DETERMINING MOTION PARAMETERS USING A PERTURBATION APPROACH*

G. B. Shaw**

COINS Technical Report 83-30

September 1983

ABSTRACT

Many workers have obtained rigid body motion parameters from observations of a few points in two successive images. These calculations have been consistently fragile in the presence of observational error. The present work provides a somewhat different basis for calculation of such motion descriptions. This approach, based on perturbation theory, identifies the source of calculational fragility and provides methods for localizing unavoidable ambiguity. That is, only those parameters which are in principle ambiguous are to be substantially affected by data errors.

This paper presents the basis for the method and results for some special cases, including determination of general rigid motion parameters from accurate position observations. A somewhat unusual choice of rigid motion parameters is employed, and its advantages for such calculations are described.

*This research was supported in part by the Defense Advanced Research Projects Agency under contract N00014-82-K-0464.

**G. B. Shaw is an Assistant Professor in the Computer and Information Science Department at the University of Oregon.

Determining Motion Parameters Using a Perturbation Approach

1.0 Introduction

Analysis of isolated images of static scenes has begun to give way to consideration of multiple images. For instance, a set of stereo images of a single scene taken from different viewpoints provides more information than any single view could. Again, a sequence of images recorded by a single camera may be analyzed to obtain motion information. The special cases of a static camera in a dynamic scene and 'egomotion' in a fixed environment have both received particular attention. (A recent book, [Huang 1982] is devoted specifically to treatment of motion).

The present paper is concerned primarily with analysis of motion. Of the two principal views of the image data employed: as the source of an 'optic flow field', or of a more conventional sequence of static 'snapshots', the latter is employed. In this approach, distinguished points or regions are tracked from frame to frame, each providing a sequence of positions and (for segments or regions) orientations due to motion.

Such treatments of motion generally suppose a scene to consist of some number of rigid 'objects', within each of which all physical distances are preserved through the complete sequence of views. (Any group of static or co-moving entities are thus 'objects' in this sense.) From this rigid-body assumption, are obtained 'structure from motion' relations [Ullman 1979] eg that four independent orthographic views of four noncoplanar points allow unique reconstruction of all obtainable motion and position parameters. Various workers, [Fennema+ 1979, Nagel 1981, Prazdny 1981, Roach+ 1980, Tsai+ 1980, Yen+ 1983] have obtained similar results, each choosing a different parametrization of motion, and with each version having its own theoretical or computational advantage. Some recent work, [Longuet-Higgins 1981, Tsai+ 1982] has been particularly elegant. From one point of view, then, the problem of motion from multiple observations is already solved. Redundantly solved. From another point of view, the situation is not so satisfactory.

What happens when one examines actual images, extracts clearly identifiable points, and applies the various formulas to obtain the underlying motion and position information? Generally, the results are very ill-behaved indeed, with small errors in observation yielding unpredictably large errors in the parameters, even when the underlying model is not violated. (Known difficulties, such as following points on highlights or occlusion contours, are a separate problem.) This situation has caused a number of workers to examine approaches applicable to special cases, such as pure translation, pure rotation, and motion confined to a plane [Lawton 1982]. For each such special case, a computational scheme could be obtained which was much more robust than any of the general methods.

The investigation reported here is an attempt to examine general motion in a way that permits image data to be utilized in an hierarchical way. This identifies some motion and depth parameters as being particularly evident from image motion, while other parameters are shown to depend on finer observation. The mathematical approach taken is to apply perturbation analysis, long used in physics and astronomy, to extract accurate approximate relations between motion parameters and observable quantities. expressions for all these parameters are obtained in a way that permits simple polynomial smoothing of observed data to be matched cleanly to the analytic results. Simple data fitting measures can thus be used as measures of the reliability of the motion parameters obtained.

2.0 Perturbation Theory, General Remarks

Perturbation theory was developed as a method for accurate calculation of the orbital parameters of celestial objects, a problem wholly unrelated to the present one, though both are attempts to relate motion parameters to observational data. In more recent years, it has been developed as a general-purpose tool for approximate solution of a variety of physical and mathematical problems (see, for example, [Morse+ 1953]). It can be applied to a problem which can be considered to be a 'small' modification of an already solved problem. ('Small' here means that the known solutions are reasonable first approximations to the solutions desired). Formally, we write the problem as an operator equation, usually involving a matrix or differential operator. There will be an equation of interest:

$$Op \; f = 0$$

and a similar equation:

$$Op^{[0]} \; f^{[0]} = 0$$

which can already be solved.

The 'perturbation trick' is to consider instead an entire class of problems parametrized by a continuous quantity x. The solvable problem and its solution correspond to x=0, while the real problem is taken to be the operator and solution at some other value, say x=1. We then try to solve the generalized problem for <u>all</u> values of x. Moreover, we assume that our general solution will be analytic in x in some interval about x=0. We can therefore expand the operator and general solution as Taylor's series in x. Combining all terms with a common power of x:

$$\sum_i \sum_k x^i \, Op^{[k]} \, f^{[i-k]} = 0$$

where

$$Op(x) = \sum_i x^i \, Op^{[i]}$$

and

$$f(x) = \sum_i x^i \, f^{[i]}$$

Since this equation is to be analytic, every coefficient of x must separately equal zero:

$$Op^{[0]} f^{[0]} = 0$$

$$Op^{[0]} f^{[1]} + Op^{[1]} f^{[0]} = 0$$

etc.

The 'zero-order' equation is just the problem already known to be solvable. Each successively higher-order equation involves the elements which appeared in the preceding equations as well as one new operator and one new 'correction' to the solution f. These equations can be significantly simpler than the original. In a typical application of the method, they are successively solved until a sufficiently accurate approximation of the desired solution f(1) is obtained. Often, only the first- or second-order equations must be solved to get useful values.

2.1 Formulations Of The Motion Equation

In order to apply this technique to motion estimation, one must first find an appropriate equation to solve and a suitable solvable related equation. A convenient form of the motion equation which we must solve is given in the next section. The related equation can be the motion generated by approximate motion parameters. In some cases, we may even use the simplest estimate: no motion.

As noted earlier, a variety of different formulations of the motion equations have appeared. All treat essentially the same situation. A set of vectors { Xi } in 3-space are given. Any rigid motion of this set of vectors to positions { Xi' } can be described uniquely as a translation T in the original reference frame followed by a rotation R about the observation point. (See [Nagel 1981] and many others)

$$Xi' = (Xi + T) * R \tag{2.1}$$

With the vectors X_i, X_i' known, this is a first-order equation for the motion parameters R_{ij} and T_i. Unfortunately, for this linear equation to represent a rigid motion, the nine elements of R must satisfy six more quadratic constraints. Further, since in our case we can observe only projections of the vectors X_i, X_i', the product contains unknown depths as well. Thus, the equations we actually solve are all quadratic, with the constraint equations quadratic in the individual unknowns. These equations are by no means intractable, but results obtained with real image data seem not totally satisfactory. Perturbation methods may be helpful in dealing with these equations directly, but in this paper a slightly different form of the motion equations will be used.

A general rotation can be specified by three parameters. Most representations use more, thus requiring additional 'constraint' equations to be satisfied, as well as the motion equations themselves. Common representations include a 3x3 orthogonal matrix, as above, a unit quaternion (4 parameters), or a magnitude w and a unit vector along the rotation axis (again, 4 parameters). A rather interesting choice is to combine these last as:

$$(Q_x, Q_y, Q_z) = \tan(w/2)\,(S_x, S_y, S_z)$$

where S and w are the unit vector and the rotation magnitude. As these are a minimal set of rotation parameters, no additional constraint equations are required. The elements of R are quadratic functions of the Qi. (Again, [Nagel 1981] described this parameter choice, though very tersely.)

If we could factor the rotation matrix into a product of linear terms, we would have a somewhat more tractable equation. Such a factorization is actually possible. The clearest demonstration uses a quaternion form of equation 2.1. While use of quaternions is not essential to the results we derive, it is still of some interest. (See [Pervin+ 1983] for some related forms). Instead of the vector quantities, Xi, Xi', Q, and T, we introduce the (four-component) quaternions:

$$\underline{Xi} = (0, xi, yi, zi) \qquad \underline{Xi}' = (0, xi', yi', zi')$$

$$\underline{Q} = (1, qx, qy, qz) \qquad \underline{T} = (0, tx, ty, tz)$$

We can then represent the motion compactly by:

$$(2.2) \qquad \underline{Xi}' = \underline{Q}+ * (\underline{Xi}) * \underline{Q} + \underline{T}$$

Here, $\underline{Q}+$ is the inverse of $\underline{Q}$, and '*' represents quaternion multiplication. This equation can be rewritten as:

$$\underline{Q} * \underline{Xi}' = (\underline{Xi} + \underline{T}) * \underline{Q} + \underline{Q} * \underline{T}$$

Rewriting this equation in terms of more conventional vector multiplication:

$$(2.3) \qquad Xi' - Xi + T = (Xi' + Xi + T) \times Q$$

This equation relates the absolute positions Xi, Xi' and the unknown (vector) motion parameters Q and T. (Of course, it can be derived without the brief excursion into quaternions, but not so briefly or so clearly).

Since the three-space positions (Xi, Xi') are not known, but only their projections, a slightly different form of the equation is easier to use. Two different (central) projections are commonly used: onto a plane or onto a unit (Gaussian) sphere. If a full solution were to be obtained, the two projections would be entirely equivalent, but as we will be obtaining approximations, they are not. While neither form seems to be greatly superior, projection on the Gaussian sphere seems to provide slightly simpler expressions. We therefore introduce unit vectors Ui and Ui' (the directions of the observed vectors Xi and Xi'), and the distances ri and ri'.

$$Xi = ri * Ui \quad \text{and} \quad Xi' = ri' * Ui'$$

Substituting into the earlier motion equation, we now get:

Eq0: ri'*Ui' - ri*Ui + T = (ri'*Ui' + ri*Ui + T)xQ

This equation seems particularly useful for our purposes. No additional constraint equations are required. It is linear in each unknown variable (allowing successive solution) and quadratic overall.

2.2 Use Of The Exact Equation For Known Motion

Though it is a digression from the main point of this paper, as the form above has not often been used a brief comment on its applicability seems in order. The primary use would seem to be when the motion is already known, ie for 'shape from motion' or 'shape from stereo'. For this case, we rewrite Eq0 as:

ri'*(Ui' - Ui'xQ) = ri*(Ui + UixQ) - T + TxQ

or,

ri'*Ai' = ri*Ai + B

with the obvious values for the vectors Ai', Ai, B.

Given exact values of Ui, Ui', these three (normally) independent linear relations between each pair of unknowns ri, ri', permit convenient solution. In some cases, the direction of view will be such that two of these equations are dependent, but only rarely will all three be dependent. In only those cases will the values for ri and ri' not be determined. In fact, with approximate Ui, Ui', the equations are likely to be inconsistent, requiring a best-fit solution of some kind. This point will be discussed briefly in Section 4.2.

3.0 Perturbation Expansion Of The Motion Equation

To apply this method to motion, we must select an appropriate perturbation parameter in which to expand the quantities above. We will first write down the most general form of the equations (so general as to be useless for calculation), then explore some special cases.

The approach here is straightforward. We assume that each of the variables appearing in Eq0 depends on the perturbation parameter. Expanding each as a power series, we see that Eq0 effectively establishes relationships which must hold among these series. We can write out these dependencies explicitly as perturbation equations of various orders. We will write, for each quantity A,

```
A = A0 + x*A1 + x^2*A2 + ...
```

Eq0 thus becomes (suppressing the index i, for compactness):

```
Eq1:  0 = r0'*U0' - r*U0 + T0 - (r0'*U0' + r0*U0 + T0)xQ0
          + x*(r1'*U0' - r1*U0 + r0'*U1' - r0*U1 + T1
             - (r1'*U0' + r1*U0 - r0'*U1' + r0*U1 + T1)xQ0)
             - (r0'*U0' + r0*U0 + T0)xQ1 + ...
          + x^n*( Yn - Z0xQn - Z1xQn-1 - ... - ZnxQ0) + ...
```

where we have used:

```
Yn = rn'*U0' +...+ r0'*Un' - rn*U0 -...- r0*Un + Tn
Zn = rn'*U0' +...+ r0'*Un' + rn*U0 +...+ r0*Un + Tn
```

Setting this equation equal to zero term by term:

```
(3.0) r0'*(U0' - U0'xQ0) - r0*(U0 + U0xQ0) + T0 - T0xQ0 = 0

(3.1)    r1'*(U0' - U0'xQ0) - r1*(U0 + U0xQ0)
       + r0'*(U1'xQ0 + U0'xQ1) - r0*(U1xQ0 + U0xQ1)
          + T0xQ1 + T1 - T1xQ0                          = 0


(3.n)    Yn - ZnxQ0 - Zn-1xQ1 - ... - Z0xQn             = 0
```

Since U, U' are always unit vectors, an additional set of equations is needed to express that constraint. To be consistent with observations, these vectors only need to be of unit length for x=0 and x=1. It is usually convenient to accomplish this by requiring that they be units for all values of the parameter, resulting in the normalization equations:

$$U0 \cdot U0 = 1$$

$$U0 \cdot U1 = 0$$

$$U0 \cdot Un + U1 \cdot Un-1 + \ldots + Un \cdot U0 = 0$$

4.0 Some Applications Of The Perturbation Equation

The above equations are a bit overwhelming. We will try taming them a bit by considering some specific problems related to observed motion. For each of these, it will be possible to truncate some of the expansions yielding more tractable forms. Two complementary questions to pose for each of these formulations are consistency and solvability. The exact equation is usually inconsistent. That is, there will commonly be no single motion exactly consistent with all the (errorful) point observations used. It is the requirement of exact satisfiability which causes the idealized equations to be ill-behaved in practice. On the other hand, something like Eq1, while consistent with any observations whatever, contains arbitrarily many unknown values for each observed point, and hence cannot be solved. An important art is to establish a set of solvable equations for which known approximation methods will obtain answers (not really solutions) which are not highly sensitive to errors.

Once gotten, the equations can be used in two different ways. One method is to take a few of the low-order equations (perhaps only the zeroth and first) and treat them as exact. We will use this approach in most of what follows. The second approach is to solve the entire system of equations. Because of the simple form of the equations, the infinite series for Q, T etc. can be formally summed, and approximate solutions for the first few terms extended to a complete solution.

4.1 Known Motion, Errorful Observations

We might first consider a problem mentioned earlier: finding 'best' values of the depths, knowing that the position measurements are inexact. The usual least squares methods treat the observations as exact, determining the depth values that create the least anomoly, while in most cases some observation error would be expected. While other regression methods can provide more equable treatment of the variables, they still fail to take into account the variation in sensitivity of the equations to different errors. So. Suppose we assume Q and T exactly known. The low order equations would then become:

(3.0') r0'*(U0' - U0'xQ) - r0*(U0 + U0xQ) + T - TxQ = 0

(3.1') r1'*(U0' - U0'xQ) - r1*(U0 + U0xQ)
+ r0'*U1'xQ - r0*U1xQ = 0

We recognize (3.0') as the usual 'exact' equation, but now it is interpreted as giving the first approximation to r, r' based on the observed U0, U0'. Equation (3.1') involves refinements of both depth and position. This equation is linear in all the new parameters, but since there are six new parameters introduced for each point and only three constraints, solution is not possible. We can get useful information, however. First, we can notice that:

$$r1' * (Q.U0') = r1 * (Q.U0)$$

so we have some idea of the relative accuracy of our two depth estimates. Second, if we have some estimate deltaU of the reliability of U and U',

$$|r1' * (U0' - U0'xQ) - r1 * (U0 - U0xQ| \approx (r0+r0') * |Q| * deltaU$$

which gives an estimate of the absolute errors possible.

Neither of these results is startling, but they might be of practical use, and certainly are more informative than the error bounds one would obtain from simple regression.

4.2 Approximately Known Motion, Accurate Observations

When we already know the motion, but not to sufficient accuracy to allow exact depth calculations to be useful, we can again simplify the perturbation equations enormously.

This allows use of additional points to get depth information and at the same time to refine the motion estimate.

For this case, assume Q = Q0 + x*Q1, T = T0 + x*T1, where Q1 and T1 are the corrections to the initial estimates Q0, T0. Further, assume U' and U to be accurately known. Eq1 is much simplified, giving:

(4.1) r0'*(U' - U'xQ0) - r0*(U + UxQ0) + T0 - T0xQ0 = 0

(4.2) r1'*(U' - U'xQ0) - r1*(U + UxQ0)
+ (r0'*U' + r0*U + T0)xQ1 + T1 - T1xQ0 = 0

etc.

Examining these equations in order, we see that (4.1) is three linear equations in the two unknowns r0', r0. Thus, we can obtain an initial approximation to the point distances using the assumed motion parameters. Proceeding to (4.2), we see that these three equations introduce eight new parameters: r1, r1', Q1 and T1. Obviously, they cannot be obtained by examining the motion of a single point. If we analyze several points sharing a common motion, however, we see that for each new point only two new depth parameters are introduced in the three equations. Thus, if we have at least six points available, we can get values for the motion corrections, Q1, T1. The remaining equations allow improved estimation of depths, as well as allowing examination of the consistency of the parameters already obtained.

4.3 Unknown Motion From A Pair Of Observations

Though the above expressions can be useful, one may be more interested in using the equations to allow estimation of the motion parameters when there is no a priori motion information available. This case suggests somewhat different perturbation choices, and yields slightly different equations.

In this case, the values of Q0 and T0 are zero: 'no motion'. Both Q and T are taken to be linear in x, as before. The initial position and distance are independent of the effects of motion. We take the second observation to be a linear function of the perturbation parameter (the simplest coherent assumption). For this case, the perturbation equations become:

(4.3) r0'*U' - r*U = 0

(4.4) r1'*U0' + r0'*U1' + T1 - 2*r*UxQ1 = 0

and for N>1, the Nth-order equation:

(4.5) Yn - Yn-1xQ1 = 0

where in this simple case, (n>1)

Yn (= Zn) = rn'*U0' + rn-1'*U1'

In addition, the requirement that U' remains a unit vector means:

$$2*U0'.U1' + U1'.U1' = 0$$

4.3.1 First-order Relations -

The zero-order equation for this case simply notes the equality of r and r0', and of U and U0'. More interesting information is contained in the first-order equation, which describes the relationship between the initial and final observations in terms of the perturbing elements Q1 and T1, and the distances to the points.

Either by examination of the equations or by determined introspection, one may realize that without knowing either the amount of displacement or the distance to some point, there is no way even in principle to determine the overall scale of the images being observed. One may, then, take some particular point to be at unit distance initially. The corresponding set of three equations is linear, with one other unknown distance and the six unknown motion parameters. Every additional observation of a point adds three equations and two distance parameters. Thus, approximate values of all motion parameters and (relative) depth maps are attainable (in principle) by observing as few as five comoving points.

In optic flow or for small displacements, (where a first-order approximation is quite good) the apparent motion of a point can be seen to result from the components of rotation and of translation perpendicular to the line-of-sight to the point. That can be made clear by the following observations: The equation can be written as:

$$(4.4') \quad r*U1' + r1'*U + T1 = 2*r*UxQ1$$

taking the scalar product of each term with U:

$$r1' = r*U \,.\, U1' - U \,.\, T1$$

(the change in depth is due to motion along the line of sight). Rewriting (4.4')

$$U1' = ((U.T1)*T1 - T1)/r + 2*UxQ1 - (U.U1')*U$$

and we can see that the apparent motion consists of three terms: one due to translation perpendicular to the line of sight, one due to rotation perpendicular to the same line, and the last an artificial term due to normalizing all observation vectors to unit length. Note that the contributions of the two motion vectors vary differently with change of observation direction, so by simply observing points lying at different visual orientations, one can establish their separate contributions. Moreover, the radial components of the vectors will be different in different viewing directions, so all six motion parameters can be reliably obtained if (and only if) widely scattered comoving points can be observed.

4.3.2 Second-order Relations -

Next, we can examine the second-order equations. Writing them out explicitly we have:

r2'*U + r1'*U1' - (r1'*U + r*U1' + T1)xQ1 = 0

We see (with some difficulty) that the second-order depth changes are generated by a term which is the product of the perpendicular component of translation with rotation about the viewing axis and another which combines line of sight translation with rotation perpendicular to the viewing axis. Just as for the first-order equation, the new parameter (r2') is a linear function of the observed data and previously obtained parameters. Higher order equations merely provide identities which hold for the case of perfect observations, and would only be of interest for measuring the accuracy of the calculated parameters. The observations would rarely be sufficiently accurate for such relations to be meaningful.

Notice that we get more from these equations than just the calculational formulae we need. From the arrangement of motions into different orders, we get some notion of the dependencies of the observed quantities on different motion parameters. Thus, (4.4") can be read as an assertion that at any one point, only the net effect of translation and rotation across the field of view is accurately observable. Further, it provides a means of determining how large an angle must separate two sets of points before this composite effect can be disambiguated.

4.4 Series Of Observations: A Brief Comment

An attractive problem for the perturbation method is determining motion from a series of observed point positions. This is perfectly feasible, and clearly useful. However, while pairs of positions can be analyzed without any reference to the dynamics of physical objects, much more is required to make sense of a series of positions.

Ideally, it would seem one would want to be able to determine an arbitrary rigid motion. No restrictions should be imposed as to the regularity of the motion from image to image. Somewhat surprisingly, this case amounts to analyzing the sequence of points a pair at a time. This is the case we have treated.

Because of the noisiness of observational data, we often want to use the consistency of the motion parameters with time as a further measure of accuracy of the calculations. For real motion, we often know that the rotation Q can change only slowly. To assume that T will be well-behaved, however, normally requires that we express the motion in some privileged coordinate system. Various choices of privileged frame are sensible, the choice depending on the physical situation: whether the observer motion is expected to be significant, and so on. It seemed best to avoid attacking these complications along with those directly posed by the perturbation formulation.

5.0 The Case Of 'pure' Motions.

While we have treated a variety of special cases, they may all seem a bit abstract. Two cases can be easily shown in more detail. We will consider the simplest special motions: pure rotation, and pure translation.

5.1 Pure Rotation

If the vector T vanishes identically, the equations above are particularly simple. Since Q = x*Q1,

Eq0' $U' = U + (U' + U) \times Q1$

where we have noted that r = r' identically.

(4.4") $U1' = 2*U \times Q1$

gives us:

$$Q1 \cdot (U \times U1') = 1/2*(U1' \cdot U1')$$

and

$$Q1 \cdot U = 0$$

leaving one projection of Q unknown.

Next, examining the higher order equations:

(4.5") Un' = Un-1'xQ1 (n>1)

we find that we can obtain similar projections of Q on all Un', allowing us to derive the (elementary) exact result:

$$Q1 . (UxU') = 1/2*(U - U').(U - U')$$

As each point observation provides us with two components of Q (and no information at all about the component along the line of sight), observing the motion of two or more points determines Q completely. Realistically, it is clear that to get accurate information about all three components, it's necessary to have at least two observations which lie in substantially different viewing directions.

6.0 Pure Translation

This case is very similar. With Q = 0 and T = x*T1:

Eq0: r'*U' + T = r*U

(4.4*) r1'*U + r*U1' + T1 = 0

(4.5*) Yn' = 0

so

$$r1' = - T1 . U$$

and T1 - (T1.U)*U = -r*U1'

Thus, because of the unknown depth parameter, we can obtain only the ratio of the components of motion orthogonal to the viewing direction from each data point. Otherwise, the expressions are much like those for pure rotations, with two points determining the direction of translation. (It has frequently been pointed out that if one finds the great circle passing through the two images of a particular point, that the solution point is given geometrically by the intersection of the great circles corresponding to two observed points). Further, (again examining an infinite sum of terms)

$$T1 - (T1.U)*U = -r*(U' - U)$$

And just as before, the first-order equation turns out to be the same as the exact one.

7.0 Using The Perturbation Equations: Point Pairs

Finally, a few explicit formulae will be given which can be applied to real images. We will take as particularly convenient the case where a pair of points has been found in an image, along with the successors in a later image. We will see that spatial motion has three main effects on observed pairs: displacing the pair, changing the separation of its points, and causing rotation about its center. The latter two effects arise primarily from the translational and rotational motions along the viewing direction, while the former is the net result of all the motions perpendicular to the viewing axis.

7.1 Dipole Lengths (apparent Angles)

For points on the unit sphere, a measure of the distance between two points is just the angle between their direction vectors. More conveniently, the cosine of this angle is given by the scalar product of the direction vectors. If the points are associated with the same rigid body, how do such distances change as the body moves?

We imagine observing two points, iX, jX --> iX', jX'. The perturbation equations derived in section 4 will apply equally to each of these points. If in fact they are associated with the same rigid body, the motion parameters 'Q' and 'T' must be identical for both iU and jU. Therefore, from the given equations we can relate the iUk and jUk to motion parameters. (The complex subscripting seems unavoidable: observations i and j involve functions of order k in x). We can take U' to be simply U + x*U1, Q = x*Q1, and T = x*T1. For these to be consistent assumptions, we must abandon unit normalization for U; instead choosing to maintain a unit projection on the original vector. That is, U'.U = 1, which implies U1.U = 0. Thus,

$$iU'.jU' = iU.jU + x*(iU1.jU + iU.jU1) + x2*iU1.jU1$$

and (4.4) yields:

```
ir1'*jU.iU + ir*jU.iU1 + jU.T1 - 2*ir*jU.(iUxQ1)        = 0
```

combining this with the same equation with indices permuted, and observing that ir1' = -T1.iU, we find:

```
iU'.jU' - iU.jU ~= iU.jU1 + iU1.jU
                 = (1 - iU.jU)*(T.iU/ir + T.jU/jr)
```

so the change in dipole separation is proportional to the separation, the amount of translation along the viewing direction, and the average inverse distance to the two physical points observed. As was noted earlier, this equation connects T and the ir, but with one undetermined scale factor. (We could choose either the length T or some particular ir to establish a unit of distance.)

7.2 Apparent Rotation

Rotation of point-pairs is a bit more complicated to analyze, but as we shall see, just as rewarding. Examining the cross-product of two dipole vectors, we see:

```
(iU'-jU') x (iU-jU) =   0
                 + x*( iU1xiU - iU1xjU - jU1xiU + jU1xjU )
                 + x2* ...
```

From the above equations we can show

$$iU1 \; x \; iU = -(T \; x \; iU)/ri + 2*(Q - (Q.iU)iU)$$

$$iU1xjU = -(TxjU - (T.iU)iUxjU)/ri + 2*((iU.jU)Q - (Q.jU)iU)$$

therefore:

$$(iU1 - jU1)x(iU - jU)$$
$$= (1/jr - 1/ir)*Tx(iU - jU)$$
$$- ((T.iU)/ir - (T.jU)/jr)*iUxjU$$
$$-2*(2*(1 - iU.jU)Q + Q.(iU - jU))*(iU - jU))$$

This expression is a bit awkward to use directly, but taking the dot product with the average direction vector of the dipole we get:

$$((iU1-jU1)x(iU-jU).(iU+jU)$$
$$= (1/rj - 1/ri) \; T.((iU-jU)x(iU+jU))$$
$$-4*(1 - iU.jU) \; Q.(iU+jU)$$

Interpreting this latter expression takes some thought. The first term is a rather complicated one, but vanishes whenever the two observed points lie at the same distance or when they both lie in the same plane as the true translation vector. For the term to be large, both of these conditions must be substantially violated. For most observed dipoles, then, we may expect its contribution to be rather small. The second term is the product of (1 - iU.jU), which depends only on the size of the dipole observed, and Q.(iU+jU), the projection of the true rotation axis on the average observation direction. If we consider the vectors with magnitude Q.(iU - jU) and direction (iU + jU), we can see that they all lie on the surface of a sphere which passes through the origin and through the desired Q vector. The center of the sphere then lies at Q / 2.

Two convenient calculational approaches suggest themselves for extracting these parameters from a set of dipole observations. The first is to do a straightforward least-squares fitting of this second-order surface. The second is to note that the surface can be described by:

$$Qx*xc + Qy*yc + Qz*zc = rc^2$$

so that the parameters can be obtained by as the best plane through the observed (xc, yc, zc, rc2) points. This linear estimation approach is especially convenient, but both approaches allow use of available values for calculating a suitably averaged result.

7.3 Apparent Motion

The final information available from motion of point pairs is the movement of the pair as a whole. This motion conveys basically the same information as the motion of a single point, and the low-order equations of motion are the same. Use of point pairs should be of most use when absolute alignment may be in doubt anyway, so only the relative measurements have been discussed here.

8.0 Discussion

This paper has tried to make a number of points. The principle technical point has been that the mathematical method of perturbation theory is a useful tool in the investigation of motion from image data. The principle philosophical point has been that it is important to extract from data, not just formulas for properties of interest, but some information about the reliability of those properties. We are all accustomed to statistical significance measures as means of getting such information. But 'sensitivity analyses' characteristic of classical numerical analysis and of the perturbation method provide an *a priori* estimate of accuracies.

In addition, a number of approximation formulas were obtained which should be of considerable practical use.

Acknowledgements

The bulk of this work was done in the COINS Dept, where Ed Riseman and Al Hanson and their group provided a highly stimulating environment. Discussions with Daryl Lawton on problems of motion have been particularly valuable.

9.0 References

[Fennema+ 1979] C.L. Fennema and W.B. Thompson, "Velocity determination in scenes containing several moving objects," CGIP, vol 11, 1979, pp. 301-327

[Huang 1981] T.S. Huang, (editor) Image Sequence Analysis, Berlin-Heidelberg: Springer-Verlag, 1981

[Lawton 1983a] D.T. Lawton, "Processing translational motion sequences", CGIP, vol 22, 1983, pp. 116-144.

[Lawton 1983b] ----------, "Processing dynamic image sequences from a moving sensor", PhD thesis, U Mass, 1983

[Longuet-Higgins 1981] H.C. Longuet-Higgings, "A computer algorithm for reconstructing a scene from two projections", Nature, vol 293, 1981, pp. 133-135

[Morse+ 1953] P.M. Morse and H. Feshbach, Methods of Theoretical Physics, vol 2, New York, McGraw-Hill, 1953, pp. 1109ff

[Nagel 1981] H.H. Nagel,"On the derivation of 3-D rigid point configurations from image sequences", Proc. of the Pattern Recognition and Image Processing Conference, Dallas, Texas 1981, pp. 103-108

[Pervin+ 1983] E. Pervin and J. Webb, "Quaternions in computer vision and robotics", Proc Computer Vision and Pattern Recognition Conf, Wash D.C., 1938, pp. 382-383

[Prazdny 1981] K. Prazdny, "Determining the instantaneous direction of motion from optical flow generated by a curvilinearly moving observer," Proc. of the Pattern Recognition and Image Processing Conference, Dallas, Texas 1981, pp. 109-114

[Roach+ 1980] J.W. Roach and J.K. Aggarwal, "Determining the movement of objects from a sequence of images," IEEE Transactionson Pattern Analysis and Machine Intelligence, vol PAMI-2, 1980, pp. 516-521

[Tsai+ 1982] R.Y. Tsai and T.S. Huang, "Uniqueness and estimation of three-dimensional motion parameters of rigid objects with curved surfaces", Proc. of the Pattern Recognition and Image Processing Conference, Las Vegas, Nevada 1982, pp. 112-118

[Ullman 1979] S. Ullman, "The interpretation of structure from motion," Proc. Royal Soc. London, vol B203, 1979, pp. 405-426

[Yen+ 1983] B.L. Yen and T.S. Huang, "Determining 3-D motion parameters of a rigid body: a vector-geometrical approach", Proc ACM Workshop, Motion: Representation and Perception, Toronto 1983, pp. 78-90