KOCH, *Computing optical flow across multiple scales: An adaptive coarse-to-fine strategy*, *Int. J. Comput. Vision*, 6 (1991), pp. 133–145.
- [6] M. BERTERO, T. A. POGGIO, AND V. TORRE, *Ill-posed problems in early vision*, *Proc. IEEE*, 76 (1988),

- pp. 869–889.
- [7] F. A. BORNEMANN AND R. KRAUSE, *Classical and cascadic multigrid—a methodological comparison*, in Domain Decomposition Methods in Sciences and Engineering, P. Bjorstad, M. Espedal, and D. Keyes, eds., DD Press, Bergen, Norway, 1998, pp. 64–71.
 - [8] A. BORZI, K. ITO, AND K. KUNISCH, *Optimal control formulation for determining optical flow*, SIAM J. Sci. Comput., 24 (2002), pp. 818–847.
 - [9] A. BRANDT, *Multi-level adaptive solutions to boundary-value problems*, Math. Comp., 31 (1977), pp. 333–390.
 - [10] W. L. BRIGGS, V. E. HENSON, AND S. F. MCCORMICK, *A Multigrid Tutorial*, 2nd ed., SIAM, Philadelphia, 2000.
 - [11] T. BROX, A. BRUHN, N. PAPPENBERG, AND J. WEICKERT, *High accuracy optical flow estimation based on a theory for warping*, in Proceedings of the European Conference on Computer Vision, Vol. 4, 2004, pp. 25–36.
 - [12] A. BRUHN, *Variational Optic Flow Computation: Accurate Modeling and Efficient Numerics*, Ph.D. thesis, Department of Mathematics and Computer Science, Saarland University, Saarbrücken, Germany, 2006.
 - [13] A. BRUHN, J. WEICKERT, C. FEDDERN, T. KOHLBERGER, AND C. SCHNORR, *Real-time optic flow computation with variational methods*, in Proceedings of the International Conference on Computer Analysis of Images and Patterns, Lecture Notes in Comput. Sci. 2756, Springer, Berlin, 2003, pp. 222–229.
 - [14] A. BRUHN, J. WEICKERT, C. FEDDERN, T. KOHLBERGER, AND C. SCHNORR, *Variational optical flow computation in real time*, IEEE Trans. Image Process., 14 (2005), pp. 608–615.
 - [15] P. Y. BURGI, *Motion estimation based on the direction of intensity gradient*, Image Vision Comput., 22 (2004), pp. 637–653.
 - [16] V. CASELLES, L. GARRIDO, AND L. IGUAL, *A contrast invariant approach to motion estimation*, in Scale Space and PDE Methods in Computer Vision, Lecture Notes in Comput. Sci. 3459, Springer-Verlag, Berlin, Heidelberg, 2005, pp. 242–253.
 - [17] I. COHEN AND I. HERLIN, *Non uniform multiresolution method for optical flow and phase portrait models: Environmental applications*, Int. J. Comput. Vision, 33 (1999), pp. 29–49.
 - [18] R. S. DEMBO AND T. STEIHAUG, *Truncated Newton algorithms for large scale unconstrained optimization*, Math. Programming, 26 (1983), pp. 190–212.
 - [19] M. DROSKE AND M. RUMPF, *A variational approach to nonrigid morphological image registration*, SIAM J. Appl. Math., 64 (2004), pp. 668–687.
 - [20] W. ENKELMANN, *Investigations of multigrid algorithms for the estimation of optical flow fields in image sequences*, Comput. Vision Graph. Image Process., 43 (1988), pp. 150–177.
 - [21] H. FARID AND E. P. SIMONCELLI, *Differentiation of discrete multidimensional signals*, IEEE Trans. Image Process., 13 (2004), pp. 496–508.
 - [22] S. GHOSAL AND P. A. VANEK, *Fast scalable algorithm for discontinuous optical flow estimation*, IEEE Trans. Pattern Anal. Mach. Intell., 18 (1996), pp. 181–194.
 - [23] F. GLAZER, *Multilevel relaxation in low-level computer vision*, in Multiresolution Image Processing and Analysis, A. Rosenfeld, ed., Springer-Verlag, New York, 1984, pp. 312–330.
 - [24] S. GRATTON, A. SARTENAER, AND P. L. TOINT, *Recursive trust-region method for multiscale nonlinear optimization*, SIAM J. Optim., 19 (2008), pp. 414–444.
 - [25] W. HACKBUSCH, *Multi-grid Methods and Applications*, Springer, Berlin, New York, 1985.
 - [26] W. HINTERBERGER, O. SCHERZER, C. SCHNÖRR, AND J. WEICKERT, *Analysis of optical flow models in the framework of the calculus of variations*, Numer. Funct. Anal. Optim., 23 (2002), pp. 69–89.
 - [27] B. HORN AND B. SCHUNK, *Determining optical flow*, Artificial Intelligence, 17 (1981), pp. 185–203.
 - [28] E. M. KALMOUN, H. KÖSTLER, AND U. RÜDE, *3D optical flow computation using a parallel variational multigrid scheme with application to cardiac C-arm CT motion*, Image Vision Comput., 25 (2007), pp. 1482–1494.
 - [29] E. M. KALMOUN AND U. RÜDE, *A variational multigrid for computing the optical flow*, in Proceedings of the Vision, Modeling, and Visualization Conference, Akademische Verlagsgesellschaft, Berlin, 2003, pp. 577–584.
 - [30] R. G. KEYS, *Cubic convolution interpolation for digital image processing*, IEEE Trans. Acoust. Speech

- Signal Process., 29 (1981), pp. 1153–1160.
- [31] R. M. LEWIS AND S. G. NASH, *Model problems for the multigrid optimization of systems governed by differential equations*, SIAM J. Sci. Comput., 26 (2005), pp. 1811–1837.
- [32] D. C. LIU AND J. NOCEDAL, *On the limited memory BFGS method for large scale optimization*, Math. Programming, 45 (1989), pp. 503–528.
- [33] B. D. LUCAS AND T. KANADE, *An iterative image registration technique with an application to stereo vision*, in Proceedings of the DARPA Image Understanding Workshop, 1981, pp. 121–130.
- [34] E. MEMIN AND P. PEREZ, *Dense estimation and object-based segmentation of the optical flow with robust techniques*, IEEE Trans. Image Process., 7 (1998), pp. 703–719.
- [35] E. MEMIN AND P. PEREZ, *A multigrid approach for hierarchical motion estimation*, in Proceedings of the Sixth International Conference on Computer Vision, IEEE Computer Society, Washington, DC, 1998, pp. 933–938.
- [36] J. J. MORÉ AND D. J. THUENTE, *Line search algorithms with guaranteed sufficient decrease*, ACM Trans. Math. Software, 20 (1994), pp. 286–307.
- [37] H.-H. NAGEL, *Constraints for the estimation of displacement vector fields from image sequences*, in Proceedings of the Eighth International Joint Conference on Artificial Intelligence, Karlsruhe, Germany, A. Bundy, ed., Morgan Kaufmann, San Francisco, 1983, pp. 945–951.
- [38] H.-H. NAGEL AND W. ENKELMANN, *An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences*, IEEE Trans. Pattern Anal. Mach. Intell., 8 (1986), pp. 565–593.
- [39] S. G. NASH, *Preconditioning of truncated-Newton methods*, SIAM J. Sci. Statist. Comput., 6 (1985), pp. 599–616.
- [40] S. G. NASH, *A multigrid approach to discretized optimization problems*, Optim. Methods Softw., 14 (2000), pp. 99–116.
- [41] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer-Verlag, New York, 1999.
- [42] J. M. ODOBEZ AND P. BOUTHEMY, *Robust multiresolution estimation of parametric motion models*, J. Vis. Commun. Image Rep., 6 (1995), pp. 348–365.
- [43] N. PAPENBERG, A. BRUHN, T. BROX, S. DIDAS, AND J. WEICKERT, *Highly accurate optic flow computation with theoretically justified warping*, Int. J. Comput. Vision, 67 (2006), pp. 141–158.
- [44] C. STILLER AND J. KONRAD, *Estimating motion in image sequences*, IEEE Signal Processing Mag., 16 (1999), pp. 70–91.
- [45] R. SZELISKI AND H. Y. SHUM, *Motion estimation with quadtree splines*, IEEE Trans. Pattern Anal. Mach. Intell., 18 (1996), pp. 1199–1210.
- [46] D. TERZOPOULOS, *Image analysis using multigrid relaxation methods*, IEEE Trans. Pattern Anal. Mach. Intell., 8 (1986), pp. 129–139.
- [47] J. WEICKERT AND C. SCHNORR, *A theoretical framework for convex regularizers in PDE-based computation of image motion*, Int. J. Comput. Vision, 45 (2001), pp. 245–264.
- [48] Z. WEN AND D. GOLDFARB, *A Line Search Multigrid Method for Large-Scale Convex Optimization*, Tech. report, Department of IEOR, Columbia University, 2007.
- [49] D. XIE AND T. SCHLICK, *Efficient implementation of the truncated-Newton algorithm for large-scale chemistry applications*, SIAM J. Optim., 10 (1999), pp. 132–154.