

Computing Optical Flow with Physical Models of Brightness Variation

Horst W. Haussecker, *Member, IEEE*, and David J. Fleet, *Member, IEEE*

Abstract—Although most optical flow techniques presume brightness constancy, it is well-known that this constraint is often violated, producing poor estimates of image motion. This paper describes a generalized formulation of optical flow estimation based on models of brightness variations that are caused by time-dependent physical processes. These include changing surface orientation with respect to a directional illuminant, motion of the illuminant, and physical models of heat transport in infrared images. With these models, we simultaneously estimate the 2D image motion and the relevant physical parameters of the brightness change model. The estimation problem is formulated using total least squares (TLS), with confidence bounds on the parameters. Experiments in four domains, with both synthetic and natural inputs, show how this formulation produces superior estimates of the 2D image motion.

Index Terms—Optical flow, physics-based brightness variation, total least squares.

1 INTRODUCTION

THIS paper examines the use of physical models of time-dependent brightness variation to estimate optical flow and physical parameters of a scene from image sequences. Although studied extensively [2], [16], reliable flow estimation remains difficult when brightness constancy is violated. The problems arise from the complex physical processes involved in scene illumination, surface reflection, and the transmission of radiation through surfaces and the atmosphere [17], [31], [32], [38], [41]. For example, brightness variations are caused by changing surface orientations with respect to a directional illuminant, by motion of an illuminant, and by thermal diffusion and decay in infrared images. The goal of this work is to formulate the problem of optical flow estimation in ways that take physical causes of brightness variation into account, leading to methods for estimating both the optical flow and the physical parameters of the brightness variation.

Like conventional formulations of optical flow estimation, our goal is to estimate an accurate approximation to 2D motion fields corresponding to the 3D motion of object surfaces [17], [38]. Many optical flow formulations assume brightness constancy, i.e., they estimate the 2D velocity of points of constant image brightness. For graylevel images $g(\mathbf{x}, t)$, where $\mathbf{x} = [x, y]^T$ denotes spatial position and t denotes time, this amounts to finding a path $\mathbf{x}(t)$ along which image brightness remains constant,

$$g(\mathbf{x}(t), t) = c, \quad (1)$$

for some constant c . The total temporal derivative of both sides of (1) yields the well-known *brightness constraint equation* [17]:

$$\frac{dg}{dt} = (\nabla g)^T \mathbf{v} + g_t = 0, \quad (2)$$

where $\mathbf{v} = [v_1, v_2]^T = [dx/dt, dy/dt]^T$ is the flow field that we wish to estimate, g_i denotes the partial derivative of g with respect to the coordinate $i \in \{x, y, t\}$, and $(\nabla g)^T = [g_x, g_y]$. Because (2) provides one constraint in two unknowns, it is common to combine constraints at pixels in a local neighborhood, assuming that the motion is smooth in the region [17], [20]. Constant and affine motion models have been used successfully [2], [3], [11], yielding linear constraints on the local flow. The resulting systems of equations can be solved using (weighted) least squares [20] or total least squares (TLS) [39], [40]. To further constrain the estimates, the neighborhoods can be extended into time if \mathbf{v} is smooth within local temporal windows (e.g., [4], [10], [29]).

If brightness is not conserved, then the optical flow field estimated from (2) can be a severely biased approximation to the underlying 2D motion field of interest [5], [6], [14], [26], [28], [31], [32], [38]. Causes of brightness variation include moving illumination envelopes, shadows cast on moving objects, changing surface orientation under directional illumination, and atmospheric influences in outdoor applications. Other instances occur in scientific applications where the quantitative analysis of dynamic processes is essential [15], [19]. For example, Figs. 1b and 1c illustrate the influence of brightness changes on optical flow estimation with two examples of physical transport processes in infrared images; although the surface translates in each case, the flow fields that conserve brightness converge or diverge.

Estimation of the parameters of the physical processes that cause brightness variation is often just as important as the accurate estimation of the motion. An example can be found in physical oceanography where the dynamics of

• The authors are with the Xerox Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304.
E-mail: {hhaussec, fleet}@parc.xerox.com.

Manuscript received 13 June 2000; revised 19 Jan. 2001; accepted 31 Jan. 2001.

Recommended for acceptance by M. Irani.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 112281.





Fig. 1. Illustration of errors in the optical flow estimation due to brightness changes. All three examples depict translational motion between two images. (a) The left-most image pair satisfies brightness constancy. The other examples exhibit exponential decay (b) and diffusion (c), which are common in infrared images and, hence, do not satisfy brightness constancy. The vectors shown depict the optical flow that would be consistent with brightness constancy in each case.

water surface heat patterns provides important information about the hydrodynamic processes underneath the water surface. As depicted in Fig. 2, although only the surface layer of the ocean may be visible to infrared sensors, by estimating the parameters of a physics-based model for brightness variations one can sometimes infer the properties of the complex 3D transport processes beneath the surface layer [15]. Let $T(\mathbf{X}, t, \mathbf{p}_T)$ denote the 3D temperature distribution at 3D location \mathbf{X} and time t , where \mathbf{p}_T is a parameter vector for the physical process. From a model of image formation, one can relate the 3D transport parameters to the parameters of the image brightness variation. This provides one with a generative model, $g(\mathbf{x}, t, \mathbf{p}_g)$, that expresses the time-varying behavior of image brightness in terms of 2D model parameters \mathbf{p}_g and image coordinates \mathbf{x} :



One goal of the current work is to estimate the 2D model parameters, \mathbf{p}_g . Ideally, one can then attempt to solve the inverse problem to find the 3D transport process, \mathbf{p}_T .

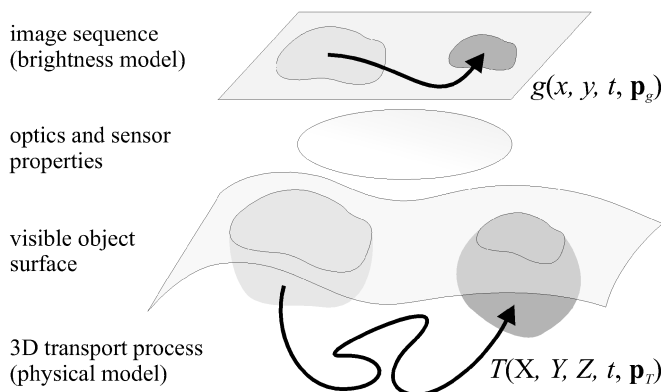


Fig. 2. The temperature patterns on the surface of a moving object, imaged by an infrared camera, are subject to complex 3D transport processes underneath the surface layer.

This paper describes a generalized framework for incorporating physical models of brightness changes into motion analysis. Brightness changes are either parameterized as time-varying analytical functions or by the differential equations that model the underlying physical processes. We only require that the brightness variation be linear in the model parameters. With this, we obtain a linear system of equations that comprise a generalization of the brightness constancy assumption.

Estimates of the model parameters can then be obtained using TLS over local spatiotemporal neighborhoods. It is important that these neighborhoods are extended in time to more than two frames. With two consecutive frames, we can only adequately constrain simple temporal models such as linear brightness change. For more complicated brightness dynamics, we need multiple frames to constrain parameter estimation. Even then not all parameters of the model are necessarily well-constrained, much like the well-known aperture problem. Accordingly, it is important that we obtain error covariance matrices for the parameter estimates to characterize parameter uncertainty [33].

After describing previous work on this topic, we outline our mathematical formulation with reference to several different sources of brightness variation, and we then describe the computational procedure. Finally, using experimental work, we show that this method produces improved optical flow estimates in several application domains. It also yields estimates of model parameters that characterize the physical processes in each case.

2 PREVIOUS WORK

Violations of brightness constancy are well-known [18], [28], [38]. One approach to coping with them has been to preprocess the image to extract image properties whose deformations through time provide reliable estimates of the flow field. One example involves band-pass filtering and contrast gain normalization to remove smooth illumination variations [1]. Another technique involves the extraction of image features, such as edges or regions, to achieve robustness with respect to large brightness variations [27], [38]. Fleet and Jepson [11], [12] proposed the tracking of band-pass phase information. Phase is stable with respect to smooth variations in illumination and smooth geometric deformations between frames. Moreover, locations of phase

instability can be detected and, therefore, ignored, making the subsequent estimation of optical flow more robust [12].

Alternatively, brightness variations have been accommodated by modifying the gradient constraint equation [6], [14], [26], [28], [32]. A general framework for this is proposed in [32], which permits changes in contrast and mean intensity. This is formulated using a multiplier and an offset field to yield constraints of the form

$$g(\mathbf{x}, t) - g(\mathbf{x} + \mathbf{v}, t + 1) = m(\mathbf{x}, t)g(\mathbf{x}, t) + c(\mathbf{x}, t). \quad (3)$$

It is certainly true that all changes between two images can be modeled by (3). Indeed, with the introduction of $m(\mathbf{x}, t)$ and $c(\mathbf{x}, t)$, there are now four unknowns at every image location (two for $\mathbf{v}(\mathbf{x}, t)$, plus $m(\mathbf{x}, t)$ and $c(\mathbf{x}, t)$), and it is, therefore, necessary to impose constraints on them, such as smoothness, in order to obtain a unique solution. However, despite the generality of (3), this formulation only models brightness changes between successive time instants; it does not allow us to discriminate different physical causes of brightness changes, or to constrain the estimation to satisfy particular physical models. With two frames, the model only captures linear brightness changes. By contrast, with the models we consider below, it is important that the neighborhoods are extended to more than two frames so higher-order brightness changes can be modeled.

In work on target tracking, Hager and Belhumeur [14] combine illumination changes and pose-dependent geometric image distortions into a parameterized model. They use robust area-based regression to fit the image to a linear combination of basis templates (eigenmodels). One disadvantage of the approach is that the basis set must be computed from the target, under varying illumination, prior to the tracking. Also, the resulting parameters specify a location in the eigenspace of training images, rather than a physical model of the brightness variation. Black et al. [5] express the change between two frames of an image sequence as a mixture of causes, including both motion and illumination effects, but they have not considered realistic time-varying physical models. Moreover, their use of mixture models, robust statistics, and the EM algorithm are computationally expensive compared to the solution developed here.

The techniques mentioned above are confined to brightness changes between two images; they do not exploit the physical nature of brightness variation over more than two frames. In contrast, the approach described in this paper generalizes the temporal brightness changes in ways governed by the underlying physical process. By confining the classes of permitted solutions to those of physical relevance, we constrain the solutions and simultaneously estimate the physical parameters of interest along with the optical flow.

Finally, this work also bears some similarity to research by Michel et al. [23] on the thermophysical interpretation of static images. Based on the principle of energy conservation, internal object composition, and object surface properties, they define object features that remain invariant under viewing and scene conditions. As such, they use physical models of heat transport to define object-specific features in




Fig. 3. Illustration of the generalized model that allows the object brightness to change within a few images. Solutions of the brightness constancy assumption are confined to the gray plane ($g(\mathbf{x}, t) - g(\mathbf{x}_0, t_0) = 0$).

infrared images. But, these models were not used for motion estimation.

Wildes et al. [41] propose a formulation for motion estimation for fluid flows and transmittant media based on a conservation of mass principle. This use of physical models is similar to the approach taken here and complements the physical models we discuss below.

3 PHYSICS-BASED BRIGHTNESS VARIATION

As a generalization of the brightness conservation equation (1), we define a path $\mathbf{x}(t)$, depicted in Fig. 3, along which brightness can change according to a parameterized function, h ; that is,

$$g(\mathbf{x}(t), t) = h(g_0, t, \mathbf{a}), \quad (4)$$

where $g_0 = g(\mathbf{x}(t_0), t_0)$ denotes the image at time 0, and $\mathbf{a} = [a_1, \dots, a_Q]^T$ denotes a Q -dimensional parameter vector for the brightness change model. Without loss of generality, we choose a parameterization such that $\mathbf{a} = \mathbf{0}$ produces the identity transformation, i.e., $h(g_0, t, \mathbf{a} = \mathbf{0}) = g_0$. While \mathbf{a} is assumed to be constant within small temporal windows, h is expressed as a function of time so that we can capture nonlinear temporal brightness changes.

The total derivative of both sides of (4) then yields a *generalized brightness change constraint equation*

$$(\nabla g)^T \mathbf{v} + g_t = f(g_0, t, \mathbf{a}), \quad (5)$$

where f is defined as

$$f(g_0, t, \mathbf{a}) = \frac{d}{dt} [h(g_0, t, \mathbf{a})]. \quad (6)$$

When brightness is conserved, as in (1), h is constant as a function of time and, therefore, $f = 0$. In this case, (5) reduces to (2).

Given constraints like that in (5), our goal is to estimate the parameters of the optical flow field \mathbf{v} , and the parameters \mathbf{a} of the physical model f . In (5), we used a constant flow model for which the flow parameters are v_1

and v_2 . But, it is straightforward to use other linear parameterized models [3], [7], [8], [13].

Finally, there are two different ways to determine the form of f . One can derive f using (6) and a known analytical form for h , or one can choose f according to the differential equations of the underlying physical process. We illustrate this in Sections 3.1, 3.2, 3.3, and 3.4 below, before returning to the general formulation in Section 3.5.

3.1 Diffusion

Infrared (IR) image sequences capture spatiotemporal thermal patterns of a scene. The heat distribution is time-varying because of object motion and because of the thermal transport processes that occur on and in the objects, such as diffusion and convection in fluid systems. Such thermal transport processes depend on physical properties of the emitting material and on the shape of the heat distribution itself at any instant. Thermal transport can cause significant variation in the brightness of infrared sequences and should be considered in order to separate the apparent variation in the heat distribution over the body from the motion of the object.

Diffusion provides a useful model for brightness changes due to thermal transport where there exist steep temperature gradients over the surface of an object. These temperature gradients diffuse according to the well-known diffusion equation:

$$f = \frac{dg}{dt} = \nabla(\mathbf{D} \nabla g), \quad \mathbf{D} = \begin{bmatrix} D_{11} & D_{12} \\ D_{12} & D_{22} \end{bmatrix}, \quad (7)$$

where \mathbf{D} is referred to as the *diffusion tensor*, a positive-definite symmetric matrix that generally depends on spatial position and time. When the diffusion tensor is shift-invariant (i.e., independent of spatial position), then (7) reduces to

$$\frac{dg}{dt} = D_{11} \frac{\partial^2 g}{\partial x^2} + 2D_{12} \frac{\partial^2 g}{\partial x \partial y} + D_{22} \frac{\partial^2 g}{\partial y^2}, \quad (8)$$

with constants D_{11} , D_{12} , and D_{22} . If the surface diffusion process is also isotropic and the object surface is frontoparallel (perpendicular to the optical axis), then the brightness variation in (8) reduces further to

$$\frac{dg}{dt} = D \left(\frac{\partial^2 g}{\partial x^2} + \frac{\partial^2 g}{\partial y^2} \right), \quad (9)$$

where D is now a scalar *diffusion constant*. As above, because f is linear in the diffusion parameters (i.e., D_{ij} and D in (8) and (9)), we have found a linear constraint equation that relates the flow \mathbf{v} , the diffusion parameters, and the image gradients.

3.2 Exponential Decay

In oceanographic applications, researchers have used lasers to induce artificial thermal distributions on the water surface. Then, from infrared image sequences, they can measure the motion of the water surface to study air-sea gas exchange [15]. However, infrared cameras only sense the thermal pattern that lies within the top 10-50 μm of the water surface. Accordingly, one must account for brightness variations that result from the transport of heat from the

water surface into deeper layers that are beyond the range of the sensors. Thermal transport from surface to deeper layers is caused by diffusion and by turbulent convection. One reason for this is that laser radiation is absorbed within only the first 10 μm of the water surface. Thus, the vertical temperature gradient is usually orders of magnitude larger than temperature gradients along the water surface. The lateral gradients are a function of the cross-sectional shape of the laser beam, which is typically about 10 mm wide. Because diffusion is a function of the heat gradient (as in (7)), vertical diffusion can dominate the thermal transport at the water surface.

Haussecker [15] showed that the net temporal temperature variation sensed by an infrared camera could be approximated by an exponential decay model. In this case, the brightness function h in (4) has the analytical form

$$h(g_0, t, \kappa) = g_0 \exp(-\kappa t). \quad (10)$$

Accordingly, the parameter vector \mathbf{a} in (4) reduces to a scalar *decay constant*, κ . From (6) and (10), it then follows that

$$f(g_0, t, \kappa) = \frac{dg}{dt} = -\kappa g_0 \exp(-\kappa t) = -\kappa g(\mathbf{x}(t), t). \quad (11)$$

This is the well-known differential equation of exponential decay. It states that the rate of change at any time is proportional to the current value. Because f is linear in κ , our brightness constraint equation is linear in the flow parameters \mathbf{v} and in the decay parameter κ .

3.3 Moving Illumination Envelope

Brightness changes caused by moving, nonuniform illumination envelopes have been considered in the 2-frame case [32]. Here, we focus on illuminants with a relatively narrow envelope, such as flashlights or spotlights. Diffuse shadows provide another case of interest when they are cast on a surface, the motion of which we wish to estimate.

We model the image as the product of an underlying surface albedo function g_c that translates with image velocity \mathbf{v} , and an illumination envelope E (surface irradiance) that translates with velocity \mathbf{u} :

$$g(\mathbf{x}, t) = g_c(\mathbf{x} - t \mathbf{v}) E(\mathbf{x} - t \mathbf{u}). \quad (12)$$

To characterize this brightness transformation, it is convenient to use a coordinate reference frame that is fixed on the underlying surface, i.e., $\mathbf{x}_r = \mathbf{x} - t \mathbf{v}$. The motion of the envelope relative to this reference frame is given by the relative motion of the envelope and the surface, $\mathbf{u}_r = \mathbf{u} - \mathbf{v}$. Then, the image brightness becomes

$$g_w(\mathbf{x}_r, t) \equiv g(\mathbf{x}_r + t \mathbf{v}, t) = g_c(\mathbf{x}_r) E(\mathbf{x}_r - t \mathbf{u}_r), \quad (13)$$

where $g_w(\mathbf{x}_r, t)$ is a warped version of $g(\mathbf{x}, t)$, for which the motion of the surface is stabilized.

To parameterize the brightness variation through time, we approximate $E(\mathbf{x}_r - t \mathbf{u}_r)$ by a Taylor series, up to second order with respect to time, about the point $(\mathbf{x}_r, 0)$:

$$E(\mathbf{x}_r - t \mathbf{u}_r) \approx E(\mathbf{x}_r) - t \nabla E^T \mathbf{u}_r + \frac{1}{2} t^2 \mathbf{u}_r^T \mathbf{H} \mathbf{u}_r, \quad (14)$$

where ∇E and \mathbf{H} are the gradient and Hessian of E at $(\mathbf{x}_r, 0)$. Substituting (14) into (13) yields the brightness function given by

$$h = g_c \left[E - t \nabla E^T \mathbf{u}_r + \frac{1}{2} t^2 \mathbf{u}_r^T \mathbf{H} \mathbf{u}_r \right], \quad (15)$$

where, for notational convenience, E , ∇E , and \mathbf{H} are evaluated at $(\mathbf{x}_r, 0)$. Then, from (6), the brightness change, f , is given by

$$f = a_1 + a_2 t, \quad (16)$$

where $a_1(\mathbf{x}_r) = -g_c \nabla E^T \mathbf{u}_r$ and $a_2(\mathbf{x}_r) = g_c \mathbf{u}_r^T \mathbf{H} \mathbf{u}_r$. The brightness change f is linear in the parameters $\mathbf{a} = [a_1, a_2]^T$.

If the moving envelope can be approximated by (14), then the quadratic time-varying model in (15) is expected to reduce the bias in optical flow estimates. By comparison, if the envelope were nearly linear in the spatial coordinates, i.e., if $|t \mathbf{u}_r^T \mathbf{H} \mathbf{u}_r| \ll |\nabla E^T \mathbf{u}_r|$, then a first-order temporal model for h and, hence, a constant model for f , would suffice. In either case, solving for the polynomial model coefficients in \mathbf{a} at a single instant of time does not allow us to separate the exact shape of E from its motion \mathbf{u}_r . However, it does provide information about the combined impact of both.

The derivation of the brightness model (16) is based on the assumption that the surface albedo function exhibits small-scale texture while the illumination envelope shows only smooth spatial variation. We have also assumed that the motion of the textured surface is the relevant motion to be estimated rather than the motion of the illuminant. The determination of both motions simultaneously is outside the scope of this paper (however, see [9]).

3.4 Changing Surface Orientation

In the last case, we address concerns over brightness variations caused by surface rotation under directional illumination. As is well-known, even Lambertian surfaces exhibit brightness changes if the angle between the surface normal, \mathbf{n} , and the direction of incident illumination, \mathbf{r} , changes with time. Although one might attempt to evenly illuminate a scene to avoid these effects, directional illumination cannot be avoided in most cases. Examples include outdoor scenes in direct sunlight, indoor illumination through a single window, and exploration of dark scenes using a collimated light source.

Given a combination of ambient illumination and a fixed, distant, point light source from direction \mathbf{r} (where $\|\mathbf{r}\| = 1$), the surface radiance from a Lambertian surface with unit normal \mathbf{n} can be expressed as $c_0 + c_1 \max(0, \mathbf{n}^T \mathbf{r})$, where c_0 is the ambient component and c_1 is proportional to the surface albedo and the intensity of the illuminant. Here, the inner product is bounded below by zero to account for cases when $\mathbf{n}^T \mathbf{r} < 0$ in which case the illuminant is not visible from the surface.

If we assume that the body is rotating about an axis \mathbf{r}_0 at a rate of ω radians per unit time, then we can write the surface normal at time t as $\mathbf{n}_t = \mathbf{R}_t \mathbf{n}_0$, where \mathbf{n}_0 is the surface normal at time 0 and \mathbf{R}_t is the 3D rotation matrix that maps the normal at time 0 to the normal at time t (through a rotation of ωt radians). Then, the time-varying

radiance becomes $c_0 + c_1 \max(0, \mathbf{n}_0^T \mathbf{R}_t^T \mathbf{r})$. One can show that the extent to which the radiance changes with time depends on the angle between the light source direction \mathbf{r} and the axis of rotation \mathbf{r}_0 . To see this, let $\mathbf{r} = \alpha \mathbf{r}_0 + \beta \mathbf{r}_1$, where $\|\mathbf{r}_0\| = \|\mathbf{r}_1\| = 1$, $\alpha = \mathbf{r}^T \mathbf{r}_0$, and $\alpha^2 + \beta^2 = 1$. With this and assuming that initially $\mathbf{n}_0^T \mathbf{r}_0 > 0$, it follows that the time varying radiance is

$$L(t) = c_0 + c_1 (\alpha \mathbf{n}_0^T \mathbf{r}_0 + \beta \max(0, \mathbf{n}_0^T \mathbf{R}_t^T \mathbf{r}_1)). \quad (17)$$

Finally, with the use of trigonometric identities on the elements of the rotation matrix and the unit normal vectors in (17), one can show that this reduces to the general form of

$$L(t) = \kappa_1 + \kappa_2 \max(0, \cos(\omega t + \phi)), \quad (18)$$

where $\kappa_1 = c_0 + c_1 \alpha \mathbf{n}_0^T \mathbf{r}_0$, $\kappa_2 = c_1 \beta$, and ω is the frequency of the temporal modulation.

One can see from the form of (18) that significant brightness oscillations occur when the angle between the surface normal and the light source direction changes significantly. For an object rotating on a level turntable and illuminated from the side, the axis of rotation is normal to the light source direction (i.e., $\alpha = 0$ and $\beta = 1$). In this case, the surfaces will move through grazing angles with respect to the light source to produce large brightness variations. Conversely, if the object on the turntable is illuminated from above, so the light source direction and the axis of rotation are parallel (i.e., $\alpha = 1$ and $\beta = 0$), then we expect surface radiance to remain invariant through time, regardless of the surface normal.

With respect to estimation, it is obvious from (18) that radiance, and, hence, image brightness, is not linear in parameters of interest $[\omega, \phi]$. However, all possible angles between visible (opaque) surfaces and the illumination direction are confined to the interval $[-\pi/2, \pi/2]$. Within this interval, the cosine can be approximated by a second-order polynomial, which provides a brightness function that is linear in its parameters, that is,

$$h(g_0, t, a_1, a_2) = (1 + a_1 t + a_2 t^2) g_0 + O(t^3), \quad (19)$$

where the parameters a_1 and a_2 are functions of ω and ϕ . Using (19) and ignoring terms higher than second-order in t , gives us the following approximation for f :

$$f \approx g_0 a_1 + 2g_0 a_2 t, \quad (20)$$

which is linear in the parameters. Once an approximate parameter set, a_1 and a_2 , has been estimated using (20) in (5), the quadratic approximation of h can be fitted to (18) to estimate the parameters $[\omega, \phi]$.

3.5 Generalized Formulation

In Sections 3.1, 3.2, and 3.3, the parametric brightness change models were linear in the parameters \mathbf{a} . For the cosinusoidal brightness change in Section 3.4, we approximated the brightness function, h , with a second-order polynomial in t . In general, all smoothly varying functions can be locally expanded by a Taylor series and approximated by a polynomial of order Q and, therefore, we can assume that h is analytic in a set of parameters $\mathbf{a} = [a_1, \dots, a_Q]^T$ without loss of generality. Accordingly,

remembering that $h(g_0, t, \mathbf{a} = 0) = g_0$, we can expand h as a Taylor series about $\mathbf{a} = 0$:

$$h(g_0, t, \mathbf{a}) = g_0 + \sum_{k=1}^Q a_k \frac{\partial h}{\partial a_k}. \quad (21)$$

Using (21), we can express f , the temporal brightness variation defined in (6), as

$$f(g_0, t, \mathbf{a}) = \frac{dh(g_0, t, \mathbf{a})}{dt} = \sum_{k=1}^Q a_k \frac{d}{dt} \frac{\partial h}{\partial a_k}, \quad (22)$$

where a_k is assumed to be constant within local windows of spatiotemporal support. As h is analytic in \mathbf{a} , we can exchange the order of differentiation to obtain the general form of our constraints:

$$f(g_0, t, \mathbf{a}) = \sum_{k=1}^Q a_k \frac{\partial f}{\partial a_k} = (\nabla_{\mathbf{a}} f)^T \mathbf{a}. \quad (23)$$

That is, f can be written as a scalar product of the parameter vector \mathbf{a} and a vector containing the partial derivatives of f with respect to the parameters a_k .

This formulation constitutes a generic extension of the brightness change constraint equation to parameterized models of brightness variation. The choice of an appropriate brightness change model remains crucial and must be designed for the application domain; in general, the physics of image formation provides the correct physical model of brightness variation. This paper does not address the problem of automatically selecting an appropriate model.

4 COMPUTATIONAL FRAMEWORK

In each of the above formulations, we obtain linear constraints that relate the optical flow parameters, the brightness change parameters, and the image measurements. The form of the constraints, assuming a translational model of optical flow, can be expressed as

$$\mathbf{c}^T \mathbf{p}_h = 0, \quad (24)$$

where

$$\mathbf{c} = [(\nabla_{\mathbf{a}} f)^T, (\nabla g)^T, g_t]^T, \quad (25)$$

$$\mathbf{p}_h = [\mathbf{p}^T, 1]^T, \quad (26)$$

$$\mathbf{p} = [-\mathbf{a}^T, \mathbf{v}^T]^T. \quad (27)$$

Here, the $(Q+2)$ -dimensional vector \mathbf{p} contains the parameters of interest, namely, the flow field parameters and the brightness parameters of h , and \mathbf{p}_h denotes its *homogeneous* counterpart (26). The $(Q+3)$ -dimensional vector \mathbf{c} combines the image derivative measurements and the gradient of f with respect to the parameters \mathbf{a} . This form of constraint is easily generalized from a constant flow model to higher-order parameterized motion models [3], [7], [8], [13].

Equation (24) provides one constraint on several unknowns. To further constrain the estimation of \mathbf{p} , we assume that \mathbf{p} remains constant within a local space-time

neighborhood. Assuming N independent constraints at pixels within the region, ignoring measurement noise, we obtain a linear system of equations

$$\mathbf{G} \mathbf{p}_h = 0, \quad (28)$$

where $\mathbf{G} = [\mathbf{c}_1, \dots, \mathbf{c}_N]^T$. Assuming isotropic Gaussian noise in all measurements (the elements of \mathbf{G}), including spatial and temporal image gradients, the maximum likelihood estimate for the unknown parameter vector is given by the *total least squares* (TLS) solution [7], [13], [33], [34], [37], [40]. The total least squares (TLS) method seeks to minimize $\|\mathbf{G} \mathbf{p}_h\|^2$, subject to the constraint that $\mathbf{p}_h^T \mathbf{p}_h = 1$ to avoid trivial solutions. As is well-known, this formulation yields a solution, $\hat{\mathbf{p}}_h$, given by the right singular direction associated with the smallest singular value of \mathbf{G} [24], [33], [37]. Moreover, one must normalize $\hat{\mathbf{p}}_h$, so that its last component is unity to obtain \mathbf{p} .

The algorithmic aspects of TLS parameter estimation have been explored in some detail [24], [25], [33], [36], [37]. As mentioned above, if the measurement noise (in each \mathbf{c}_i) is independent, isotropic and Gaussian, then TLS yields an unbiased, maximum-likelihood estimate [37], [33], and the smallest singular value provides an estimate of the noise variance. In practice, it is generally the case that errors in the measurement of temporal image derivatives are larger than those in the measurement of spatial derivatives. Similarly, derivative measurement noise at neighboring pixels in an image is often correlated [30]. Therefore, the noise might not be i.i.d. (independent and identically distributed). In this case, the direction of the singular vectors is biased and, in severe cases, the order of the singular values can be swapped. This problem can be handled by *renormalization* or *equilibration* techniques, as discussed in [21], [25], or with the approach proposed by [22]. In this paper, we have not performed a careful examination of noise variances, nor have we applied renormalization.

Finally, as in conventional optical flow techniques, our confidence in the parameter estimates depends on the distribution of constraints and the magnitude of the noise. Poor conditioning of the solution, resulting from near degeneracy of the constraints, is often referred to as the aperture problem. Here, we capture estimator reliability with an approximation to the error covariance matrix (Cramer-Rao lower bound). Nestares et al. [33] derived a Cramer-Rao bound for total least squares; beginning with an error-in-variables formulation of the likelihood [37], they integrate out the true values of the measurements, given a conditional Gaussian prior over the noiseless measurements. This yields the TLS likelihood function, from which the Cramer-Rao lower bound can be determined from the Hessian of the log-likelihood, \mathbf{H} . Following [33], we approximate the error covariance matrix by $\Sigma = \mathbf{H}^{-1}$, where the Hessian is evaluated at the TLS estimate $\hat{\mathbf{p}}_h = [\hat{\mathbf{p}}^T, 1]^T$:

$$\mathbf{H} = \frac{\gamma}{\sigma_n^2 \|\hat{\mathbf{p}}_h\|^2} \left(\mathbf{M} - \frac{1}{\|\hat{\mathbf{p}}_h\|^2} (\hat{\mathbf{p}}_h^T \mathbf{C} \hat{\mathbf{p}}_h) \mathbf{I}_{Q+2} + \frac{4}{\|\hat{\mathbf{p}}_h\|^4} \left[(\hat{\mathbf{p}}_h^T \mathbf{C} \hat{\mathbf{p}}_h) \hat{\mathbf{p}} - \|\hat{\mathbf{p}}_h\|^2 (\mathbf{M} \hat{\mathbf{p}} + \mathbf{A}^T \mathbf{b}) \right] \hat{\mathbf{p}}^T \right), \quad (29)$$




Fig. 4. Rotating sphere under directional illumination. (a) and (b) Frames 1 and 5. (c) and (d) Difference between the ground truth image motion and the optical flow estimated using brightness constancy (c) and the quadratic temporal brightness model (d). Also, shown are the corresponding uncertainty ellipses in each case.

where $\mathbf{C} \equiv \mathbf{G}^T \mathbf{G}$, \mathbf{A} contains all but the last column of \mathbf{G} , \mathbf{b} is the last column of \mathbf{G} , $\mathbf{M} \equiv \mathbf{A}^T \mathbf{A}$, and \mathbf{I}_{Q+2} denotes a $(Q+2) \times (Q+2)$ identity matrix. Also, $\gamma = \sigma^2 / (\sigma_n^2 + \sigma^2)$, where σ^2 denotes the signal power in the measurements (i.e., the expected magnitude of the measurements in \mathbf{G}) and $\sigma_n^2 \mathbf{I}_{Q+3}$ is the covariance of the IID Gaussian noise.

5 EXPERIMENTAL RESULTS

We have applied the method described above to synthetic image sequences for which ground truth is available and to natural image sequences. The application domains include scientific applications with infrared images and computer vision applications with natural lighting. In all cases, estimation was based on space-time neighborhoods of support with Gaussian weights centered about the center of the space-time support window. The temporal support extended between 5 and 9 frames, and the spatial support was varied between 5×5 pixels and 32×32 pixels for different experiments. As spatial and temporal gradient operators, we used optimized differential operators proposed in [35].

The experimental results include estimates of optical flow and of the brightness change parameters. Error covariances (29) are also estimated. Below, we show optical flow estimates for natural images and error vector fields for synthetic sequences for which ground truth is available. In each case, we also show error ellipses that satisfy $\mathbf{e}^T \Sigma_v^{-1} \mathbf{e} = 4.6$, where Σ_v is the 2D error covariance submatrix for v_1 and v_2 . These ellipses capture approximately 90 percent of the expected errors. Finally, for display convenience, we don't show estimates at those image locations for which the L_2 norm of the error covariance is extremely large.

5.1 Changing Surface Orientation

Figs. 4a and 4b show two frames from a computer generated image sequence of a randomly textured 3D sphere under directional illumination. The sphere was rendered to be illuminated under an angle of 45° with respect to viewing direction, and it was rotating about a vertical axis through its center. The angular velocity of the sphere was varied in several experiments, staying within the spatiotemporal sampling limits imposed by the scale of the spatial texture (this allowed us to avoid the need for a coarse-to-fine

estimation strategy in the current experiments). In some cases, we rotated the light source in one direction while rotating the sphere in the other. This allowed us to increase the rate of brightness change while keeping image velocities reasonable small. The temporal brightness function was modeled using either brightness constancy (2), a linear brightness change model, or the quadratic approximation to the sinusoidal relationship in (19).

As expected, we found that the constant brightness model performed poorly compared to the linear and quadratic models which account for brightness variations. For slow rotational speeds, the linear and quadratic models produced similar results. With faster rotational velocities (or with the light source and sphere rotated in opposite directions), larger brightness changes are produced which exaggerate differences between the linear and quadratic models. Figs. 4c and 4d show optical flow results with uncertainty ellipses for the constant and the quadratic brightness model, respectively. Because the synthetic sequence provides ground truth, we plot the difference between the estimated flow and ground truth. One can see that the optical flow field estimated with the quadratic model is more accurate than that found with the constant model. The errors are largest where the illumination is near grazing angles.

Another example of brightness variation caused by changing surface orientation is shown in Fig. 5. Here, a human arm was illuminated by natural light from a window in a workplace environment. The arm is rotating about its main axis towards the left while slightly translating to the right, so that the resultant flow field is close to stationary. Fig. 5e shows the optical flow field estimated with the quadratic model. Fig. 5d shows optical flow estimates obtained with the brightness constancy model. By comparison, one can see by inspection that the quadratic model produces more plausible estimates than the constant model. Smoothness in Figs. 5d and 5e is in part due to the fact that we estimate motion at every pixel while using filters with spatial support of 32×32 pixels.

5.2 Moving Illuminant

Our next experiment involved moving illuminants. To generate a synthetic example, we multiplied a sample of smoothed white noise with a Gaussian envelope. The Gaussian simulates the illumination envelope, and it is the motion of the noise signal that we wish to estimate.

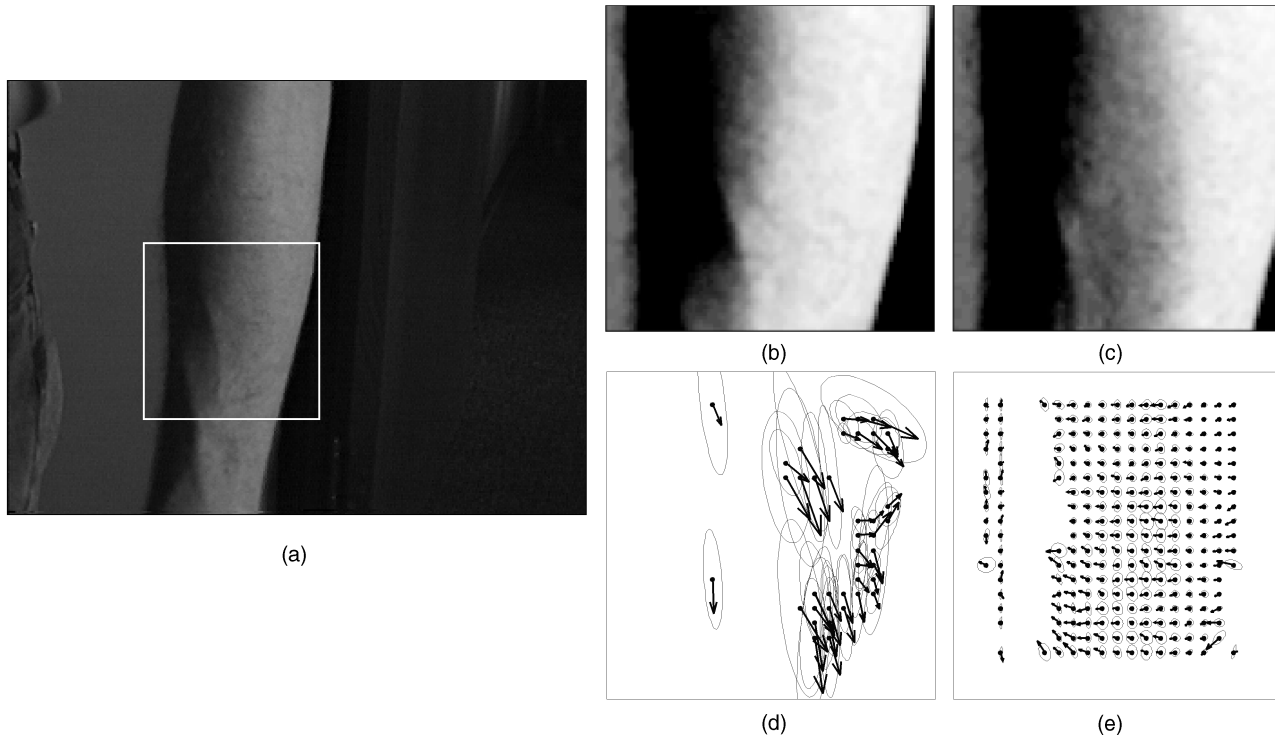


Fig. 5. A human arm, under directional illumination, is rotating about its main axis towards the left while slightly translating to the right. (a) Full image indicating the area displayed in the right images. (b) and (c) Frames 1 and 3. (d) and (e) Optical flow estimates and uncertainty ellipses for the constant and quadratic temporal brightness models. The estimates are shown only at those locations where the L_2 norm of the error covariance matrix was below a threshold.

Following the formulation in Section 3.3, both signals translated with constant velocity.

Fig. 6 shows results for the different temporal brightness models, namely, the constant (2), linear (3), and quadratic (15) models. As above, with the synthetic sequence, we plot the difference between the estimated flow and the ground truth. Due to the shape of the illumination envelope, no part of the moving pattern remains at constant brightness along its path. As a consequence, the constant brightness model fails to predict the true velocity (Fig. 6c). The linear model correctly accounts for brightness changes in regions where the instantaneous motion of the illumination envelope is mainly parallel to its level contours (Fig. 6d). In these regions, the parameter a_1 in (16) exceeds the parameter a_2 , and the temporal brightness changes are nearly linear. However, the linear model fails in regions of high positive or negative values of the combined motion/curvature parameter a_2 , which correspond to the dark and bright regions in Fig. 6f. The quadratic brightness change model allows us to accurately estimate the motion of the pattern (Fig. 6e).

The next experiment comprised an individual walking over a dark textured carpet, where the only illumination was a flashlight carried by the subject. Two frames from the sequence are shown in Figs. 7a and 7b. Negahdaripour [32] investigated an analogous instance of a nonuniform illuminant in an underwater scene. He showed that the linear brightness change model (3) performs well in estimating the correct optical flow as compared to the bias that occurs with the constant brightness model. However, in his case, the illuminant was stationary with respect to the

camera. We find the same result; for small $\mathbf{u}_r \approx -\mathbf{v}$, the linear model performs as well as a quadratic model. For faster motions $\|\mathbf{u}_r\| \gg \|\mathbf{v}\|$, the quadratic terms in (15) become more significant, so the quadratic model yields better results (compare Figs. 7d and 7e).

5.3 Exponential Decay

The next two examples involve thermal transport processes in infrared image sequences. Initial experiments with exponential decay were conducted with a synthetic image sequence of a noisy, Gaussian-shaped brightness distribution that is subject to exponential decay while moving with a velocity $\mathbf{v} = [-1, 0]^T$ pixels/frame. The decay constant was $\kappa = 0.3$ frames⁻¹. Parameter estimation was performed first with the brightness constancy model and then with the exponential decay model (11). Results are shown in Fig. 8.

As expected, the constant brightness model performs poorly. The error vector field, \mathbf{e} , in Fig. 8d shows that the constant brightness model interprets the decaying heat spot as a convergent flow field. The most severe errors occur in regions of highest brightness values where the decay rate is highest according to (11).

When we used the correct physical transport model, the optical flow field is estimated more accurately as shown by the difference between the true flow field and the estimated flow in Fig. 8e. In addition to the improved accuracy in the flow estimates, the exponential decay model yields an estimate of the decay constant κ (Fig. 8f). Within the image region where the L_2 norm of the error covariance matrix shows the problem to be reasonably well conditioned (i.e.,




Fig. 6. Synthetic moving illumination envelope. The underlying texture and the envelope move with velocities $\mathbf{v} = [1, 1]^T$ and $\mathbf{u}_E = [-2, -2]^T$ pixels/frame, respectively. (a) and (b) Frames 2 and 6. (c), (d), and (e) Difference between the true and estimated flow, with the uncertainty ellipses, for the constant, linear, and quadratic temporal brightness models, respectively. (f) Estimated parameter a_2 of the combined curvature and motion of the envelope (see (16)).

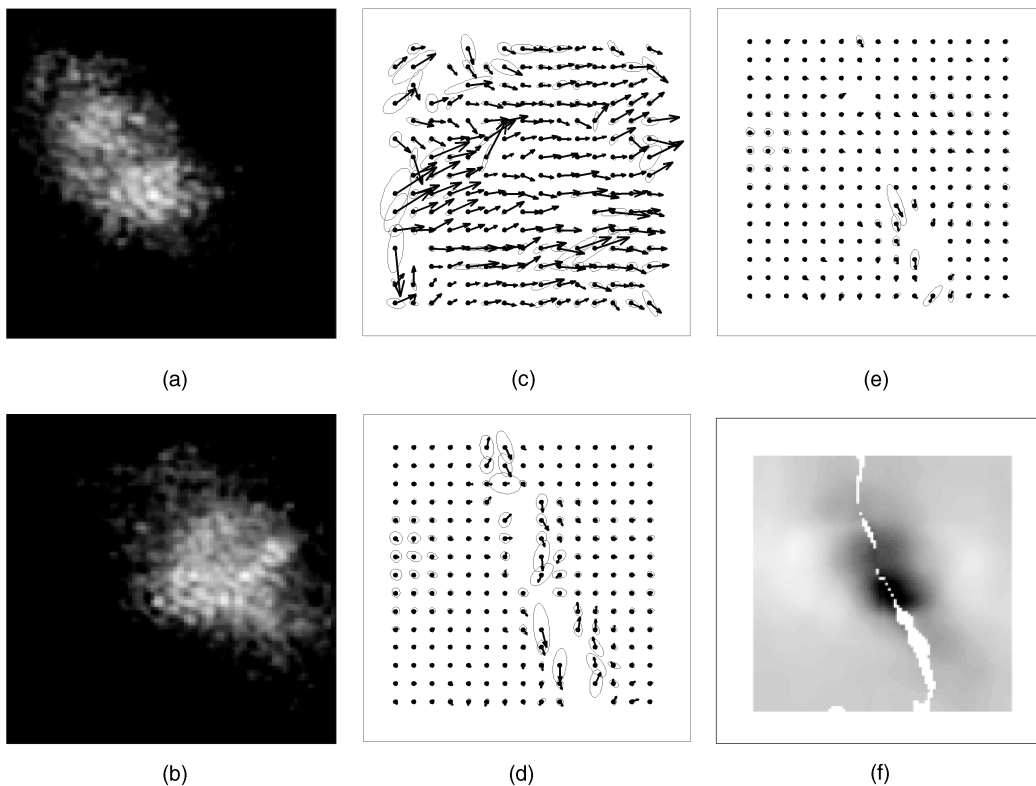


Fig. 7. Moving flashlight illuminates a carpet. The envelope moves down and to the right while the carpet texture remains almost stationary. (a) and (b) Frames 2 and 6. (c) and (d) Optical flow estimates and uncertainty ellipses for constant and linear brightness models. (e) and (f) Flow estimates with the quadratic model and an image of estimates of parameter a_2 which depends on the curvature and motion of the envelope.




Fig. 8. Synthetic test sequence of an object subject to exponential decay. A Gaussian shaped object is moving with the velocity $\mathbf{v} = [-1, 0]^T$ pixels/frame and simultaneously exponentially decaying with the decay rate $\kappa = 0.3$ frames $^{-1}$. (a), (b), and (c) First, third, and fifth frame of the sequence. (d) Difference vector field $\mathbf{e} = \mathbf{v} - \mathbf{v}'$ between the true and the estimated velocity field and uncertainty ellipses estimated with the constant brightness assumption. (e) Difference vector field and uncertainty ellipses estimated with the exponential decay model. (f) Relative error in the decay rate κ , ranging from 0 percent (black) to 25 percent (white) threshold by the confidence measure. The arrows in (d) and (e) are drawn to the same scale in order to compare the differences in the estimated error vector fields.

locations with estimates in Fig. 8f), the relative error in the estimates of the decay constant is well below 20 percent.

Comparing the accuracy estimates of the two optical flow fields (Figs. 8d and 8e) reveals an interesting feature of the parameter estimation. While the error covariance ellipses for the exponential decay model correctly capture the residual errors in the optical flow estimate, the constant brightness model seems to severely underestimate the errors. The constant model forces the optical flow estimates to directions that conserve brightness within the spatio-temporal neighborhood. As these directions can often be found with high confidence, and the number of parameters estimated is small, the error covariances are sometimes quite small. However, these directions do not correctly capture the underlying optical flow field, which helps to underline the importance of using the correct physical model of the brightness variation.

Fig. 9 shows an application from physical oceanography where heat spots are artificially generated by a laser so that the water surface can be tracked with an infrared camera. One of the scientific tasks was to estimate the decay rate of heat spots on the water surface in a wind/wave tank. The size of the depicted region in Fig. 9 is approximately 5×2.5 cm. The heat decay rate is directly related to the transfer velocity of heat across the air/water interface, which is to be estimated. In addition to the exponential decay, the thermal pattern deforms according to the underlying turbulent flow field, which we also wish to estimate.

If the brightness is assumed to remain constant (Fig. 9d), the estimated flow, especially the convergent flow in the center, is unrealistic. In fact, the heat spot in this example is sheared and elongated from one frame to the next. Using an exponential decay model, by comparison, the estimated flow field is much more realistic (Fig. 9e), as are the estimates of the decay rate κ . An estimate $\bar{\kappa}$ is obtained by averaging $\kappa(\mathbf{x})$ over all image points where the estimation was reasonably well conditioned (where the L_2 norm of the error covariance matrix Σ lies below a threshold as discussed above). For the example depicted in Fig. 9, the average decay rate is $\bar{\kappa} = 4.02$ s $^{-1}$. This value is in good agreement with the expected value at the same wind speed [15].

5.4 Heat Diffusion

Our last experiments concern thermal diffusion in infrared images. As above, we first generated a synthetic test sequence of a noisy Gaussian-shaped brightness distribution that is subject to diffusion while moving with a velocity $\mathbf{v} = [-1, 0]^T$ pixels/frame. The diffusion constant was set as $D = 2.5$ pixels 2 frames $^{-1}$. Estimation was done first with the brightness constancy model and then with the diffusion model (9). The results are shown in Fig. 10.

As expected, the constant brightness model performs poorly compared to the diffusion model. The error vector field in Fig. 10d shows that the constant brightness model yields convergent and divergent optical flow in regions of negative or positive image curvature. The most severe errors occur in the center of the image where the brightness




Fig. 9. Application example of an exponentially decaying heat spot on a wavy water surface in a laboratory wind-wave tank. (a), (b), and (c) First, third, and fifth frame of the sequence. (d) Optical flow field \mathbf{v} and uncertainty ellipses estimated with the constant brightness assumption. (e) Optical flow field and uncertainty ellipses estimated with an exponential decay model.




Fig. 10. Synthetic test sequence of an object subject to diffusion. A Gaussian shaped object is moving with the velocity $\mathbf{v} = [-1, 0]^T$ pixels/frame and simultaneously diffusing with the diffusion constant $D = 2.5$ pixels²/frame. (a), (b), and (c) first, second, and third frame of the sequence. (d) Difference vector field $\mathbf{e} = \mathbf{v} - \mathbf{v}'$ between the true velocity field \mathbf{v} and the estimated velocity field \mathbf{v}' , and uncertainty ellipses estimated with the constant brightness assumption. (e) Difference vector field and uncertainty ellipses estimated with a diffusion model. (f) relative error in the diffusion constant D , ranging from 0 percent (black) to 25 percent (white) threshold by the confidence measure. The arrows in (d) and (e) are drawn to the same scale in order to compare the differences in the estimated error vector fields.

variation, according to the diffusion equation (9), is greatest. By comparison, with the correct physical transport model the estimated flow field is more accurate; Fig. 10e shows the corresponding errors. In addition to the superior flow estimates, the diffusion model yields an estimate of the diffusion constant D (Fig. 10f). Within the region containing pixels at which the problem was reasonably well-conditioned, the relative error in the estimated diffusion constant is below 25 percent.

The last example shows a natural infrared image sequence (Fig. 11) of a localized heat spot induced by a laser on a thin object. The underlying material, with the heat spot, is translating towards the top of the image. As it translates the heat spot is simultaneously diffusing along the target surface. Figs. 11d and 11e show the optical flow fields computed with the constant brightness model and with the diffusion model. The flow field computed with the constant brightness model shows convergence about the




Fig. 11. Heat spot on a target surface. While the target is moving in upward direction, the heat distribution on the target surface is subject to diffusion. (a), (b), and (c) first, third, and fifth frame of the sequence. (d) Optical flow field v and uncertainty ellipses estimated with the constant brightness assumption. (e) Optical flow field and uncertainty ellipses estimated with a diffusion model.

center of the heat spot, while the diffusion model yields a translational flow field within the same region. Also, as in the synthetic case, use of the diffusion model allows one to estimate the diffusion constant of the heat spot, which was estimated to $D = 1.34 \cdot 10^{-3} \text{cm}^2 \text{s}^{-1}$. The estimates obtained were similar to those obtained using the same experimental setup, but without motion so that brightness change was known to be solely a function of thermal diffusion.

6 CONCLUSIONS

This paper presents a new approach to quantitatively estimating motion and physical parameters of image sequences. We use physical models of brightness change to facilitate the estimation of both optical flow and physical parameters of the scene. Previous approaches have accommodated violations of brightness constancy with the use of robust statistics or with generalized brightness constancy constraints that allow generic types of contrast change. Here, we consider models of brightness variation that have time-dependent physical causes, including changing surface orientation with respect to a directional illuminant, motion of the illuminant, and physical models of heat transport (diffusion and decay) in infrared images.

The new technique is a straightforward extension of the standard brightness change constraint equation to incorporate the spatiotemporal signature of particular dynamic processes. With our formulation, the resulting problems result in linear systems of constraint equations that can be solved by standard numerical techniques, such as total-least-squares. The method provides both accurate optical

flow estimates and accurate estimates of the relevant physical parameters. The usual sensitivity of total least squares to measurement noise and conditioning [40] is, in part, mitigated with the use of the error covariance.

ACKNOWLEDGMENTS

Portions of this work were performed while H.W. Haussecker was with the Interdisciplinary Center for Scientific Computing (IWR), Heidelberg University, Germany. Also, portions of this work were supported by the "Deutsche Forschungsgemeinschaft," DFG, within the research unit "Image Sequence Analysis to Investigate Dynamic Processes." The authors would like to thank C. Garbe for providing data and helping with the implementation, O. Nestares for his work on covariance estimates, and M. Black, B. Jähne, A. Jepson, and H. Spies for helpful discussions.

REFERENCES

- [1] P. Anandan, "A Computational Framework and an Algorithm for the Measurement of Visual Motion," *Int'l J. Computer Vision*, vol. 2, pp. 283–310, 1989.
- [2] J.L. Barron, D.J. Fleet, and S.S. Beauchemin, "Performance of Optical Flow Techniques," *Int'l J. Computer Vision*, vol. 12, no. 1, pp. 43–77, 1994.
- [3] J. Bergen, P. Anandan, K. Hanna, and R. Hingorani, "Hierarchical Model-Based Motion Estimation," *Proc. European Conf. Computer Vision*, pp. 237–252, 1992.
- [4] M.J. Black and P. Anandan, "Robust Dynamic Motion Estimation Over Time," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 296–302, June 1991.

- [5] M.J. Black, D.J. Fleet, and Y. Yacoob, "Robustly Estimating Changes in Image Appearance," *Computer Vision and Image Understanding*, vol. 78, no. 1, pp. 8–31, 2000.
- [6] N. Cornelius and T. Kanade, "Adapting Optical Flow to Measure Object Motion in Reflectance and X-Ray Image Sequences," *ACM Workshop Motion: Representation and Perception*, pp. 50–58, 1983.
- [7] B. Duc, "Feature Design: Applications to Motion Analysis and Identity Verification," PhD thesis, École Polytechnique Fédérale de Lausanne, 1997.
- [8] D.J. Fleet, M.J. Black, Y. Yacoob, and A.D. Jepson, "Design and Use of Linear Models for Image Motion Analysis," *Int'l J. Computer Vision*, vol. 36, no. 3, pp. 171–193, 2000.
- [9] D.J. Fleet and K. Langley, "Computational Analysis of Nonfourier Motion," *Vision Research*, vol. 22, pp. 3057–3079, 1994.
- [10] D.J. Fleet and K. Langley, "Recursive Filters for Optical Flow," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 17, pp. 61–67, 1995.
- [11] D.J. Fleet and A.D. Jepson, "Computation of Component Image Velocity from Local Phase Information," *Int'l J. Computer Vision*, vol. 5, pp. 77–104, 1990.
- [12] D.J. Fleet and A.D. Jepson, "Stability of Phase Information," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, pp. 1253–1268, 1993.
- [13] L. Florac, W. Niessen, and M. Nielsen, "The Intrinsic Structure of Optical Flow Incorporating Measurement Duality," *Int'l J. Computer Vision*, vol. 27, no. 3, pp. 263–286, 1998.
- [14] G.D. Hager and P.N. Belhumeur, "Efficient Region Tracking with Parametric Models of Geometry and Illumination," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 27, no. 10, pp. 1025–1039, 1998.
- [15] H.W. Haussecker, "Physics from IR Sequences: Quantitative Analysis of Transport Models and Parameters of Air-Sea Gas Transfer," *Proc. Intl. Symp. Gas Transfer at Water Surfaces*, 2000.
- [16] H.W. Haussecker and H. Spies, "Motion," *Handbook of Computer Vision and Applications*, B. Jähne, H.W. Haussecker, and P. Geissler, eds., pp. 309–396, 1999.
- [17] B.K. Horn and B.G. Schunck, "Determining Optical Flow," *Artificial Intelligence*, vol. 17, pp. 185–204, 1981.
- [18] B.K. Horn, *Robot Vision*. Cambridge, Mass.: MIT Press, 1986.
- [19] B. Jähne, H.W. Haussecker, H. Spies, D. Schmundt, and U. Schurr, "Study of Dynamical Processes with Tensor-Based Spatiotemporal Image Processing Techniques," *Proc. European Conf. Computer Vision*, vol. II, H. Burkhardt and B. Neumann, eds., pp. 322–335, 1998.
- [20] B. Lucas and T. Kanade, "An Iterative Image Registration Technique with an Application to Stereo Vision," *Proc. DARPA Image Understanding Workshop*, pp. 121–130, 1981.
- [21] W.J. MacLean, "Removal of Translation Bias When Using Subspace Methods," *Proc. IEEE Int'l Conf. Computer Vision*, vol. II, pp. 753–758, 1999.
- [22] B. Matei and P. Meer, "A General Method for Errors-in-Variables Problems in Computer Vision," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. II, pp. 18–25, 2000.
- [23] J. Michel, N. Nandhakumar, and V. Velten, "Thermophysical Algebraic Invariants from Infrared Imagery for Object Recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 19, no. 1, pp. 41–51, 1997.
- [24] M. Mühlich and R. Mester, "The Role of Total Least Squares in Motion Analysis," *Proc. European Conf. Computer Vision*, vol. II, H. Burkhardt and B. Neumann, eds., pp. 305–321, 1998.
- [25] M. Mühlich and R. Mester, "A Considerable Improvement in Pure Parameter Estimation Using TLS and Equilibration," *Pattern Recognition Letters*, 2000.
- [26] N. Mukawa, "Estimation of Shape, Reflection Coefficients and Illumination Direction from Image Sequences," *Proc. IEEE Int'l Conf. Computer Vision*, pp. 507–512, 1990.
- [27] D.W. Murray and B.F. Buxton, *Experiments in the Machine Interpretation of Visual Motion*. Cambridge, Mass.: MIT Press, 1990.
- [28] H.-H. Nagel, "On a Constraint Equation for the Estimation of Displacement Rates in Image Sequences," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 11, no. 1, pp. 13–30, 1989.
- [29] H.-H. Nagel, "Extending the 'Oriented Smoothness Constraint' into the Temporal Domain and the Estimation of Derivatives of Optical Flow," *Proc. European Conf. Computer Vision*, O. Faugeras, ed., pp. 139–148, 1990.
- [30] H.-H. Nagel, "Optical-Flow Estimation and the Interaction between Measurement Errors at Adjacent Pixel Positions," *Int'l J. Computer Vision*, vol. 15, no. 3, pp. 271–288, July 1995.
- [31] S.K. Nayar and G. Narasimhan, "Vision in Bad Weather," *Proc. IEEE Int'l Conf. Computer Vision*, vol. II, pp. 820–827, 1999.
- [32] S. Negahdaripour, "Revised Definition of Optical Flow: Integration of Radiometric and Geometric Clues for Dynamic Scene Analysis," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 20, no. 9, pp. 961–979, Sept. 1998.
- [33] O. Nestares, D.J. Fleet, and D.J. Heeger, "Likelihood Functions and Confidence Bounds for Total-Least-Squares Estimation," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. I, pp. 523–530, 2000.
- [34] N. Ohta, "Optical Flow Detection Using a General Noise Model," *IEICE Trans. Information and Systems*, vol. 79, no. 7, pp. 951–957, 1996.
- [35] H. Scharr, S. Körkel, and B. Jähne, "Numerische Isotropieoptimierung von FIR-Filtern mittels Querglättung," *Mustererkennung 1997*, E. Paulus and F.M. Wahl, eds., pp. 367–374, 1997.
- [36] G.W. Stewart, "Errors in Variables for Numerical Analysts," *Recent Advances in Total Least Squares Techniques*, S. Van Huffel, ed., pp. 3–10, 1997.
- [37] S. Van Huffel and S. Vandewalle, *The Total Least Squares Problem: Computational Aspects and Analysis*. Philadelphia, Penn.: SIAM Press, 1991.
- [38] A. Verri and T. Poggio, "Motion Field and Optical Flow: Qualitative Properties," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 11, no. 5, pp. 490–498, 1989.
- [39] S. Wang, V. Markandey, and A. Reid, "Total Least Squares Fitting Spatiotemporal Derivatives to Smooth Optical Flow Fields," *Proc. SPIE: Signal and Data Processing of Small Targets*, pp. 42–55, 1992.
- [40] J. Weber and J. Malik, "Robust Computation of Optical Flow in a Multiscale Differential Framework," *Int'l J. Computer Vision*, vol. 15, no. 1, pp. 67–81, 1995.
- [41] R.P. Wildes, M.J. Amabile, A. Lanzillotto, and T. Leu, "Recovering Estimates of Fluid Flow from Image Sequence Data," *Computer Vision and Image Understanding*, vol. 80, pp. 246–266, 2000.



Horst W. Haussecker received his doctoral degree in physics from Heidelberg University, Germany, in 1996. He was visiting scientist at the Scripps Institution of Oceanography, University of California, San Diego, in 1994. From 1996 to 1999, he was researcher in computer vision at the Interdisciplinary Center for Scientific Computing (IWR), Heidelberg University, where he also lectured on optical flow computation. Since 1999, he has been a member of the research staff at Xerox's Palo Alto Research Center (PARC). His research interests include computer vision, optical flow computation, infrared thermography, physical transport processes, and application of computer vision as quantitative instrument in science and technology. He is coeditor and the main contributing author of two books and has authored or coauthored more than 50 technical articles. He is a member of the IEEE and the IEEE Computer Society.



David J. Fleet received the PhD degree in computer science from the University of Toronto in 1991. From 1991 to 2000, he held a faculty position at Queen's University, in computer science, psychology, and electrical engineering. In 1999, he joined the Xerox Palo Alto Research Center (PARC) where he currently manages the Digital Video Analysis Project and the Perceptual Documents Project. His general research interests include computer vision, image processing, visual perception, and visual neuroscience. He has published one book and research articles on various topics including the estimation of optical flow and stereoscopic disparity, probabilistic methods in motion analysis, modeling appearance changes in image sequences, non-Fourier motion and stereo perception, and the neural basis of stereo vision. He is a member of IEEE, the IEEE Computer Society, and an associate editor of *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

► For further information on this or any computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.