What is wrong with so-called 'linear' photogrammetric methods?

Berthold K.P. Horn

1999 August 12

Several photogrammetric problem lead to sets of coupled non-linear equations. In some cases, claims have been made that the solutions can be found by solving some linear equations instead. These claims invariable depend on use of a transformation from the world to the image that does not reflect physical reality — i.e. perspective projection. Essentially, to obtain a set of linear equations, some additional — and unnatural — degrees of freedom are 'added' to perspective projection — such as anistropic scaling or skew of what should be perpendicular axes. While the 'linear' methods may produce the correct result with perfect data, they produce physically meaningless solutions when measurement errors are present. Further, because of the added degrees of freedom, additional contraint is required. Importantly, these methods also are more sensitive to noise than the correct photogrammetric methods. In the case of the central photogrammetric problems in 3-D this can be hard to demonstarte analytically or numerically. As a result, it is difficult to make this important point clearly and convincingly. In this regard, it can be instructive instead to first consider equivalent 2-D problems, where the derivations are short and transparent,

Two-dimensional analog of the "Location Determination Problem"

In the three dimensional "Location Determination Problem" (LDP), the "3-point" method leads to three coupled sets of quadratic equations. The "4-point" method instead leads to linear equations, hence is easier to implement, but since it does not model the physical reality correctly, is less accurate (despite using more measurements). With perfectly accurate measurements the 3-point and 4-point methods produce the same result. So the question is one of how they differ in sensitivity to measurement errors. It may be hard to see the fundamental differences between these two methods in three dimensions.

One way to gain intuitive insight is to consider an analoguous but simpler two-dimensional problem. Consider the problem of finding the position and orientation of a flat part on a conveyor belt from a direct overhead view. Position and orientation are relative to a standard position and orientation of the part. If we do not accurately know the height of the camera above the conveyor belt, then the scale of the image of the object is also unknown. Overall then we have four unknown quantities: rotation of the part relative to its reference position (θ say), translation of the part relative to its reference position ((x_o, y_o) say), and scale (s say).

We can determine these four parameters by measuring where known reference points on the object appear. Let P_1 be an identifiable point on the object with coordinates (x_1, y_1) in its reference position, and measured coordinates (x_1', y_1'). Since a measurement of a point in the image only gives us two numbers ((x_1', y_1')), we cannot expect to recover all four parameters of the coordinate transformation. Intuitively, if we only know where one point on the object goes to, then we can rotate the object about this point and shrink or expand it relative to this point without changing where that particular point lies. So the transformation is not fully constrained by a single measurement.

Two-dimensional 2-point method (Non-linear)

If, in addition, we measure a second point P_2 with reference coordinates (x_2, y_2) and measured coordinates (x'_2, y'_2), we have four numbers, which is enough to solve for the four unknowns of the transformation (clearly then two points are the minimum required to obtain a solution). In fact, the transformation is

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = s \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} x_o \\ y_o \end{pmatrix}$$

where the translation (x_o, y_o) is given by

$$\begin{aligned} x_o &= \tfrac{1}{2}(x'_1 + x'_2) - \tfrac{1}{2}(x_1 + x_2) \\ y_o &= \tfrac{1}{2}(y'_1 + y'_2) - \tfrac{1}{2}(y_1 + y_2) \end{aligned}$$

the scale $s = r'/r$ where

$$r = \sqrt{\delta x^2 + \delta y^2} \quad r' = \sqrt{\delta x'^2 + \delta y'^2}$$

and

$$\begin{aligned} \delta x &= x_2 - x_1 & \delta x' &= x'_2 - x'_1 \\ \delta y &= y_2 - y_1 & \delta y' &= y'_2 - y'_1 \end{aligned}$$

while the rotation matrix components are given by

$$\begin{aligned} \cos\theta &= \frac{1}{rr'}(\delta x\, \delta x' + \delta y\, \delta y') \\ \sin\theta &= \frac{1}{rr'}(\delta x\, \delta y' - \delta y\, \delta x') \end{aligned}$$

This transformation maps P_1 exactly into P'_1, and P_2 exactly into P'_2 — as can be verified by substituting the coordinates into the equation. If there are no measurement errors the same transformation also maps all other points of the object into their correct rotated, translated and scaled position.

Two-dimensional 3-point method (Linear)

The above solution of the problem requires use of square-roots (arising from the solution of a pair of quadratic equations). One may wish to find a method that only involves linear equations, assuming this is possible. We can do this by relaxing the constraint on the transformation and "generalizing" it to the following linear transformation:

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} e \\ f \end{pmatrix}$$

with six unknowns a, b, c, d, e, and f. Note that the actual transformation, discussed above, is a special case of this with the two constraints

$$a^2 + b^2 = c^2 + d^2 \quad \text{and} \quad ac + bd = 0$$

since $\sin^2\theta + \cos^2\theta = 1$ and $\cos\theta\sin\theta - \cos\theta\sin\theta = 0$.

With six unknowns, we need three measurements to obtain enough constraint (since there are two coordinates per measurement). Each such measurements yields two equations of the form:

$$\begin{aligned} x'_i &= a x_i + b y_i + e \\ y'_i &= c x_i + d y_i + f \end{aligned}$$

for $i = 1, 2, \ldots 3$. We can view these as six linear equations in the six unknowns a, b, c, d, e, and f. These can be conveniently split into two independent groups of three linear equations:

$$\begin{pmatrix} x_1 & y_1 & 1 \\ x_2 & y_2 & 1 \\ x_3 & y_3 & 1 \end{pmatrix} \begin{pmatrix} a \\ b \\ e \end{pmatrix} = \begin{pmatrix} x_1' \\ x_2' \\ x_3' \end{pmatrix}$$

$$\begin{pmatrix} x_1 & y_1 & 1 \\ x_2 & y_2 & 1 \\ x_3 & y_3 & 1 \end{pmatrix} \begin{pmatrix} c \\ d \\ f \end{pmatrix} = \begin{pmatrix} y_1' \\ y_2' \\ y_3' \end{pmatrix}$$

Conveniently, the 3×3 coefficent matrix is the same in the two groups of equations. The determinant of the coefficient matrix is

$$x_1 y_2 - x_1 y_3 + x_3 y_1 - x_2 y_1 + x_2 y_3 - x_3 y_2$$

which is zero when the three points lie on a straight line — in this case the 3-point method fails entirely. If the three points are approximately aligned, the determinant will be small, and the solution unstable, since then one will be dividing by a very small determinant, with the potential for great amplification of measurement error. The 2-point method — despite the fact that it uses fewer points — has no such problem, of course.

Numerical Example

Suppose $(x_1, y_1) = (-1, 0)$, $(x_2, y_2) = (0, \epsilon)$, and $(x_3, y_3) = (1, 0)$, with ϵ small ($\epsilon \ll 1$). The coefficient matrix is in this case

$$\begin{pmatrix} -1 & 0 & 1 \\ 0 & \epsilon & 1 \\ 1 & 0 & 1 \end{pmatrix}$$

The determinant is -2ϵ, so can be made as small as we please by making ϵ small. The inverse of the coefficient matrix is

$$\frac{1}{2} \begin{pmatrix} -1 & 0 & 1 \\ -1/\epsilon & 2/\epsilon & -1/\epsilon \\ 1 & 0 & 1 \end{pmatrix}$$

The coefficients of the solution that are particularly sensitive to small measurement errors in this case are

$$\begin{aligned} b &= -(x_1' - 2x_2' + x_3')/(2\epsilon) \\ d &= -(y_1' - 2y_2' + y_3')/(2\epsilon) \end{aligned}$$

The “second differences” that appear in these two formulae will be very small when the three points are nearly colinear and evenly spaced, as in this example. So b and d are given as the ratio of two very small numbers.

If ϵ is small, $1/\epsilon$ is large, and so small errors in measuring any of the six individual coordinate components x_1, y_1, x_2, y_2, x_3, and y_3, can induce a large error in the coefficients b and d. For example, if ϵ happens to be 0.001 then measurement errors are amplified one thousand fold in the computation of b and d. This in turn will affect how other points away from the measured points are mapped by the transformation.

Summary

The two dimensional problem can be solved using methods for linear equations, but at the cost of not enforcing the known physical constraint. If we solve the problem with six unknown parameters, it will in general not yield a solution that corresponds to physical reality. That is, the 2×2 matrix that appears in the transformation will not be the product of a scale factor and an orthonormal rotation matrix unless measurements are perfect. In general, the matrix will not satisfy the constraint that its two rows be orthogonal, and of the same magnitude.

With perfect measurement accuracy, the recovered transformation will be correct either way. So the question is what happens when there are errors in the measurements. The transformation recovered will always fit the three measured points exactly — but large errors will be apparent when we move away from the points for which we have measurements.

Three point measurements overconstrain the problem since we know that only two points are needed. The extra information can be used effectively to obtain a “best fit” or “least squares” solution (we don’t show details of this here, but in the two dimensional case there is a closed form solution to this least squares problem) which will be better. If we are willing to measure more points, we should use them optimally, not merely to make possible a simple linear solution method that discards some of the real constraints and reduces accuracy.

The same argument applies to the three dimensional “Location Determination Problem” (a.k.a. Exterior Orientation in photogrammetry). In this case, three point measurements are enough, while four measurements make it possible to use a simpler linear method, but at the cost of throwing away information.

A more famous analoguous situation can be found in relation to the problem referred to as "Relative Orientation” by photogrammetrists, where the task is to recover the translation and rotation of one camera relative to another using only measurements form the images taken by the two cameras. Solution of this problem is fundamental to the reduction of aerial photography in the creation of topographic maps. The same problem also arises in “motion vision” where the translational and rotational motion of a vehicle is to be recovered from “before” and “after” images of the same scene.

In this case, the minimum number of points is five, leading to a coupled system of five quadratic equations. It has been proposed that one can solve this problem using instead eight measurements and selectively discarding constraints in order to obtain a set of linear equations which can be easily solved. There have been numerous papers pointing out the sensitivity of the linear “eight-point” method to measurement error. These are backed up by calculations based on large numbers of simulated measurements.

In photogrammtry, the problems of “interior orientation,” “exterior orientation,” “relative orientation,” and “absolute orientation” are central. Since photogrammetrists are committed to make the best possible use of the (in their situation expensive) raw data, they never considered anything but least squares methods using many more measurements than the minimal set required. They do not resort to methods that throw away some constraint in order to simplify the equations, since this introduces errors that they do not need to accept.