

Optical Flow from 1-D Correlation: Application to a Simple Time-to-Crash Detector

NICOLA ANCONA

Tecnopolis CSATA Novus Ortus, Robotics & Automation Laboratory, 70010 Valenzano, Bari, Italy

TOMASO POGGIO

Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 545 Technology Square, Cambridge, MA 02139

Received February 15, 1993; Revised January 20, 1994.

Abstract. In the first part of this article we show that a new technique exploiting 1-D correlation of 2-D or even 1-D patches between successive frames may be sufficient to compute a satisfactory estimation of the optical flow field. The algorithm is well suited to VLSI implementations. The sparse measurements provided by the technique can be used to compute qualitative properties of the flow for a number of different visual tasks. In particular, the second part of the article shows how to combine our 1-D correlation technique with a scheme for detecting expansion or rotation (Poggio et al. 1991) in a simple algorithm which also suggests interesting biological implications. The algorithm provides a rough estimate of time-to-crash. It was tested on real-image sequences. We show its performance and compare the results to previous approaches.

1 Introduction

The problem of how to compute efficiently estimates of the optical flow at sparse locations is of critical importance for practical implementations in a number of different tasks. A specific example is the detection of expansion of the visual field with a rough estimate of *time-to-crash* (TTC). The question has also interesting relations with biology, as we will discuss later. Here, we propose an efficient algorithm for computing the optical flow, which performs well in a number of experiments with sequences of real images and is well suited to a VLSI implementation.

Optical-flow algorithms based on patchwise correlation of filtered images perform in a satisfactory way (Little et al. 1988), and better in practice than most other approaches (see Christof et al. 1991). Their main drawback is computational complexity that forbids at present useful VLSI implementations. In this article, we show that 1-D patchwise correlation may provide a sufficiently accurate estimate of the optical flow.¹ We will then show with experiments on real-image sequences how to apply this technique to measure time-to-crash, by exploiting a recently proposed scheme (Poggio et al. 1991). The latter scheme, which is robust and invariant to the position of the focus of expansion or the center of rotation, relies on sparse measurements

of either the normal or the tangential component of the optical flow (relative to a closed contour). We will also discuss some broad implications of this work for the practical computation of the optical flow and for biology, in particular its relation to Reichardt-type models.

There are two main and quite separate contributions in this article:

1. an efficient 1-D correlation scheme to estimate the optical flow along a desired direction
2. the experimental demonstration that a previously proposed algorithm for estimating time-to-crash performs satisfactorily in a series of experiments with real images in which the elementary measurements of the flow are obtained by the new 1-D correlation scheme.

2 Computing the Optical Flow Along a Direction

How can the component of the optical flow be measured efficiently along a certain desired direction? As argued by Verri and Poggio (1989), a qualitative estimate is often sufficient for many visual tasks. For the task of detecting a potential crash, for instance, it has been suggested (Poggio et al. 1991) that a *precise* measurement of the normal component of the flow may not be necessary, since the precise definition of the optical

flow is itself somewhat arbitrary: it is sufficient that the estimate be qualitatively consistent with the values of the perspective 2-D projection of the “true” 3-D velocity field *for the particular stimulus*. In other words, even estimates that don’t really measure image-plane velocity (like Reichardt’s correlation model or equivalent energy models), since they also depend on spatial structure of the image, may be acceptable for several visual tasks, if their estimates are consistent over the visual field. Certain uses of a crash detector are good examples. It turns out that even a rough estimate of time-to-crash (TTC) is possible using approximate estimates of the optical flow field. Flies and other insects rely for landing on what appears to be a qualitative estimate of the time-to-crash!

2.1 1-D Correlation of 2-D Patches

A possible approach for an estimate of the optical flow is to use a *1-D correlation scheme* between two successive frames, instead of 2-D correlation, as in [2]. The basic idea underlying the full 2-D correlation technique—we call it the 2-D–2-D technique (it is also called the *winner-take-all* method)—is to measure, for each desired location, the (x, y) shift that maximizes the correlation between 2-D patches centered around the desired location in successive frames. The patchwise correlation between the image at time t and at time $t + \delta t$ is defined as

$$\begin{aligned} \Phi(\delta x, \delta y; t) &\equiv I^w \otimes I \\ &= \int I^w(\xi, \eta; t) I(\delta x + \xi, \delta y \\ &\quad + \eta; t + \delta t) d\xi d\eta \end{aligned} \quad (1)$$

where $I^w(\xi, \eta; t)$ is the image at time t windowed to the patch of interest and set to 0 outside it. The L_2 distance has very similar properties to the correlation measure.² In the context of this article, minimizing the L_2 distance is exactly equivalent to maximizing the correlation (the observation is due to F. Girosi). As noticed before (Little et al. 1988), the previous idea can be regarded as an approximation of a regularization solution to the problem of computing the optical flow.³ Usually, one does not use grey values directly but rather some filtered version of the image, for instance through a Laplacian-of-a-Gaussian filter [see Little et al. (1988)], possibly at different resolutions.

Let us call $D(\delta_x, \delta_y)$ the L_2 distance between 2 patches in 2 frames at location (x, y) as a function of the shift vector (δ_x, δ_y) . The “winner-take-all” scheme




Fig. 1. The search space for the 1-D–2-D scheme used for the computation of the x and y components of the optical flow.

finds $\mathbf{s}^* = (\delta_x^*, \delta_y^*)$ that minimizes D [or maximizes the correlation function $\Phi(\delta_x, \delta_y)$] and assumes that the optical flow estimate is $\mathbf{u}^* = \mathbf{s}^*/\Delta t$, where Δt is the interframe interval.

It is natural to consider whether the component of \mathbf{u}^* along a given direction, for instance x , may be estimated in a satisfactory way simply by computing the δx that minimizes $D(\delta x, 0)$, that is the patchwise correlation as a function of x shifts only. We have found in our experiments that 1-D correlation of a 2-D patch provides estimates of δx^* are very close to the estimates obtained from the 2-D–2-D technique. We label this technique *1-D–2-D*, since it involves one-dimensional correlations on 2-D patches.

If we combine horizontal and vertical motion detectors of our 1-D, winner-take-all type (see Fig. 1), we obtain an appealing scheme to estimate the optical flow field at one point. The optical flow in one point is the vector sum of the x and y components computed by using such motion detectors. The key aspect of this approach is its reduction of the complexity of the problem, while maintaining a good estimation of the flow field: a complete two-dimensional search required in the winner-take-all scheme (Little et al. 1988) is reduced to two one-dimensional searches. Let us call v_{\max} the maximum velocity expected on the image plane. In (Little et al. 1988), the search space size to scan is $(2v_{\max} + 1)^2$ for each point; in our approach, its size is limited to $2(2v_{\max} + 1)$.

2 1-D Correlation of 1-D Patches

So far we have discussed that 1-D correlation of 2-D patches gives a satisfactory estimate of the optical flow between two successive frames, reducing the search space of corresponding points. This is equivalent to saying that the

$$\max_{\delta x} \Phi(\delta x, 0) \quad \text{and} \quad \max_{\delta y} \Phi(0, \delta y)$$

give a satisfactory estimate of

$$\max_{\delta x, \delta y} \Phi(\delta x, \delta y)$$

This suggests a further simplification: instead of $\Phi(\delta x, 0)$, consider a projection on x of $\Phi(\delta x, \delta y)$ obtained by some form of averaging operation on y , that is

$$\Phi * h_2$$

where h_2 is a 2-D filter such as a Gaussian elongated in the y direction and $*$ stands for the convolution operator. By well-known properties of the Gaussian function, h_2 can always be written as

$$h_2 = h * h$$

where h are Gaussian functions of appropriate variance. Assuming that we can neglect the patch size in the definition of Φ , we can write

$$\Phi * h_2 = (I_t * h) \otimes (I_{t+\delta t} * h) \quad (2)$$

where $I_t = I(x, y, t)$.

Thus, in the approximation of a large patch size, projecting the correlation function is equivalent to appropriately filtering the two images before correlation. Since it is usually better to discount the average intensity as well as small gradients through a high-pass filtering operation, in order to estimate the x -component of \mathbf{u} , we just perform a Gaussian smoothing in the y direction, as shown in Eq. (2), and then perform an additional convolution with the first or second derivative of a Gaussian function elongated in the x direction. Therefore the intensity function that is used in practice in the correlation operation is

$$I_t = (G_{\sigma_y}(y) * I_t) * G''_{\sigma_x}(x) \quad (3)$$

where σ_x and σ_y define the receptive field of such an elementary motion detector. After this filtering step, it is sufficient to evaluate the maximum of the correlation function only on 1-D patches to obtain an estimate of the x component of the flow. The previous argument

does not strictly apply to the L_2 distance measure we have used in our experiments. The very close similarity between correlation and distance, however, suggests a very similar behavior in both cases. We label this technique the *1-D-1-D* scheme since it involves 1-D correlations of 1-D patches.

3 A Crash Detector: The Green Theorem Scheme

As described by Poggio, Verri, and Torre (1991; see also Cipolla and Blake 1992), the divergence of the optical field $\nabla \mathbf{u}(x, y)$ is a differential measure of the local expansion

$$\nabla \mathbf{u}(x, y) = \frac{\partial u_x(x, y)}{\partial x} + \frac{\partial u_y(x, y)}{\partial y}$$

For a linear field (i.e., $\mathbf{u}(\mathbf{x}) = \mathbf{A}\mathbf{x}$), the divergence of \mathbf{u} is the same everywhere. In the case of linear fields (and all fields can be approximated by linear fields close to the singularity), the integral of the divergence over an area is invariant with respect to the position of the center of expansion. Green's theorems show that the integral over a surface patch S of the divergence of a field \mathbf{u} is equal to the integral along the patch boundary of the component of the field which is normal to the boundary ($\mathbf{u} \cdot \mathbf{n}$). In formula

$$\int_S \nabla \mathbf{u}(x, y) dx dy = \int_C \mathbf{u} \cdot \mathbf{n} dl \quad (4)$$

Therefore, since for a linear field $\nabla \mathbf{u} = 2/\tau$ where τ is the time to crash (TTC), a TTC detector that exploits the Green theorem just needs to sum over a closed contour, say a circle, the normal component of the flow measured at n points along the contour. We assume that the task is to compute *time to crash* (TTC) for pure translational motion. Possibly the simplest TTC detector of this type, shown in Fig. 3, is composed of just 4 elementary motion detectors. In this case, we have to sum the x component of \mathbf{u} for the horizontal detectors and the y component of \mathbf{u} for the vertical ones, with the correct sign.

Due to the invariance with respect to the position of the focus of expansion (or contraction) we can in principle arrange a certain number of TTC detectors, on the image plane (see Fig. 4). Our simulations suggest that one detector with a large radius (Fig. 3) is better than several smaller detectors (Fig. 4) in situations in which the whole visual field expands, probably because of better numerical stability of the estimates. Of course, a




Fig. 2. A TTC detector consisting of elementary motion detectors (see Fig. 1) at several locations along a closed contour. Each of the elementary motion detectors could be replaced by a single detector normal to the circle.




Fig. 4. A possible arrangement of TTC detectors in the image plane that is not as efficient as a single TTC detector with greater radius but has higher spatial resolution.

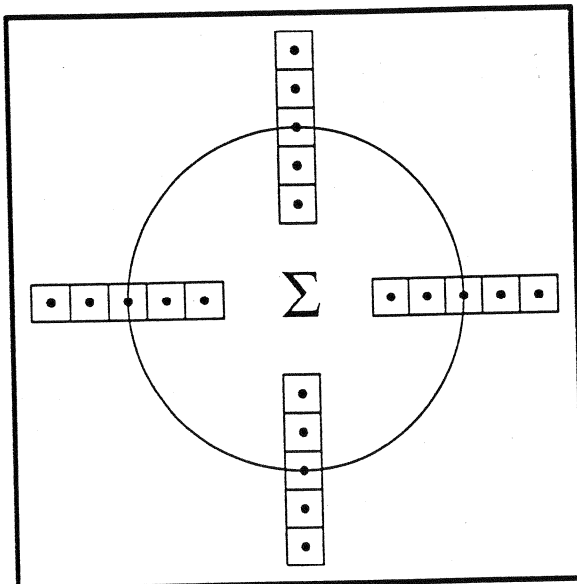


Fig. 3. Time-to-crash detector that exploits the Green theorem.

“large” detector has a poorer spatial resolution and this may be a problem in some applications (but not ours).

We have discussed so far schemes for detecting expansion. Similar arguments hold for rotation. The Green theorem relevant to this case is usually called

Stoke’s theorem and takes the form

$$\int_S \nabla \wedge \mathbf{u}(x, y) \cdot d\mathbf{S} = \int_C \mathbf{u} \cdot d\mathbf{r} \quad (5)$$

which says that the total flux of the differential measure of *rotationality* of the field $\nabla \wedge \mathbf{u}$ across the surface path S is equal to the integral along the boundary C of the surface patch of the component of the field that is tangential to the boundary. As described by Poggio et al. (1991), each elementary detector evaluates the tangential flow component at the contour of the receptive field (see Fig. 5). In this case, a detector has to compute the component of \mathbf{u} along the tangential direction at the contour.

4 Experimental Results

4.1 The 1-D-2-D Scheme

We have extensively tested our approach on real-image sequences. Each sequence was acquired from a camera mounted on a mobile platform moving at constant velocity. In all experiments the movement of the vehicle was a forward translation along a straight trajectory. We have verified the results obtained from our 1-D-2-D approach with the standard winner-take-all (2-D-2-D) scheme (Little et al. 1988; Christof et al. 1991).

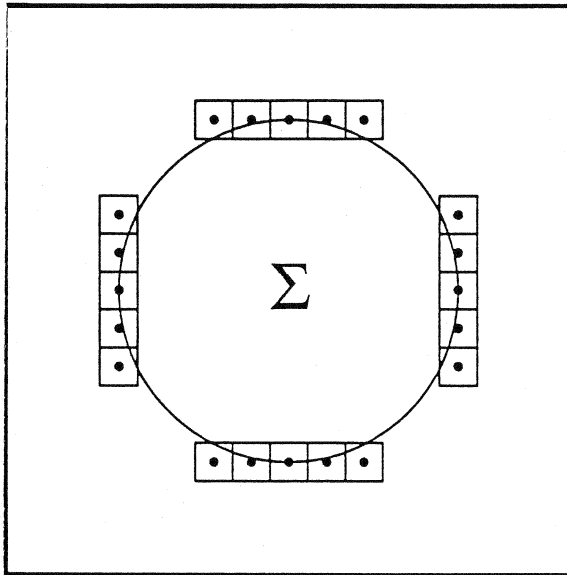


Fig. 5. Motion detector that exploits Stokes' theorem.

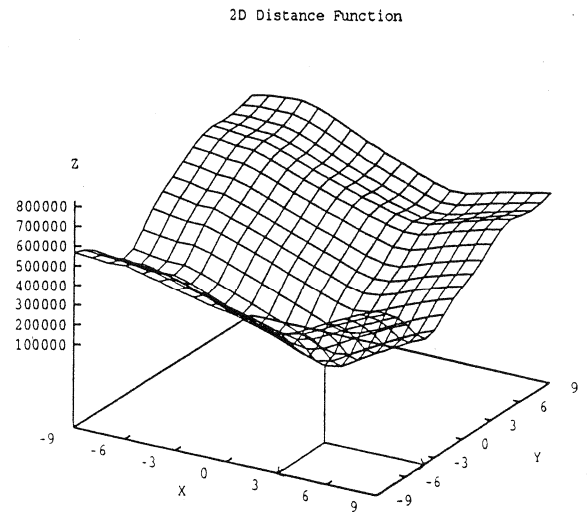


Fig. 6. (a) First and (b) last image of the sequence.

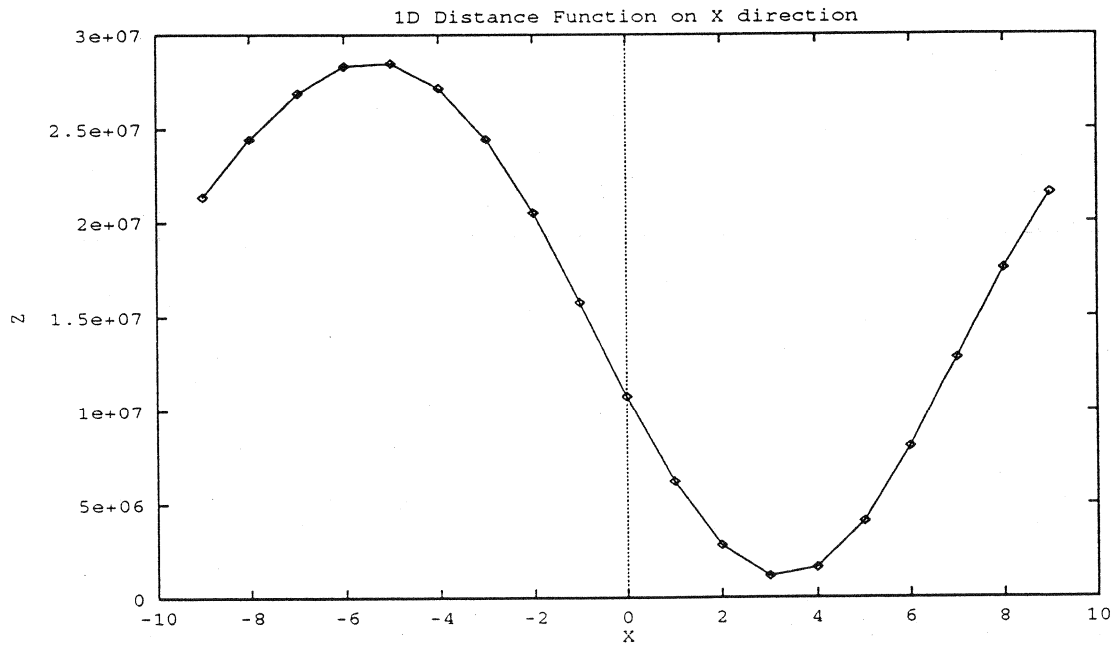


Fig. 7. An example of optical flows computed by using (a) 2-D-2-D, (b) 1-D-2-D, (c) 1-D-1-D. In most frame pairs in a sequence the three flows are much more similar to each other.

Figure 9 shows the first and last image of a sequence composed of 100 frames. Each image of the sequence is first convolved with a Gaussian filter having $\sigma = 0.5$. In both the algorithms we have used $v_{\max} = 9$ and $v = 20$ pixels, where v_{\max} is the maximum expected velocity of the points on the image plane and v is the

ray of the patch used for the evaluation of Φ . In other words, the correlation window used for the optical-flow computation is 41×41 pixels and the search space used is 19×19 by 2-D-2-D and $19 + 19$ by 1-D-2-D. Figure 10 shows the optical flows computed by the two methods using two successive images of the

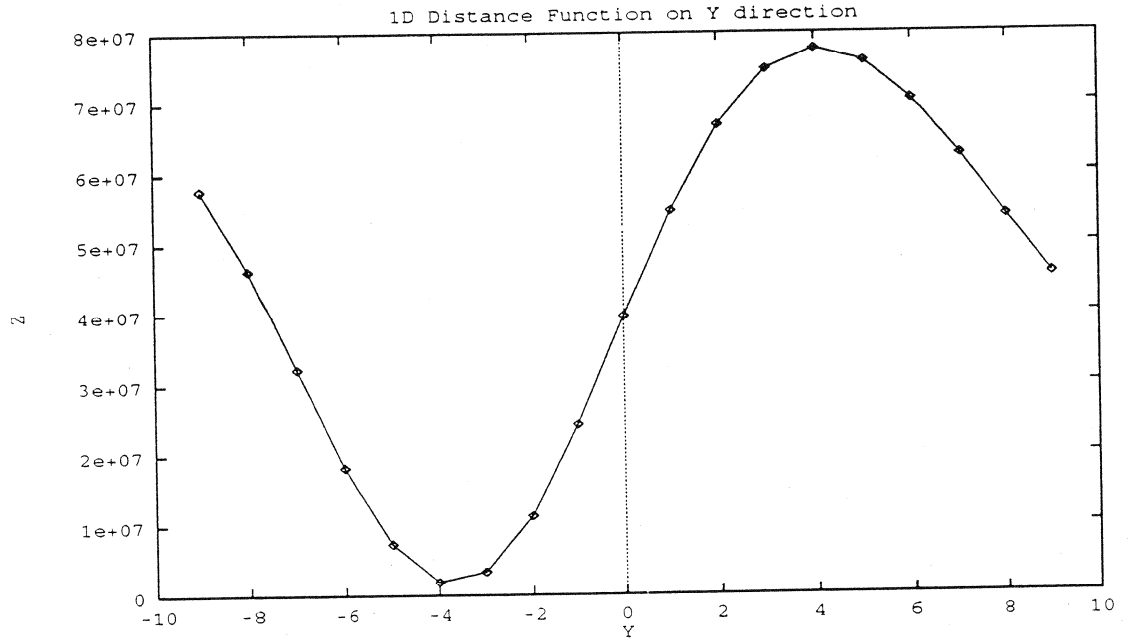


Fig. 8. TTC measurements comparison by using 2-D-2-D and 1-D-2-D. In this and the following figures the abscissa gives the time in terms of elapsed frames; the ordinate gives the estimate of the time to crash in frame units.

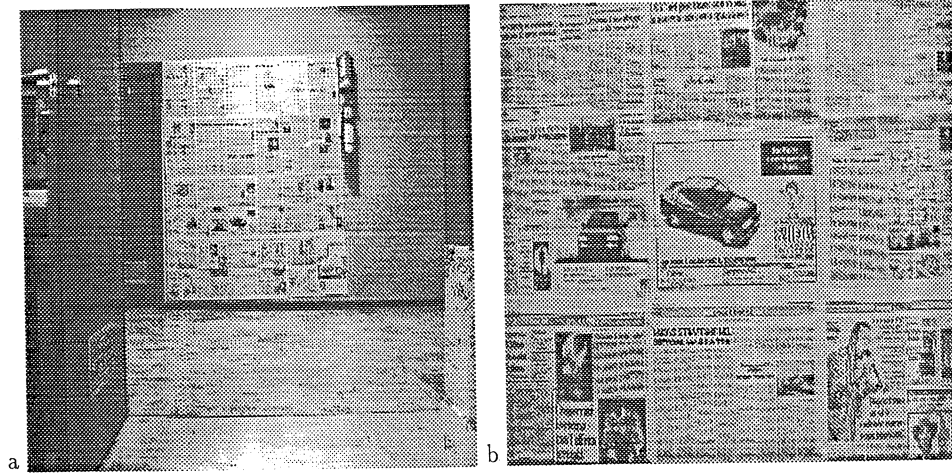


Fig. 9. 2-D distance function. The arrows indicate the position of the minimum.

sequence. The position of the focus of expansion was computed by using the approach described by Verri and Campani (1990).

We have used the methods described by Verri and Campani (1990) and Poggio et al. (1991) to verify the TTC estimation. To compute the TTC at a point by using the first method, we used an area of 81×81 pixels around that point. The points were 10

pixels apart. To compute TTC by using the method described by Poggio et al. (1991), we used a lattice of overlapping motion detectors. The distance between two points on the lattice was 10 pixels. Each detector had a receptive field of ray $r = 40$ pixels. In Fig. 11, we compare the results obtained by using the 2-D-2-D estimation of the optical flow with the 1-D-2-D one, by using the two different methods in the first stage of the

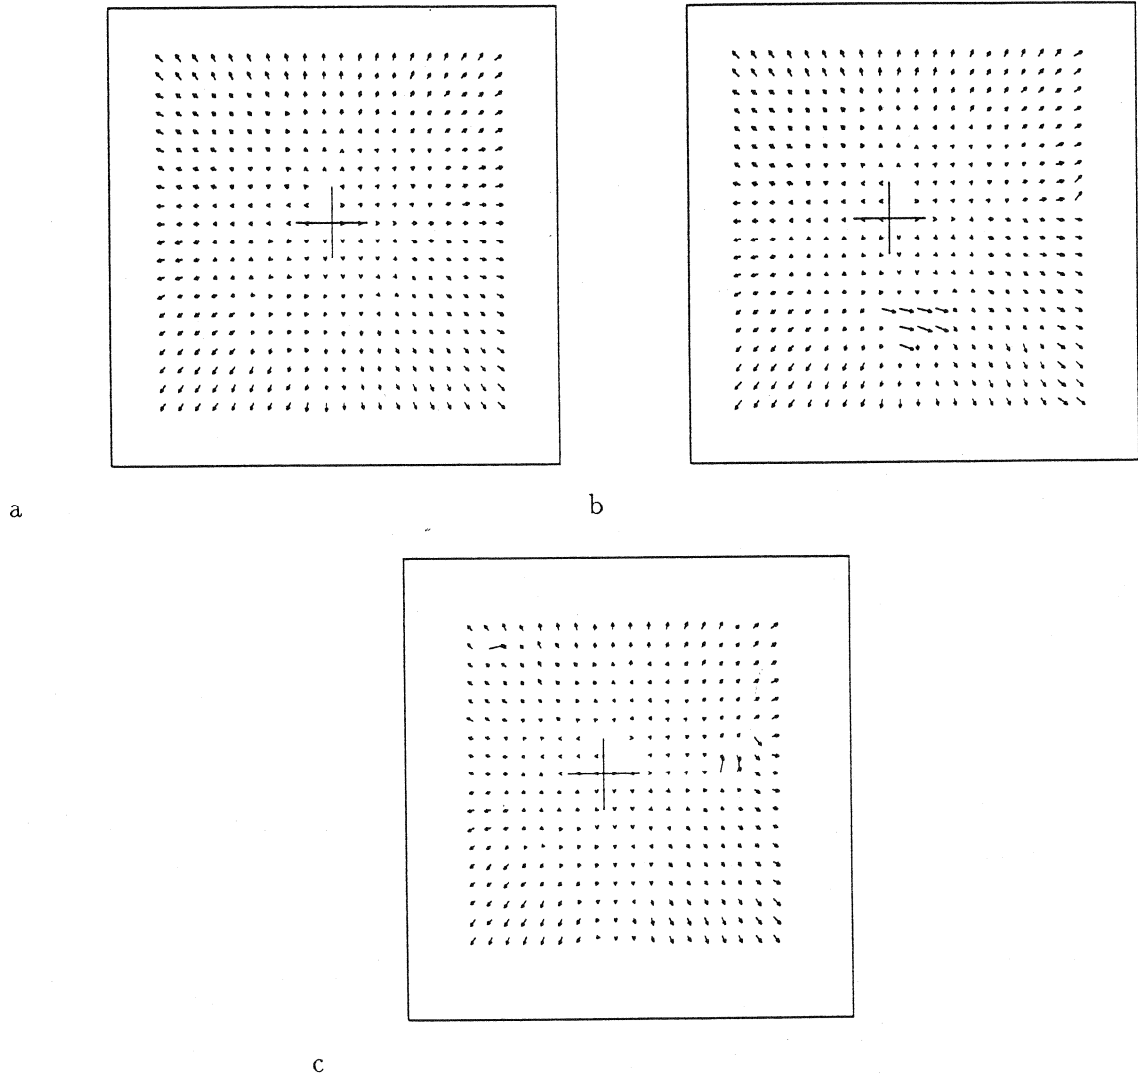


Fig. 10. 1-D distance function computed in the x direction by using $1-D-1-D$.

TTC. Performing a linear best fit on the TTC measurements, we obtain a slope of $m = -1.036$ by using the optical flows computed by $2-D-2-D$ and the method described by Verri and Campani (1990), and $m = -1.139$ by using the optical flows computed by $1-D-2-D$ and the method described by Poggio et al. (1991). Comparing the *true* TTC (straight line in Fig. 11) with the TTC measures obtained by using the second method, we estimate an absolute error in the mean of 2.63, with a standard deviation of 3.35 frame unit. In terms of rel-

ative units, the error in the mean is 5.7% with a standard deviation of 6.1%.

4.2 The $1-D-1-D$ Scheme

In this section we compare the results obtained by using $2-D-2-D$ and $1-D-1-D$. In both techniques, we have used $v_{\max} = 9$ and $v = 20$ pixels. In other words, the correlation window used for the optical flow estimation

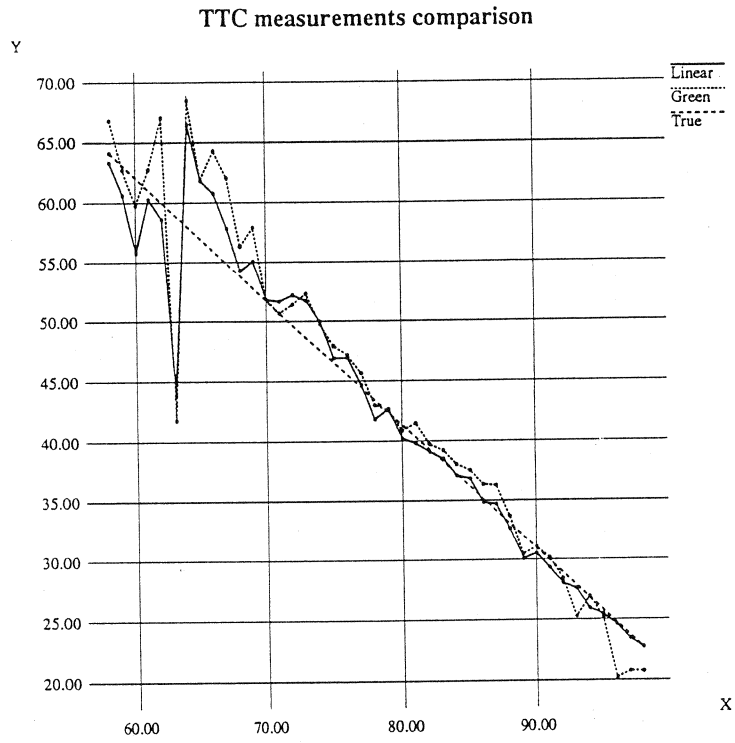


Fig. 11. 1-D distance function computed in the y direction by using 1-D-1-D.




Fig. 12. TTC measurements by using 1-D-1-D and a TTC detector with four elementary motion detectors.




Fig. 13. (a) First and (b) last image of the sequence.




Fig. 14. An example of a flow field computed at one point in time in the above sequence (Fig. 13) obtained by using (a) 2-D-2-D and (b) 1-D-1-D.

is 41×41 pixels for 2-D-2-D and 41 pixels for 1-D-1-D. In the filtering step, we have used $\sigma_y = 6$ and $\sigma_x = 3$ pixels for computing the x component of the optical flow. These values of σ produce a receptive field of an elementary motion detector equal to that used by 2-D-2-D (1681 pixels). Figure 9 shows a plot of the 2-D correlation function used in 2-D-2-D over a 2-D search space and a 2-D integration area. Figs. 10 and 11 show a plot of 1-D correlation functions used by the 1-D-1-D technique to estimate both components of the optical flow. In this case we used two 1-D search

spaces (in x and y directions respectively) and a 1-D integration area. Notice that this approach is capable of computing a reliable estimation of the flow vectors, while reducing the complexity of the problem.

Figures 13, 17, 21 show the first and the last image of three sequences acquired from a camera mounted on a mobile platform moving at constant velocity, along a straight trajectory. Figures 14, 18, 22 show the optical flows computed by the two methods, by using two successive frames of the sequences. The mean (continuous line) and the standard deviation (dashed line)




Fig. 15. Mean (dotted line) and standard deviation (dashed line) of the optical flow estimation.




Fig. 16. TTC estimation by using one TTC detector, with receptive field of $r = 80$ and 32 elementary motion detectors. The slope of the true TTC, computed by using the optical flows obtained by 2-D-2-D, is $m = -0.672$. The slope of the straight line, computed by using the TTC measures obtained by 1-D-1-D, is $m = -0.64$. A comparison of the TTC measures obtained by 1-D-1-D with the true TTC yields a mean absolute error of 9.02, with a standard deviation of 9.54. The relative error in the mean is 10.79% with a standard deviation of 9.49%. In order to evaluate the error in the time-to-crash estimation, the following steps have been performed. The true time to crash was estimated from a linear best fit of the TTC measures obtained by using the 2-D-2-D scheme for the optical flow estimation. The figures show the straight line that represents the theoretical behavior of the TTC. A linear best fit of the TTC measures obtained by using the 1-D-1-D scheme for the optical flow estimation was then performed in order to evaluate the slopes of the two straight lines. The absolute and relative error between the "true" TTC and the one measured by the 1-D-1-D scheme was then estimated. Let us call τ^* the true TTC. The absolute error is $E_a = |\tau^* - \tau|$ and the relative error is $E_r = |\tau^* - \tau|/|\tau|$.




Fig. 17. (a) First and (b) last image of the sequence.




Fig. 18. An example of a flow field obtained by using (a) 2-D-2-D and (b) 1-D-1-D.




Fig. 19. Mean (dotted line) and standard deviation (dashed line) of the error relative to optical flow estimation.




Fig. 20. TTC estimation by using one TTC detector, with receptive field of $r = 80$ and 32 elementary motion detectors. The slope of the *true* TTC, computed by using the optical flows obtained by $2-D-2-D$, is $m = -0.77$. The slope of the straight line, computed by using the TTC measures obtained by $1-D-1-D$, is $m = -0.83$. Comparing the TTC measures obtained by $1-D-1-D$ with the true TTC, we had a mean absolute error of 8.02, with a standard deviation of 8.97. With respect to the relative error we had a mean of 10.9% and a standard deviation of 9.72%.




Fig. 21. (a) First and (b) last image of the sequence.

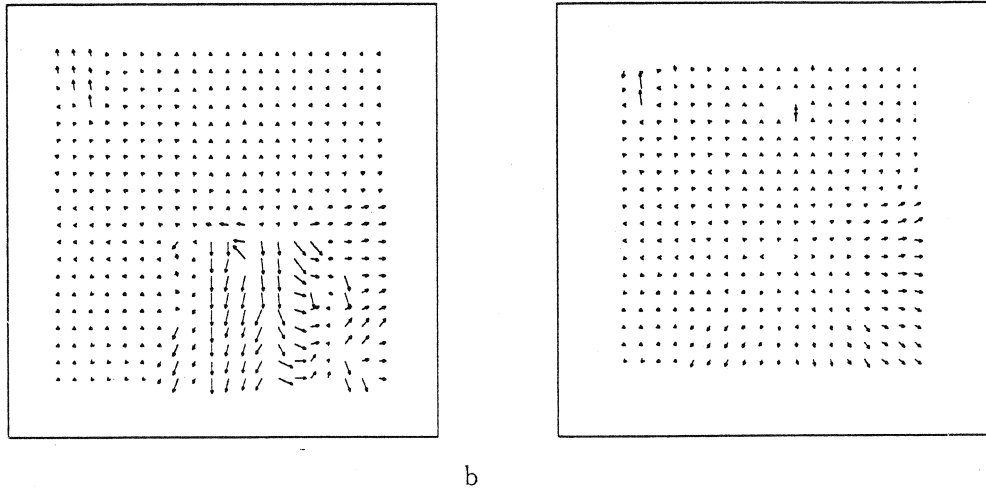


Fig. 22. Flow field obtained by using (a) 2-D-2-D and (b) 1-D-1-D.




Fig. 23. Mean (dotted line) and standard deviation (dashed line) of the error relative to optical flow estimation.




Fig. 24. TTC estimation by using one TTC detector, with receptive field of $r = 80$ and 32 elementary motion detectors. The slope of the true TTC, computed by using the optical flows obtained by *wta-2-D*, is $m = -1.24$. The slope of the straight line, computed by using the TTC measures obtained by *1-D-1-D*, is $m = -1.14$. Comparing the TTC measures obtained by *1-D-1-D* with the true TTC, we had a mean absolute error of 7.6, with a standard deviation of 7.9. With respect to the relative error, we had a mean of 11.4% and a standard deviation of 10.3%.

of the error on the optical flow estimation is shown in Figs. 15, 19, 23. Figures 16, 20, 24 show the TTC estimation by using the two different methods. In each experiment, we have used only one TTC detector, with receptive field of $r = 80$ pixel, composed by 32 elementary motion detectors (see Fig. 2).

5 Conclusions

5.1 Extensions of the Optical Flow Algorithm

There are several directions in which we plan to improve and extend our scheme:

- It may be possible to reduce further the number of sample points for D_p (i.e., the number of shifts) by using techniques for learning from examples such as the RBF technique (Poggio and Girosi 1990) to approximate $D_p(\delta x)$ as $D_p(\delta x) = \sum c_n G(x - t_n)$, and then find the minimum of D_p in terms of the dynamical system $dx/dt = -\epsilon \sum c_n G'(x - t_n)$. An alternative strategy is to try to learn directly the

function $\min D_p(x)$ from the samples of D_p , using a few examples of D_p typical for the specific situation. The conjecture is that the RBF technique may be able to learn the mapping $\min D_p(x)$ from examples of functions of the same class (compare Poggio and Vetter 1992). A similar idea is to try to learn how to sample the correlation function as a function of past sampled values. Again, the training examples would be functions of the same class. This would provide at each t an estimate of the most appropriate correlation shifts to try.

- Instead of simply measuring $D_{i,i-1}$, that is the distance between frame i and frame $i - 1$, we could measure in addition also $D_{i,i-2}$, $D_{i,i-3}$, ... and combine them in an estimate of the optical flow component by taking the average of $D_{i,i-1}/\Delta t$, $D_{i,i-2}/2\Delta t$, $D_{i,i-3}/3\Delta t$, etc. This technique may be improved further by using a Kalman filter.
- The basic scheme of Fig. 1 may be used to compute horizontal and vertical disparities among the two frames of a stereo pair.
- Confidence measures will be developed to further improve the performance of the technique.

5.2 Biological Implications of Our 1-D Technique

Poggio et al. (1991) conjectured that "the specific type of elementary motion detectors that are used to provide the estimate of the normal component of the flow is probably not critical. Radially oriented (for expansion and contraction), two input elementary motion detectors such as the correlation model (Reichardt 1961, Poggio and Reichardt 1973; van Santen and Sperling, 1984)—or approximations of it—are likely to be adequate. The critical property is that they should measure motion with the correct sign." Our results confirming their conjecture—since they suggest that 1-D correlation (or L_2 distance estimation) is sufficient for an adequate estimate of qualitative properties of the optical flow—have interesting implications for biology. Consider a 2-D array of Reichardt's detectors (for motion in the x direction) with spacing Δx and also detectors with spacings $2\Delta x$ etc. Take the sum of all detectors with the same spacing over a 2-D patch. Perform a winner-take-all operation on these sums. Select the set with optimal spacing as the one corresponding to the present estimate of optical flow. This scheme is analog

in time but otherwise equivalent to the one we have implemented. In formulae

$$\sum [I_i(t) - I_{i+k}(t - \Delta t)]^2$$

where Δt is the interframe interval in our implementation and is the delay in Reichardt's model,⁴ k represents the shift in our computation of D and represents the separation between the inputs to Reichardt's modules, $I_i(t)$ is the image value (in general spatially and temporally filtered) at location i and time t and the sum \sum is taken over the 2-D patch of detectors of the same type.

Thus an array of Reichardt's models with different spacings of the 2 inputs (in x) could be used in a plausible way to estimate the optical-flow component along the direction of the two-inputs detectors. Notice that a plausible implementation in terms of Reichardt's detectors of the 2-D correlation-based algorithm would be much harder, since it would effectively require detectors with all possible 2-D spacings. This seems implausible and contrary to experimental evidence where only a small number of separations and directions (as small as 3) seem to be present.

The above description is equivalent to our $1-D-2-D$ scheme and involves the summation over x and y "patches" of elementary 1-D motion detectors. In the fly this is plausible, given the known summation properties of specific wide field lobula plate cells.⁵ Our $1-D-1-D$ scheme on the other hand would require a summation over the x dimension only (in our example), but an oriented filtering of the image with receptive fields elongated in y before, the elementary motion detectors. It is possible that this second scheme may be used in the fly by different summation cells with smaller receptive fields. It is also possible that the wide field lobula plate cells effectively implement a scheme between the $1-D-2-D$ and the $1-D-1-D$ by using some oriented filtering before motion detection and limited y integration of the output of the elementary motion detectors. Similar considerations may apply to some of the motion-selective cortical cells.

5.3 The Time-to-Crash Detector

The TTC detector we have simulated is not the only possible scheme. Others possible schemes (see for instance Cipolla and Blake 1992) take into account more complex motions than just a frontal approach to a horizontal surface.

It is also conceivable that the scheme we suggest may be simplified even further in certain situations.

For instance, it may be sufficient in the summation stage to use the value of the correlation for a fixed (and reasonable) shift—instead of an estimate of the optical flow, that is the shift that maximizes correlation. This is equivalent to the direct use of the output of Reichardt's correlation nets instead of using the result of a winner-take-all operation on a set of Reichardt's nets with different spacings (or delays).

Another related idea is to continuously adjust the correlation shifts in order to track as closely as possible the maximum of the correlation (or the minimum of the distance): in this way it may be possible to reduce the computation of the correlation to just a few shifts, especially if time-filtering techniques are also used.

Acknowledgments

We are grateful to John Harris and Federico Girosi for many discussions and very useful comments.

Notes

1. In this article we use mainly the L_2 distance rather than the correlation itself. Since the two measures are equivalent for the purposes of this article, we will often use the terms correlation and distance in an interchangeable way.
2. The L_2 distance is in this case the square root of the sum of the squares of the differences between values of corresponding pixels. Other "robust" distance metric may be used, such as the sum of absolute values.
3. And in turn several definitions of the optical flow such as Horn and Schunk's, can be shown to be approximations of the correlation technique Poggio 1998.
4. We have written here the quadratic version of Reichardt's model; the same argument carries over to the standard model with multiplication: for the basic equivalence of the the quadratic and multiplication version (see Poggio and Reichardt 1973).
5. The patch would be very large and would correspond to the receptive field of the cell, that is its integration domain.

References

- Christof, K., Moore, A., Bair, W., Horiuchi, T., Bishofberger, B., and Lazzaro, J. 1991. Computing motion using analog vlsi vision chips: An experimental comparison among four approaches, *IEEE Workshop on Visual Motion*.
- Cipolla, R. and Blake, A. 1992. Surface orientation and time-to-contact from image divergence and deformation, *2nd Europ. Conf. Comput. Vis.* Santa Margherita Ligure-Italy, pp. 187-202.
- Little, J., Bulthoff H., and Poggio, T. 1988. Parallel optical flow using local voting, *Proc. 2nd Intern. Conf. Comput. Vis.*, Tampa, FL.

- Poggio, T. and Girosi, F. 1990. Networks for approximation and learning, *Proc. IEEE*, 78(9).
- Poggio, T., and Reichardt, W. 1973. Considerations on models of movement detection, *Kybernetik*, 13:223-227.
- Poggio, T., Verri, A., and Torre, V. 1991. Green theorems and qualitative properties of the optical flow, Tech. Rep. A.I. Memo No. 1289, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA.
- Poggio, T., Yang, W., and Torre, V. 1988. Optical flow: Computational properties and networks, biological and analog. In *The Neuron as a Computational Units Proceedings*, Scientific Applications International Corporation: Cambridge, UK, June.
- Reichardt, W. 1961. Autocorrelation, a principle for evaluation of sensory information by the central nervous system. In W.A. Rosenblith, Ed., *Principles of Sensory Communication*, Wiley: New York, pp. 303-317.
- Reichardt, W. 1987. Evaluation of optical motion information by movement detectors, *J. Comput. Physics*, 161:533-547.
- van Santen, J.P.H. and Sperling, G. 1984. Temporal covariance model of human motion perception. *J. Opt. Soc. Amer. A*, 1:451-473.
- Verri, A. and Campani M. 1990. An algebraic procedure for the computation of optical flow from first order derivatives of the image brightness, *Proc. IEEE Workshop on Robust Computer Vision*, Seattle.
- Verri, A. and Poggio, T. 1989. Motion field and optical flow: Qualitative properties, *IEEE Trans. Patt. Anal. Mach. Intell.*, 11:490-498.