

Computing Optical Flow with Physical Models of Brightness Variation *

Horst W. Haussecker David J. Fleet
 Xerox Palo Alto Research Center,
 3333 Coyote Hill Road, Palo Alto, CA 94304
 {hhaussec,fleet}@parc.xerox.com

Abstract

This paper exploits physical models of time-varying brightness in image sequences to estimate optical flow and physical parameters of the scene. Previous approaches handled violations of brightness constancy with the use of robust statistics or with generalized brightness constancy constraints that allow generic types of contrast and illumination changes. Here, we consider models of brightness variation that have time-dependent physical causes, namely, changing surface orientation with respect to a directional illuminant, motion of the illuminant, and physical models of heat transport in infrared images. We simultaneously estimate the optical flow and the relevant physical parameters. The estimation problem is formulated using total least squares (TLS), with confidence bounds on the parameters.

1. Introduction

This paper uses physical models of time-dependent brightness variation in image sequences to estimate optical flow and physical parameters of the scene. Physical causes of brightness variation include changing surface orientation with respect to a directional illuminant, motion of the illuminant, and physical models of heat transport in infrared images such as diffusion and decay. We wish to estimate optical flow and the relevant physical parameters.

Many computer vision applications require accurate estimates of the optical flow field. Although studied extensively [1, 11], reliable optical flow computation still remains difficult in many cases. Problems arise from the complex physical processes involved in scene illumination, surface reflection, and the transmission of radiation through surfaces and the atmosphere [12, 24, 19]. Without a model of image formation it is not possible to unambiguously relate spatiotemporal brightness to motion.

*Part of this work was performed while HWH was with the Interdisciplinary Center for Scientific Computing, Heidelberg University, Germany. Portions of this work were supported by the 'Deutsche Forschungsgemeinschaft' (Image Sequence Analysis to Investigate Dynamic Processes). The authors thank M. Black, B. Jähne, A. Jepson, and O. Nestares for helpful discussions, and C. Garbe for providing the rotating sphere sequence.




Figure 1. Illustration of errors in the optical flow estimation due to brightness changes. (a) constant brightness (correct flow field), (b) exponential decay, (c) diffusion.

Common optical flow techniques assume brightness constancy. For graylevel images $g(\mathbf{x}, t)$, where $\mathbf{x} = (x, y)$, the tracking of points of constant brightness amounts to finding a path $\mathbf{x}(t)$ along which image brightness is constant, i.e.,

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

for some constant c . Taking the total temporal derivative of both sides of (1) yields the well-known *brightness change constraint equation* (BCCE) [12]:

$$\frac{dg}{dt} = g_x v_1 + g_y v_2 + g_t = 0, \quad (2)$$

where $\mathbf{v} = [v_1, v_2]^T = [dx/dt, dy/dt]^T$ is the *optical flow* that we wish to estimate, and g_i denotes the partial derivative of g with respect to the coordinate $i \in \{x, y, t\}$. Because (2) provides one constraint in two unknowns it is common to combine constraints at neighboring pixels, assuming that the optical flow field is locally constant or affine [1, 2]. This produces a system of linear equations that can be solved using standard (weighted) least squares [14], or total least squares (TLS) [25, 26]. To further constrain the estimates, the regions can be extended into time if \mathbf{v} is smooth within local temporal windows [7, 27]. For the enhanced brightness change models (Section 3) we consider here, it is very important that the neighborhoods are extended to more than two frames. With only two frames, one can only model brightness changes that are linear in time.

If brightness is not conserved, then the optical flow field estimated from (2) can be severely biased [4, 17, 18, 24, 3, 20, 9, 19]. Causes of brightness variation include moving illumination envelopes, changing orientation of surfaces under directional illumination, and atmospheric influences in outdoor applications. Other instances occur in scientific applications that quantitatively investigate dynamic processes [13]. Figure 1 illustrates the influence of brightness changes on optical flow estimation with two examples of physical transport processes in infrared images, namely, exponential decay and diffusion; although the surface translates in each case, the flow field that conserves brightness may converge or diverge.

This paper describes a generalized framework for incorporating brightness changes into motion analysis using physical models. 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, not in the image brightness or in the spatiotemporal coordinates. With this, we obtain a linear system of equations that constitutes a straightforward generalization of the brightness constancy assumption.

Estimates of the parameters can then be obtained using TLS. We show that this produces improved optical flow estimates, and it allows us to estimate additional information that characterizes the physical processes. TLS error covariance matrices [21] are used to quantify the accuracy of the optical flow and the brightness change parameters.

2. Previous Work

Brightness variations have been modeled by [4, 17, 18, 24, 3, 20, 9]. A general framework is proposed in [20] where the brightness change between two frames consists of a multiplier and an offset field:

$$g(\mathbf{r} + \delta\mathbf{r}) - g(\mathbf{r}) = m(\mathbf{r})g(\mathbf{r}) + c(\mathbf{r}), \quad (3)$$

where, for notational convenience, $\mathbf{r} = [\mathbf{x}^T, t]^T$ denotes a space-time 3D vector. It is certainly true that all changes between two images can be modeled according to (3). However, this approximation only yields the instantaneous brightness change, which does not allow us to discriminate different physical causes of brightness changes, or to constrain the estimation to satisfy particular physical models.

In related work on target tracking, Hager and Belhumeur [9] 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




Figure 2. 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 $(g(\mathbf{x}, t) - g(\mathbf{x}_0, t_0) = 0)$.

physical model of the brightness variation. Black et al. [3] 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 linear 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. Our approach 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 parameters of interest.

3. Physics-Based Brightness Variation

Given a physical model, constraints on brightness variations can be specified by a functional relation or, more directly, by differential equations that relate temporal brightness changes to the spatial image structure. As a generalization of brightness conservation (1), shown in Fig. 2, we define a temporal trajectory $\mathbf{x}(t)$ along which brightness can change according to a parameterized function, h :

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

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, $h(g_0, t, \mathbf{a} = \mathbf{0}) = g_0$. While \mathbf{a} is assumed to be constant within small temporal windows, h is expressed as a function of time to be able to capture nonlinear temporal brightness changes.

Taking the total derivative of both sides of (4) yields a

generalized brightness change constraint equation:

$$g_x v_1 + g_y v_2 + g_t = f(g_0, t, \mathbf{a}), \quad (5)$$

where f is defined as:

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

With the constant brightness model (1), $f = 0$, and (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 such as affine motion [6].

Finally, there are two different ways to specify the form of f . One can compute f by (6) using a known analytical form of h , or one can choose f according to the differential equations of the underlying physical processes. We will illustrate this below in Sections 3.1 - 3.4, before returning to the general formulation in Section 3.5.

3.1. Exponential Decay

In certain instances, brightness in infrared images is well modeled using exponential decay. This occurs in applications where heat is removed from the thin layer on an object surface that infrared cameras are imaging. In many such cases the brightness function h in (4) has the analytical form

$$h(g_0, t, \kappa) = g_0 \exp(-\kappa t). \quad (7)$$

In this case, the parameter vector \mathbf{a} reduces to a scalar *decay constant*, $a = \kappa$. Therefore, from (6) it follows that

$$f(g_0, t, \kappa) = -\kappa g_0 \exp(-\kappa t) = -\kappa g(\mathbf{x}(t), t). \quad (8)$$

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 , we can estimate \mathbf{v} and κ using linear methods (see Section 4).

3.2. Diffusion

Another elementary physical transport process, used to model brightness changes in infrared image sequences is *diffusion*. If the brightness change can be modeled as isotropic diffusion in the image coordinates, then it depends on the spatial image structure according to the well known diffusion equation:

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

where D is the scalar *diffusion constant*. Because f is linear D , we can use linear methods to estimate \mathbf{v} and D (Sec. 4).

3.3. Moving Illumination Envelope

Brightness changes caused by moving, non-uniform illumination envelopes have been considered in the 2-frame case [20]. 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} :

$$g(\mathbf{x}, t) = g_c(\mathbf{x} - t\mathbf{v}) E(\mathbf{x} - t\mathbf{u}). \quad (10)$$

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

$$g_w(\mathbf{x}_r, t) \equiv g(\mathbf{x}_r + t\mathbf{v}, t) = g_c(\mathbf{x}_r) E(\mathbf{x}_r - t\mathbf{u}_r), \quad (11)$$

where $g_w(\mathbf{x}_r, t)$ is a warped version of $g(\mathbf{x}, t)$ for which the motion of the surface reflectance 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 $(\mathbf{x}_r, 0)$:

$$E(\mathbf{x}_r - t\mathbf{u}_r) \approx E(\mathbf{x}_r) - t \nabla E^T \mathbf{u}_r + \frac{1}{2} t^2 \mathbf{u}_r^T \mathbf{H} \mathbf{u}_r, \quad (12)$$

where ∇E and \mathbf{H} are the gradient and Hessian of E at $(\mathbf{x}_r, 0)$. Substituting (12) into (11) yields a brightness function:

$$h = g_c \left[E - t \nabla E^T \mathbf{u}_r + \frac{1}{2} t^2 \mathbf{u}_r^T \mathbf{H} \mathbf{u}_r \right]. \quad (13)$$

Then, from (6), the brightness change, f , is given by

$$f = a_1 + a_2 t, \quad (14)$$

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

When the moving envelope can be approximated by (12) the quadratic time-varying model in (13) reduces the bias in optical flow estimates. By comparison, if the envelope were nearly linear, a first-order temporal model for h , and hence a constant model for f , would suffice. In either case, solving for the polynomial coefficients of the model in \mathbf{a} 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.

3.4. Changing Surface Orientation

The last case we address here concerns 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{l} 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. In some applications one might intentionally use a directional source to enhance edges while simultaneously tracking surface properties.

Given a combination of ambient illumination and a fixed, distant, point light source from direction \mathbf{l} (where $\|\mathbf{l}\| = 1$), the surface radiance from a Lambertian surface with unit normal \mathbf{n} can be expressed as $c_0 + c_1 \mathbf{n}^T \mathbf{l}$, where c_0 is the ambient component and c_1 is proportional to surface albedo. If we assume a rotating body, then we can write the surface normal at time t as $\mathbf{n}_t = \mathbf{R}_t \mathbf{n}_0$, where \mathbf{R}_t is a 3D rotation matrix and \mathbf{n}_0 is the normal at time 0. Then, the time-varying radiance becomes $c_0 + c_1 \mathbf{n}_0^T \mathbf{R}_t^T \mathbf{l}$.

The extent to which the radiance changes with time depends on the angle between the light source direction \mathbf{l} and the axis of rotation \mathbf{l}_0 . To see this, let $\mathbf{l} = \alpha \mathbf{l}_0 + \beta \mathbf{l}_1$, where $\|\mathbf{l}_0\| = \|\mathbf{l}_1\| = 1$, $\alpha = \mathbf{l}^T \mathbf{l}_0$, and $\alpha^2 + \beta^2 = 1$. With this, one can show that the time varying radiance becomes

$$L(t) = c_0 + c_1 \left(\alpha \mathbf{n}_0^T \mathbf{l}_0 + \beta \mathbf{n}_0^T \mathbf{R}_t^T \mathbf{l}_1 \right). \quad (15)$$

Finally, with some manipulation of the second term, one can show that this reduces to the general form of

$$L(t) = \kappa_1 + \kappa_2 \cos(\omega t + \phi), \quad (16)$$

where $\kappa_1 = c_0 + c_1 \alpha \mathbf{n}_0^T \mathbf{l}_0$, $\kappa_2 = c_1 \beta$, and ω denotes the frequency of the temporal modulation which depends directly on the rate of object rotation.

It is obvious from (16) 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 our brightness function that is linear in its parameters:

$$h(g_0, t, a_1, a_2) \approx g_0 (1 + a_1 t + a_2 t^2), \quad (17)$$

where the parameters a_1 and a_2 are functions of ω and ϕ . Using (17), f can be approximated by

$$f = g_0 a_1 + 2g_0 a_2 t, \quad (18)$$

which is linear in the parameters. Once an approximate parameter set, a_1 and a_2 , is estimated using (18) in (5), the quadratic approximation of h can be fitted to (16) to estimate the parameters ω and ϕ .

3.5. Generalized Formulation

In Sections 3.1 - 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 $\mathbf{a} = 0$:

$$h(g_0, t, \mathbf{a}) = g_0 + \sum_{k=1}^Q a_k \frac{\partial h}{\partial a_k}. \quad (19)$$

Using (19) we can express f , the temporal brightness variation defined in (6), as

$$f(g_0, t, \mathbf{a}) = \frac{dh(g_0, t, \mathbf{a})}{dt} = \sum_{k=1}^Q a_k \frac{d}{dt} \frac{\partial h}{\partial a_k}, \quad (20)$$

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

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

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

4. Computational Framework

In each of the above formulations we obtain linear constraints that relate the variables of interest and noisy measurements. The general form of the constraints, assuming a constant optical flow model, can be expressed as

$$\begin{aligned} \mathbf{d}^T \mathbf{p}_h &= 0, \quad \text{with } \mathbf{d} = [(\nabla_{\mathbf{a}} f)^T, \nabla g^T, g_t]^T \\ \text{and } \mathbf{p}_h &= [\mathbf{p}^T, 1]^T, \quad \mathbf{p} = [-\mathbf{a}^T, \mathbf{v}^T]^T, \end{aligned} \quad (22)$$

where \mathbf{p} contains the parameters of interest (the flow field parameters and the brightness parameters of h), and \mathbf{p}_h denotes the *homogeneous* counterpart of \mathbf{p} . The $(Q + 3)$ -dimensional vector \mathbf{d} 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 model to higher-order parameterized motion models [2, 6], such as an affine model.

Equation (22) is just one constraint in several unknowns. To further constrain the parameters, \mathbf{p} , we assume that \mathbf{p} is constant within a local space-time region. We then use a

collection of N such constraints at neighboring pixels in the region to obtain a linear system:

$$\mathbf{G} \mathbf{p}_h = 0, \quad \text{with} \quad \mathbf{G} = [\mathbf{d}_1, \dots, \mathbf{d}_N]^T. \quad (23)$$

Assuming IID Gaussian noise in the measurement matrix \mathbf{G} , a maximum likelihood estimate of \mathbf{p} is given by the TLS estimator [21, 23], often formulated as the minimum of

$$\frac{\mathbf{p}_h^T \mathbf{G}^T \mathbf{G} \mathbf{p}_h}{\mathbf{p}_h^T \mathbf{p}_h}. \quad (24)$$

In practice, errors in measuring temporal image derivatives are often larger than errors in measuring the spatial derivatives. Also, derivative measurements at adjacent pixels are often correlated. Thus, one must renormalize the constraints before using TLS to avoid bias in the estimates [15, 16]. Although not always referred to as such, many recent approaches to optical flow computation (e.g., [5, 8, 22, 18, 26]) are based on TLS or related techniques.

To quantify the measurement error in terms of an error covariance matrix, we use the Hessian of the negative log-likelihood evaluated at the TLS estimate, $\hat{\mathbf{p}}_h = [\hat{\mathbf{p}}^T, 1]^T$. In [21], this is shown to be

$$\mathbf{H} = \frac{\gamma}{\sigma_n^2 \|\hat{\mathbf{p}}_h\|^2} \left(\mathbf{M} - \frac{1}{\|\hat{\mathbf{p}}_h\|^2} (\hat{\mathbf{p}}_h^T \mathbf{D} \hat{\mathbf{p}}_h) \mathbf{I}_{Q+2} + \frac{4}{\|\hat{\mathbf{p}}_h\|^4} \left[(\hat{\mathbf{p}}_h^T \mathbf{D} \hat{\mathbf{p}}_h) \hat{\mathbf{p}} - \|\hat{\mathbf{p}}_h\|^2 (\mathbf{M} \hat{\mathbf{p}} + \mathbf{A}^T \mathbf{b}) \right] \hat{\mathbf{p}}^T \right), \quad (25)$$

where $\mathbf{D} \equiv \mathbf{G}^T \mathbf{G}$ is a $(Q+3) \times (Q+3)$ matrix, \mathbf{A} contains the first $(Q+2)$ columns 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. The factor γ is defined as $\gamma = \sigma^2 / (\sigma_n^2 + \sigma^2)$, where σ^2 denotes the variance of the expected distribution of gradients in \mathbf{G} , and $\sigma_n^2 \mathbf{I}_{Q+3}$ is the covariance of the IID Gaussian noise. If the signal to noise ratio (SNR) is high ($\sigma^2 \gg \sigma_n^2$), then $\gamma \approx 1$. The error covariance matrix suggested by [21], is then given by $\Sigma = \mathbf{H}^{-1}$.

5. Experimental Results

We have applied the technique to both synthetic and natural images sequences. These include scientific applications with infrared image sequences, as well as more conventional computer vision applications. In each case we measure the optical flow and the brightness change parameters. We also compute error covariance matrices, given above in Sec. 4, that serve as confidence bounds on the estimates [21].

5.1. Changing Surface Orientation

Figures 3(a,b) 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




Figure 3. Rotating sphere under directional illumination. (a, b) Frames 1 and 5. (c, d) Optical flow estimates and uncertainty ellipses for the constant and quadratic temporal brightness models.

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). The temporal brightness function was modeled with the quadratic approximation to the sinusoidal relationship in (17), as described in Section 3.4. For comparison, we also obtained estimates using a local linear approximation (3) and the brightness constancy model (2).

In this and other experiments, we found that the constant brightness model performed poorly compared to the linear and quadratic models. For slow rotations the linear and quadratic temporal change models produced very similar results. Fast motions produce faster brightness changes, which then show differences between the linear and quadratic models. To illustrate this while obeying sampling limits, we let the sphere rotate in one direction while the light source rotated in the other direction.

Figure 3(c,d) shows the optical flow estimates with uncertainty ellipses. The error ellipses satisfy $\mathbf{e}^T \mathbf{C}^{-1} \mathbf{e} = 4.6$ to capture 90% of the expected errors, where \mathbf{C} is the 2D error covariance matrix for v_1 and v_2 , and \mathbf{e} is the optical flow error. For convenience, we only display flow estimates when the norm of Σ is less than a liberal threshold. It is easy to see in Figure 3(b,c) that the flow field estimated with the

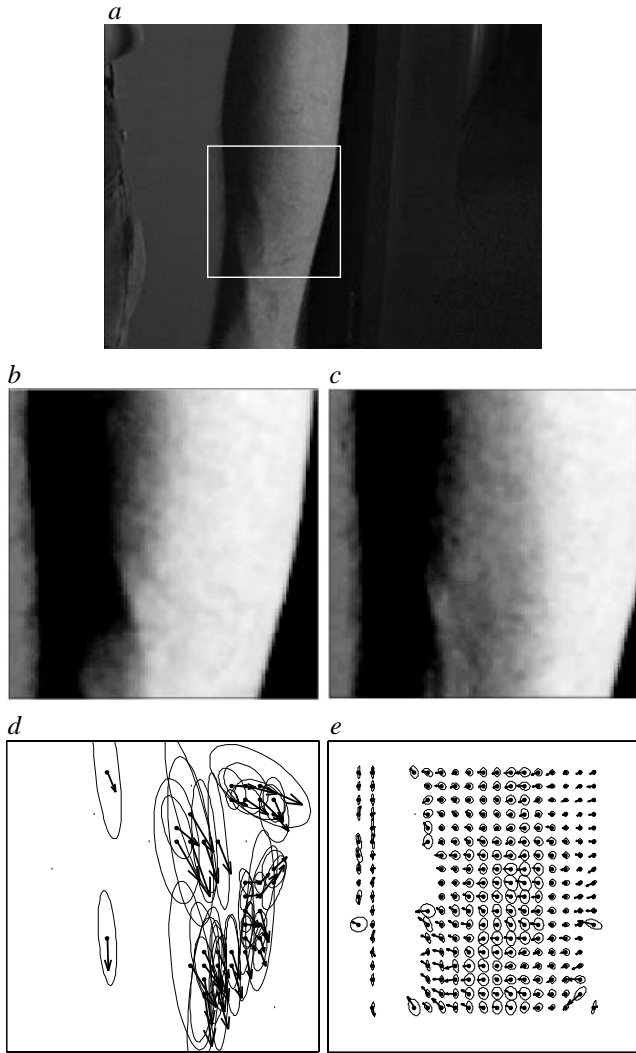


Figure 4. A human arm, under directional illumination, is rotating about its main axis towards the left while slightly translating to the right. (a) Full image indicating the area displayed in the lower images. (b, c) Frames 1 and 3. (d, e) Optical flow estimates and uncertainty ellipses for the constant and quadratic temporal brightness models.

quadratic model is more accurate than that found with the constant model.

Another example of brightness variation caused by changing surface orientation is shown in Figure 4. Here, a human arm was illuminated by light through a window in a workplace environment. The arm was turning with respect to the direction of the illumination. Due to the complex structure of non-rigid motion of the skin, the surface normal exhibited fast changes leading to brightness changes. Figure 4 shows the optical flow estimates based on the quadratic model described in Section 3.4. By comparison one can see that the estimates with the brightness constancy model are severely biased.

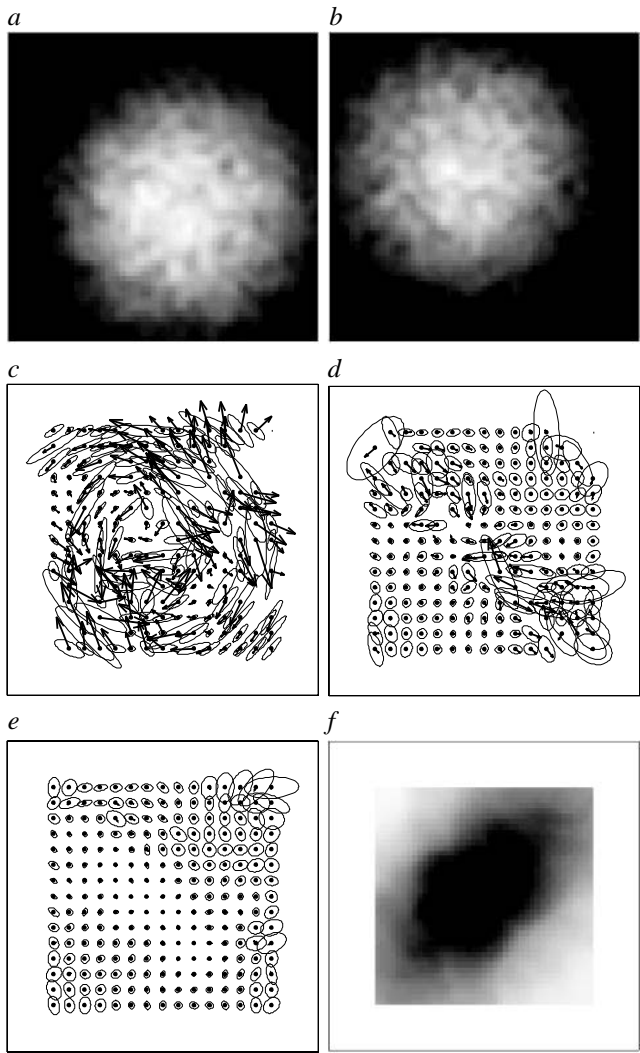


Figure 5. Synthetic moving illumination envelope. The underlying texture and the envelope move with velocities $\mathbf{v} = [1, 1]^T$ and $\mathbf{u}_E = [-2, -2]^T$ pixels/frame respectively. (a, b) Frames 2 and 6. (c, d, e) Difference between the true and estimated flow, with the uncertainty ellipses, 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 Eqn. (14)).

5.2. Moving Illuminant

To generate a synthetic example of a moving illuminant, we simply multiplied a sample of smoothed white noise with a Gaussian envelope. As described in Section 3.3, both signals translate with constant velocity. The Gaussian simulates the illumination envelope, and it is the motion of the noise signal that we wish to estimate.

Figure 5 shows results for different temporal brightness models, namely, the constant (2), linear (3), and quadratic (13) models. Because the synthetic sequence provides ground truth, we plot the difference vectors \mathbf{e} between the




Figure 6. Moving flashlight illuminates a carpet. The envelope moves right while the carpet texture remains almost stationary. (a,b) Frames 2 and 6. (c, d) Optical flow estimates and uncertainty ellipses for constant and linear brightness models. (e, 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.

flow estimates and ground truth. Due to the shape of the illumination envelope, no part of the moving pattern remains at constant brightness along its path. Consequently, the constant brightness model fails to predict the true velocity over the entire image (Fig. 5c). 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. 5d). In these regions the parameter a_1 in (14) 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 (dark and




Figure 7. Exponentially decaying heat spot on a wavy water surface. (a, b, c): Frames 1, 3 and 5. (d) Decay rate κ threshold by the confidence measure. (e, f) Optical flow estimates and uncertainty ellipses estimated with the constant brightness model and the exponential decay model.

bright regions in Figure 5f. The quadratic brightness change model allows us to accurately estimate the motion of the pattern (Fig. 5e).

The next example shows an image sequence taken as a moving observer shone a flashlight on a dark textured carpet (Fig. 6). *Negahdaripour* [20] investigated a similar example of a non-uniform illumination pattern 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 (13) become more significant, so the quadratic model yields better results (compare Figs. 6 d,e). This is clearest in regions with high envelope curvature (Fig. 6 f).

5.3. Heat Transport in Infrared Images

Figure 7 shows an application example from physical oceanography. The scientific task was to estimate the decay rate of an exponentially decaying heat spot on the water surface in a wind/wave tank. The broader goal is to estimate the transfer velocity of heat across the air/water interface, which is known to be related to the heat decay rate on the water surface. In addition to the exponential decay, the image is expected to deform due to the underlying turbulent flow field, another important parameter of air/sea interac-

tion. Thus, both the decay rate and the flow field need to be estimated.

In Figure 7, the images show an area of about 5×5 cm. If we assume brightness constancy then the estimated flow (Fig. 7 e), especially the convergent flow in the center, is unrealistic. In fact the heat spot is sheared and elongated from one frame to the next. Using the exponential brightness change model the flow (Fig. 7 f) is accurately estimated together with the thermal decay rate κ (Fig. 7 d).

Further examples of heat transport in infrared images, including further examples of heat decay and examples of heat diffusion can be found in [10].

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 have linear solutions using total-least-squares. With the use of an error covariance, we show that 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 is mitigated with the use of the error covariance.

References

- [1] J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *IJCV*, 12(1):43–77, 1994.
- [2] J. Bergen, P. Anandan, K. Hanna and R. Hingorani. Hierarchical model-based motion estimation. *Proc. ECCV*, pp. 237–252. Springer-Verlag, 1992.
- [3] M. J. Black, D. J. Fleet, and Y. Yacoob. Robustly estimating changes in image appearance. *CVIU* 78:8–31, 2000.
- [4] N. Cornelius and T. Kanade. Adapting optical flow to measure object motion in reflectance and x-ray image sequences. In *ACM Workshop on Motion*, pp. 50–58, Toronto, 1983.
- [5] B. Duc. *Feature design: applications to motion analysis and identity verification*. PhD thesis, École Polytechnique Fédérale de Lausanne, 1997.
- [6] D. J. Fleet, M. J. Black, Y. Yacoob, and A. D. Jepson. Design and use of linear models for image motion analysis. *IJCV* 36:169–191, 2000.
- [7] D. J. Fleet and K. Langley. Recursive filters for optical flow. *IEEE Trans. PAMI*, 17:61–67, 1995.
- [8] L. Florac, W. Niessen, and M. Nielsen. The intrinsic structure of optical flow incorporating measurement duality. *IJCV*, 27(3):263–286, 1998.
- [9] G. D. Hager and P. N. Belhumeur. Efficient region tracking with parametric models of geometry and illumination. *IEEE Trans. PAMI*, 27(10):1025–1039, 1998.
- [10] H. Haussecker. Simultaneous estimation of optical flow and heat transport in infrared image sequences. *Proc. IEEE Workshop on Computer Vision Beyond the Visible Spectrum*, (submitted), Hilton Head, 2000.
- [11] H. Haussecker and H. Spies. Motion. In *Handbook of Computer Vision and Applications*. B. Jähne, H. Haussecker, and P. Geissler (eds), Academic Press, pp. 309–396. 1999.
- [12] B. K. Horn and B. G. Schunck. Determining optical flow. *AI*, 17:185–204, 1981.
- [13] B. Jähne, H. Haussecker, H. Spies, D. Schmundt, and U. Schurr. Study of dynamical processes with tensor-based spatiotemporal image processing techniques. *Proc. ECCV*, pp. 322–335. Springer-Verlag, 1998.
- [14] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. *Proc. DARPA IU Workshop*, pp. 121–130, 1981.
- [15] W. J. MacLean. Removal of translation bias when using subspace methods. *Proc. ICCV*, pp. 753–758, Corfu, 1999.
- [16] M. Mühlich and R. Mester. A considerable improvement in pure parameter estimation using TLS and equilibration. *Proc. Conf. Stat. Methods for Image Proc.*, 1999.
- [17] N. Mukawa. Estimation of shape, reflection coefficients and illumination direction from image sequences. *Proc. IEEE ICCV*, pp. 507–512, Osaka, 1990.
- [18] H.-H. Nagel. On a constraint equation for the estimation of displacement rates in image sequences. *IEEE Trans. PAMI*, 11(1):13–30, 1989.
- [19] S. K. Nayar and G. Narasimhan. Vision in bad weather. *Proc. IEEE ICCV*, pp. 820–827, Corfu, 1999.
- [20] S. Negahdaripour. Revised definition of optical flow: integration of radiometric and geometric clues for dynamic scene analysis. *IEEE Trans. PAMI*, 20(9):961–979, 1998.
- [21] O. Nestares, D. J. Fleet, and D. J. Heeger. Likelihood functions and confidence bounds for total-least-squares estimation. *Proc. IEEE CVPR*, (this volume), Hilton Head, 2000.
- [22] N. Ohta. Optical flow detection using a general noise model. *IEICE Trans. Inf. & Syst.*, E79-D(7):951–957, 1996.
- [23] S. Van Huffel and S. Vandewalle. *The Total Least Squares Problem: Computational Aspects and Analysis*. SIAM, Philadelphia, 1991.
- [24] A. Verri and T. Poggio. Motion field and optical flow: qualitative properties. *IEEE Trans. PAMI*, 11(5):490–498, 1989.
- [25] S. Wang, V. Markandey, and A. Reid. Total least squares fitting spatiotemporal derivatives to smooth optical flow fields. In *SPIE: Signal and Data Processing of Small Targets*, pp. 42–55, 1992.
- [26] J. Weber and J. Malik. Robust computation of optical flow in a multiscale differential framework. *IJCV*, 15:67–81, 1995.
- [27] L. Zelnik-Manor and M. Irani. Multi-frame alignment of planes. *Proc. IEEE CVPR*, Fort Collins, pp. 151–156, 2000.