



Image Deformations Are Better than Optical Flow

C. TOMASI

Computer Science Department, Stanford University
Stanford, CA 94305, U.S.A.

J. SHI

Computer Science Division, University of California at Berkeley
Berkeley, CA 94720, U.S.A.

Abstract—In many computer vision applications it is necessary to compute the direction of heading of a moving camera from the images it produces. Traditionally, this computation has been based on the optical flow, that is, on the motion of point features in the field of view. We show that the differential changes in the angles between the projection rays of pairs of point features are a better input for this purpose. These angles, the *image deformations*, do not depend on viewer rotation, so the key problem of separating the effects of rotation from those of translation is solved at the input. Experiments show both the feasibility of the method on real images and the advantages of using deformations rather than optical flow.

Keywords—Computer vision, Optical flow, Image motion analysis, Egomotion, Motion parallax.

1. INTRODUCTION

As we walk down a hallway, the moving images on our retinas convey enough information to determine our direction of heading. Several researchers have investigated how this direction could be computed, either in the human visual system, or by a computer processing images from a moving camera.

The main difficulty of this computation is to separate the effects of viewer rotation from those of viewer translation. In fact, with only translation the task would be quite simple; features in the image move toward or away from a single point, the focus of expansion, which points toward the direction of heading (see Figure 1). If the viewer also rotates, however, the focus of expansion vanishes. Unfortunately, the rotational component is often dominant; the effects of translation on the images are small, being inversely proportional to the usually large distance to the scene, while the effects of rotation are independent of distance. This has always been known to movie directors, who use expensive dollies to move their cameras with as little vibration as possible.

In this paper, we show that the effects of rotation can be eliminated from the image measurements right at the input of the computation. Specifically, we observe that the angle between the projection rays of two features in the scene does not depend on the viewer's rotation, and we use the changes in these angles as the input to our algorithm. In other words, rather than measuring how the image points *move* in the field of view, the traditional flow-based approach, we measure how the image *deforms* over time.

Of course, our approach uses the same data as the flow-based methods do, since we compute deformations from feature positions in successive images. However, our method differs in how

We thank the anonymous reviewers for their useful comments that improved the quality of this paper. This research was supported under NSF Grants IRI-9201751 and IRI-9496205.




Figure 1. With pure camera translation, the trajectories of all points pass through one image point, the *focus of expansion*. As the white lines suggest, the streaks in this slow-exposure picture all converge to a point near the middle of the picture.

those data are used; computing deformations removes rotation at the outset, rather than at some later stage of the computation. With noisy data, canceling rotation by computing deformations has two advantages.

1. The direction of heading is computed from deformations by minimizing a residual function with a deeper minimum than the one based on optical flow. Even very small rotations flatten the minimum of the flow-based residual function considerably, leading to a minimization that is more sensitive to noise.
2. The magnitude of deformations with respect to image noise is a direct indication of their reliability for the estimation of the direction of heading. This is not true for optical flow; a large flow may be just caused by viewer rotation, and the differences between flow vectors may be too small with respect to noise to be used for the computation of heading. Deformations, on the other hand, can be monitored as the viewer moves, and heading computed only once they are large enough.

With n feature points there are $n(n - 1)/2$ pairs of features, and therefore, $O(n^2)$ angles between features. However, only $2n - 3$ of these angles are independent. We propose a method for selecting a sufficient set of $O(n)$ points in time $O(n \log n)$.

In the next section, we discuss previous work on the subject. Then, in Section 3, we define image deformations and derive an equation that links them to the direction of heading. In Section 4, we combine these equations into a single system of equations for a sufficient set of feature pairs. In the same section, we also show how the distances to the feature points in space can be eliminated from the equations, leading to a minimization problem in two variables, which we solve in Section 5. In Section 6, we explain why local minima are shallower than global ones, an important point for finding the correct solution. We then test our solution on both simulated and real images (Sections 7 and 8), and discuss applications and complexity issues in Section 9.

2. RELATION WITH PREVIOUS WORK

We now compare our method with others regarding how the effects of rotation are eliminated from the images. We do not discuss methods that require space-continuous velocity fields or second derivatives of image motion [1–3], since we use as input the instantaneous image velocities of a set of discrete image points.

The methods presented in [4–6] are based on the following observation: if a point \mathbf{p} in the first image moves to \mathbf{q} in the second, then the vectors \mathbf{q} , $R\mathbf{p}$, \mathbf{T} , where R is rotation and \mathbf{T} is translation, are coplanar

$$\mathbf{q} \cdot (R\mathbf{p} \times \mathbf{T}) = 0.$$

The three methods above enforce this constraint over several points to compute R and \mathbf{T} . The effect of noise on these nonlinear equations depends on R . For instance, if the two vectors $R\mathbf{p}$ and \mathbf{T} happen to be nearly parallel to each other, small perturbations of \mathbf{p} can change the direction of $R\mathbf{p} \times \mathbf{T}$ considerably.

When the translation of the viewer is large, the three methods just cited are preferable to ours, since they make no assumption as to the distance traveled by the viewer. In contrast, we measure differential changes in the angles between features. On the other hand, when the viewer moves little, the methods above can fail altogether when a certain quadratic equation has no real solution [4,6]. Our method degrades more gracefully; as image deformations become smaller and smaller relative to noise, the uncertainty in the direction of heading grows, but no outright failure occurs.

An observation by Helmholtz [7] has been used in [8–11]; the vector difference in velocity between two points that are nearby in the image but at different depths is nearly independent of rotation (and exactly so, when the points are at the same image location). Changes in this difference, called *motion parallax*, supply sufficient constraints to recover the direction of heading. Our method is also based on motion parallax; our image deformations are the magnitude (in degrees of visual angle) of the vector difference used in those papers. However, by ignoring the direction of the parallax and only considering its magnitude we make parallax independent of rotations exactly, regardless of the image positions of the two feature points. The two hard problems of determining pairs of image features along depth boundaries and measuring their image velocities (given the interference of the boundary) are thereby avoided.

Our approach is similar to those in [12–15] in that we minimize some residual over the measurements in the least squares sense. Like Heeger and Jepson, we reduce minimization to that of a function of two variables (the parameters for the direction of heading), but we use rotation-independent image deformations rather than image flow, with the advantages mentioned in the introduction.

3. IMAGE DEFORMATIONS

Consider two points \mathbf{P} and \mathbf{Q} in space, as in Figure 2. As the viewer moves from \mathbf{C} to \mathbf{C}' , the magnitude α of the angle \mathbf{PCQ} formed by the projection rays changes to $\mathbf{PC'Q}$. The *image deformation* is $\dot{\alpha}$, the time derivative of α . The angle α is given by

$$\alpha = \arccos(\mathbf{p}^T \mathbf{q}) \quad (1)$$

where \mathbf{p} and \mathbf{q} are two unit vectors from the viewer center to the points \mathbf{P} and \mathbf{Q} .




Figure 2. As the viewer moves, the angle α between projection rays \mathbf{CP} , \mathbf{CQ} varies. The time derivative of this variation is the *image deformation*.

To find $\dot{\alpha}$, we first determine the derivative of α with respect to viewer position when \mathbf{C} moves along three special directions: \mathbf{p} , \mathbf{q} , and the direction

$$\mathbf{r} = \frac{\mathbf{p} \times \mathbf{q}}{|\mathbf{p} \times \mathbf{q}|},$$

orthogonal to \mathbf{p} and \mathbf{q} . If p, q, r are the amounts of viewer motion along $\mathbf{p}, \mathbf{q}, \mathbf{r}$, we find that

$$\frac{\partial \alpha}{\partial p} = \frac{\sin \alpha}{|\mathbf{Q}|}, \quad \frac{\partial \alpha}{\partial q} = \frac{\sin \alpha}{|\mathbf{P}|}, \quad \frac{\partial \alpha}{\partial r} = 0.$$

If the viewer motion is expressed instead in an orthogonal reference system through its components x, y, z , the chain rule for differentiation yields

$$\begin{bmatrix} \frac{\partial \alpha}{\partial x} \\ \frac{\partial \alpha}{\partial y} \\ \frac{\partial \alpha}{\partial z} \end{bmatrix} = J^T \begin{bmatrix} \frac{\partial \alpha}{\partial p} \\ \frac{\partial \alpha}{\partial q} \\ \frac{\partial \alpha}{\partial r} \end{bmatrix},$$

where the Jacobian J is given by

$$J = [\mathbf{p} \ \mathbf{q} \ \mathbf{r}]^{-1}. \quad (2)$$

Finally, we apply the chain rule once more to compute the derivative of α with respect to time, given the three time derivatives \mathbf{t} of the viewer position \mathbf{C} :

$$\dot{\alpha} = \sin \alpha \mathbf{t}^T J^T \begin{bmatrix} |\mathbf{Q}|^{-1} \\ |\mathbf{P}|^{-1} \\ 0 \end{bmatrix}. \quad (3)$$

We can rewrite equation (3) in a more compact form to obtain the following fundamental measurement equation:

$$b = \mathbf{t}^T A \mathbf{d}, \quad (4)$$

where the scalar

$$b = \frac{\dot{\alpha}}{\sin \alpha} \quad (5)$$

is a quantity that can be measured from two or more images, the vector

$$\mathbf{d} = \begin{bmatrix} d_P \\ d_Q \end{bmatrix} = \begin{bmatrix} |\mathbf{P}|^{-1} \\ |\mathbf{Q}|^{-1} \end{bmatrix}$$

collects the reciprocals of the unknown depth values, the columns of the 3×2 matrix A are the known second and first row of J , and the vector \mathbf{t} is the unknown viewer velocity.

4. COMBINING MULTIPLE MEASUREMENTS

With n features instead of two, we can write $n(n-1)/2$ equations like equation (4), one for every pair of features. However, only $2n - 3$ of these equations provide independent measurements.

In fact, if two points $\mathbf{p}_1, \mathbf{p}_2$ are picked as reference, the positions of all n points is identified up to a mirror flip by the distance between \mathbf{p}_1 and \mathbf{p}_2 and by the $2(n-2)$ pairs of angles that \mathbf{p}_1 and \mathbf{p}_2 form with the other points \mathbf{p}_i for $i = 3, \dots, n$, for a total of $1 + 2(n-2) = 2n - 3$ parameters. This construction, however, is only useful as a way to count independent equations. In fact, an arbitrary choice of \mathbf{p}_1 and \mathbf{p}_2 can lead to long and narrow triangles $(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_i)$ yielding a numerically poor system of equations.

Instead, we measure angles α for pairs of image points that are connected by the edges of a Delaunay triangulation [16] of all the points (see Figure 7). This triangulation is a planar graph, and thus, has $O(n)$ edges. Every vertex is connected to at least two others, leading to a sufficient set of measurement equations. Edges connect nearest neighbors, yielding large ratios in equation (5) for reduced noise sensitivity. Finally, points strictly inside the convex hull belong to triangles that are as close as possible to equilateral, yielding maximally independent measurement equations.

We can then let

$$\mathbf{d} = [d_1 \ \cdots \ d_n]^\top,$$

$$\mathbf{b} = [b_1 \ \cdots \ b_m]^\top,$$

where b_k is the left-hand side of equation (4) for edge k , and m is the number of edges in the Delaunay triangulation. If edge k connects points i and j , the matrix A for the k^{th} measurement equation (4) is formed from the Jacobian $[\mathbf{p}_i \ \mathbf{p}_j \ \mathbf{r}_{ij}]^{-1}$ of equation (2). Because the matrix A for pair (i, j) is equal to that for pair (j, i) , but with its columns switched, we can consistently define

$$A_k = [\mathbf{a}_{ji} \ \mathbf{a}_{ij}],$$

and define the $m \times n$ matrix $A(\mathbf{t})$ whose entry k, l is

$$t_{kl} = \begin{cases} \mathbf{t}^\top \mathbf{a}_{ij}, & \text{for } l = i, \\ \mathbf{t}^\top \mathbf{a}_{ji}, & \text{for } l = j, \\ 0, & \text{otherwise.} \end{cases}$$

With these definitions, the m equations (4) can be collected into the following bilinear system:

$$\mathbf{b} = A(\mathbf{t})\mathbf{d}. \quad (6)$$

If \mathbf{t}, \mathbf{d} is a solution to equation (6), so is $c\mathbf{t}, \mathbf{d}/c$ for any nonzero c , consistently with the fact that absolute scale cannot be recovered from images alone [17]. With the additional constraint that the translation \mathbf{t} have unit norm, \mathbf{t} is the viewer's direction of heading.

5. SOLVING FOR THE DIRECTION OF HEADING

In the presence of noise, equation (6) will only be satisfied approximately, so we need a measure for how close \mathbf{b} is to the column space of $A(\mathbf{t})$. Conceptually, a measure for this distance is obtained by replacing \mathbf{d} in equation (6) by its solution \mathbf{d}^+ in terms of the pseudoinverse,

$$\mathbf{d}^+ = (A(\mathbf{t})^\top A(\mathbf{t}))^{-1} A(\mathbf{t})^\top \mathbf{b},$$

and then measuring the residual

$$\rho(\mathbf{t}) = |A(\mathbf{t})\mathbf{d}^+ - \mathbf{b}|. \quad (7)$$

This residual does not depend on \mathbf{d} , and can be minimized with respect to \mathbf{t} . Of course, from a numerical standpoint, more efficient algorithms for computing (7) should be used than computing the pseudoinverse of $A(\mathbf{t})$ for given values of \mathbf{t} : the reason for introducing this pseudoinverse is only to show that conceptually a residue can be built that does not depend on \mathbf{d} . In practice, we use a variable projection method [18,19] on the unit sphere $|\mathbf{t}| = 1$ to minimize the residual (7) with respect to \mathbf{d} and \mathbf{t} . Local minima do exist, but the more points are available in the field of view, the smoother the residual (7) turns out to be. In our experiments, we usually find one or two local minima, with the correct minimum considerably deeper than the other (see Sections 7 and 8 for examples, and Section 6 for a discussion of this point). One can then use local minimization methods starting at a few random points on the unit sphere and choose the convergence point of smallest residual as the solution. At the end of the computation, the depth vector \mathbf{d} is simply discarded. In fact, it has been shown [20] that if the translation is small and rotation is nonzero the depth results are unreliable no matter what algorithm is used. However, one could conceivably estimate better depth results from a long sequence of depth measurements from consecutive pairs of images, using, for instance, a Kalman filter.

6. GLOBAL AND LOCAL MINIMA

In this section, we give an intuitive explanation for why local minima of the residual function (7) are deeper than the global minimum, a fact that we have invariably noticed in our experiments and simulations. First, the camera motion is known [5] to be uniquely determined by optical flow in the absence of noise. If noise is added, the global minimum is raised away from zero. Local minima, on the other hand, are away from zero already in the ideal case. This by itself does not imply that they become even higher when noise is added. It just says that without noise the global minimum is deepest.




Figure 3. The correct (crosses and solid arrow) depth and motion interpretation of image deformations yields the global minimum of the residual (7). The crosses on the left, together with the circles and the dashed arrow, yield a local minimum.

With noise, we can gain some understanding of what happens by considering the example of Figure 3, in which the camera sees two clouds of points in front of itself. Crosses denote the true position of the points in the world, and the solid arrow points along the true direction of motion of the camera. As the camera moves ahead, two phenomena can be observed in the image. The dominant one is that the two clouds of points become more distant from each other as they move toward the image periphery. The secondary effect is that the points within each cloud are spread further apart from each other. The relative extent of these two effects depends on the actual geometry of the scene. The minimization algorithm then tries to “explain” the image deformations by determining the positions of the points in the world and the motion of the camera. A good explanation corresponds to a deep minimum.

The correct explanation is given in Figure 3 by the crosses and the solid arrow. However, a different explanation is possible, represented by the left cloud of crosses, the cloud of circles, and the dashed arrow for camera motion. The algorithm has no notion of “in front” or “behind” the camera, although this would be easy to check after depths have been computed. Consequently, the two explanations are equally good in terms of the data, since the points behind the camera

(the circles) project onto the same image locations as the corresponding points in front of the camera (the crosses on the right).

The wrong explanation accounts for the increasing image separation of the two clouds as the camera moves along the dashed arrow, and thereby, creates a local minimum in the residual function. On the other hand, the points within each of the two clouds would essentially only translate together if the camera were actually moving along the dashed arrow. Consequently, the secondary effect of the diverging points within a cloud, which is actually observed in the data, is not explained by this incorrect interpretation, and the corresponding minimum is shallower.

This argument also suggests that fewer local minima could be obtained if the minimization of the residual (7) were constrained to positive depth values, an approach that we did not pursue.

7. A SIMULATION EXPERIMENT

Figure 4 shows a contour plot of the residual $\rho(\mathbf{t})$ on the hemisphere of heading directions \mathbf{t} corresponding to a forward moving viewer. Because $A(\mathbf{t})\mathbf{d} = A(-\mathbf{t})(-\mathbf{d})$ (see equation (6)), the residual function is the same on the opposite hemisphere ($\rho(\mathbf{t}) = \rho(-\mathbf{t})$). For this simulation we use 30 feature points, spanning a visual angle of about 120 degrees and distributed in depth between one and ten units away from the camera. Absolute depths are irrelevant because scale does not influence the results. The camera translation is one hundredth of the average distance to the scene, and the camera rotation is 3 degrees around a vertical axis. Random uniform noise is added to the second image. The width of the noise distribution is half a pixel for a 500×500 pixel image.




Figure 4. Contour plot of the residual $\rho(\mathbf{t})$ on the forward hemisphere. The two loops slightly above and to the left of the center are minima; the bigger loop is around the global minimum. Other loops are maxima.

The true direction, randomly generated, was the unit vector $(-0.306, -0.066, 0.950)^\top$. Our method computed $\mathbf{t} = (-0.302, -0.126, 0.945)^\top$, corresponding to a heading error of 3.45 degrees.

Figure 5 shows the direction error versus feature position uncertainty for the same situation. Each point on the graph is the average heading error for ten runs with the same noise distribution but different noise samples. For increasing noise the results become more and more erratic. However, the algorithm never fails completely, but degrades gracefully.

Figure 6 compares our deformation-based residual,

$$|\mathbf{b} - \mathbf{t}^\top \mathbf{A} \mathbf{d}|$$

(see equation (4)) with the more traditional residual based on optical flow,

$$|\mathbf{u} - (\delta \mathbf{A}_1 \mathbf{t} + \mathbf{A}_2 \boldsymbol{\omega})|,$$




Figure 5. Direction of heading error versus feature position uncertainty.

where \mathbf{u} is the optical flow, ω is a vector representing the camera rotation, A_1 and A_2 are matrices that depend only on image position, δ is inverse depth along the optical axis, and \mathbf{t} is the direction of heading (see for instance [15]). In both cases, the overall residual is the root mean square of the residuals for the individual features. All plots were obtained for a viewer translating exactly forward and no image noise. The two top diagrams show the residuals for pure translation; the widths and shapes of the two minima are essentially the same. When the camera rotates even by just one degree, the minimum of the flow-based residual (bottom right plot) becomes elongated and shallow, leading to a more noise-sensitive minimization. The deformation residual, on the other hand, remains unaltered (bottom left).

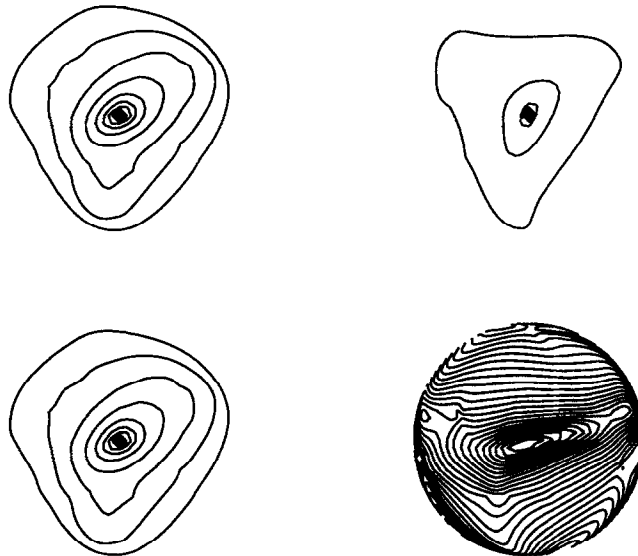


Figure 6. Contour plots of the deformation-based residual $\rho(\mathbf{t})$ (left plots), and the traditional flow-based residual (right), for pure translation (top), and with an added rotation of one degree (bottom).

To summarize, deformations are preferable to direct optical flow when the viewer rotates by even a small amount. On the other hand, regardless of whether flow or deformations are used, Figure 5 shows that recovering the direction of heading from a pair of images requires cameras and tracking systems of good quality. In the next section, we show that our method can be used with real images.

8. AN EXPERIMENT ON REAL IMAGES

Figure 7 shows the first of a sequence of images taken in our lab, with the Delaunay triangulation of the tracked features superimposed. The viewing angle is about thirty degrees and the objects in the scene are between about 50 and 100 cm away.




Figure 7. The first frame used in the experiment, with the Delaunay triangulation of the selected features.

About 230 features were automatically selected in the first frame and tracked using the algorithm described in [21]. Of those, 44 were handpicked to provide a roughly uniform distribution over the image. The camera was moved by a Puma arm proceeding in small steps, first along a constant reference direction (roughly towards the pencils), then along a direction at an angle of 30 degrees from the reference direction. For every new frame, features were tracked and the root mean square image deformation from start to current frame was determined. As soon as the rms deformation exceeded one pixel, the direction of heading was computed from the deformations between the start frame and the current frame. The current frame then became the new start frame for the next measurement. This procedure guarantees that the deformations are substantially greater than the feature position uncertainty (about 0.1 pixels), leading to a reliable computation of the direction of heading.

This input monitoring method would not be possible with optical flow. In fact, the magnitude of flow says nothing about its reliability for the computation of the heading direction, since there is no direct relation between the amount of flow and the amount of camera translation. Deformations, on the other hand, are nonzero only in the presence of translation, and their size with respect to the uncertainty of feature positions is a direct indication of how reliable deformations are for the computation of translation.

Three sufficiently distant images were obtained with this procedure, with the second image at the turning point of the camera path. The residual functions for the two image pairs are shown in Figure 8. The error on the angle between the two directions of heading, as computed by our method, was about 8 degrees. Since we do not know the accuracy of the Puma arm for very small motions and our camera was not calibrated, this error is only a rough indication of the accuracy of our method. We are planning more accurate experiments. Each residual function has a clean and deep global minimum, even with the small motion (1 cm) and narrow field of view (30 degrees) of our experiment. There is also one local minimum in each residual, but this is much more shallow and created no problem for our minimization procedure.




Figure 8. Contour plots of $\rho(\mathbf{t})$ on the positive hemisphere for the two real image pairs. Global minima are in the first quadrant, local minima in the third.

9. APPLICATIONS AND COMPUTATIONAL ISSUES

The deformation-based method proposed in this paper makes it possible to compute the direction of heading of a camera from its images in a more reliable way than previous approaches based on optical flow. If our method can be implemented in real time, it can be applied to many tasks. For instance, in hand-eye coordination, camera motion can now be measured in the camera's frame of reference, as opposed to the robot's. This makes hand-eye calibration straight-forward. In robot navigation, a new motion measurement becomes available that is immune to problems like wheel-slippage or bumpy terrain that severely degrade mechanical motion measurements. In satellite docking or ballistic target homing applications, the vehicle's motion can be continuously compared to the image of the target for accurate guidance even in the presence of rotations.

The question of computational complexity is therefore of primary importance. The core of our method is the solution of the bilinear system (6). The variable-projection methods we use for this purpose (see Section 5) are iterative, and it currently takes a few minutes on a Sparc2 workstation to solve for the direction of heading. The matrix A in equation (6) is formed from a $m \times n \times 3$ tensor of numbers, where n is the number of feature points and m is the ($O(n)$) number of feature pairs considered. While implementing the entire computation in one thirtieth of a second would require substantial hardware, we notice the following points.

- Measurements are used only when the deformation-monitoring system described in Section 8 deems the data to be reliable with respect to image noise. It may take several frames before a new computation is necessary.
- Complete convergence is not necessary at every computation. One can stop the iterations when new data become available, and start a new minimization using the current estimate as a starting point. This amounts to doing gradient descent on a slowly varying landscape.
- A much faster, although coarser, estimation of the heading direction \mathbf{t} can be obtained by sampling the unit sphere $|\mathbf{t}| = 1$, and computing the residual (7) everywhere in parallel. The minimum can then be found by a comparison tree in time that is logarithmic in the number of sampling points.

10. CONCLUSION

Computing the direction of heading from image deformations is an interesting alternative to using the image flow field directly, because it removes the effects of rotation right at the outset in a clearly understandable and straight-forward way. The minimum of the residual function computed from deformations is deeper than that of the flow-based residual, leading to a more reliable solution. Furthermore, the magnitudes of the deformations, when compared to the noise level in the images, are a direct indication of their reliability for the computation of the direction of heading. Finally, our minimization method degrades gracefully with feature position uncertainty.

Simulations confirm that a wide angle of view and accurate image measurements are necessary for good heading estimates, and our experiments with real images indicates that these requirements are realistic. However, it is obvious from the quantitative data presented in this paper

that determining the camera's direction of heading from closely-spaced images is a delicate task, and only explicit attention to the sensitivity of the problem, as well as to camera calibration and to the numerical aspects of the computation, can lead to accurate motion estimates.

REFERENCES

1. J.J. Koenderink and A.J. Van Doorn, Invariant properties of the motion parallax field due to the movement of rigid bodies relative to an observer, *Optica Acta* **22** (9), 773–791 (1975).
2. A.M. Waxman, B. Kamgar-Parsi, and M. Subbarao, Closed-form solution to image flow equations for 3D structure and motion, *International Journal of Computer Vision* **1**, 239–258 (1987).
3. A. Verri, F. Girosi, and V. Torre, Mathematical properties of the two-dimensional motion field: From singular points to motion parameters, *Journal of the Optical Society of America A* **6** (5), 698–712 (1989).
4. H.C. Longuet-Higgins, A computer algorithm for reconstructing a scene from two projections, *Nature* **293**, 133–135 (1981).
5. R.Y. Tsai and T.S. Huang, Uniqueness and estimation of three-dimensional motion parameters of rigid objects with curved surfaces, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *PAMI* **6** (1), 13–27 (1984).
6. J. Weng, T.S. Huang and N. Ahuja, Motion and structure from two perspective views: Algorithms, error analysis, and error estimation, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11** (5), 451–476 (1989).
7. H. Helmholtz, *Treatise on Physiological Optics*, (German original 1856–67), Dover, New York, (1925).
8. H.C. Longuet-Higgins and K. Pradny, The interpretation of a moving retinal image, *Proceedings of the Royal Society of London*, Series B **208**, 385–397 (1980).
9. J.H. Rieger and D.T. Lawton, Processing differential image motion, *Journal of the Optical Society of America A* (2), 254–360 (1985).
10. J.E. Cutting, *Perception with an Eye towards Motion*, MIT Press, Cambridge, MA, (1986).
11. E.C. Hildreth, Recovering heading for visually-guided navigation, *Vision Research* **32** (6), 1177–1192 (1992).
12. A.R. Bruss and B.K.P. Horn, Passive navigation, *Computer Vision, Graphics, and Image Processing* **21**, 3–20 (1983).
13. G. Adiv, Determining three-dimensional motion and structure from optical flow generated by several moving objects, *IEEE Pattern Analysis and Machine Intelligence* **7**, 384–401 (1985).
14. S.J. Maybank, A Theoretical Study of Optical Flow, Ph.D. Thesis, University of London, (1987).
15. D.J. Heeger and A.D. Jepson, Subspace methods for recovering rigid motion I: Algorithm and implementation, *International Journal of Computer Vision* **7** (2), 95–118 (1992).
16. F.P. Preparata and M.I. Shamos, *Computational Geometry*, Springer-Verlag, New York, (1985).
17. E.H. Thompson, A rational algebraic formulation of the problem of relative orientation, *Photogrammetric Record* **3** (14), 152–159 (1959).
18. G.H. Golub and V. Pereyra, The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate, *SIAM Journal on Numerical Analysis* **10** (2), 413–432 (1973).
19. A. Ruhe and P.Å. Wedin, Algorithms for separable nonlinear least squares problems, *SIAM Review* **22** (3), 318–337 (1980).
20. J.J. Koenderink and A.J. van Doorn, Facts on optic flow, *Biological Cybernetics* **56**, 247–255 (1987).
21. C. Tomasi and T. Kanade, Shape and motion from image streams: A factorization method—3. Detection and tracking of point features, Technical Report CMU-CS-91-132, Carnegie Mellon University, Pittsburgh, (1991).