Introducing a Smoothness Constraint in a Matching Approach for the Computation of Displacement Fields

P. Anandan and Richard Weiss

COINS Technical Report 85-38**
December, 1985

Abstract

Correlation matching techniques for computing displacement fields from successive frames in a dynamic image sequence are known to be error-prone. Our previous work [Anan84] has been concerned with identifying the sources of these errors and computing a confidence measure. Here we formulate a smoothness constraint that is useful for improving unreliable matches in an image based on the reliable ones. We note the relationship between our formulation and the gradient based formulation of the smoothness constraint. We provide a hierarchical matching algorithm that includes our smoothness constraint and show preliminary results of applying our algorithm to real images.

** The report is sponsored by Allen R. Hanson and Edward M. Riseman. It appeared in DARPA IU Workshop Proceedings, Miami Beach, Fl., 1985. This research was supported by DARPA under grant N00014-82-K-0464.

1 Introduction

Intensity based techniques for the computation of displacements of image points between a pair of image frames rely on the similarity of the light intensity reflected from a scene location in the two frames. Such techniques include *gradient techniques* and *correlation techniques.*

In the traditional formulation of the gradient techniques, an intensity constancy assumption is used to provide a partial constraint on the displacement vector at each point in the image. A smoothness constraint on the displacement field is then used to uniquely determine the displacements [Horn81,Hild85] of all the points. The correlation techniques usually provide a unique displacement vector based on local matches, without incorporating any smoothness constraint. Hence, obvious local errors in the displacement fields computed by the correlation techniques go undetected.

In our previous work [Anan84], we identified some common sources of matching-errors during correlation and provided a scalar confidence measure with the estimated displacement at an image point. This measure, which indicated the reliability of the associated displacement estimate, used information already available during the correlation process and hence did not require significant additional computation.

Although our previous measure was useful in determining the reliability of a displacement estimate, we indicated that since the displacement is a vector quantity, a vector valued confidence measure may be more appropriate. This paper provides such a measure and formulates a smoothness constraint on the displacement field that uses our new confidence measure.

Intuitively, we regard a displacement field to be "smooth" in an area of the image if its variation over the area is small. An example of a measure of the spatial variation of a displacement field is

$$E_{smooth}(U) = \int\int |\nabla U|^2 dx dy$$

where $U = (u, v)$ is the displacement vector at (x, y) (with u and v as its components in x and y directions respectively), and the domain of the integral is an area of the image (possibly the

whole image). This is due to Horn and Schunck [Horn81], who use this measure in a gradient-based approach for the computation of optical flow. Another example of such a measure is provided by Terzopoulos [Terz84], which we will describe in detail later in this document.

The formulation provided here is the minimization of the sum of two *errors*, E_{smooth} and E_{approx}. Given a displacement field, E_{smooth} measures its spatial variation and E_{approx} measures its deviation from the initial displacement vectors computed by the matching process.

Our choice of the approximation error is based on the results of a matching process. Each displacement vector is represented in a convenient local ortho-normal basis (e_{max}, e_{min}), which are usually not parallel to (x, y). For a given displacement field U, the approximation error is a weighted sum of the deviations of the components (along these basis vectors) of the displacement vectors in the field from the corresponding components of known initial values (provided by the matching process). The weights are the components of the vector-valued confidence measures along the basis vectors.

Intuitively, the basis vectors e_{max} and e_{min} and the weights c_{max} and c_{min} can be understood as follows: At a point along an edge in the image, the basis vector e_{max} will be approximately oriented in the direction normal to the edge, and e_{min} will be oriented parallel to edge. At such a point, the weight c_{max} will be large and the weight c_{min} will be small. On the other hand, in an area of the image with small intensity variations, both the weights will be small, whereas at a point along contour with high curvature, both the weights will be high. The precise definition of the basis vectors the weights are provided in section 5.

The precise form of the approximation error is

$$\begin{aligned} E_{approx}(U) &= \sum_{x,y} c_{max}(U \cdot e_{max} - D \cdot e_{max})^2 \\ &+ \sum_{x,y} c_{min}(U \cdot e_{min} - D \cdot e_{min})^2 \end{aligned}$$

where c_{max} and c_{min} are the weights, U is the displacement vector and D is the initial displacement vector provided by the correlation matching algorithm, at location (x, y) in the image. Here c_{max} and c_{min} indicate the confidence in the components of the displacement vector D in the directions

e_{max} and e_{min} respectively.

In the rest of this paper, we develop this idea and formulate it as a relaxation process. We extend the analysis of Terzopoulos [Terz84] for the surface-reconstruction problem to our two dimensional minimization problem. We follow a finite-element approach to solving the problem, primarily because this approach enables us to deal with arbitrary, known discontinuities in the displacement field.

We also incorporate the smoothness process within the framework of a hierarchical algorithm for the computation of displacement fields. Our hierarchical algorithm is similar to the multi-frequency coarse-fine techniques described in [Glaz83,Anan84]. At each level of the hierarchy, we include the smoothness process after computing displacement estimates by correlation a matching process.

The relationship of this minimization formulation to other formulations is described in section 2. In section 3, the conditions for the existence of a solution to this problem are explained, and in section 4 the finite element approach is used to provide a discrete relaxation algorithm. The source of the weights c_{max} and c_{min}, and the basis vectors e_{max} and e_{min} is discussed in section 5. The incorporation of this smoothness constraint in the hierarchical algorithm is described in Section 6. Finally, some experimental results are provided in section 7 and the scope for future work is discussed in section 8.

2 Relationship to other work

The minimization problem posed in the previous section has its roots in other similar work in computer vision. In particular, the formulation of the intensity and smoothness constraints for the computation of optic flow fields by Horn and Schunck [Horn81], the related work by Nagel [Nage83,Nage84] and Hildreth [Hild85], and the surface reconstruction problem posed by Terzopoulos [Terz84] are closely related to our formulation.

First, we explain the relationship of our work to that of Horn and Schunck. Our smoothness

constraint can be chosen to be the same as theirs. The intensity constraint they use is

$$\nabla I \cdot V + \frac{\partial I}{\partial t} = 0 \tag{1}$$

where ∇I is the spatial gradient of the image intensity I at a point (x, y) in the image and V is the image-velocity at that point. Under the assumption that the time interval δt between successive image frames is small, we can approximate V by $D/\delta t$, where D is the displacement of an image-point, and $\partial I/\partial t$ by $\Delta I/\delta t$. Then, we can rewrite equation 1 as

$$\nabla I \cdot D + \Delta I = 0 \tag{2}$$

Based on this, we can define an error

$$E_{approx} = \sum_{x,y} (\nabla I \cdot U + \Delta I)^2$$

which is zero when $U = D$ where D is any value that satisfies the equation 2 above.

This error can be rewritten as

$$E_{approx} = \sum_{x,y} |\nabla I| (U \cdot e_{\nabla I} + \frac{\Delta I}{|\nabla I|})^2$$

where $e_{\nabla I}$ is the unit vector in the direction of ∇I.

From equation 2, we see that $D \cdot e_{\nabla I} = -\Delta I/|\nabla I|$. Considering our E_{approx} term described in section 1, if we set $c_{min} = 0$, $c_{max} = |\nabla I|^2$, $e_{max} = e_{\nabla I}$, our approximation error is the same as that of Horn and Schunck. Indeed, their intensity constraint simply provides values for one component of the displacement vector at any point in the image, viz., the component in the direction of the intensity gradient. This is because only the first order image intensity variations are considered.

Nagel [Nage83] points out that at certain areas in the image it may be possible to locally obtain values for both components of the displacement vector. This is usually true at points of high curvature along image-contours, or in textured areas, where the second order intensity variations

are large and can be useful for obtaining the unique displacement vector. In section 5, we explain how a correlation matching algorithm behaves at such areas of the image.

Both our formulation and the formulation of Horn and Schunck can be regarded as a two dimensional version of the smoothness process presented by Terzopoulos [Terz84]. Terzopoulos is interested in visual surface reconstruction. This leads to the problem of smoothing a scalar variable u, the depth of the visible surface along each viewing direction. Approximate values for the depth and associated confidence measures are assumed to be known at some locations in the image. Similarly, approximate orientation of the surface and associated confidence measures are also assumed to be known at selected locations.

The problem is that of minimizing the sum of three errors $E_{smoothness}$ and E_D, and E_O where $E_{smoothness}$ is due to a surface smoothness constraint, E_D is due to the known approximate depth values, and E_O is due to known approximate orientations of the surface. Terzopoulos considers two possible measures of the spatial-variation of the depth values. One is based on what is called a "membrane" model of the visual surface and the other is based on a "thin-plate" model of the surface. From our point of view, the two models are different simply with respect to the functional they minimize, which we will describe in detail in the next section. In attempting to extend this model to the two dimensional displacement fields, we ignore the orientation constraint, because we do not have any such information regarding our vector fields. For details, we refer the reader to [Terz84], Chapter 4.

In what follows, we provide a mathematical analysis of our problem and derive a method to iteratively solve it. Our mathematical analysis borrows heavily from that of Terzopoulos (see [Terz84] Chapters 5 and 6), with suitable modifications to deal with the two dimensional vector field.

3 Mathematical formulation

The problem of finding a smooth displacement field which approximates the estimated displacements on a discrete subset can be formulated as a minimization problem. That is, we are trying to find a vector field $U = (u, v)$ which minimizes a quadratic functional $E(U)$ where $E(U) = E_{smooth}(U) + E_{approx}(U)$. As mentioned above, we consider two choices for E_{smooth}. E_{approx} measures how well U approximates the data given at a set of grid points (x, y) in D.

The two functionals which we consider have the property that under certain weak conditions there will always exist a unique minimum solution. The proof of this is based on the following theorem:

If $E(U) = B(U,U) - f(U)$, where B is a symmetric, bilinear form on a Banach space and f is a linear form, then E has a unique minimum if B is positive definite (in that case, E is said to be elliptic).

We will derive conditions on E_{approx} for each of the two smoothness constraints so that the positive definiteness condition is satisfied.

3.1 Smoothness constraint for thin plate model

This functional, which is minimized by an ideal thin plate, can be extended to a vector field $U = (u, v)$

$$\begin{aligned} E_{smooth}(U) &= \int\int (u_{xx}^2 + 2u_{xy}^2 + u_{yy}^2) \\ &+ (v_{xx}^2 + 2v_{xy}^2 + v_{yy}^2) dx dy \end{aligned} \tag{3}$$

In order to apply the theorem, we decompose E into quadratic and linear terms. E_{smooth} is purely quadratic and E_{approx} has both linear and quadratic terms, so

$$B = B_{smooth} + B_{approx} \tag{4}$$

and

$$\begin{aligned} E_{approx}(U) &= \sum c_{max}(U \cdot e_{max} - D \cdot e_{max})^2 \\ &+ c_{min}(U \cdot e_{Min} - D \cdot e_{min})^2 \end{aligned} \tag{5}$$

So

$$B_{approx}(U) = \sum c_{max}(U \cdot e_{max})^2 + \sum c_{min}(U \cdot e_{min})^2 \tag{6}$$

We have confidence measurements c_{max} and c_{min} at each point, such that c_{max} is the confidence in direction e_{max} and c_{min} is the confidence in direction e_{min}, where e_{max} and e_{min} form an orthonormal basis for the (x, y)-plane, and $0 \leq c_{min} \leq c_{max} \leq 1$.

B is positive definite if $B(U, U) = 0$ implies that $U = 0$. So we need to find conditions such that if $E_{smooth}(U) = 0$ and $B_{approx}(U) = 0$ then $U = 0$.

The smoothness constraint in particular implies that all of the second partial derivatives of u and v are 0, so u and v are linear

$$U = (ax + by + c, dx + ey + f) \tag{7}$$

We would like to show that all of these constants are zero. For example, if there are three non-collinear grid points which are corner points, then c_{max} and $c_{min} > 0$, and we have 6 independent equations for the coefficients. If on the other hand e_{max} always points in the same direction and c_{min} is always 0, then we cannot get independent equations, and we know that the answer will not be unique. Let $e_{max}(x, y) = (cos\theta(x, y), sin\theta(x, y))$ then $e_{min}(x, y) = (-sin\theta(x, y), cos\theta(x, y))$

The equations are of two types :

$$axcos\theta + bycos\theta + ccos\theta + dxsin\theta + eysin\theta + fsin\theta = 0 \tag{8}$$

or if $c_{min} > 0$

$$-axsin\theta - bysin\theta - csin\theta + dxcos\theta + eycos\theta + fcos\theta = 0 \tag{9}$$

The condition that will guarantee a unique solution is that one can find 6 independent equations from the set of equations defined by the points at which the vector field has been approximated. In particular, three non-collinear points with $c_{min} > 0$ or 6 points with $c_{max} > 0$ and which satisfy the following two conditions

1. e_{max} points in the same direction for no more than three points.

2. if e_{max} lies in the same direction for any subset of three points then they are not collinear.

3.2 Smoothness constraint for the membrane model

The smoothness constraint for this model is given by

$$E_{smooth}(u,v) = \int\int |\nabla u|^2 + |\nabla v|^2 dxdy \tag{10}$$

If $E_{smooth}(U) = 0$ then all of the first partial derivatives are 0, so U must be constant, i.e. $U = (a,b)$ Thus if two independent equations can be derived from B_{approx}, then U must be 0. This happens if there is a corner point or there are two points with different e_{max} vectors. Thus, this condition implies that B is positive definite and therefore E has a unique minimum.

4 Solving the variational problem

In the previous section we showed that there is a unique smooth vector field which minimizes the constraints; in this section we present methods to obtain a discrete approximation to this solution. The choice of domain and basis functions here is not one of the standard ones and is based on the approach of Terzopoulos. We have extended his results to two dimensional flow fields, and we present the masks which represent the solution. We intend to investigate a more standard choice of basis for the finite element technique.

We take as our discrete solutions piecewise polynomials; each polynomial being defined on a domain D in the plane. For the membrane model, we choose a triangular finite element as shown in figure 1. For the thin plate, we choose a domain which consists of a set of six points on a square grid (see figure 2).

There are two requirements for the finite element approach; that a unique solution exist for each particular grid, and that the solutions converge as the mesh size goes to zero. Our choice of the finite-elements follows that of Terzopoulos. He verifies that these elements satisfy these two requirements.

4.1 Computation of masks

In order to solve the discrete minimization problem, linear equations in the values at the grid-locations are derived. These equations are used to update the values $U = (u, v)$ at a point in terms of its neighbors. This is commonly done by convolution with a mask; the masks for each of the two models are given below:

For the membrane model, the convolution mask is

$$\frac{1}{4} \times \begin{array}{ccc} & 1 & \\ 1 & 0 & 1 \\ & 1 & \end{array}$$

For the thin-plate model it is,

$$\frac{1}{20} \times \begin{array}{ccccc} & & -1 & & \\ & -2 & 8 & -2 & \\ -1 & 8 & 0 & 8 & -1 \\ & -2 & 8 & -2 & \\ & & -1 & & \end{array}$$

Let us denote the convolution mask as A. Also, let $\bar{U} = A * U$. Then, solving the discrete problem is the same as solving the following system of coupled equations:

$$\begin{aligned} (U - \bar{U}) \; &+ \; c_{max}(U \cdot e_{max} - D \cdot e_{max})e_{max} \\ &+ \; c_{min}(U \cdot e_{min} - D \cdot e_{min})e_{min} \; = \; \vec{0} \end{aligned} \tag{11}$$

for each point on the image grid.

4.2 Relaxation algorithm

There are a number of numerical-methods for solving the system of coupled linear equations described above. One of the simplest methods is the Gauss-Seidel relaxation algorithm. This is an iterative process, where during each iteration the value of U at each point in the image is solved in terms of the values of its neighbors.

In our case, we have

$$\begin{aligned} U^{n+1} = U^n &+ \frac{c_{max}}{c_{max}+1}((D-U^n)\cdot e_{max})e_{max} \\ &+ \frac{c_{min}}{c_{min}+1}((D-U^n)\cdot e_{min})e_{min} \end{aligned} \tag{12}$$

where the superscripts denote the number of the iteration.

From a geometric point of view, this updating scheme can be regarded as choosing a point in the (u,v) space that is a combination U^n and D. We illustrate this idea in figure 3. For convenience, we have chosen to represent the displacements in a cartesian coordinate system with its axes parallel to (e_{max}, e_{min}). The two key parameters are $c_{max}/(1+c_{max})$ and $c_{min}/(1+c_{min})$, which vary between 0 and 1, as c_{max} and c_{min} vary between 0 and ∞. Since $c_{max} \geq c_{min}$, the location of U^{n+1} will always be on or above the line joining D and U^n in figure 3. In particular, it can be seen that the location of U^{n+1} will be always within the triangle that is shown in that figure, moving towards the line joining D and U^n as c_{min} gets closer to c_{max}.

Finally, we note that although the Gauss-Seidel relaxation scheme is simple to implement, there are more efficient techniques to solve a sparse system of equations. Studying these and choosing an efficient, parallel scheme will be a part of our future work.

5 The basis vectors and the confidence measures

One of the key elements of our formulation is the approximation constraint. The proper choice of the ortho-normal basis vectors e_{max} and e_{min} and the confidence measures c_{max} and c_{min} are crucial to our model and to our algorithm. In this section we discuss how these measures can be obtained.

In attempting to choose these measures, it is useful to formalize our efforts in the form of a set of design considerations. These considerations are,

1. If both components of the two dimensional displacement vector are reliably known at an image location, both c_{max} and c_{min} should be large. If one component is reliably known and

the other is not, then the c_{max} should be large and c_{min} small, and e_{max} should be oriented along the direction of the reliable component. If both are unreliable then both c_{max} and c_{min} should be low.

2. The computation of c_{max} and c_{min} should take minimum effort in addition to what is already necessary for the computation of the displacement vectors.

In our previous study on the behavior of *error* (sum of squared-differences) surfaces [Anan84], we noted that the shape of these surfaces contained useful information regarding the reliability of the displacement vectors. Below, we explain what we mean by these terms, and briefly summarize our major observations.

The process of matching by discrete correlation consists of choosing a window around a point of interest in one frame (the first window), and correlating the intensity values in that window with those in the candidate match windows in the second frame (the second window). The term *correlation* is used in a generic sense here. The measure that we prefer (for reasons explained in [Anan84]) is the sum of squared differences between corresponding pixel values (SSD).

By the term *error surface*, we mean the surface whose height is the SSD value corresponding to each possible displacement. Hence, for each location of interest in the first frame, we have available an error surface. The best match is, indeed, the displacement corresponding to the minimum height of the error surface (henceforth called the *pit*).

In our study of the error surfaces [Anan84], we also noted that the curvature of the error surface around the pit, taken along any direction seems to provide information regarding the reliability of the match along that direction. In particular, we observed the following :

1. At a corner point, the pit is sharp in all directions; correspondingly all the directional curvatures are high.

2. At a point along an edge, the error surface shows a long valley like structure, and the orientation of the valley corresponds to that of the edge. This means that the curvature in the

direction parallel to the edge is low, whereas that in the direction perpendicular to the edge is high.

3. At a point in a homogeneous area all curvatures are low, and the error surface looks rather flat.

Figures 4, 5, and 6 show examples of error surfaces at a corner-point, an edge-point, and a homogeneous point respectively. The surfaces are shown inverted, in order to enhance visibility. Note that in a small area around the point with minimum error (the peak in these figures), the surfaces demonstrate the properties described above.

The directional curvatures of the error surface seem to satisfy the first of our design criteria. Further, the computation of the error surface is a necessary part of the matching process, and the computation of the curvatures requires only minimal additional effort. Hence, this meets both our design criteria. Based on these, our method of choosing c_{max} and c_{min} are as follows :

1. Compute the principal curvatures and the directions of the principal axes around the pit of the error surface. Let C_{max} and C_{min} be the principal curvatures and E_{max} and E_{min} be the unit vectors along the directions of the principal axes.

2. Compute the error corresponding to the location of best match. Let this be SSD_{min}.

3. We choose $e_{max} = E_{max}$ and $e_{min} = E_{min}$. For c_{max} and c_{min}, we have a number of choices:

   (a) We can simply take $c_{max} = C_{max}$, and $c_{min} = C_{min}$.

   (b) We can normalize the principal curvatures to be within some desired range and to be scaled by a desired factor, and choose $c_{max} = NORM(C_{max})$ and $c_{min} = NORM(C_{min})$. Here $NORM(.)$ is the normalizing function of the curvature.

   (c) We can normalize the principal curvatures using SSD_{min}. We can choose

$$c_{max} = \frac{C_{max}}{k_1 C_{max} + k_2 SSD_{min} + k_3}$$

and

$$c_{min} = \frac{C_{min}}{k_1 C_{min} + k_2 SSD_{min} + k_3}$$

This has the effect of making the confidence inversely proportional to the error corresponding to the best match. Usually this error is high when the local physical structure represented in the image changes, due to expansion, rotation, etc. of the image, or if the image has been corrupted by noise. Strictly speaking, correlation (or SSD) provides truly reliable matches only when these effects are absent. Hence, this type of normalization takes into account problems that are unexpected, but influence the error measure.

The choice of the normalization of the principal curvatures to obtain the confidence measures is a crucial factor. For the experiments described in this paper, we used a normalization of the type discussed in 3(c) above, choosing $k_1 = 0, k_2 = 1$, and $k_3 = 100$. These choices were made on an empirical basis, although we observed that varying them did not significantly affect our results. This issue is open for further research.

6 A hierarchical matching algorithm with smoothing

The foregoing discussion regarding matching and smoothing and the associated confidence measures is based on the strong assumption that the displacements are small in magnitude compared to the scale of the image and the spatial rate of image intensity variations. It has been noted by other researchers [Burt83,Glaz83] that in order to realistically deal with general images with larger displacements, it is desirable to use a hierarchical matching approach. This has the advantage of achieving both savings in computation and the use of image intensity variation at a scale appropriate for the amount of motion. information Our technique here is based on the hierarchical correlation algorithm described in [Glaz83] and [Anan84].

Briefly, the hierarchical matching/smoothing algorithm is as follows:

- Each of the two images are processed with a set of band-pass filters and reduced in resolution. The band-pass filters are isotropic difference of Gaussian filters, one octave wide and one octave apart from each other. These are implemented based on the ideas suggested by Burt ([Burt81].

- The matching begins at a level of resolution corresponding to a displacement of less than 1 pixel. The process of correlation is described in [Glaz83,Anan84]. Sum of squared-differences is used as the match measure. For the experiments of this paper, we chose square windows 5 pixels wide.

- After the matching at one level, confidence measures are computed as described in the previous section and smoothing is performed. For the experiments of this paper, we used a membrane model and finite element method for smoothing.

- The smoothed vector field is projected to the next finer-level image, based on the modified projection strategy described in [Anan84]. These are used as initial values for the matching/smoothing process at this finer level.

- This process is repeated at all levels up to the level of the image.

7 Experiments

We performed two experiments to illustrate our algorithm. These experiments are meant to be illustrative of the ideas presented in this paper. In this sense these results should be regarded as preliminary. However, it will be easy to see from these examples that the smoothness process reduces the errors in the initial match estimates.

The first experiment demonstrates the effect of applying our smoothness algorithm at a single level of resolution and computing a dense displacement field. The second experiment demonstrates the use of our confidence measure in a hierarchical matching/smoothing algorithm. In both cases, we used the membrane model for the smoothness constraint.

The first experiment was performed by using a single real image and generating a second image by digitally translating the first image. This enabled us to have "ground-truth" displacement data for comparison. The second experiment was performed using a pair of images from a *real* image sequence called *road-scene image sequence.*

7.1 Testing the single level smoothing algorithm

For this experiment, we used the image shown in figure 7 as the first of the pair of images. The second image was obtained by rotating the first image clockwise by 4 degrees on the plane of the image about the center of the image. The initial results of the matching process, prior to smoothing are shown in figure 8. The results after 100 iterations of the smoothness process are shown in figure 9. Although these results are useful to gain a qualitative understanding of the effect of the smoothness process, it is more useful to consider the statistics of the error vectors (i.e, the difference at each pixel between the true flow vector and the computed vector). These are shown in the table in figure 10.

Note that we have shown component-wise error statistics as well as the statistics of the length of the error vector. The component-wise statistics are based on a local decomposition of the error vectors in the direction of e_{max} and e_{min}. We use these in order to illustrate the point that most of the error is usually in the direction of e_{min}.

7.2 Testing the hierarchical algorithm

Our second experiment involved performing the hierarchical matching algorithm with smoothing on a pair of images from the *road scene image sequence*, which is a sequence of images of a road scene obtained from a translating camera. The two images are shown in figures 11 and 12.

The displacement of any point between these two images is less than 10 pixels. This implies that four levels of the hierarchy are sufficient for processing. At each level, we performed 10 iterations of the relaxation algorithm. This number was chosen arbitrarily.

In figures 13 and 14, we display the displacement fields at the four levels of the hierarchy. The dramatic improvement in the final results are the results of the integration of the matching and the smoothing process in the hierarchy. Smoothing at one level of resolution will not provide such improvement.

In figure 15 we display the confidence measures c_{max} and c_{min} as intensity images. The brightness of a point in these images are proportional to the confidence measures. We have also superimposed the unit vectors e_{max} on the figure containing c_{max}, in order to demonstrate that it is usually perpendicular to the edges in the image.

For the purposes of a closer scrutiny, in figure 16 we display the results only at the image-resolution.

8 Future Work

In this paper, we provided an outline of a technique for incorporating a smoothness constraint in a matching algorithm. We also indicated how this may be incorporated in a hierarchical technique. Although we demonstrated these ideas with some experiments, many details of the algorithm are yet to be worked out. In particular, we feel the following areas need greater analysis and closer scrutiny.

1. The choice of the smoothness criterion. The membrane model as well as the thin plate model are heuristics for the type of variation we expect the displacement field to possess. It may be possible to limit the choice of the smoothness criterion based on an understanding of the geometric structure of continuous flow fields. Such analyses abound in the literature (e.g., [Adiv85,Waxm84]).

2. The normalization of the confidence measures. One of the important factors that affect the outcome of the process is the validity of the confidence measure. At the outset, a conservative measure seems more appropriate, i.e., one that is more prone to doubt the reliability of the local match estimate. This will help us avoid the process from being misguided by incorrect

values with high confidence. In general, a more thorough investigation is critical to make this process useful in real images.

3. One of the motivations behind choosing the finite element approach for solving the minimization problem is that this approach allows us to deal with known motion and occlusion boundaries. Therefore, the recognition of such boundaries, and the inclusion of that information in the smoothness process is one of goals of future research. We seek an understanding of what information is available for recognizing such events and how they may be utilized.

We believe that the hierarchical matching/smoothing algorithm is a useful technique for the computation of dense, reliable displacement fields in real images. Many issues remain to be solved. However, the approach takes into consideration the limits of its various components, and uses them appropriately. This provides us the motivation for further investigation along this line.

Acknowledgments

We wish to thank Dr. George Reynolds for his comments on the manuscript, and Frank Glazer for many valuable discussions. Thanks are also due to Brian Burns and Bill Guazzo for their help in obtaining the figures for the paper.

References

[Adiv85] Adiv, Gilad, Determining 3-D Motion and Structure from Optical Flow Generated by Several Moving Objects, *IEEE PAMI*, Vol. PAMI-7, July 1985, pp. 384–401.

[Anan84] Anandan P., Computing Dense Displacement Fields with Confidence Measures in Scenes Containing Occlusion, *SPIE Intelligent Robots and Computer Vision Conference*, Vol. 521, pp 184-194, 1984, also *COINS Technical Report 84-32*, University of Massachusetts, December 1984.

[Burt81] Burt, P. J., Fast Filter Transforms for Image Processing, CGIP vol. 16, pp 20-51, 1981.

[Burt83] Burt, P. J., Yen C. and Xu X., Multi-Resolution Flow-Through Motion Analysis, *IEEE CVPR Conference Proceedings*, June 1983, pp. 246-252.

[Glaz83] Glazer, F., Reynolds, G. and Anandan, P., Scene Matching by Hierarchical Correlation, *IEEE CVPR conference*, June 1983, pp. 432-441.

[Hild85] Hildreth, E. C., The Measurement of Visual Motion, *PhD dissertation*, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, Ma., 1983.

[Horn81] Horn, B. K. P. and Schunck B. G., Determining Optical Flow, *Artificial Intelligence* Vol. 17, pp. 185-203

[Nage83] Nagel H. H., Displacement Vectors Derived from Second-Order Intensity Variations in Image Sequences, *Computer Vision, Graphics, and Image Processing*, 21, pp 85-117, 1983.

[Nage84] Nagel H. H., Constraints for the Estimation of Displacement Vector Fields from Image Sequences, *IJCAI-83*, Karlsruhe, W. Germany, pp 945-951, 1983.

[Terz84] Terzopoulos, D., Mutiresolution Computation of Visible-Surface Representations, *Phd Dissertation*, Massachusetts Institute of Technology, Jan. 1984.

[Waxm84] Waxman A. and Wohn K., Contour Evaluation, Neighborhood Deformation and Global Image Flow: Planar Surfaces in Motion, *CS-TR-1394* University of Maryland, April 1984.

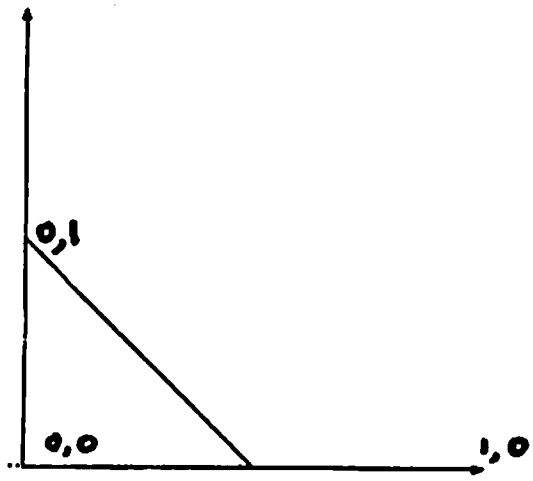

Figure 1: The finite-element domain for the membrane model

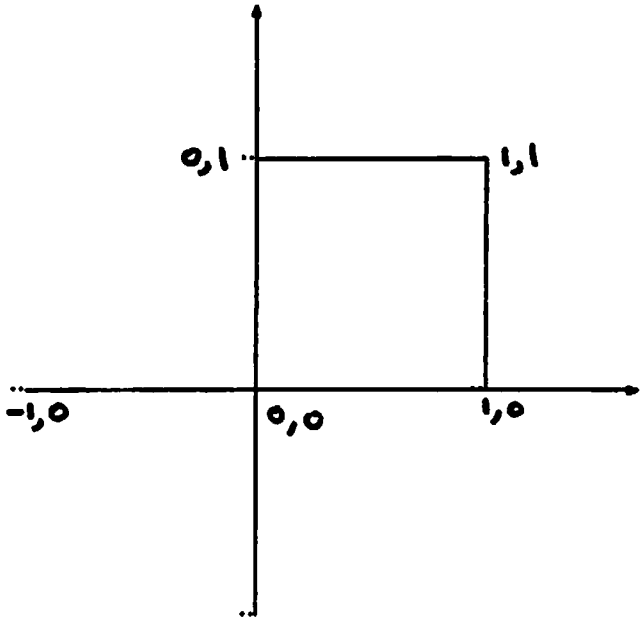

Figure 2: The finite-element domain for the plate model

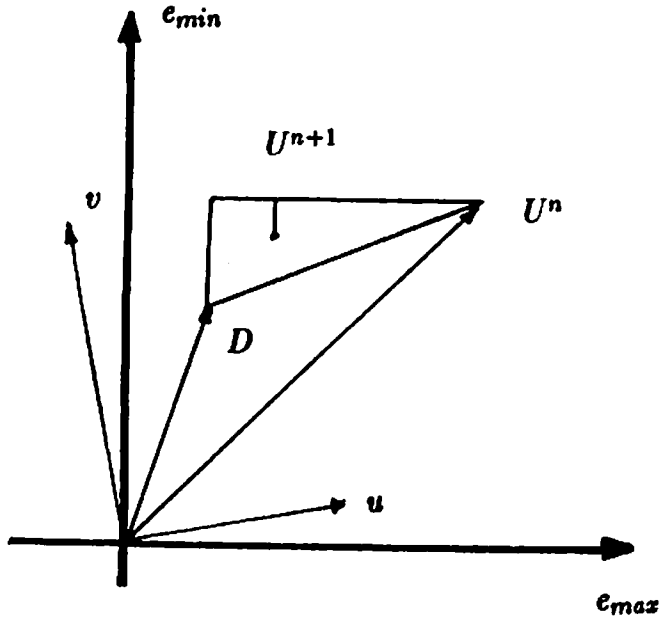

Figure 3: A geometrical illustration of the relaxation process

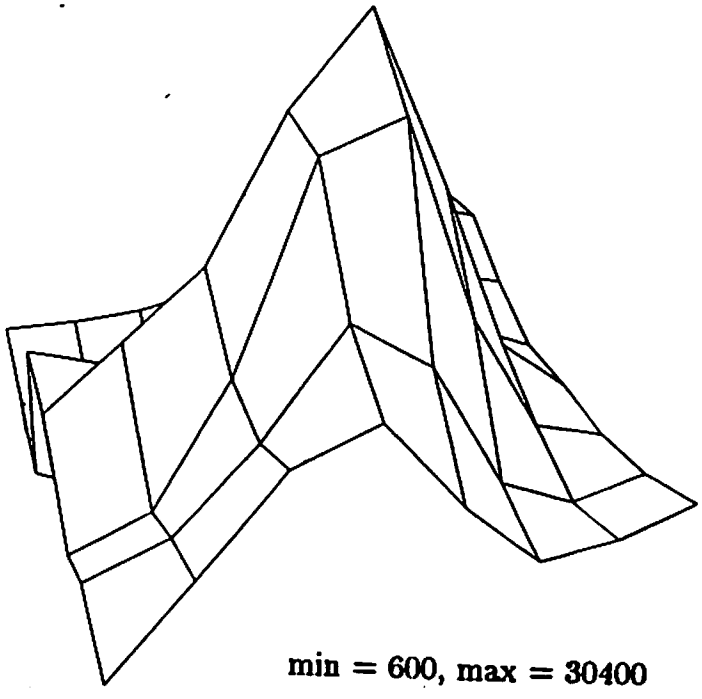

Figure 4: The error surface at a corner point

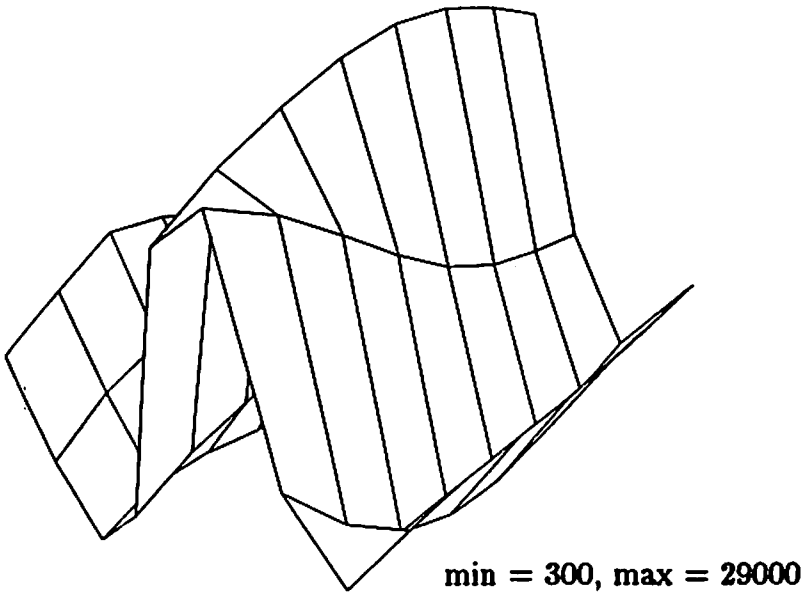

Figure 5: The error surface at an edge point

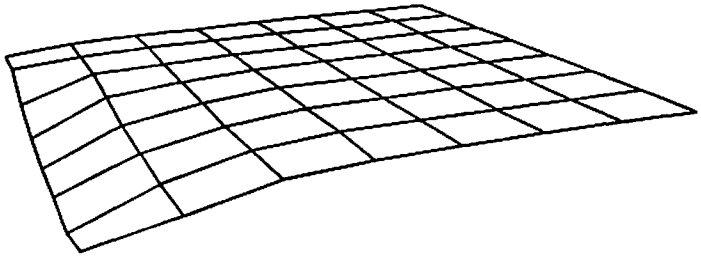

Figure 6: The error surface at a homogeneous point

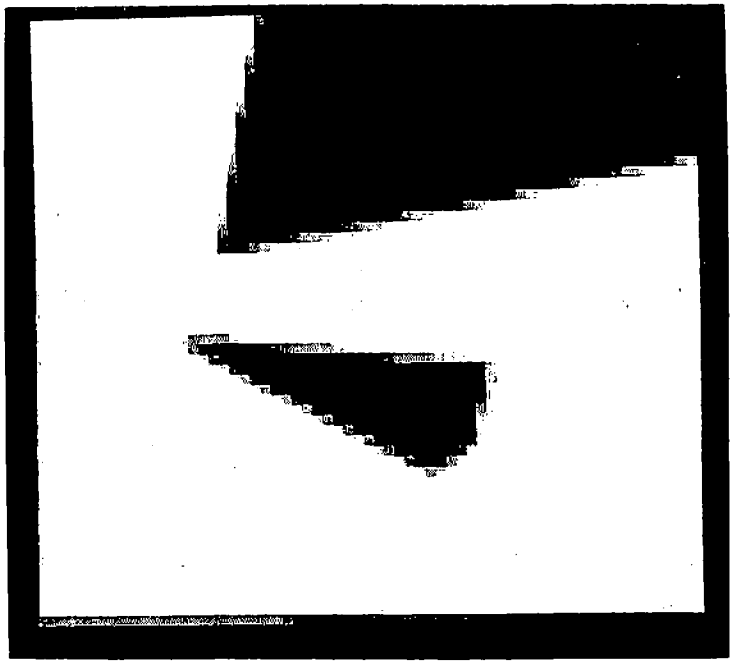

Figure 7: The input image for the single level experiment

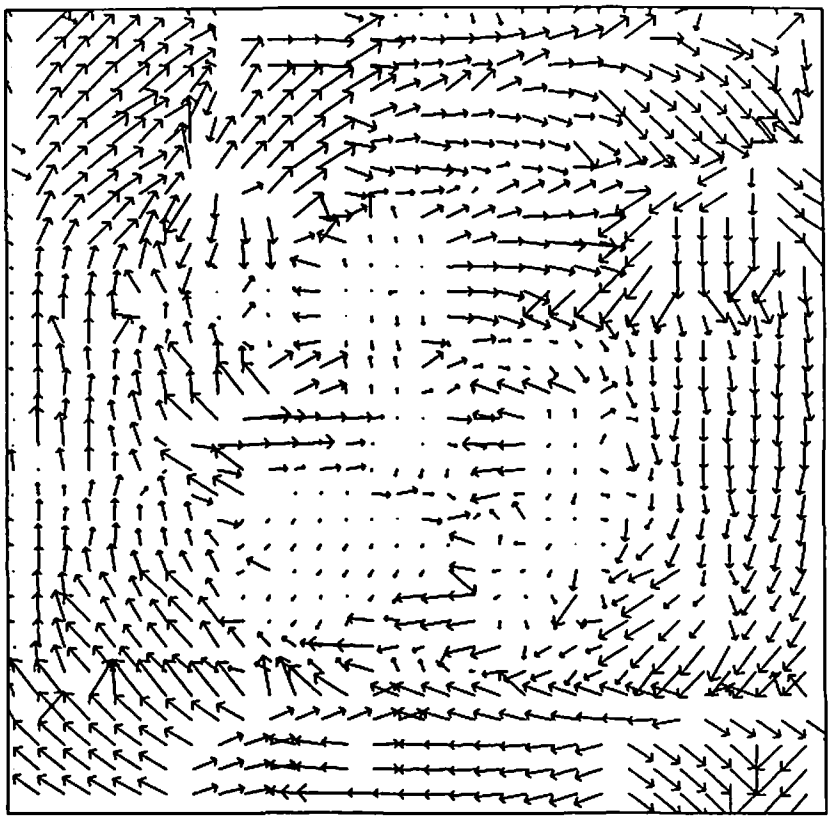

Figure 8: The initial displacement estimates provided by the matching process. Only a quarter of the displacement vectors are shown, in order to enhance visibility

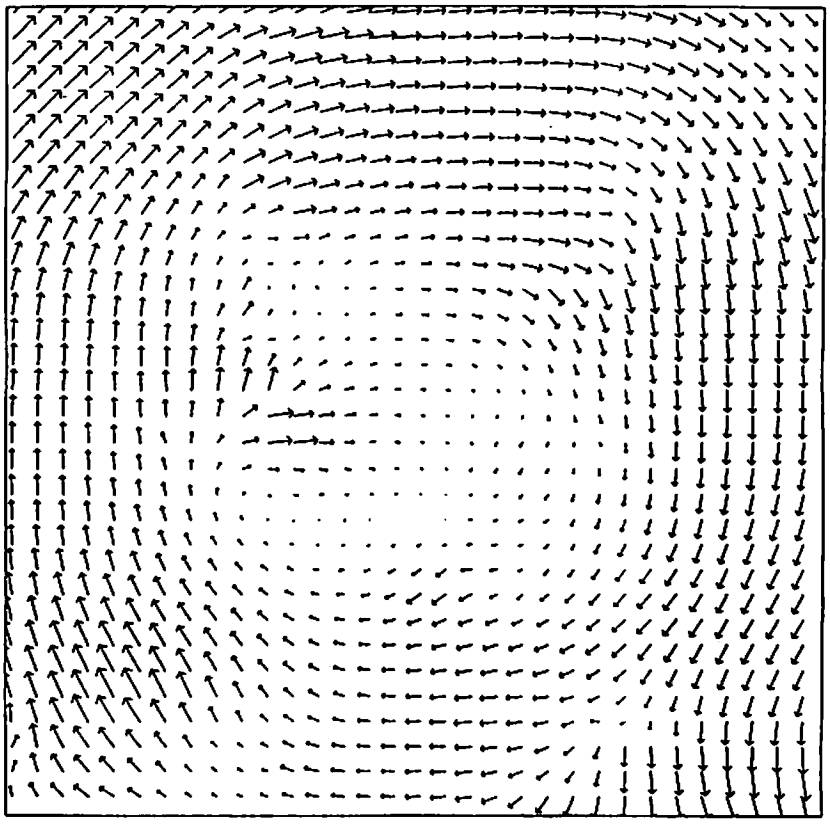

Figure 9: The displacement field after 100 iterations of smoothing using the membrane model

	err. along e_{max}		err. along e_{min}		err. mag	
	mean	std. dev	mean	std. dev	mean	std dev
initial disp.	0.1571	0.1964	0.4066	0.3341	0.4693	0.3464
membrane mod.	0.1191	0.1264	0.1609	0.1418	0.2263	0.1580
plate model	0.1323	0.1465	0.2179	0.1857	0.2845	0.1799

Figure 10: Table of error statistics for the single level experiment The table shows the statistics of the absolute values of the error component along e_{max} and e_{min} as well as those of the magnitude of the error vectors

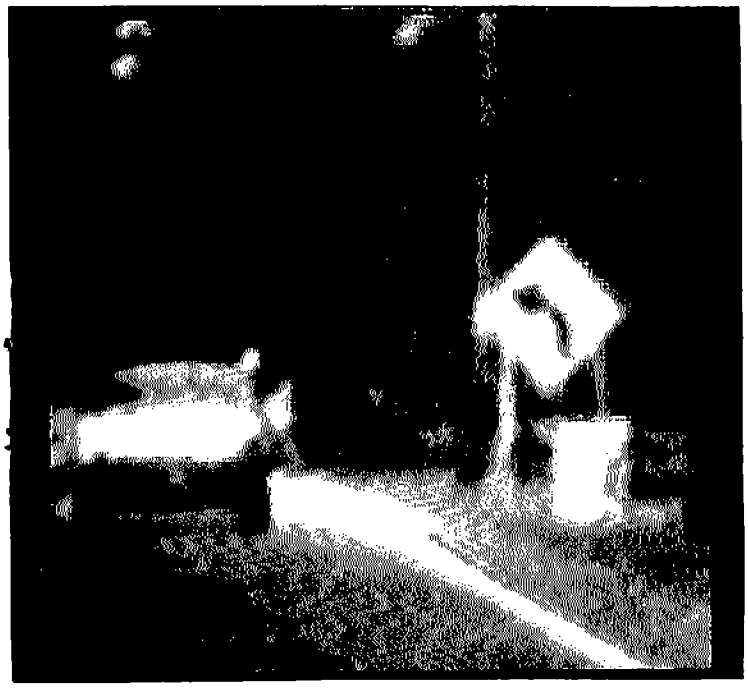

Figure 11: The first image of the road-scene image pair

Figure 12: The second image of the road-scene image pair

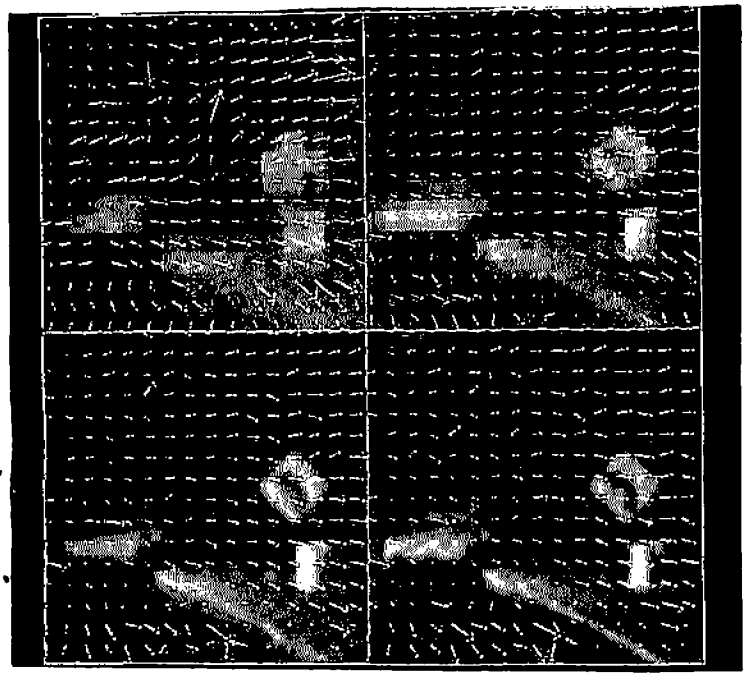

Figure 13: The results of the hierarchical matching algorithm without the smoothing process when applied to the road-scene image-pair.

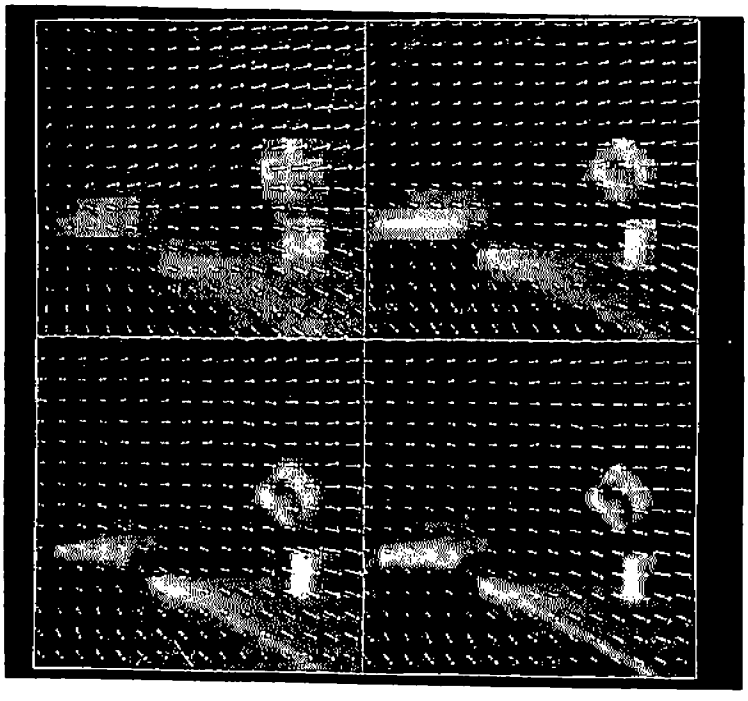

Figure 14: The results of the hierarchical matching/smoothing algorithm applied to the road-scene image-pair

In both the figures above, results at four levels of the hierarchy are shown. The background images are the low-pass filtered versions of the first road image. The image resolution is 128×128 (shown in the bottom-right quadrant). At each level, except the 16×16 resolution in the top-left quadrant, we have shown only a sample of the displacement vectors. This was done so as to enhance visibility.

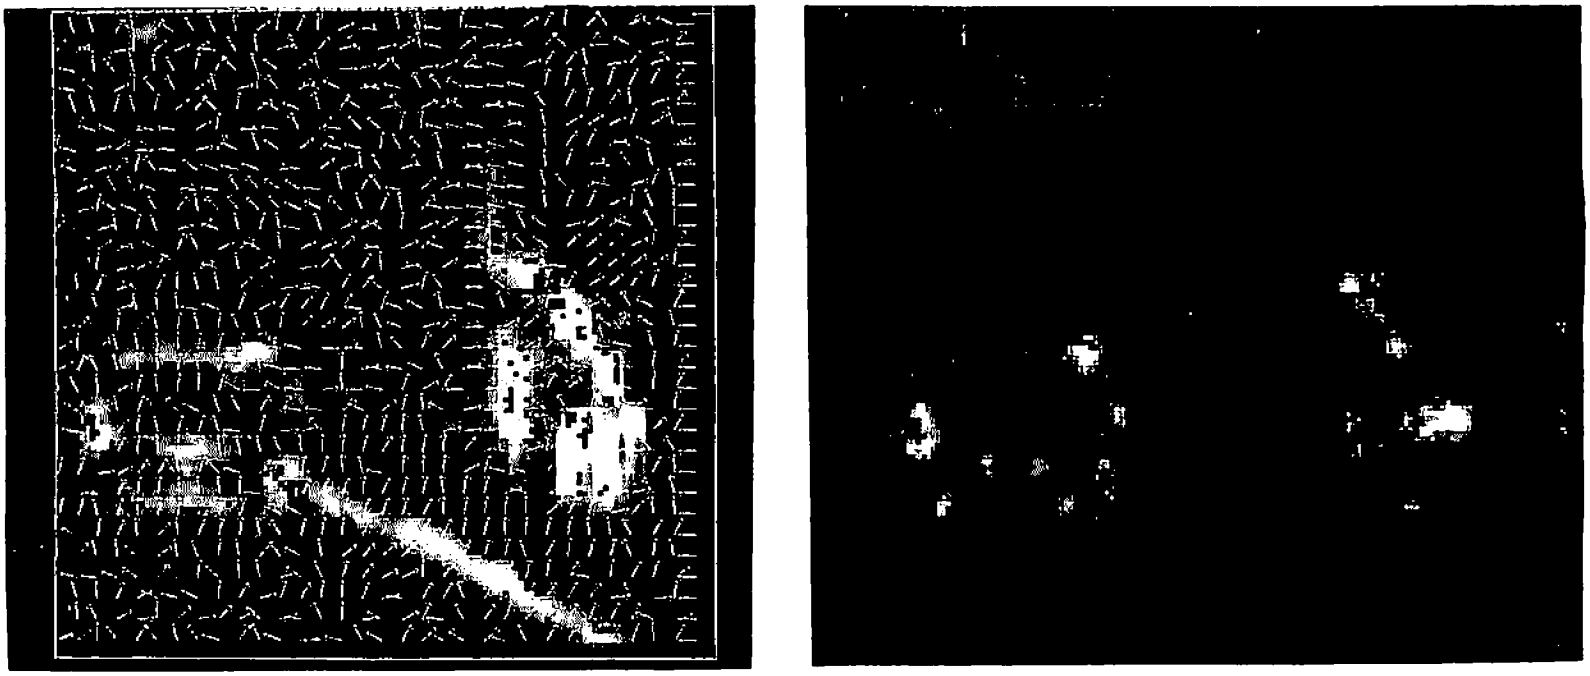

Figure 15: The confidence measures c_{max} (on the left) and c_{min} (on the right) at the finest-level of the hierarchy. These measures are displayed here as intensity images. The left figure also contains the unit-vectors e_{max}

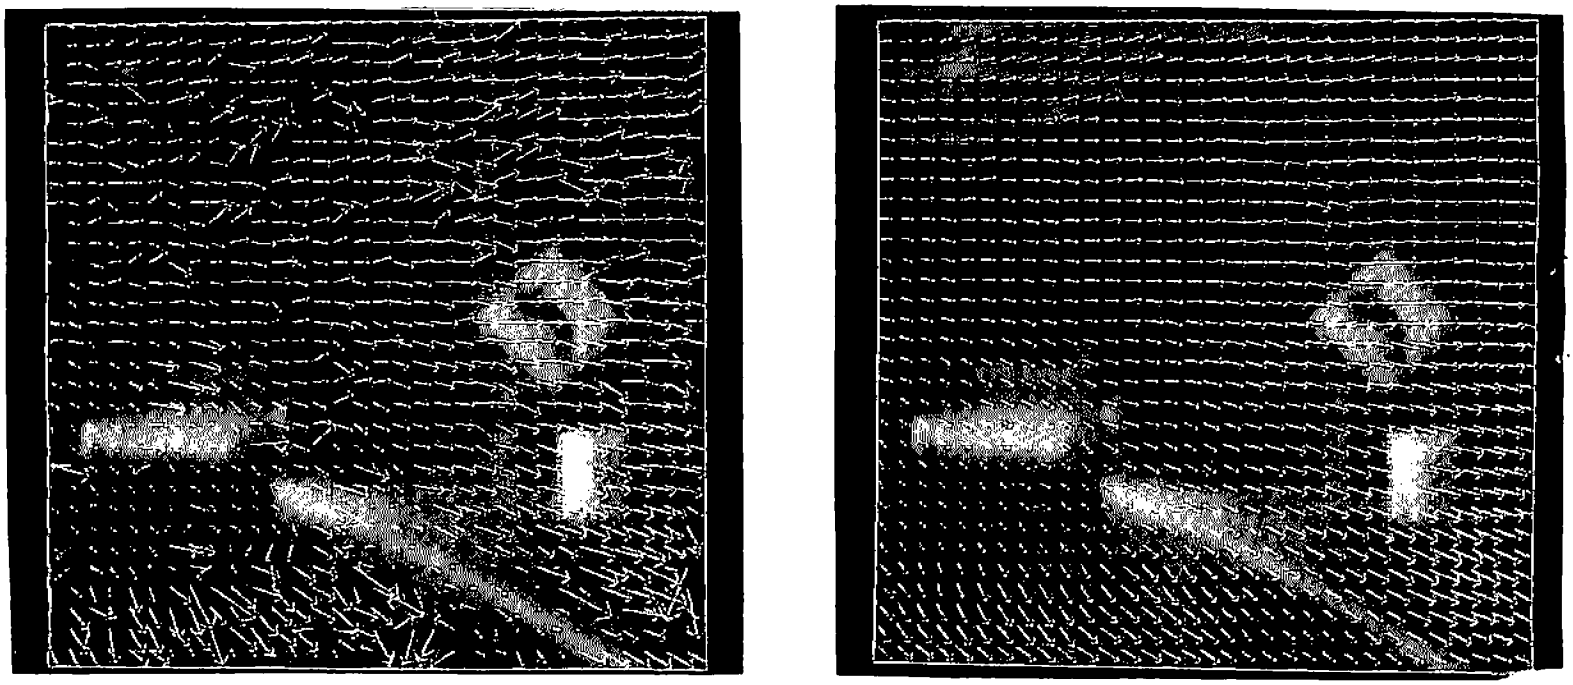

Figure 16: The results at image-resolution for the road-scene image-pair