Dense Estimation of Fluid Flows

Thomas Corpetti, Étienne Mémin, and Patrick Pérez

Abstract—In this paper, we address the problem of estimating and analyzing the motion of fluids in image sequences. Due to the great deal of spatial and temporal distortions that intensity patterns exhibit in images of fluids, the standard techniques from Computer Vision, originally designed for quasi-rigid motions with stable salient features, are not well adapted in this context. We thus investigate a dedicated minimization-based motion estimator. The cost function to be minimized includes a novel data term relying on an integrated version of the continuity equation of fluid mechanics, which is compatible with large displacements. This term is associated with an original second-order div-curl regularization which prevents the washing out of the salient vorticity and divergence structures. The performance of the resulting fluid flow estimator is demonstrated on meteorological satellite images. In addition, we show how the sequences of dense motion fields we estimate can be reliably used to reconstruct trajectories and to extract the regions of high vorticity and divergence.

Index Terms—Fluid motion, continuity equation, div-curl regularization, nonconvex minimization, trajectories, vorticity, and divergence concentration.

✦

1 INTRODUCTION

IN a number of scientific domains, the analysis of image sequences involving fluids is of the highest importance. In environmental sciences such as oceanography, meteorology, and climatology, the sequences of satellite images are one of the main sources of information. Their analysis allows, for example, the tracking of fish eggs, pollutants, or ice floes in oceanography [12], [40] and of depressions or convective cells in meteorology [30], [36], [37]. In medicine, tomography sequences are used to monitor blood flow in which vortices can indicate pathologies [29], [35]. In the field of fluid mechanics, experimental imaging techniques now routinely produce various types of videos which constitute a unique source of information for both applied and theoretical studies [1], [43], [46].

In all these domains, cameras provide in a versatile and nonintrusive way huge amounts of spatio-temporal data, as opposed to *in situ* measurement techniques (e.g., physical drifter in oceanography, mechanical anemometry in meteorology, thermal anemometry in fluid mechanics), which are complex, expensive, and supply only sparse information. Unfortunately, unlike dedicated probes, images only give indirect access to the physical quantities of interest. With videos or image series, one must face the complex task of extracting these quantities from the intensity patterns recorded in the successive images. As compared to more classical motion analysis problems addressed by the computer vision community, the analysis of fluid motion from images is particularly challenging due to the great deal of spatial and temporal distortions that the intensity patterns exhibit.

In this specific context, a number of methods have been proposed to measure *dense displacement fields* [4], [7], [11], [13], [14], [15], [18], [21], [28], [33], [36], [40], [47], [48]. The obtained displacement fields already provide valuable information in their own right. They are, for instance, intended to capture the winds in meteorology applications [36] or they can be used directly to feed the numerical models of weather prediction [36], ocean circulation [11], or fluid mechanics [46]. These motion fields are also useful in deriving other entities of interest. The numerical derivation of these fields provides estimates of the divergence and vorticity [43], whereas numerical integration provides stream functions, velocity potentials [39], [40], and point trajectories [40]. Based on the estimated motion fields, the so-called critical points of the flow can also be recovered and classified [11], [16], [17], [18], [33], [35]. Whatever the domain of application, these points are important to understand, predict, and compactly represent the flows of interest.

Apart from a few cases to be discussed later, the techniques of dense motion estimation used in all these studies are similar to those classically used in computer vision, where the motions under consideration are quasi-rigid and stable salient features are present in the scene. Unfortunately, there is no such thing as an "object" in most imaged fluid flows. There are mainly deformable and transient brightness patterns, as illustrated in Fig. 1 with a sequence of meteorological satellite images. On such sequences, techniques based on standard computer vision ingredients are thus intrinsically limited. The design of alternative approaches dedicated to fluid motion thus remains a widely open domain of research. Our work is a contribution in that direction.

2 BACKGROUND AND CONTRIBUTION

2.1 Standard Dense Motion Estimation and Its Limitations

The most accurate techniques to address the generic problem of estimating the apparent motion from image sequences are based on the seminal work of Horn and
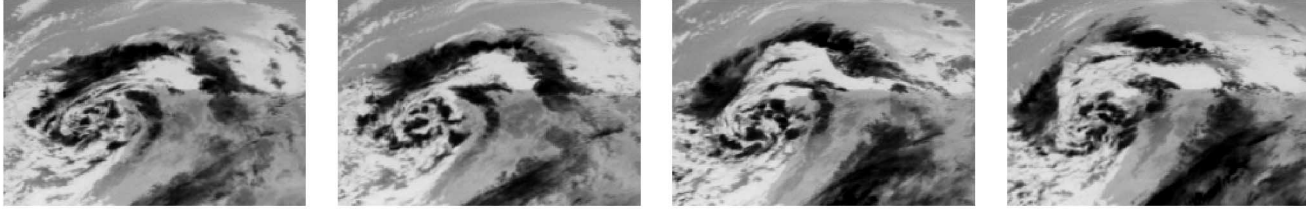

- *T. Corpetti and E. Mémin are with IRISA/Université de Rennes I, 35042 Rennes Cedex, France. E-mail: {tcorpett, memin}@irisa.fr.*
- *P. Pérez is with Microsoft Research, Cambridge, CB2 3NH, United Kingdom. E-mail: pperez@microsoft.com.*

Fig. 1. **Example of complex imaged fluid flow** sequence of four infrared Meteosat images showing a spiraling depression.

Schunck [23]. See for instance [2], [9], [11], [25], [26], [27], [31], [44], [45] for the description of such state-of-the-art techniques and [5], [6] for comprehensive comparisons with completely different approaches.

These techniques are based on the minimization of a global cost function of the form:

$$\int_{\Omega} f_1\left[\nabla E(\boldsymbol{x},t) \cdot \boldsymbol{v}(\boldsymbol{x},t) + \frac{\partial E(x,t)}{\partial t}\right] + \alpha f_2[|\nabla \boldsymbol{v}(\boldsymbol{x},t)|], \quad (1)$$

where $\boldsymbol{v}(\boldsymbol{x},t)$ is the unknown velocity at time t and location $\boldsymbol{x} = (x,y)$ in the image plane Ω, $E(\boldsymbol{x},t)$ is the image brightness, α is a positive parameter, and f_1 and f_2 are two penalty functions. The first part of this cost function relies on the assumption that the visible points conserve roughly their intensity in the course of a small displacement, that is:

$$\frac{dE}{dt} = \nabla E \cdot \boldsymbol{v} + \frac{\partial E}{\partial t} \approx 0. \quad (2)$$

The associated penalty function f_1 is often the L_2 norm. However, better estimates are usually obtained by choosing a "softer" penalty function [9], [31]. Such functions, arising from robust statistics [24], limit the impact of the many locations where the brightness constancy assumption does not hold, such as on occlusion boundaries.

The second term of the cost function (1) favors smooth solutions. It leads to a well-posed optimization problem, while incorporating a sensible prior on the velocity to be recovered. As with the penalty function in the data term, the penalty function f_2 was taken as a quadratic in early studies, but a softer penalty is now preferred in order not to smooth out the natural discontinuities of the velocity field [9], [11], [26], [31].
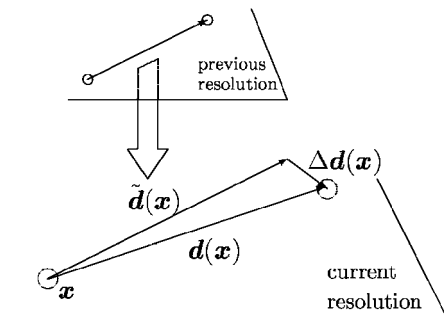

Fig. 2. **Incremental multiresolution setting** the displacement estimated at a previous image resolution is refined by estimating a displacement increment at the current resolution.

Based on (1), the estimation of motion presents a major flaw: Due to its differential nature, the expression (2) of the brightness conservation is no longer usable when large displacements occur between consecutive frames of the sequence. To circumvent this problem, the brightness conservation must be expressed in an integrated way:

$$E(\boldsymbol{x} + \boldsymbol{d}(\boldsymbol{x}), t + \Delta t) - E(\boldsymbol{x}, t) \approx 0, \quad (3)$$

where Δt is the temporal sampling rate and $\boldsymbol{d}(\boldsymbol{x})$ is the *displacement* from time t to $t + \Delta t$ of the point located at position $\boldsymbol{x}$ at time t. This new form is highly nonlinear in the unknown vector field. As a consequence, almost all studies resort to a succession of linearizations embedded within a multiresolution scheme. Given a previous estimate $\tilde{\boldsymbol{d}}$ of the displacement field obtained at a coarser resolution, a first-order expansion of the first term in (3) is performed around $(\boldsymbol{x} + \tilde{\boldsymbol{d}}(\boldsymbol{x}), t + \Delta t)$, yielding the following cost function:

$$\int_{\Omega} f_1\left[\nabla E(\boldsymbol{x} + \tilde{\boldsymbol{d}}(\boldsymbol{x}), t + \Delta t) \cdot \Delta \boldsymbol{d}(\boldsymbol{x}) + E(\boldsymbol{x} + \tilde{\boldsymbol{d}}(\boldsymbol{x}), \\ t + \Delta t) - E(\boldsymbol{x}, t)\right] + \alpha f_2[|\nabla(\tilde{\boldsymbol{d}} + \Delta \boldsymbol{d})(\boldsymbol{x})|] \quad (4)$$

to be minimized with respect to the displacement increment $\Delta \boldsymbol{d} = \boldsymbol{d} - \tilde{\boldsymbol{d}}$ (Fig. 2).

Although motion estimators based on a cost function of the type in (4) have proven useful for fluids [4], [11], [28], [33], [43], the two basic assumptions—brightness constancy and first-order smoothness—seem inappropriate for the complex deformable motions encountered in this context. First, the deformable nature of the fluid motions, the complexity of the imaging processes, and the possible variations of temperature and pressure of the moving fluid all contribute to variations of the observed brightness for a given element of fluid.

In addition to these breakdowns of the brightness constancy assumption, optical-flow estimators equipped with a first-order regularization, as in (4), have difficulty recovering most complex fluid motions. Even with a soft penalty function f_2, a smoothing of this type results in estimated motion fields with low divergence and vorticity —and these two quantities are known to be large in a number of locations for most imaged fluid flows. This over-smoothing is all the more problematic since it occurs usually around the critical points, which are of primary importance to describe and understand the whole flow.

2.2 Related Works and Contribution

The classical approaches based on a cost function of the type in (4) see the weakness of their two underlying assumptions dramatically exaggerated in the peculiar case of fluids. We propose here a new minimization-based

technique where both assumptions and, consequently, both parts of the cost function are changed to better suit the specificities of fluid images.

Following several authors, we first suggest to base the data term on a *continuity equation* applied to image brightness. This idea was originally suggested by Schunck [38] for general motion estimation. Since then, it has been acknowledged as an especially appealing alternative to the brightness constancy assumption in the context of moving fluids [3], [7], [13], [14], [15], [47], [48]. The recorded image intensity is indeed related, to a degree of complexity that depends on the type of imagery, either to the density of the fluid, or to the concentration of a passive marker transported by it. The rationale is then that the continuity equation, which these physical quantities obey for the tridimensional velocity field (at least in absence of matter/ marker creation or disappearance), can be applied to the image brightness itself for the apparent two-dimensional velocity. Although more physically-grounded, this remains an *ad-hoc* constraint on the image brightness (except in specific cases [14], [47]) since the precise relationship between the brightness and the observed physical quantities is beyond reach in most cases, including satellite images. Due to this limitation, we will adopt a robust formulation, in contrast to all other studies making use of the continuity equation.

As we shall see, the resulting brightness constraint is similar to (2), relying on the apparent velocity. It is thus not directly applicable to large displacements. We address its limitation by deriving an integrated continuity equation, resulting in the novel term that is presented in Section 3.

Beside this new data term, we investigate a new regularization approach that is presented in Section 4. In order to better preserve the concentrations of both divergence and vorticity, which are smoothed out by a standard regularization, we resort to an original second-order "div-curl" regularization. The "div-curl" regularization has been used by Suter [41] and Gupta and Prince [20] for the estimation of motion. We depart from these earlier studies in two ways. First we turn the second-order div-curl regularization into two interleaved first-order regularizations based on two auxiliary scalar fields which approximate the divergence and the vorticity of the flow. This reduces the complexity of the associated minimization problem. From a modeling point of view, this also offers the opportunity to incorporate physical priors directly on the divergence and the vorticity, as well as to assimilate physical measurements of these quantities. Second, we make use of soft penalty functions in this regularizer for an even better preservation of the divergence and vorticity structures of the flow.

The problem of minimizing the resulting nonconvex cost function is addressed in Section 5. Comparative results of motion field estimation are reported and discussed in Section 6. Finally, based on two different uses of dense fluid motion fields, namely the reconstruction of long individual trajectories and the extraction of vorticity/divergence structures, we demonstrate in Section 7 that the motion fields provided by our approach constitute reliable inputs for further computations of real interest.

3 INTEGRATED CONTINUITY EQUATION

In most image sequences, the classic assumption of brightness constancy along the motion trajectories is known to be violated in a number of locations in the image plane. For example, brightness constancy does not apply at occluding contours, where pieces of background appear or disappear at each time step. When an explicit modeling of pathological situations is not possible, their impact on the rest of the motion estimation can be limited by using nonquadratic penalty functions stemming from robust statistics.

In fluid imagery, the brightness constancy assumption is, unfortunately, even less appropriate. Indeed, image sequences involving fluids often exhibit dramatic temporal changes of brightness. A number of factors, depending on the imaging technology and on the nature of the flow, can be responsible for such a behavior. First of all, imaged flows, like those in meteorology, might be *compressible*. In this case, the fluids experience changes of volume during displacement, which is incompatible with the brightness constancy assumption.

Often, the recorded irradiance either corresponds to a slice of the flow (like in particle velocimetry experiments using a laser sheet, [1], [46]) or results from an integration process over a certain layer of the flow (like in tomography [13], [29], angiography [35] or transmittance imagery [14], [47]). In the former case, tridimensional motions nonparallel to the visualization plane cause fluid elements to enter or exit the imaged slice. In the latter case, the integration process might result in apparent displacements for which the brightness constancy assumption is not valid even if there is an actual Lagrangian invariant—i.e., a certain physical quantity that is conserved along trajectories—in the tridimensional flow.

Finally, the recorded brightness is often a function of the thermodynamic conditions (temperature and pressure). Hence, the brightness of a given element of fluid changes whenever these conditions change along the trajectory. This is true in infrared satellite imagery, for example, where the image brightness is a direct function of the temperature [48] and, thus, depends, among other factors, on the altitude of clouds visible in the image, as well as on the temperature of the ground or sea beneath.

Based on these considerations, an alternative brightness constraint, which would be better suited to the physics of fluids, has been sought. A few authors have suggested the use of the continuity equation, as a more physically-grounded constraint, and they have demonstrated that it is indeed an appealing alternative [3], [7], [13], [14], [15], [47], [48]. It turns out that the image irradiance for a fluid is usually related to the density of a physical quantity—e.g., mass in transmittance imagery [14], [47], dye, smoke, or particle concentration in particle image velocimetry [43], heat in infrared imagery—transported by the flow under a global conservation constraint. This density ρ obeys the continuity equation:

$$\frac{\partial \rho}{\partial t} + \mathrm{div}(\rho \mathbf{V}) = 0, \tag{5}$$

where $\mathbf{V}$ is the three-dimensional velocity field. This equation is derived from the global conservation assumption by stating that the temporal variation of the quantity

under consideration within an infinitesimal volume amounts exactly to the flux of this quantity through the boundary surface of the volume.

One can then assume by analogy that the two-dimensional image brightness E and apparent velocity v satisfy:

$$\frac{\partial E}{\partial t} + \text{div}(E\boldsymbol{v}) = 0. \qquad (6)$$

For incompressible fluids such as water, the three-dimensional flow is divergence free. Assuming the resulting apparent bidimensional flow is divergence free as well, the bidimensional continuity equation above amounts exactly to the brightness constancy constraint (2), since $\text{div}(E\boldsymbol{v}) = \boldsymbol{v} \cdot \nabla E + E\text{div}\boldsymbol{v}$. In other cases, i.e., when flows are compressible such as in meteorological satellite imagery, the brightness constraint expressed by (6) differs from the standard one, (2), by the additional term $E\text{div}\boldsymbol{v}$.

The actual validity of this new brightness constraint has been formally demonstrated by Fitzpatrick [14] in the particular case of transmittance images with certain known boundary conditions. In the same context, it was recently generalized by Wildes et al. [47], under the simple assumption of null normal flow on boundaries.

In other contexts, such as the one of atmospheric satellite images on which our experiments are conducted, both the image formation model and the boundary conditions are insufficiently accessible (the boundaries of the image have nothing to do with the boundaries of the fluid) to obtain such a strong result. Consequently, even if the constraint has experimentally proven appealing in previous studies [7], [48], only qualitative justifications can be provided.

In the case of the Meteosat images on which our experiments are conducted, simple models of atmospheric mechanics and thermodynamics can be mobilized in favor of the continuity equation. For images in both infrared and water vapor channels, the image irradiance can be converted linearly into brightness temperature based on parameters reactualized every three months by the European consortium Eumetsat in charged of the Meteosat images. Considering a homogeneous temperature profile of the troposphere and using a linearization of the so-called Laplace formula [42], the temperature can finally be converted linearly into pressure P, making the original brightness approximately linear in the pressure.

Based on the hypothesis of atmospheric static equilibrium, the vertical component of pressure forces and the weight cancel out, yielding:

$$\frac{\partial P}{\partial z} = -g\rho, \qquad (7)$$

where z denotes the altitude, g the gravity, assumed constant, and ρ the mass density. Hence, the image brightness is approximately proportional to the vertical integral of the mass density:

$$E(x, y, t) \propto \int_{z_0}^{z_1} \rho(x, y, z, t)dz. \qquad (8)$$

This image formation model is similar to the one used in transmittance imagery. Following Fitzpatrick [14], the integration of the continuity equation (5) with respect to z and the use of (8) finally provide:

$$\frac{\partial E}{\partial t} + \text{div}(E\boldsymbol{v}) + [E\boldsymbol{n}_z \cdot \mathbf{V}]_{z_0}^{z_1} = 0, \qquad (9)$$

with the apparent bidimensional velocity v defined as

$$\boldsymbol{v} = \frac{\int_{z_0}^{z_1} \rho \mathbf{V}_h}{\int_{z_0}^{z_1} \rho},$$

with $\mathbf{V}_h$ the vector made up with the two first components of $\mathbf{V}$ and $\boldsymbol{n}_z^T = [0\ 0\ 1]$.

Making the assumption, common in meteorology, that the vertical component $\boldsymbol{n}_z \cdot \mathbf{V}$ of the instantaneous atmospheric velocity is negligible, we end up with the image brightness respecting a bidimensional continuity equation for the bidimensional "apparent" atmospheric velocity v.

For flows with nonzero divergence in general and those present in infrared and water vapor Meteosat images in particular, we thus believe that the constraint (6) is more principled than the brightness constancy assumption (2). The comparative results we shall provide on the latter type of images support this belief.

In contrast to previous studies, however, we think that the direct use of the constraint in (6) is plagued in exactly the same way as the classical constraint (2)—i.e., by its differential nature—and is not suitable to estimate large displacements. This is particularly problematic when the imaged flow is very rapid (like in a number of fluid experiments), or when the temporal sampling rate is very low (as with satellite images). To circumvent this limitation, we need to derive an integrated version of the constraint, in the same way as displacement-based brightness constancy (3) is derived from the velocity-based version (2).

To this end we first rewrite (6) using the identity $\text{div}(E\boldsymbol{v}) = E\text{div}\boldsymbol{v} + \nabla E \cdot \boldsymbol{v}$ and the definition $\frac{dE}{dt} = \frac{\partial E}{\partial t} + \nabla E \cdot \boldsymbol{v}$ of the total derivative:

$$\frac{dE}{dt} + E\text{div}\boldsymbol{v} = 0. \qquad (10)$$

Now, consider an infinitesimal imaged fluid element with trajectory $\{\boldsymbol{x}(u), u \in (t, t + \Delta t)\}$ joining location $\boldsymbol{x}$ at time t to location $\boldsymbol{x} + \boldsymbol{d}(\boldsymbol{x})$ at time $t + \Delta t$. We assume that its velocity is constant along this piece of trajectory, implying that $\boldsymbol{d}(\boldsymbol{x}) = \Delta t \boldsymbol{v}(\boldsymbol{x}, t)$ and

$$\text{div}\boldsymbol{v}(\boldsymbol{x}(u), u) = \Delta t^{-1}\text{div}\boldsymbol{d}(\boldsymbol{x}), \ \forall u \in (t, t + \Delta t).$$

The ordinary differential equation (10) is then readily integrated:

$$\left.\frac{dE(\boldsymbol{x}(t), t)}{dt}\right|_{t=u} = -E(\boldsymbol{x}(u), u)\text{div}\boldsymbol{v}(\boldsymbol{x}(u), u), \qquad (11)$$

$$\forall u \in (t, t + \Delta t)$$

$$\Rightarrow E(\boldsymbol{x} + \boldsymbol{d}(\boldsymbol{x}), t + \Delta t) = E(\boldsymbol{x}, t)\exp(-\text{div}\boldsymbol{d}(\boldsymbol{x})). \qquad (12)$$

According to this constraint, the brightness is scaled by the factor $\exp(-\text{div}\boldsymbol{d}(\boldsymbol{x}))$. It decreases (respectively, increases) for motions with positive (respectively, negative) divergence. When the divergence is zero, this constraint amounts exactly to the brightness constancy constraint (3).
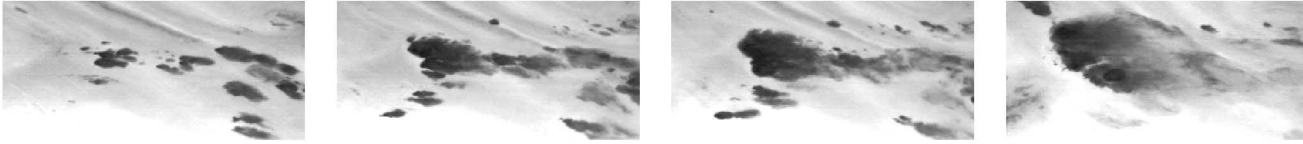
Fig. 3. **Example of divergence concentrations** in these four consecutive water vapor Meteosat images, the dark spots correspond to exploding convective cells whose motion in the image plane is highly divergent.

This integrated continuity equation has the advantage of dealing explicitly with displacements instead of velocities, but, as a consequence, the dependency on the unknown vector field d is now highly nonlinear. As with standard incremental motion estimation techniques described in Section 2.1, we cope with this problem by using a coarse-to-fine succession of linearizations. Given a previous estimate $\tilde{d}$ of the displacement field obtained at a coarser resolution, a first-order expansion of the left-hand-side of (12) is performed around $(x + \tilde{d}(x), t + \Delta t)$. Dropping the time indices of the intensity function for sake of clarity, we end up with the following data term for our dedicated cost function:

$$F_1(\Delta d) = \int_\Omega f_1 \left[\exp(\mathrm{div}\tilde{d}(x)) \big((\tilde{E}(x) \nabla \mathrm{div}\tilde{d}(x) + \nabla \tilde{E}(x)) \cdot \triangle d(x) + \tilde{E}(x) \big) - E(x) \right] dx, \quad (13)$$

where $\triangle d = d - \tilde{d}$ and we introduce the compact notation $\tilde{E}(x) = E(x + \tilde{d}(x), t + \Delta t)$.

Although the continuity equation seems more appropriate than the brightness constancy constraint in the context of compressible fluids, it remains an ad-hoc choice. As we previously discussed, no one has given to date a rigorous proof of its validity, except in the specific case of transmittance imagery [14], [47]. Moreover, the underlying assumption of global conservation of some physical quantity, such as mass, is likely to be locally violated for physical reasons (presence of actual sources or sinks in the flow), as well as for image-based reasons. Concerning the latter issue, three-dimensional motions nonparallel to the image plane, such as convective motions of clouds, can result in the appearance or disappearance of matter in the imaged volume. Similarly, the motion at the boundary of the image plane is not parallel to this boundary in general. This again results in the entrance and the exit of fluid elements in the imaged area.

For all these reasons, we would like a soft penalty function in (13), which will limit the impact of the various violations of the continuity equation on the overall motion estimate.

A whole range of such soft penalty functions have been proposed in the literature (see for instance [10]). Some are convex, which is convenient from a minimization point of view. However, we have found that a bounded penalty is more effective. We chose the Leclerc function defined as $f_1(x) = 1 - \exp(-\tau_1 x^2)$, where τ_1 is a positive parameter. The hierarchical minimization scheme we describe in Section 5 proven effective in dealing with the resulting nonconvexity.

We now turn to the definition of the smoothness prior to be used in conjunction with (13).

4 ADAPTED DIV-CURL REGULARIZATION

The first-order regularization used in the standard cost function (4) penalizes the spatial variations of the displacement field $d = (u, v)$. Such a regularization thus results in estimated vector fields with rather small divergence and vorticity. This can be made more visible when the penalty function f_2 is quadratic by noticing that the Euler-Lagrange equations associated with

$$\int_\Omega |\nabla u|^2 + |\nabla v|^2 \quad (14)$$

are the same as those associated with the so-called first-order div-curl regularizer [41]

$$\int_\Omega \mathrm{div}^2 d + \mathrm{curl}^2 d, \quad (15)$$

where $\mathrm{div}\, d = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y}$ and $\mathrm{curl}\, d = \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}$ are, respectively, the divergence and the vorticity of the bidimensional vector field d. From a minimization point of view, the two first-order regularizers, (14) and (15), are thus equivalent. This equivalence does not hold any longer for a softer penalty function f_2. But, it is easy to see that the discontinuities authorized by the use of such a penalty function in (4) will not remove the global bias towards low divergence and vorticity.

For fluid motion estimation, such a first-order smoothing does not seem appropriate since the flows of interest are expected to exhibit large "concentrations" of vorticity and divergence. In fluid mechanics experiments, for example, turbulent flows develop vortices which are interacting rotational structures characterized by large vorticity. In meteorological satellite images, the spiraling motion of depressions (Fig. 1) possesses both substantial vorticity and divergence. Important divergent bidimensional motions are also associated to tridimensional motions that are nonparallel to the image plane, such as those shown in Fig. 3.

The key information present in these concentrations of divergence and vorticity will then be partly ignored by standard motion estimation techniques based on a first-order regularization. Besides, in the prospect of using the data cost term introduced in the previous section, which makes explicit use of the divergence, an underestimation of the divergence field would be critical. It would neutralize the advantage of the new data term by making it tend to the standard one based on the brightness constancy constraint.

Hence, we propose to use a second-order regularization. Since the divergence and the vorticity of the flow are much more physically meaningful than the spatial gradient, the second-order div-curl regularizer introduced by Suter [41]

$$\int_\Omega |\nabla \mathrm{div}\boldsymbol{d}|^2 + |\nabla\mathrm{curl}\boldsymbol{d}|^2 \qquad (16)$$

is particularly appealing. However, we modify it in two ways. First, we substitute a soft penalty function f_2 for the quadratic one. Second, we introduce two auxiliary scalar fields, ξ and ζ, which will constitute direct estimates of the divergence and vorticity, respectively. The new regularizer is given by:

$$\int_\Omega |\mathrm{div}\boldsymbol{d} - \xi|^2 + \lambda f_2(|\nabla\xi|) + \int_\Omega |\mathrm{curl}\boldsymbol{d} - \zeta|^2 + \lambda f_2(|\nabla\zeta|).$$
$$(17)$$

The first part of each integral encourages the displacement to comply with the current divergence and vorticity estimates, ξ and ζ, through a quadratic goodness-of-fit. The second part equips the divergence and the vorticity estimates with a robust first-order regularization favoring piece-wise smooth configurations.

Although it is not investigated in this paper, the introduction of the two scalar fields, ξ and ζ, permits even more sophisticated priors on the divergence and the vorticity of the imaged flow. In [33], we investigate, for instance, the use of a parametric description of these two quantities in the neighborhood of the critical points of the flow. In the same vein, the explicit manipulation of div and curl estimates should simplify the assimilation of physical measures provided by dedicated probes, such as sparse vorticity measurements obtained by thermal anemometry [43].

The augmented formulation also provides computational advantages. The second-order regularization is replaced by two interleaved first-order regularizations. From the Euler-Lagrange point-of-view, this amounts to replacing the fourth-order PDE associated with (16), by a coupled PDE of order two. In the next section, we shall see that our minimization, although not posed in terms of PDEs, turns out to be a discrete problem of reasonable complexity as assessed by the locality of the interactions among the unknown variables (or, equivalently, by the bandwidth of the nonlinear sparse system to be solved).

The soft penalty function f_2 is important to allow for abrupt variations of divergence and vorticity. The vorticity and divergence of complex flows are actually known to concentrate in aggregates with sharp borders. These boundaries will consequently be better recovered by a loose penalty function. We chose again the Leclerc penalty function, $f_2(x) = 1 - \exp(-\tau_2 x^2)$.

To embed the regularization term (17) in the incremental hierarchical framework briefly presented in the previous section, we must express it in terms of the displacement increment $\Delta\boldsymbol{d} = \tilde{\boldsymbol{d}} - \boldsymbol{d}$, where $\tilde{\boldsymbol{d}}$ is a crude estimate of the displacement field. The second term of our global cost function, to be combined with (13), is then:

$$F_2(\Delta\boldsymbol{d}, \xi, \zeta) = \int_\Omega |\mathrm{div}(\tilde{\boldsymbol{d}} + \triangle\boldsymbol{d}) - \xi|^2 + \lambda f_2(|\nabla\xi|) +$$
$$\int_\Omega |\mathrm{curl}(\tilde{\boldsymbol{d}} + \Delta\boldsymbol{d}) - \zeta|^2 + \lambda f_2(|\nabla\zeta|).$$
$$(18)$$

We now detail the minimization of the total cost function.

5   MINIMIZATION

The incremental estimation of the dense displacement field is conducted as follows: From each image in the sequence, lower resolution images (three in our experiments) are derived by Gaussian smoothing and subsampling. This multiresolution structure is then used in a coarse-to-fine way. At a given time in the sequence and at a given resolution, the displacement field $\tilde{\boldsymbol{d}}$ obtained from the previous resolution is refined by solving

$$\min_{\Delta\boldsymbol{d}, \xi, \zeta} F_1(\triangle\boldsymbol{d}) + \alpha F_2(\triangle\boldsymbol{d}, \xi, \zeta), \qquad (19)$$

where F_1 and F_2 are defined in (13) and (18). The approximate solution obtained by the technique described hereafter is added to $\tilde{\boldsymbol{d}}$ and passed to the next resolution (Fig. 2). At the coarsest resolution, $\tilde{\boldsymbol{d}} \equiv 0$.

5.1   Discretization

The discretization of the minimization problem (19) can be done in several ways. It can be implemented at the level of the Euler-Lagrange PDE that the minimizer must satisfy or it can be performed right away at the level of the cost functional. We prefer the second solution, reminiscent of MRF-based approaches, for its simplicity. With the unknown fields $\Delta\boldsymbol{d}$, ξ, and ζ now restricted to the finite grid S of pixel centers, the discrete counterpart of cost terms F_1 and F_2 is:

$$F_1(\Delta\boldsymbol{d}) = \sum_{s \in S} f_1\Big[\exp(\mathrm{div}\tilde{\boldsymbol{d}}(s))\big((\tilde{E}(s)\nabla\mathrm{div}\tilde{\boldsymbol{d}}(s) +$$
$$\nabla\tilde{E}(s)) \cdot \Delta\boldsymbol{d}(s) + \tilde{E}(s)\big) - E(s)\Big]$$
$$(20)$$

$$F_2(\Delta\boldsymbol{d}, \xi, \zeta) = \sum_{s \in S}[\mathrm{div}\tilde{\boldsymbol{d}}(s) + \mathrm{div}\Delta\boldsymbol{d}(s) - \xi_s]^2 +$$
$$\lambda \sum_{\langle s,r\rangle \in C} f_2(|\xi_s - \xi_r|) +$$
$$\sum_{s \in S}[\mathrm{curl}\tilde{\boldsymbol{d}}(s) + \mathrm{curl}\Delta\boldsymbol{d}(s) - \zeta_s]^2 +$$
$$\lambda \sum_{\langle s,r\rangle \in C} f_2(|\zeta_s - \zeta_r|),$$
$$(21)$$

where C is the set of adjacent pixel pairs (in the nearest-neighbor sense), f_1 and f_2 are Leclerc penalty functions with parameters τ_1 and τ_2, respectively, and div (respectively, curl) is now a discrete divergence (respectively, vorticity) operator.

Using a standard central finite difference discretization, the divergence and the vorticity can be defined at location $s = (i, j)$ as

$$\begin{cases} \mathrm{div}\boldsymbol{d}(i,j) = \frac{1}{2}[(u_{i,j+1} - u_{i,j-1}) + (v_{i+1,j} - v_{i-1,j})] \\ \mathrm{curl}\boldsymbol{d}(i,j) = \frac{1}{2}[(v_{i,j+1} - v_{i,j-1}) - (u_{i+1,j} - u_{i-1,j})], \end{cases} \quad (22)$$

where $\boldsymbol{d}_{i,j} = (u_{i,j}, v_{i,j})$. However, because the central point (i, j) does not appear in this discretization, the set of pixels is split into subsets (the red and black of a checkerboard) that only weakly interact through the data term F_1. This results in a badly conditioned system for which the resolution has been found problematic, with the appearance of numerical instabilities. Hence, we prefer to use a nonsymmetric discretization of the partial derivatives so

as to include the central point. The resulting discrete div and curl operators are:

$$
\begin{cases}
\text{div}\boldsymbol{d}(i,j) = \frac{1}{6}[(3u_{i,j} + u_{i,j-2} - 6u_{i,j-1} + 2u_{i,j+1}) + \\
\qquad\qquad\quad (3v_{i,j} + v_{i-2,j} - 6v_{i-1,j} + 2v_{i+1,j})] \\
\text{curl}\boldsymbol{d}(i,j) = \frac{1}{6}[(3v_{i,j} + v_{i,j-2} - 6v_{i,j-1} + 2v_{i,j+1}) - \\
\qquad\qquad\quad (3u_{i,j} + u_{i-2,j} - 6u_{i-1,j} + 2u_{i+1,j})].
\end{cases}
\tag{23}
$$

The discretization of the x-derivatives used here arises from solving the linear system

$$
\begin{aligned}
z_{i+h,j} = z_{i,j} &+ h \frac{\partial z}{\partial x}\Big|_{i,j} + \frac{h^2}{2}\frac{\partial^2 z}{\partial x^2}\Big|_{i,j} + \\
&\frac{h^3}{6}\frac{\partial^3 z}{\partial x^3}\Big|_{i,j} + 0(h^4),\ h = -1, 1, 2,
\end{aligned}
$$

for the three derivatives at location (i,j), where z denotes either u or v. The discretized y-derivatives are similarly obtained.

5.2 Semiquadratic Alternate Minimization

The global discrete cost function $F = F_1 + \alpha F_2$ must be minimized with respect to the three sets of variables $\boldsymbol{d} = \{\triangle\boldsymbol{d}_s\}_{s\in S}$, $\xi = \{\xi_s\}_{s\in S}$, $\zeta = \{\zeta_s\}_{s\in S}$. The minimization is conducted alternately with respect to one of the three sets at a time. If the penalty functions, f_1 and f_2, are quadratic, then each of the three optimization subproblems amounts to a large, sparse, linear system for which efficient solvers (e.g., multigrid Gauss-Seidel, preconditioned conjugate gradient) exist. The choice of softer penalty functions, however, makes the problem nonquadratic and even nonconvex in the case of the bounded Leclerc function. Although general nonconvex minimization techniques are available, the specific minimization problem we face is classically turned into an augmented half-quadratic minimization problem [22]. Indeed, the Leclerc penalty shares, with all penalty functions f such that $f(\sqrt{.})$ is concave, the property that:

$$
f(x) = \min_{z\in(m,M)} zx^2 + \psi(z),
\tag{24}
$$

where $M = \lim_{0^+} \frac{f'(x)}{2x}$, $m = \lim_{+\infty}\frac{f'(x)}{2x}$, and ψ, for which an expression can be found in [10], [19], is such that the minimizer on the right-hand-side is given by $z = \frac{f'(x)}{2x}$. For the Leclerc penalty with parameter τ, we have $m = 0$, $M = \tau$, and $\frac{f'(x)}{2x} = \tau\exp -\tau x^2$.

Using (24), each minimization problem of the generic form $\min_x \sum_k f(g_k(x))$ can be replaced by the auxiliary problem $\min_{x,\{z_k\}} \sum_k z_k g_k^2(x) + \psi(z_k)$, which can be solved by iteratively reweighted least squares (IRLS) [22]: For fixed auxiliary (weight) variables $z_k \in (m, M)$, one faces a least-squares problem; for fixed x, the optimal value for each weight is known in closed form as $\frac{f'(g_k(x))}{2g_k(x)}$.

The sum constituting F_1, as well as the sums in F_2 where penalty function f_2 appears, can be rewritten this way. Three sets of weights, $\{z_s^d\}_{s\in S}$, $\{z_{sr}^\xi\}_{\langle s,r\rangle \in C}$, and $\{z_{sr}^\zeta\}_{\langle s,r\rangle \in C}$ are thus introduced. Each of the three minimizations of F, with respect to $\triangle\boldsymbol{d}$, ξ, and ζ, respectively, is then handled by

IRLS. As suggested in [34], however, we use a nonclassic version. Whereas the chosen least-squares solver is usually applied to each least-squares problem at fixed weights until a convergence criterion is met, we only apply one iteration (i.e., one sweep over the pixel grid) of the Gauss-Seidel solver in between two weight updates. This is proven to leave the theoretical convergence properties of the algorithm unchanged—namely, convergence towards a local minimum in nonconvex cases—and yields in practice a faster convergence towards a better local minimum [34].

Finally, the key minimization with respect to the increment field $\triangle\boldsymbol{d}$ is improved, both in terms of result quality and speed, by applying the coarse-to-fine technique introduced in [32]. For sake of completeness, we provide a brief description of this hierarchical minimization technique in Appendix I.

The complete algorithm is summarized below:

- $\tilde{\boldsymbol{d}} \equiv 0$
- From coarsest to finest resolution do
  - $\triangle\boldsymbol{d} \equiv 0$, $\xi \equiv 0$, $\zeta \equiv 0$, $z^d \equiv \tau_1$, $z^\xi \equiv \tau_2$, $z^\zeta \equiv \tau_2$
  - Until convergence do$^{a\ b\ c}$

$$
\begin{aligned}
\triangle\boldsymbol{d} \ \leftarrow\ & GS\Big(\sum_s z_s^d g^2(\triangle\boldsymbol{d})_s \\
& +\alpha[\text{div}\tilde{\boldsymbol{d}}(s) + \text{div}\triangle\boldsymbol{d}(s) - \xi_s]^2 \\
& +\alpha[\text{curl}\tilde{\boldsymbol{d}}(s) + \text{curl}\triangle\boldsymbol{d}(s) - \zeta_s]^2 \Big)
\end{aligned}
\tag{25}
$$

$$
z_s^d \ \leftarrow\ \tau_1 \exp -\tau_1 g^2(\triangle\boldsymbol{d})_s,\ \forall s \in S
$$

$$
\xi \ \leftarrow\ GS\Big(\sum_s [\text{div}\tilde{\boldsymbol{d}}(s) + \text{div}\triangle\boldsymbol{d}(s) - \xi_s]^2 \\
+\lambda \sum_{\langle s,r\rangle} z_{sr}^\xi |\xi_s - \xi_r|^2 \Big)
\tag{26}
$$

$$
z_{sr}^\xi \ \leftarrow\ \tau_2 \exp -\tau_2 |\xi_s - \xi_r|^2,\ \forall \langle s,r\rangle \in C
$$

$$
\zeta \ \leftarrow\ GS\Big(\sum_s [\text{curl}\tilde{\boldsymbol{d}}(s) + \text{curl}\triangle\boldsymbol{d}(s) - \zeta_s]^2 \\
+\lambda \sum_{\langle s,r\rangle} z_{sr}^\zeta |\zeta_s - \zeta_r|^2 \Big)
\tag{27}
$$

$$
z_{sr}^\zeta \ \leftarrow\ \tau_2 \exp -\tau_2 |\zeta_s - \zeta_r|^2,\ \forall \langle s,r\rangle \in C
$$

- If not finest resolution, pass $\tilde{\boldsymbol{d}} \leftarrow \tilde{\boldsymbol{d}}+\triangle\boldsymbol{d}$ to next resolution.
- If finest resolution, $\boldsymbol{d} \leftarrow \tilde{\boldsymbol{d}} + \triangle\boldsymbol{d}$.

---

^{a}Convergence is reached when the relative change in the L_2-norm of $\triangle\boldsymbol{d}$ is below 3%.
$^b g(\triangle\boldsymbol{d})_s$ denotes $\exp(\text{div}\tilde{\boldsymbol{d}}(s))\big((\tilde{E}(s)\nabla\text{div}\tilde{\boldsymbol{d}}(s) + \nabla\tilde{E}(s))\cdot \triangle\boldsymbol{d}(s) + \tilde{E}(s)\big) - E(s)$.
^{c}"$GS()$" denotes one Gauss-Seidel iteration on the quadratic form under consideration.

6 Experimental Results

In this section, we compare a variety of dense displacement fields to assess the respective merits of the data cost term F_1, (13), introduced in Section 3 and of the div-curl regularizer F_2, (18), introduced in Section 4. Their classic counterparts, reviewed in Section 2 [where they constitute global cost function (4)] will be referred to as G_1 and G_2, respectively, in what follows. The first is discretized in the same way as F_1 in (20). The discretization of the first-order spatial derivatives in G_2 is similar to those of ζ and ξ in (21).
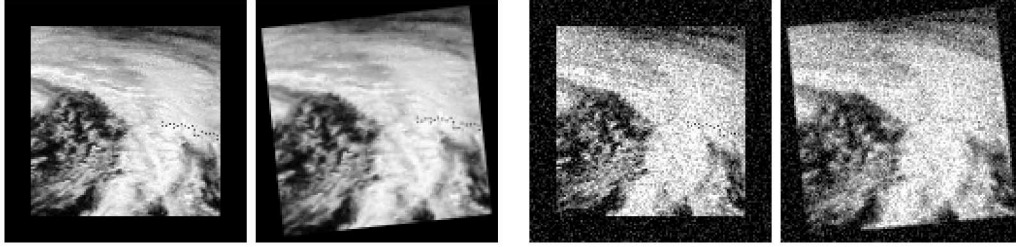
Fig. 4. **Synthetic spiral motion under the brightness constancy assumption** the motion with $\mathrm{curl} = 0.2$ and $\mathrm{div} = 0.2$ is applied to a real Meteosat image from the visible channel to get (Left) a noise free image pair and (Right) a noisy image pair.

To ease the reading of the next sections, one might think of F as in "fluid" and of G as in "general." Also, each of the four cost functions involves a penalty function. When this penalty function is chosen to be quadratic, a superscript q is used and when the bounded Leclerc penalty function is used instead, the superscript L appears. For instance, the cost function $G_1^q + \alpha G_2^q$ amounts to an incremental version of the original Horn and Schunck functional and $G_1^L + \alpha G_2^L$ is similar to the robust costs used in [9] and [31]. When using the Leclerc penalty function, an additional parameter τ must be set.

Whatever the combination of data term and regularizer under consideration, the minimization procedure employed is the straightforward adaptation of the one described in Section 5 for $F_1^L + \alpha F_2^L$, with an incremental estimation based on a multiresolution representation of the images and a hierarchical technique to speed up and improve the minimization at each resolution.

6.1 Synthetic Motions

In order to assess the relevance of the second-order div-curl regularizer F_2 (in its simple quadratic form F_2^q), we start with comparisons on two image pairs obtained by applying the same synthetic spiral motion to a real Meteosat image (Fig. 4). This motion is the sum of a zoom-in (constant divergence of 0.2) and a counter-clockwise rotation (constant vorticity of 0.2). In the first pair, the second image is built from the first real image in order to satisfy the brightness constancy constraint. In the second pair, both images are then corrupted by white Gaussian noise at 10dB.

On this benchmark we compared $G_1^q + \alpha G_2^q$, $G_1^q + \alpha G_2^L$, and $G_1^q + \alpha F_2^q$. The quadratic data term G_1^q is employed since the image pairs respect the brightness constancy constraint for the true field by construction (up to an

additive Gaussian noise, for which the quadratic estimator is the optimal one, in the second pair). The first cost has been tested with 10 different values of regularization parameter $\alpha \in [50, 500]$. For the second one, 100 different parameter combinations $(\alpha, \tau_2) \in [50, 500] \times [0.4, 2.0]$ have been tried. Finally, the third one has been used with 100 different parameter combinations

$$(\alpha, \lambda) \in [50, 500] \times [50, 500].$$

The true displacement field being available, we can quantify the global correctness of the various estimates according to the angular discrepancy criterion of Barron et al. [5]. For a given estimated motion field, the mean μ and the standard deviation σ of the angular discrepancy over the whole pixel grid are computed. For each method, we average these quantities over the estimates obtained with the different parameter sets, getting $\langle \mu \rangle$ and $\langle \sigma \rangle$. These quantities are reported in Table 1, along with the extremal values, (μ_{min}, μ_{max}) and $(\sigma_{min}, \sigma_{max})$, obtained within each set of trials. We also report the total number of image sweeps taken by each algorithm to converge with its best set of parameters.

The average values of the mean discrepancy and its standard deviation over the multiple trials are lower for the regularizer F_2^q in both examples. The much tighter range of variations of μ and σ for this regularizer also indicates that this superiority holds consistently over the large range of parameter values we have explored.

These first simple experiments show that the second-order div-curl regularizer, even in quadratic form, is better suited than a standard regularizer, robust or not, for rotational and divergent motions, which are key components of fluid flows. From a complexity point of view, this new regularizer results in a computation increase of roughly 15 percent.

TABLE 1
Results for the Two Image Pairs in Fig. 4

Cost function	$\langle \mu \rangle$	μ_{min}	μ_{max}	$\langle \sigma \rangle$	σ_{min}	σ_{max}	image sweeps
$G_1^q + \alpha G_2^q$	$3.26°$	$2.63°$	$8.34°$	$3.27°$	$2.90°$	$6.32°$	936
$G_1^q + \alpha G_2^L$	$3.40°$	$2.42°$	$18.93°$	$2.80°$	$1.11°$	$26.71°$	1136
$G_1^q + \alpha F_2^q$	$\mathbf{2.64°}$	$\mathbf{2.35°}$	$\mathbf{4.96°}$	$\mathbf{1.44°}$	$\mathbf{1.04°}$	$\mathbf{3.67°}$	1377
$G_1^q + \alpha G_2^q$	$10.59°$	$8.55°$	$15.38°$	$11.48°$	$8.28°$	$18.14°$	1103
$G_1^q + \alpha G_2^L$	$10.36°$	$5.99°$	$25.13°$	$10.81°$	$6.05°$	$28.41°$	1204
$G_1^q + \alpha F_2^q$	$\mathbf{7.90°}$	$\mathbf{5.89°}$	$\mathbf{11.97°}$	$\mathbf{8.51°}$	$\mathbf{5.72°}$	$\mathbf{14.26°}$	1384

Average and extremal values of the mean (respectively, standard deviation) of the angular discrepancy to ground-truth for three different global cost functions. (Top) Noise free image pair. (Bottom) Noisy image pair.
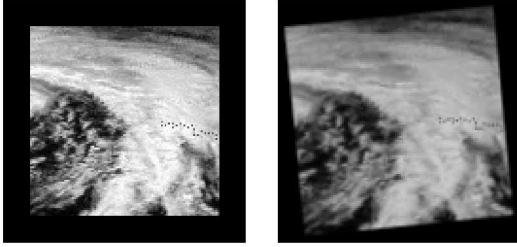
Fig. 5. **Synthetic spiral motion under the continuity equation assumption** the motion with $\mathrm{curl} = 0.2$ and $\mathrm{div} = 0.2$ is applied to a real Meteosat image from the visible channel, so as to respect the integrated continuity equation, i.e., the brightness of a displaced point is multiplied by $\exp -0.2$.

The results on the noisy image pair also demonstrate that the second-order div-curl regularizer is more robust to noise. One of the reasons might be that a higher regularization parameter α is required to limit the noise influence, which in turn exaggerates even more the tendency of first-order smoothing to underestimate divergent and rotational motion components.

It is worth noting that the image pairs are only related to fluids via the photometric content of the first real image. The apparent motion we have generated is that which any fronto-parallel plane rigid object undergoing a rotation while closing in on the camera would exhibit. Although the use of this new regularizer in other contexts falls beyond the scope of this paper, we think that these simple experiments already indicate that it is an appealing alternative to standard robust first-order smoothing as soon as important divergent and vorticity motions are present, irrespective of the nature of the imaged scene.

Continuing these synthetic image experiments, we assessed the consistency of the data cost F_1 by considering a new image pair (Fig. 5), where the same spiral motion has been applied to the same satellite image, but accompanied this time by brightness changes compliant with the integrated continuity equation (12). More precisely, the divergence of the motion being 0.2, the luminance of each displaced pixel is scaled down by $\exp(-\mathrm{div}\boldsymbol{d}) = 0.818$.

On this image pair we compared the four combinations of data costs (G_1 and F_1) and regularizers (G_2 and F_2). As the brightness constancy is, by construction, not satisfied, we have considered only the robust version G_1^L of the data cost relying on it. The value of the associated parameter τ_1 has been tested within the range [1,30]. Conversely, because of the compliance with the continuity equation, only the quadratic data cost F_1^q is considered. The regularizers were

taken to be G_2^L and F_2^q, as in previous experiments, with their parameters set to the best combinations found in the previous experiments namely, $(\alpha, \tau_2) = (400, 1.6)$ for the first-order robust smoothing and $(\alpha, \lambda) = (450, 200)$ for the quadratic div-curl regularization. The results are summarized in Table 2. Note that, since F_1^q does not involve parameters, only one parameter combination was finally considered for the two global costs using this data term (resulting in blank entries for the corresponding averaged quantities in the table).

The dramatic difference between the standard model ($G_1^L + \alpha G_2^L$), whose convergence is difficult on this example and the complete new one ($F_1^q + \alpha F_2^q$) can be visually appreciated by examining the plots in Fig. 6.

Because of the compliance of the data with the integrated continuity equation, no conclusion can be drawn from the drastic superiority of the data term F_1 which relies on it. However, two interesting observations can be made. First, it is striking to note the extent to which a modest intensity loss can fool an estimation based on brightness constancy, even if a soft penalty function is used. The second remark is that the superiority of the div-curl regularization already noticed in previous experiments is even more apparent in this case. The use of a first-order smoothing with the continuity equation leads indeed to an inherent contradiction in the presence of brightness changes: The data term favors a nonnull divergence to account for the brightness changes, whereas the first-order smoothing encourages a low divergence. Here, the estimated mean divergence with $F_1^q + \alpha G_2^L$ exhibits a relative error of 21 percent, whereas this error drops to 1 percent with the full dedicated cost $F_1^q + \alpha F_2^q$.

6.2 Real World Sequences

We now turn to qualitative comparisons on two different types of real meteorological image sequences.

The first sequence, issued from the infrared (IR) channel of Meteosat (Fig. 7), was acquired 21 January 1998 and shows a large spiraling depression in the lower left part of the image, together with a cold (dark) cloud structure expanding in the upper right part of the image.

In Fig. 8, we compare, for two consecutive images of the sequence, the displacement fields estimated with the five cost functions $G_1^L + \alpha G_2^L$, $G_1^L + \alpha F_2^L$, $F_1^L + \alpha G_2^L$, $F_1^L + \alpha F_2^q$, and $F_1^L + \alpha F_2^L$, as well as the associated vorticity and divergence.

The minimization of each cost function provides a plausible displacement field, where both the counterclockwise spiral of the depression and the downward motion of

TABLE 2
Results on the Image Pair in Fig. 5

Cost	$\langle \mu \rangle$	μ_{min}	μ_{max}	$\langle \sigma \rangle$	σ_{min}	σ_{max}	image sweeps
$G_1^L + \alpha G_2^L$	59.49°	43.75°	75.65°	22.14°	12.66°	27.67°	28365
$G_1^L + \alpha F_2^q$	57.60°	45.36°	73.46°	10.84°	4.69°	22.44°	1625
$F_1^q + \alpha G_2^L$	8.91°	-	-	15.46°	-	-	1225
$F_1^q + \alpha F_2^q$	**2.12°**	-	-	**1.54°**	-	-	3733

Average and extremal values of the mean (respectively, standard deviation) of the angular discrepancy to ground-truth for four different global cost functions.
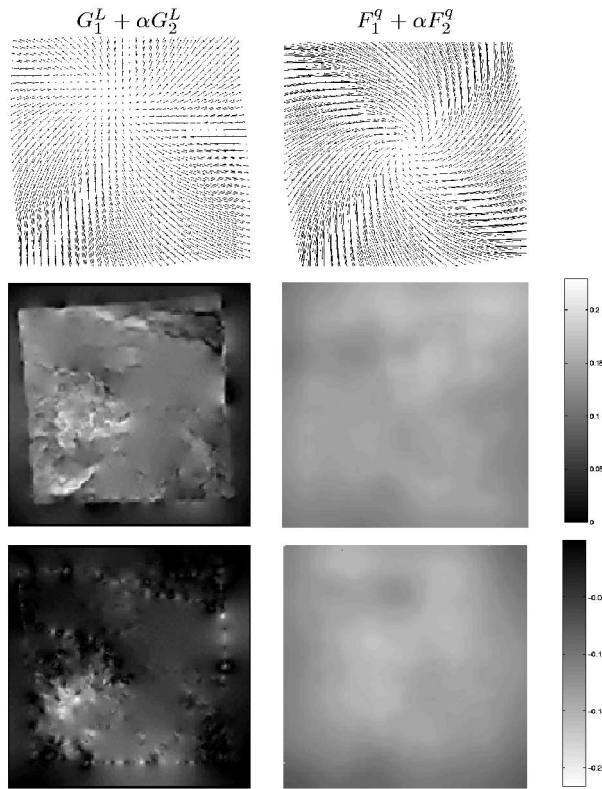
$G_1^L + \alpha G_2^L$ $F_1^q + \alpha F_2^q$

Fig. 6. The robust generic model $(G_1^L + \alpha G_2^L)$ as compared to the quadratic dedicated model $(F_1^q + \alpha F_2^q)$ on the spiral motion with intensity loss of Fig. 5. (Top) Subsampled displacement fields. (Middle) Corresponding divergence. (Bottom) Corresponding vorticity.

cold clouds are captured. Nevertheless, the complete generic estimator underestimates the different components of this complex motion, providing an estimated field with smooth and small divergence and vorticity. Replacing the robust first-order regularizer by the novel div-curl regularizer already reduces this tendency. Whereas the impact of the novel data-term is not noticeable when used in conjunction with the first-order smoothing (remember that, for low divergence, this data term gets close to the standard data-term), the full benefit of this data-term appears clearly with the second-order regularization. Indeed, the two versions $F_1^L + \alpha F_2^q$ and $F_1^L + \alpha F_2^L$ of the complete dedicated cost function provide fields which capture the full extent of the spiraling motion of the depression motion and recover better the diverging motion of the cold cloud structure. At the divergence and vorticity level, these various motion components estimated by the dedicated approach are associated with a large concentration of vorticity in the middle of the depression and several divergence concentrations at the periphery of the depression and in the middle of the cloud structure entering at the top right of the image. The visualization of these scalar fields finally demonstrates the impact of the robust cost penalty in the div-curl regularizer. As compared to the quadratic version F_2^q, one notes that F_2^L provides an even better recovery of the divergence and vorticity concentrations.

The importance of using an integrated version of the continuity equation to recover incrementally large divergent displacements has been outlined in Section 3. We can

illustrate this by leading comparative experiments on real divergent sequences at various frame rates. We applied our approach on one convective cell in the three first images of the water vapor sequence in Fig. 3, using the data term based on the integrated continuity equation, as well as its counterpart, based on the plain differential continuity equation. If we denote d^{12}, d^{23}, and d^{13} the displacement fields estimated between images 1 and 2, 2 and 3, and 1 and 3, respectively, one should obtain $d^{12} + d^{23} = d^{13}$. Focusing on the divergence only, we noted a much larger loss of accuracy without the integrated continuity equation when the frame rate is divided by two (thus increasing the amplitude of the divergent motion). Indeed, with the integrated continuity equation $\mathrm{div}d^{13}$ is only 3.3 percent lower than $\mathrm{div}(d^{12} + d^{23})$ in average, whereas this discrepancy is 10.9 percent with the standard differential continuity equation.

Obviously, the true value of all the estimated quantities (motion components, divergence, vorticity) are unknown. The good quality of the results obtained by the dedicated approach, as compared to the generic approach, has however been assessed by meteorologists in the context of a cooperation with the European consortium Eumetsat. Also, the superiority of the dedicated approach on the whole sequence will be made clearer via trajectory reconstruction in the next section.[1]

The second satellite sequence (Fig. 9) is a water vapor (WP) image sequence acquired 4 August 1995. This sequence represents a depression in the left part of the imaged area and a set of active convective cells in the right part. Such water vapor images contain valuable information for meteorologists. However, automatic analysis of these images remains a challenge due to their low photometric contrast.

Two sequences of displacement fields obtained on this sequence by the generic and the dedicated cost functions are shown in Fig. 10. It is important to point out that, to ensure a minimum of temporal coherence, the estimation procedure at the coarsest resolution was initialized in both cases by the displacement field estimated at the preceding time index. At the very beginning of the sequence, the procedure was initialized by a null field as explained in Section 5.

On this test sequence the difference in behavior of the two cost functions is even more drastic than on the IR sequence. The generic approach smoothes the displacement field in such a way that the spiraling motion of the depression is completely lost. In contrast, the dedicated approach demonstrates its capacity to recover this important motion component, even in the very adverse conditions imposed by the poor contrast of the images. Furthermore, the observation of the consecutive displacement fields indicates that the good instantaneous estimation of the large vorticity displacement by the dedicated method does not preclude a satisfactory stability in time. The spatio-temporal quality of the motion fields thus obtained will be further confirmed in the next section, where entire trajectories are recovered through the complete test sequence.

The vorticity and divergence of the estimated motion fields obtained with the dedicated approach are shown in
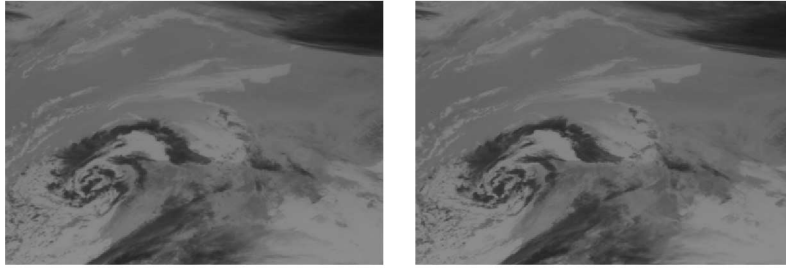
Fig. 7. **IR Meteosat images** two consecutive images from the test sequence, with a large spiraling depression in the lower left quarter of the observed region and a cold (dark) cloud structure expanding in the upper right corner.
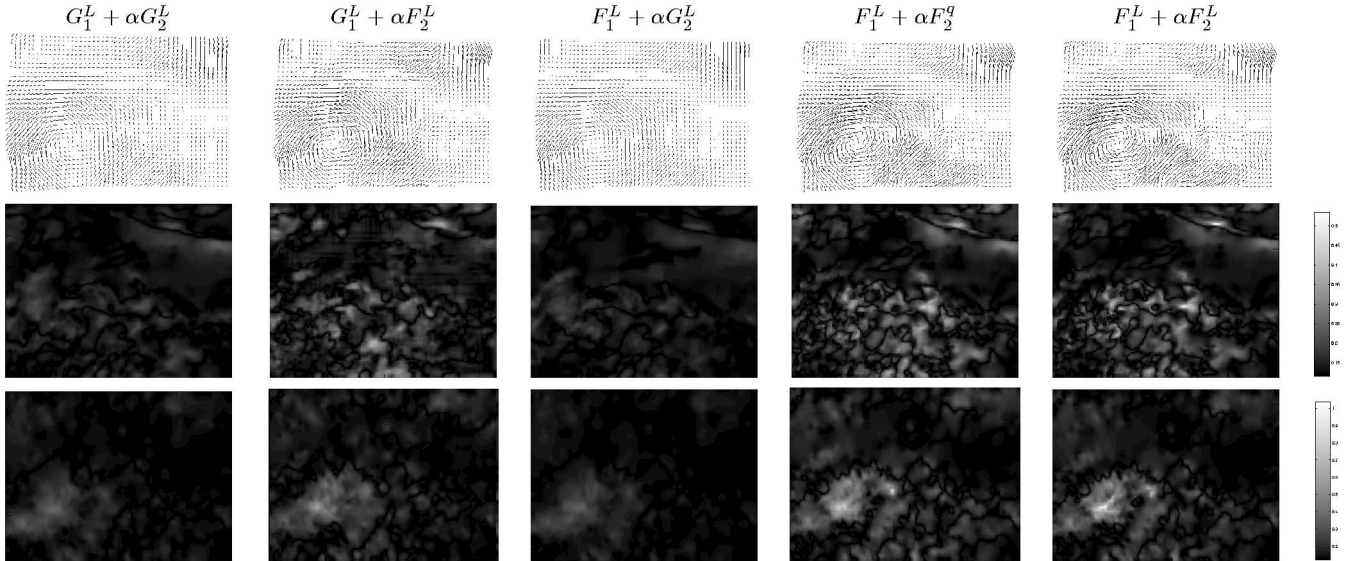

Fig. 8. **The dedicated data-model and regularizer as compared to standard data-model and regularizer on the IR Meteosat sequence** (Top) Subsampled displacement fields estimated on the image pair in Fig. 7. (Middle) Corresponding divergence. (Bottom) Corresponding vorticity.

Fig. 11. A stable clear-cut concentration of vorticity characterizes the central part of the depression, whereas small and transient divergent blobs appear in the depression and at the level of the various convective cells.

7 TWO APPLICATIONS

7.1 Trajectory Reconstruction

Reconstructing the entire trajectories of some elements of fluid is a problem of practical interest in its own right. It indeed provides numerical Lagrangian drifters which can complement or even replace complicated, expensive, and sparse physical drifters in a number of applications ranging from the tracking of fish eggs to the tracking of atmospheric pollutants. Apart from this applicative interest, the reconstruction of trajectories provides a unique visualization tool to assess the quality of estimated motion fields. As suggested in [40], trajectory reconstruction is performed here using a fourth-order Runge-Kutta integration method (see Appendix II).

We present results obtained for the complete IR and WP Meteosat sequences partly shown in Figs. 7 and 9. We compare various trajectories reconstructed in these two sequences, based on the displacement fields obtained with the generic robust cost $G_1^L + \alpha G_2^L$ and the dedicated robust cost $F_1^L + \alpha F_2^L$, respectively. These trajectories are superimposed on the final image of the sequence in Figs. 12 and 13.

The starting points were either picked manually or taken to be the nodes of a regular grid.

Again, no ground truth is available to assess the quality of these trajectories. However, the global flow evolution pictured by the trajectories stemming from a regular grid of starting points is visibly more plausible when based on the dedicated approach. In both sequences, it is clear, for instance, that these trajectories are in better agreement with the spiraling motion of depressions than those based on the generic model. The latter model fails to capture this motion at all in the WP sequence, where the photometric information is rather poor.

This first global impression can be made more precise by having a closer look at just a few sparse trajectories. For each trajectory, one can visually assess the relevance of the obtained endpoint given the selected starting point.[2]

The quality of these trajectories demonstrates the relevance of the instantaneous motion fields estimated by our method, as well as their temporal stability within a long time period, despite the absence of an explicit temporal link within the estimator (apart from the use of the previous estimate as an initial guess).

2. The complete sequences and the reconstructed trajectories can be viewed at http://www.irisa.fr/vista/Themes/Demos/MouvementFluide/fluide.english.html.
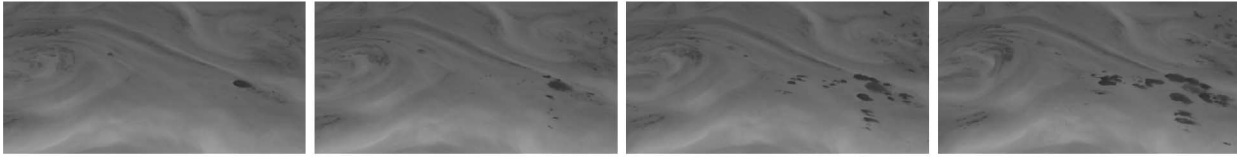
Fig. 9. **WP Meteosat images** four consecutive images from the test sequence, with a large spiraling depression in the left part and convective cells (dark) rapidly expanding in the right part.
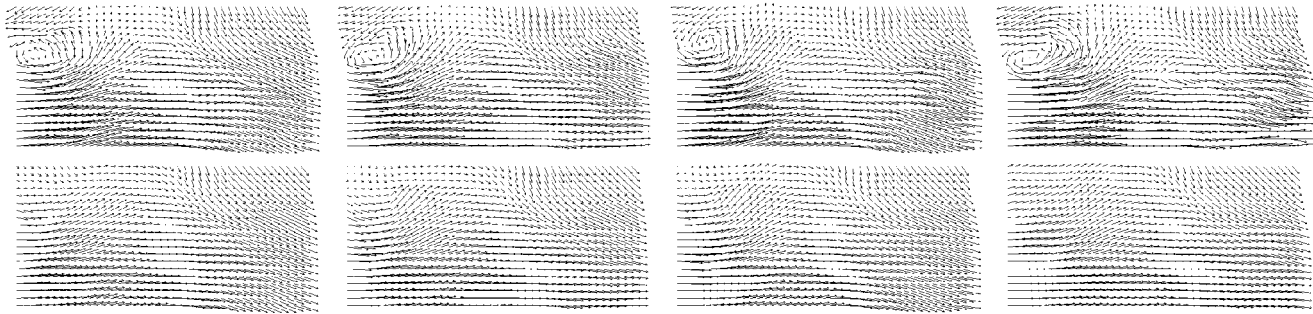

Fig. 10. **The dedicated model as compared to the robust generic model on the WP Meteosat sequence**: Subsampled displacement fields estimated from the WP images in Fig. 9 using (Top) the dedicated cost function $F_1^L + \alpha F_2^L$ and (Bottom) the generic cost function $G_1^L + \alpha G_2^L$.
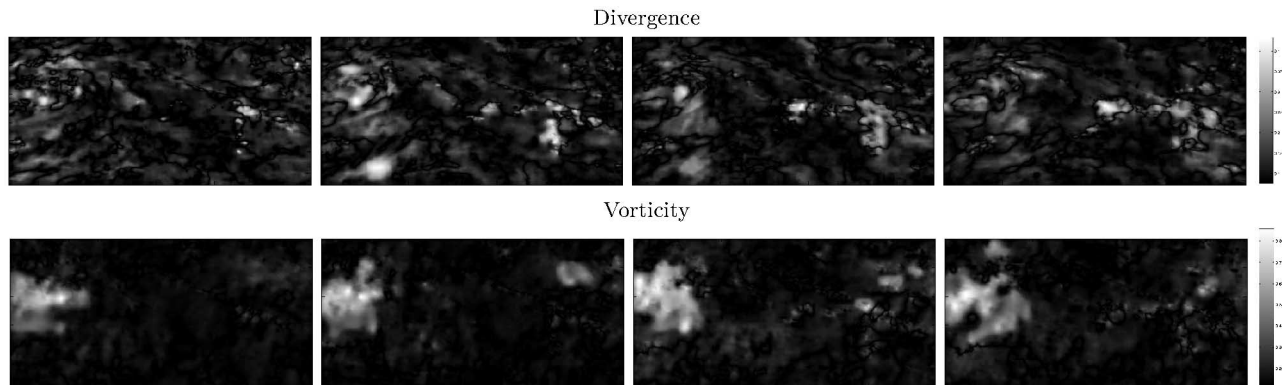
Divergence

Vorticity

Fig. 11. **Divergence and vorticity sequences with the dedicated approach** as derived from the displacement fields in the top row of Fig. 10. (Top) divergence, (Bottom) vorticity. The stable vorticity structure of the depression and the transient divergent structures of the convective cells emerge clearly.
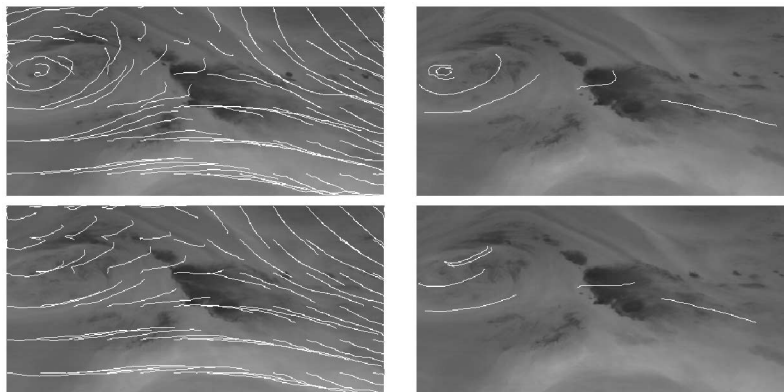

Fig. 12. **Trajectories reconstructed in the WP images** based on the displacement fields estimated with (Top) the robust dedicated cost and (Bottom) the robust generic cost. The trajectories are superimposed on the final image of the sequence.

7.2 Extraction of Main Divergence and Vorticity Structures

As we have repeatedly mentioned, the vorticity and divergence structures are key entities for describing, understanding, and predicting, imaged fluid flows. Based on a dense displacement field previously estimated, such structures can be identified via the extraction and analysis of the critical points (i.e., where the speed vanishes locally) of the flow [11], [16], [17], [18], [33], [35]. This usually requires the fitting of parametric motion models and gives access only to the structure of interest modulo the laminar (i.e., divergence and vorticity free) component of the flow: If a vortex is transported by a global laminar motion, say a translation, its center, being nonstationary, will not be a
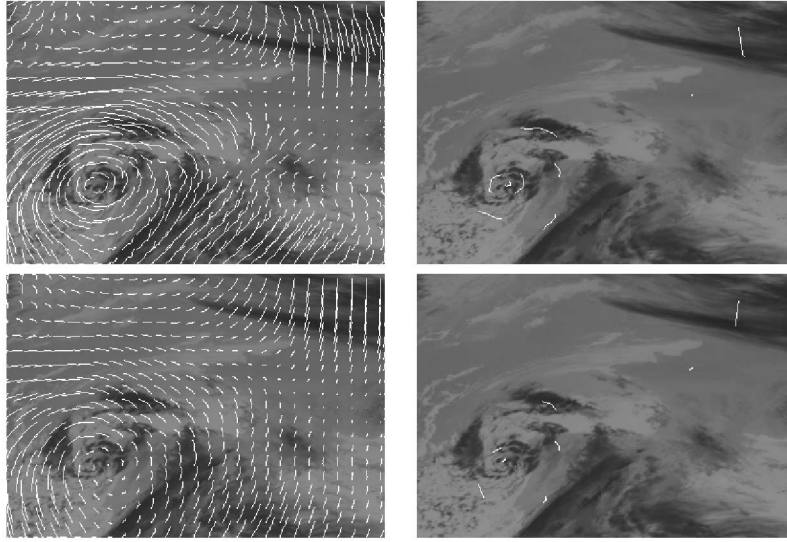
Fig. 13. **Trajectories reconstructed in the IR images** based on the displacement fields estimated with (Top) the robust dedicated cost and (Bottom) the robust generic cost. The trajectories are superimposed on the final image of the sequence.
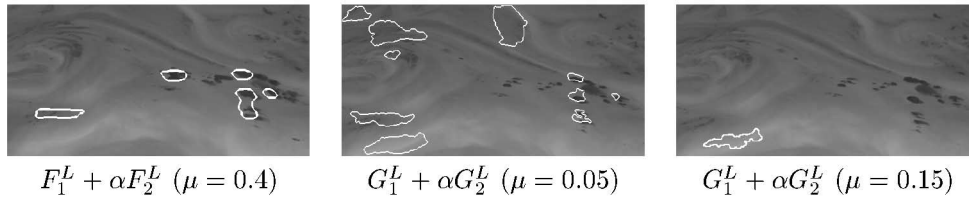

$F_1^L + \alpha F_2^L \ (\mu = 0.4)$ $\qquad$ $G_1^L + \alpha G_2^L \ (\mu = 0.05)$ $\qquad$ $G_1^L + \alpha G_2^L \ (\mu = 0.15)$

Fig. 14. **Main divergent structures** extracted from estimated flows on WP Meteosat images.

$F_1^L + \alpha F_2^L \ (\mu = 0.4)$ $\qquad$ $G_1^L + \alpha G_2^L \ (\mu = 0.05)$ $\qquad$ $G_1^L + \alpha G_2^L \ (\mu = 0.15)$
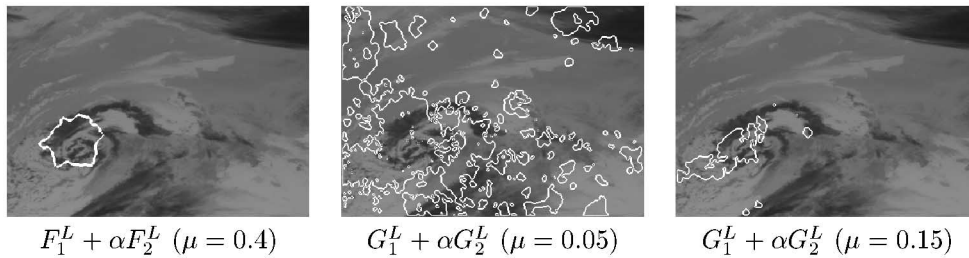
Fig. 15. **Main rotational structures** extracted from estimated flows on IR Meteosat images.

singular point of the flow, although it is the real location of interest. This typical situation arises in meteorological images. The depressions are indeed spiraling structures whose center is often in motion with respect to the earth.

We propose here a simple way to extract the divergence and vorticity structures from our dense displacement fields, while avoiding the above-mentioned problem. Based on the characterization of these structures in terms of concentrations of large divergence and vorticity, we pose the problem as one of extracting automatically such concentrations in the divergence and vorticity fields computed numerically from the estimated flow. We thus face two identical binary labelling problems in which a soft thresholding must be combined with a compactness prior on the regions to be extracted. This combination is classically achieved by casting the problem as the minimization of a twofold cost function (see Appendix III for the details).

In Figs. 14 and 15, we provide examples of divergence and vorticity concentrations extracted by this method on the WP and IR Meteosat images used in previous sections. In the two sequences, the simple labelling technique, applied to the motion fields estimated with our dedicated motion estimator, provides a sensible detection of the main divergent and rotational structures of the flow. See the convective cells in the right part of WP images in Fig. 14 and the large spiraling depression in the IR images in Fig. 15. In contrast, the application of the same detection technique on the fields provided by the generic motion estimator is doomed to fail because of the underestimation of the main divergence and vorticity components of the motion. With low values of the threshold parameter μ involved in the minimization-based detection the detected regions are much too wide with hardly any isolated blob to be identified as a structure of interest. If increasing this parameter suffices to reduce the number and span of detected regions, this might not be sufficient to recover a correct estimation of the structures. In the right image of Fig. 15 for instance, a single rotational structure is extracted, but it does not capture correctly the shape and location of the depression, which the extraction based on our dedicated model does (left image).

Hence, the divergence and the vorticity derived from the displacement fields estimated by our approach can be used simply and reliably to identify key structures of imaged fluid flows.

8   CONCLUSION

In this paper, we have presented a new method for estimating fluid flows from image sequences. This method is an extension of the standard minimization-based approaches, where a two-fold robust objective function is minimized. The two parts (i.e, the data term and the regularizer) of the novel cost function have been specifically designed to suit image sequences of fluid flows.

The data term is based on a continuity equation, as a more physically-grounded alternative to the brightness constancy assumption. In contrast to previous uses of such an alternative brightness constraint, we propose to use it in an integrated form. This makes it compatible with the standard multiresolution incremental set-up that any large displacement situation requires. Such situations occur when the imaged flow is fast (like in a number of fluid experiments) or when the temporal sampling rate is low (as with satellite images).

Concerning the regularization, we argue that only a second-order regularizer is able to preserve completely the divergence and vorticity structures of the flows. Based on the div-curl formalism, we thus introduced a robust second-order regularizer via auxiliary scalar fields which capture the divergence and vorticity of the unknown flow.

On both synthetic images and real satellite images of different natures, the merit of the two ingredients have been demonstrated. We showed in particular on the Meteosat images that the dedicated objective function made up of these two terms allows the estimation of complex flows with multiple divergence and vorticity structures. These structures, which are of primary importance for analyzing and describing fluid flows, are much better recovered by the proposed technique than by generic robust techniques which work very well on other types of images.

Based on the displacement fields provided by the proposed technique, by-products of interest to experts (e.g., meteorologists) can be obtained easily and reliably. We provided two examples of such post-processing with the detection of the main divergence and vorticity concentrations and the reconstruction of long trajectories. Similar tasks, but with a higher semantics, can now be envisaged. The automatic detection, classification, and tracking of various types of clouds in meteorology is such an example.

In view of the number of domains where the analysis of imaged fluids is important, we believe that the motion estimation technique we have presented constitutes a valuable tool. Further improvements probably lie in the incorporation of more physics. When more precise information is available, such as a model of the image formation process, an experimental knowledge of the boundary conditions, a theoretical prior on the aspect and behavior of the divergence and vorticity structures, or a set of *in situ* measurements, the cost function should be modified to capture it.

APPENDIX I

HIERACHICAL MINIMIZATION WITH REGARDS TO THE DISPLACEMENT INCREMENT FIELD

For better and faster convergence, the mimimization with respect to the incremental displacement field $\triangle d$ at a given resolution is performed using the coarse-to-fine technique introduced in [32]. As the alternate minimization iterations proceed at a given resolution, the increment field is constrained to be piecewise-parametric relatively to an image partition which becomes finer and finer. More precisely, at "grid" level ℓ, the pixel grid is partitionned in blocks B_n^ℓ of size $2^\ell \times 2^\ell$ and the increment field is constrained to satisfy

$$\Delta d_s = P^\ell(s)\theta_n^\ell, \ \forall n, \ \forall s \in B_n^\ell, \tag{28}$$

where $P^\ell(s)$ is a matrix depending on the chosen parameterization at level ℓ and θ_n^ℓ is the parameter vector for block B_n^ℓ. Until the penultimate grid level, the affine parameterization is chosen, i.e:

$$\forall \ell > 0, P^\ell(s = (x,y)) = \begin{bmatrix} 1 & x & y & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & x & y \end{bmatrix} \text{and } \theta_n^\ell \in \mathbb{R}^6. \tag{29}$$

The last level $\ell = 0$ corresponds to the pixel grid itself: When it is finally reached, no parametric constraint applies anymore.

At level ℓ, the cost terms to be minimized with respect to Δd must be rewritten in function of the parameter vectors (θ_n^ℓ) which defines Δd via (28). This symbolic rewriting is explained in [32]. It yields a quadratic cost $\frac{1}{2}\theta^{\ell T} A^\ell \theta^\ell - b^{\ell T}\theta^\ell$, where the matrix A^ℓ and the vector b^ℓ are function of the three sets of variables z^d, ξ, and ζ.

Using this technique, the synopsis presented at the end of Section 5 is replaced by:

- $\tilde{d} \equiv 0$
- From coarsest to finest resolution do
  - From $\ell = \ell_{\max}$ to 0 do
  - Until convergence do

  $$\theta^\ell \leftarrow GS\left(\frac{1}{2}\theta^{\ell T} A^\ell \theta^\ell - b^{\ell T}\theta^\ell\right)$$

  Δd is defined by (28)
  $z^d, \ \xi, \ z^\xi, \ \zeta, \ z^\zeta$ are updated as in original synopsis

- If not finest resolution, pass $\tilde{d} \leftarrow \tilde{d} + \Delta d$ to next resolution.
- If finest resolution, $d \leftarrow \tilde{d} + \Delta d$.

APPENDIX II

NUMERICAL INTEGRATION FOR TRAJECTORY RECONSTRUCTION

Knowing the dense displacement fields $d(t_0 + n\Delta t) = \{d(x, t_0 + n\Delta t), x \in \Omega\}, \ n = 0, \cdots, N-1$ for a sequence of $N + 1$ images, we aim to reconstruct the trajectory $x(t)$ of the point located at x_0 in the first image. Approximating the velocity $\frac{dx(t)}{dt}$ by $\frac{d(x(t),t)}{\Delta t}$, we have to solve the following ordinary differential equation with initial condition:

$$\begin{cases} \frac{d\boldsymbol{x}(t)}{dt} = \frac{\boldsymbol{d}(\boldsymbol{x}(t),t)}{\Delta t} \\ \boldsymbol{x}(t_0) = \boldsymbol{x}_0. \end{cases} \qquad (30)$$

Runge-Kutta methods are standard numerical techniques to solve this problem in a robust way. They work by approximating the integral

$$\underbrace{\boldsymbol{x}(t_0 + (n+1)\Delta t)}_{\boldsymbol{x}_{n+1}} - \underbrace{\boldsymbol{x}(t_0 + n\Delta t)}_{\boldsymbol{x}_n} =$$
$$\int_{t_0+n\Delta t}^{t_0+(n+1)\Delta t} \frac{d\boldsymbol{x}(t)}{dt} dt = \int_{t_0+n\Delta t}^{t_0+(n+1)\Delta t} \frac{\boldsymbol{d}(\boldsymbol{x}(t),t)}{\Delta t} dt. \qquad (31)$$

A fourth-order approximation gives the classical result:

$$\boldsymbol{x}_{n+1} = \boldsymbol{x}_n + \frac{1}{6}\left(\boldsymbol{d}_n^{(1)} + 2\boldsymbol{d}_n^{(2)} + 2\boldsymbol{d}_n^{(3)} + \boldsymbol{d}_n^{(4)} \right), \qquad (32)$$

where

$$\begin{cases} \boldsymbol{d}_n^{(1)} &= \boldsymbol{d}(\boldsymbol{x}_n, t_0 + n\Delta t) \\ \boldsymbol{d}_n^{(2)} &= \boldsymbol{d}(\boldsymbol{x}_n + \frac{\boldsymbol{d}_n^{(1)}}{2}, t_0 + n\Delta t + \frac{\Delta t}{2}) \\ \boldsymbol{d}_n^{(3)} &= \boldsymbol{d}(\boldsymbol{x}_n + \frac{\boldsymbol{d}_n^{(2)}}{2}, t_0 + n\Delta t + \frac{\Delta t}{2}) \\ \boldsymbol{d}_n^{(4)} &= \boldsymbol{d}(\boldsymbol{x}_n + \boldsymbol{d}_n^{(3)}, t_0 + n\Delta t + \Delta t). \end{cases}$$

Since displacement fields $\boldsymbol{d}$ are only available at pixel locations and at times $t_0 + n\Delta t$, we have to resort to a spatial and temporal interpolation, which we chose bilinear.

APPENDIX III

MINIMIZATION-BASED DIVERGENCE AND CURL SEGMENTATION

Given the estimated displacement field $\boldsymbol{d}$, we can numerically compute its divergence and vorticity using (22) or (23). When using the second-order regularizer we introduced, regularized versions of these two scalar fields are made available at the same time as the vector field itself.

In either case, denote ξ and ζ the scalar fields from which the significant concentrations of divergence and vorticity have to be identified. We now explain briefly how we deal with this detection problem in the case of the divergence. The technique for the vorticity is the same, with ζ replacing ξ in what follows.

We must assign to each pixel $s \in S$ a binary label σ_s indicating whether it belongs to a divergence structure ($\sigma_s = +1$) or not ($\sigma_s = -1$). This decision should be based on two criteria: The absolute value of ξ should be large in divergent structures and these structures should be compact and at least a few pixels in size. This two-fold requirement is classically captured by designing an ad hoc Markov-type cost function to be minimized. We used:

$$\sum_{s \in S} \frac{\sigma_s(\mu - |\xi_s|)}{\mu + |\xi_s|} - \beta \sum_{<s,r> \in C} \sigma_s \sigma_r, \qquad (33)$$

where μ is a "soft" threshold on the divergence and β is a positive parameter which controls the relative importance of the Ising regularizer.

The minimization of this global cost is conducted iteratively according to the deterministic Iterated Conditional Modes (ICM) algorithm [8]: All pixels are visited in turn a number of times and at each step, only the assignment of the current pixel s is updated so as to maximize the decrease of the cost function. This local update only involves the neighborhood $N(s)$ of the current pixel. Simple calculations provide the following update rule:

$$\sigma_s = \begin{cases} +1 & \text{if } \frac{|\xi_s| - \mu}{|\xi_s| + \mu} + \beta \sum_{r \in N(s)} \sigma_r > 0 \\ -1 & \text{otherwise.} \end{cases} \qquad (34)$$

The initialization is provided by the minimizer of the first sum in (33) only, i.e., the binary labelling associated with the thresholding of $|\xi|$ at level μ.

REFERENCES

[1] R. Adrian, "Particle Imaging Techniques for Experimental Fluid Mechanics," *Annal Rev. Fluid Mechanics,* vol. 23, pp. 261-304, 1991.
[2] L. Alvarez, J. Weickert, and J. Sànchez, "Reliable Estimation of Dense Optical Flow Fields with Large Displacements," *Int'l J. Computer Vision,* vol. 39, no. 1, pp. 41-56, 2000.
[3] A. Amini, "A Scalar Function Formulation for Optical Flow," *Proc. European Conf. Computer Vision,* pp. 125-131, 1994.
[4] L. Bannehr, R. Rohn, and G. Warnecke, "A Functionnal Analytic Method to Derive Displacement Vector Fields from Satellite Image Sequences," *Int'l J. Remote Sensing* vol. 17, no. 2, pp. 383-392, 1996.
[5] J. Barron, D. Fleet, and S. Beauchemin, "Performance of Optical Flow Techniques," *Int'l J. Computer Vision,* vol. 12, no. 1, pp. 43-77, 1994.
[6] S. Beauchemin and J. Barron"The Computation of Optical Flow," *ACM Computing Survey,* vol. 27, no. 3, pp. 433-467, 1995.
[7] D. Béréziat, I. Herlin, and L. Younes, "A Generalized Optical Flow Constraint and its Physical Interpretation," *Proc. Conf. Computer Vision Pattern Recognition,* vol. 2, pp. 487-492, 2000.
[8] J. Besag, "On the Statistical Analysis of Dirty Pictures," *J. Royal Statistical Soc.,* vol. 48, no. 3, pp. 259-302, 1986.
[9] M. Black, "Recursive Nonlinear Estimation of Discontinuous Flow Fields," *Proc. European Conf. Computer Vision,* pp. 138-145, 1994.
[10] M. Black and A. Rangarajan, "On the Unification of Line Processes, Outlier Rejection and Robust Statistics with Applications in Early Vision," *Int'l J. Computer Vision,* vol. 19, no. 1, pp. 75-104, 1996.
[11] I. Cohen and I. Herlin, "Non Uniform Multiresolution Method for Optical Flow and Phase Portrait Models: Environmental Applications," *Int'l J. Computer Vision,* vol. 33, no. 1, pp. 29-49, 1999.
[12] S. Das Peddada and R. McDevitt, "Least Average Residual Algorithm (LARA) for Tracking the Motion of Arctic Sea Ice," *IEEE Trans. Geoscience and Remote Sensing,* vol. 34, no. 4, pp. 915-926, 1996.
[13] J.M. Fitzpatrick, "A Method for calculating Velocity in Time Dependent Images Based on the Continuity Equation," *Proc. Conf. Computer Vision Pattern Recognition,* pp. 78-81, 1985.
[14] J.M. Fitzpatrick, "The Existence of Geometrical Density-Image Transformations Corresponding to Object Motion," *Computer Vision, Graphics, Image Proc.,* vol. 44, no. 2, pp. 155-174, 1988.
[15] J.M. Fitzpatrick and C.A. Pederson, "A Method for Calculating Fluid Flow in Time Dependant Density Images," *Electronic Imaging,* vol. 1, pp. 347-352, 1988.
[16] R.M. Ford, "Critical Point Detection in Fluid Flow Images Using Dynamical System Properties," *Pattern Recognition,* vol. 30, no. 12, pp. 1991-2000, 1997.
[17] R.M. Ford and R. Strickland, "Representing and Visualizing Fluid Flow Images and Velocimetry Data by Nonlinear Dynamical Systems," *Graphics Model Image Proc.,* vol. 57, no. 6, pp. 462-482, 1995.
[18] R.M. Ford, R. Strickland, and B. Thomas, "Image Models for 2D Flow Visualization and Compression," *Graphics Model Image Proc.,* vol. 56, no. 1, pp. 75-93, 1994.
[19] D. Geman and G. Reynolds, "Constrained Restoration and the Recovery of Discontinuities," *IEEE Trans. Pattern Analysis Machine Intelligence,* vol. 14, no. 3, pp. 367-383, 1992.
[20] S. Gupta and J. Prince, "Stochastic Models for Div-Curl Optical Flow Methods," *Signal Proc. Letters,* vol. 3, no. 2, pp. 32-34, 1996.
[21] H. Hausseker and D. Fleet, "Computing Optical Flow with Physical Models of Brightness Variation," *IEEE Trans. Pattern Analysis Machine Intelligence,* vol. 23, no. 6, pp. 661-673, 2001.

[22] P. Holland and R. Welsch, "Robust Regression Using Iteratively Reweighted Least-Squares," *Comm. Statisics—Theory Method,* vol. 6, no. 9, pp. 813-827, 1977.

[23] B. Horn and B. Schunck, "Determining Optical Flow," *Artificial Intelligence,* vol. 17, pp. 185-203, 1981.

[24] P. Huber, *Robust Statistics.* John Wiley & Sons, 1981.

[25] X. Ju, M.J. Black, and A.D. Jepson, "Skin and Bones: Multi-Layer, Locally Affine, Optical Flow and Regularization with Transparency," *Proc. Conf. Computer Vision Pattern Recognition,* pp. 307-314, 1996.

[26] P. Kornprobst, R. Deriche, and G. Aubert, "Image Sequence Analysis via Partial Differential Equations," *J. Math. Imaging and Vision,* vol. 11, no. 1, pp. 5-26, 1999.

[27] S. Lai and B. Vemuri, "Reliable and Efficient Computation of Optical Flow," *Int'l J. Computer Vision,* vol. 29, no. 2, pp. 87-105, 1998.

[28] R. Larsen, K. Conradsen, and B.K. Ersboll, "Estimation of Dense Image Flow Fields in Fluids," *IEEE Trans. Geoscience and Remote Sensing,* vol. 36, no. 1, pp. 256-264, 1998.

[29] H. Mass, A. Stefanidis, and A. Gruen, "Feature Tracking in 3-D Fluid Tomography Sequences," *Proc. Int'l Conf. Image Processing,* pp. 530-534, Sept. 1992.

[30] M. Maurizot, P. Bouthemy, and B. Delyon, "2D Fluid Motion Analysis from a Single Image," *IEEE Int'l Conf. Computer Vision and Pattern Recognition,* pp. 184-189, 1998.

[31] E. Mémin and P. Pérez, "Dense Estimation and Object-Based Segmentation of the Optical Flow with Robust Techniques," *IEEE Trans. Image Processing,* vol. 7, no. 5, pp. 703-719, 1998.

[32] E. Mémin and P. Pérez., "A Multigrid Approach for Hierarchical Motion Estimation," *Proc. Int'l Conf. Computer Vision,* pp. 933-938, 1998.

[33] E. Mémin and P. Pérez, "Fluid Motion Recovery by Coupling Dense and Parametric Motion Fields," *Proc. Int'l Conf. Computer Vision,* vol. 3, pp. 732-736, 1999.

[34] E. Mémin and T. Risset, "VLSI Design Methodology for Edge-Preserving Image Reconstruction," *Real-Time Imaging,* vol. 7, no. 1, pp. 109-126, 2001.

[35] H. Nogawa, Y. Nakajima, and Y. Sato, "Acquisition of Symbolic Description from Flow Fields: A New Approach Based on a Fluid Model," *IEEE Trans. Pattern Analysis Machine Intelligence,* vol. 19, no. 1, pp. 58-63, 1997.

[36] A. Ottenbacher, M. Tomasini, K. Holmund, and J. Schmetz, "Low-Level Cloud Motion Winds from Meteosat High-Resolution Visible Imagery," *Weather and Forecasting,* vol. 12, no. 1, pp. 175-184, 1997.

[37] C. Papin, P. Bouthemy, E. Mémin, and G. Rochard, "Tracking and Characterization of Highly Deformable Cloud Structures," *Proc. European Conf. Computer Vision,* vol. 2, pp. 428-442, 2000.

[38] B.G. Schunk, "The Motion Constraint Equation for Optical Flow," *Proc. Int'l Conf. Pattern Recognition,* pp. 20-22, 1984.

[39] J. Shukla and R. Saha, "Computation of Nondivergent Streamfunction and Irrotational Velocity Potential from the Observed Winds," *Monthly Weather Rev.* vol. 102, pp. 419-425, 1974.

[40] J. Simpson and J. Gobat, "Robust Velocity Estimates, Stream Functions and Simulated Lagrangian Drifters from Sequential Spacecraft Data," *IEEE Trans. Geosciences and Remote Sensing,* vol. 32, no. 3, pp. 479-492, 1994.

[41] D. Suter, "Motion Estimation and Vector Splines," *Proc. Conf. Computer Vision Pattern Recognition,* pp. 939-942, June 1994.

[42] J.P. Triplet and G. Roche, *Météorologie Générale.* Ecole Nationale de la Météorologie, 1971.

[43] J. Wallace and J. Foss, "The Measurement of Vorticity in Turbulent Flows," *Ann. Rev. Fluid Mechanics,* vol. 27, pp. 469-514, 1995.

[44] J. Weickert and C. Schnörr, "A Theoretical Framework for Convex Regularizers in PDE-Based Computation of Image Motion," Technical Report 13/2000 Computer Sciences Series, Univ. Mannheim, 2000.

[45] J. Weickert and C. Schnörr, "Variational Optic Flow Computation with a Spatio-Temporal Smoothness Constraint," *J. Math. Imaging and Vision,* vol. 14, no. 3, pp. 245-255, 2001.

[46] P. Wernert, W. Geissler, M Raffel, and J. Kompenhans, "Experimental and Numerical Investigations of Dynamic Stall on a Pitching Airfoil," *Am. Inst. of Aeronautics and Astronautics J.,* vol. 34, no. 5, pp. 982-989, 1996.

[47] R. Wildes, M. Amabile, A.-M. Lanzillotto, and T.-S. Leu, "Recovering Estimates of Fluid Flows From Image Sequence Data," *Computer Vision and Image Understanding,* vol. 80, pp. 246-266, 2000.

[48] L. Zhou, C. Kambhamettu, and D. Goldgof, "Fluid Structure and Motion Analysis from Multi-Spectrum 2D Cloud Images Sequences," *Proc. Conf. Computer Vision Pattern Recognition,* vol. 2, pp. 744-751, 2000.

Thomas Corpetti graduated from the National Institute of Applied Sciences (INSA) of Rennes, France, in electrical engineering, in 1999. He has been preparing a PhD degree in signal processing and telecommunications from the University of Rennes, France, since October 1999. His main research interests are in fluid and deformable motion analysis.

Étienne Mémin received the PhD degree in computer science from the University of Rennes, France, in 1993. He was an assistant professor at the University of Bretagne Sud, France from 1993 to 1999. He now holds a position at the University of Rennes, France. His research interests include computer vision, statistical models for image (sequence) analysis, fluid and medical images, and parallel algorithms for computer vision.

Patrick Pérez received the PhD degree in signal processing and telecommunications from the University of Rennes, France, in 1993. After one year as an Inria post-doctoral fellow at Brown University (Department of Applied Mathematics), he was appointed as a full-time Inria researcher. In 2000, he joined Microsoft Research, Cambridge, United Kingdom. His research interests include probabilistic models for image understanding, high dimensional inverse problems in image analysis, analysis of motion, and tracking in image sequences.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** http://computer.org/publications/dilb.