Robust Software for Computing Camera Motion Parameters

Robert Schreiber Zeyu Li Harlyn Baker
HP Labs, 1501 Page Mill Road, Palo Alto, California 94304
{rob.schreiber,zeyu.li,harlyn.baker}@hp.com

March 5, 2007

Abstract

We revisit the method of Tsai, Huang, and Zhu for the computation of camera motion parameters in computer vision. We elucidate some spectral properties of the homography matrices that arise, which are rank-one perturbations of rotation matrices. We show how to correct for noise by finding the rank-one perturbation of a rotation closest to a give matrix. We illustrate some of the inaccuracies and computational failures that can arise when using the formulas given by Tsai, and we propose new formulas that avoid these pitfalls. A computational experiment shows that the new methods are indeed quite robust.

1 Introduction

A 3×3 homography matrix H maps the image of a plane taken by one camera to the image of the same plane taken by a second camera. The relative position of these two cameras can be described by a three dimensional rotation and a translation that takes the position of one of the cameras to the position of the other. The rotation may be described by a 3×3 rotation matrix, an orthogonal matrix whose determinant is $+1$. A suitably scaled homography H satisfies

$$H = R - xy^T \tag{1}$$

where R is a rotation matrix, x is a scaling of the translation, the difference between the coordinates of the camera centers, and y is a scaling of the normal of the imaged plane. (Clearly x and y are determined by their outer product xy^T only up to a mutual scaling.) We call any matrix of the form (1) a *rank-one perturbation of a rotation* or ROPR. The basic problem that we consider is, given the ROPR matrix H, to compute a triple R, x, and y, which are the motion parameters, so that (1) holds. In all discussions of ROPRs we shall normalize y so that $||y|| = 1$, where $|| \cdot ||$ denotes the 2-norm, so as to pin down the mutual rescaling. This still leaves open the signs, since changing the signs of both x and y changes nothing, but we won't specify the sign. A *unit vector* is a vector z satisfying $z^T z = ||z||^2 = 1$.

The basic problem was solved by Tsai, Huang, and Zhu [5]. They explained the mathematics using geometric arguments, and derived formulas for the computation of the motion parameters. Some things, however, were left for us to consider. In this paper we shall state and solve three problems that, in their solution and in the context of the Tsai paper, complete a full theory of the computation of the motion parameters. First we prove properties, relevant to the computation, of a ROPR. We discuss the existence and uniqueness of solutions of (1) and derive some new results. Second, based on a theorem of Nievergelt, we show how to find the ROPR closest to a given matrix. Finally, we discuss inaccuracies in the computed results that can arise from roundoff error when the formulas of Tsai are used, and we present alternatives that are at once simpler and more accurate.

2 Properties of a rank-one perturbation of a rotation

In this section we shall determine the existence and uniqueness of solutions to the basic problem. Our discussion hinges on the multiplicity of the singular values of H. We shall prove that the second singular value of any ROPR is one. The first singular value may be greater than or equal to one and the third may be less than or equal to one. In the general case all three of the singular values are different; the special cases can arise when the rotation of the image-plane normal vector is colinear with the translation vector, in other words when x is a scalar multiple of Ry. We shall show that the basic problem has infinitely many solutions when H has three singular values equal to one, has a unique solution when two of the singular values of H are equal to one, and has exactly two solutions when the singular values of H are distinct.

Every matrix H has a singular value decomposition

$$H = U\Sigma V^T. \tag{2}$$

Here U and V are orthogonal matrices, $U^TU = V^TV = I$, and

$$\Sigma = \begin{pmatrix} \sigma_1 & 0 & 0 \\ 0 & \sigma_2 & 0 \\ 0 & 0 & \sigma_3 \end{pmatrix} \tag{3}$$

is a diagonal matrix whose diagonal elements, the singular values of H, are given in nonincreasing order: $\sigma_1 \geq \sigma_2 \geq \sigma_3$. We will use the notation $D(a, b, c)$ for the diagonal matrix of order 3 having $a, b,$, and c on the diagonal; so $\Sigma = D(\sigma_1, \sigma_2, \sigma_3)$. The singular values of H are the positive square roots of the eigenvalues of HH^T, the columns of U are eigenvectors of HH^T and the columns of V eigenvectors of H^TH.

We introduce the vector $w = Ry$ and the matrix $B \equiv HH^T - I$. It is elementary that the matrix H is orthogonal if and only if $B = 0$, if and only if $\Sigma = I$.

If H is a ROPR, then

$$\begin{aligned} B &= xx^T - wx^T - xw^T \\ &= (x - w)(x - w)^T - ww^T. \end{aligned} \tag{4}$$

We shall also make reference to $C \equiv H^TH - I$, which when H is a ROPR satisfies

$$C = (x^Tx)yy^T - (R^Tx)y^T - y(R^Tx)^T\ . \tag{5}$$

As addition of I does not change eigenvectors, we see that U consists of eigenvectors of B and V of eigenvectors of C. If H is a ROPR then because the columns of B are all linear combinations of x and Ry, it cannot have full rank: $\text{rank}(B) \leq 2$. Moreover, $\text{rank}(C) = \text{rank}(B)$, since the columns of C are linear combinations of R^Tx and R^TRy. We deal in turn with the three possibilities, $\text{rank}(B) = 0, 1,$ or 2, below.

Lemma 1 *If z is a unit vector then the matrix $X \equiv I - 2zz^T$ is symmetric and orthogonal, and* $\det(X) = -1$.

Proof. It is straigtforward that $X^TX = I$, so that X is symmetric and orthogonal. Extend z to an orthonormal basis let Z be the orthogonal matrix with this basis as its columns. Then $XZ = Z - 2zz^TZ = ZD(-1, 1, 1)$ whence $\det(X) = \det(Z^TXZ) = -1$. **QED**

Matrices of the form $I - 2zz^T$ with unit-length z are called *elementary reflectors* or *Householder transformations* after Alston Householder, who pioneered their use in matrix computation [3].

Thanks to the SVD, the basic problem can be recast as the problem of finding a rank-one perturbation of a diagonal matrix that is orthogonal and has determinant of the correct sign, namely

$$\Delta \equiv \det(U)\det(V) \; ,$$

Lemma 2 *The triple R, x, and y solves (1) if and only if the matrix*

$$Q \equiv \Sigma + (U^T x)(V^T y)^T$$

is orthogonal and $\det(Q) = \Delta$.

Proof. If Q has the specified properties then $R = UQV^T = H + xy^T$ is a rotation and is a rank-one perturbation of H. If R, x, and y satisfy (1) and R is a rotation then $U^T RV = U^T(H + xy^T)V = Q$ is orthogonal and has determinant equal to Δ. **QED**

We consider first the case in which $\mathrm{rank}(B) = 0$.

Lemma 3 *The ROPR H is orthogonal (and $\Sigma = I$, and $B = 0$) if and only if either $x = 0$ or $x = 2Ry$. In the former case,* $\det(H) = 1$ *and in the latter,* $\det(H) = -1$.

Proof. By (4), $B = 0$ if and only if $x - w = \pm w$ if and only if $x = 0$ or $x = 2w$. When $x = 0$, $H = R$ is a rotation and its determinant is 1. When $x = 2w = 2Ry$, $H = R - 2Ryy^T = R(I - 2yy^T)$ and hence $\det(H) = \det(I - 2yy^T) = -1$ by Lemma 1. **QED**

We now completely understand the case $B = 0$, $\Sigma = I$, and H is orthogonal. If the determinant of H is one, then the displacement vector $x = 0$ and there is no way to recover the image plane normal vector y from H; there are infintely many solution triples. Indeed $R = H$, $x = 0$, y satisfy (1) for any unit-length y. On the other hand, if the determinant of H is -1, there are again infinitely many solutions. Let z be any unit vector, and $X = I - 2zz^T$ the corresponding elementary reflector. Let $R = HX = H(I - 2zz^T) = H - (2Hz)z^T$. Clearly, $\det(R) = \det(H)\det(X) = 1$, so R is a rotation. The triple R, $x = -2Hz$, and $y = z$ satisfies (1), and is clearly distinct for each unit vector z.

We summarize this in the following.

Theorem 1 *Let H be orthogonal. Then H is a ROPR. If* $\det(H) = 1$ *then $R = H$, $x = 0$, and $y = z$ satisfy (1) for any unit vector z. If* $\det(H) = -1$ *then $R = HX$, $x = -2Hz$, and $y = z$ satisfy (1) for any unit vector z and corresponding elementary reflector $X = I - 2zz^T$.*

The case $\mathrm{rank}(B) = 0$ now being settled, we consider the next possibility, in which $\mathrm{rank}(B) = 1$, or equivalently the case in which two of the singular values of H are unity. As shown next, this case arises when $x = \alpha Ry$ with $\alpha \notin \{0, 2\}$.

Lemma 4 *If $x = \alpha Ry$ then the symmetric matrix B in (4) has*

(i) one negative eigenvalue if $0 < \alpha < 2$,

(ii) no nonzero eigenvalue (that is, $B = 0$) if $\alpha = 0$ or $\alpha = 2$,

(iii) one positive eigenvalue if $\alpha \notin [0, 2]$.

Proof. If $x = \alpha R y$ then $B = (\alpha^2 - 2\alpha) w w^T$. The vector $w = Ry$ is nonzero. Thus, $\text{rank}(B) \leq 1$, at least two of the eigenvalues of B are zero, and the third is given by $\alpha^2 - 2\alpha$. **QED**

Suppose we have a solution R, x, y to the basic problem. From Lemma 2 we see that with the definitions $u \equiv U^T x$ and $v \equiv V^T y$ we have that $Q = U^T R V = \Sigma + u v^T$ is orthogonal. We compute that

$$I = QQ^T = \Sigma^2 + (\Sigma v) u^T + u (\Sigma v)^T + u u^T$$

where we have used $1 = y^T y = v^T (V^T V) v = v^T v$. Suppose without loss of generality that it is the first singular value σ_1 that is different from, indeed greater than, unity. Then we have that

$$\begin{aligned} I - \Sigma^2 &= \begin{pmatrix} 1 - \sigma_1^2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \\ &= (\Sigma v) u^T + u (\Sigma v)^T + u u^T \\ &\equiv M \ . \end{aligned}$$

The matrix M evidently has rank one, and we can therefore conclude that Σv is a scalar multiple αu, and that M is therefore $(1 + 2\alpha) u u^T$. The form of M then implies that only the first element of u is nonzero, and then this is also true of Σv, and hence of v, and finally the normalization of v leads to the conclusion $v = (1,\, 0,\, 0)^T$ and $u = (u_1,\, 0,\, 0)^T$ where u_1 is a root of the quadratic $u_1^2 + 2\sigma_1 u_1 + (\sigma_1^2 - 1) = 0$. The solution $u_1 = 1 - \sigma_1$ leads to $Q = I$ and is appropriate when $\Delta = 1$. The other solution, $u_1 = -1 - \sigma_1$ leads to $Q = D(-1,\, 1,\, 1)$ and is correct when $\Delta = -1$.

When it is the third rather than the first singular value that differs from one, then $v = (0,\, 0,\, 1)$ and $u = (0,\, 0,\, u_3)$, but nothing essential changes. This proves

Theorem 2 *Let H have two singular values equal to one, so that* $\text{rank}(B) = 1$. *Then H is a ROPR and there is a unique solution to (1). When B has a positive eigenvalue then x is a multiple of the first column of U and y is the first column of V. When B has a negative eigenvalue then x is a multiple of the third column of U and y is the third column of V.*

In Tsai, the possibilities in the rank-deficient cases are further limited, perhaps by what is physically realizable in a two-camera situation. Thus, the fully degenerate case ($\Sigma = I$) is identified with a translation x equal to zero; the other possibility is ignored. In the partly degenerate case in which w and x are linearly dependent, it is assumed or perhaps shown via geometric arguments that $\sigma_1 = \sigma_2 > \sigma_3$, in other words that B is negative semidefinite. We think there is considerable value to considering the fully general algebraic rather than geometric problem of reconstructing the constituent terms making up a otherwise completely arbitrary ROPR. So we consider all the possibilities here, even if they cannot be realized with physical cameras.

Moving ahead, we consider for the remainder of this section the generic case, in which $\text{rank}(B) = 2$ and the singular values of H are distinct.

Lemma 5 *Let H be a ROPR and let B be given by (4). B has rank 2 if and only if x and Ry are linearly independent. In that case, the singular values of H satisfy*

$$\sigma_1 > \sigma_2 = 1 > \sigma_3 \, . \tag{6}$$

Furthermore, x is not an eigenvector of B, and y is not an eigenvector of the matrix C of (5).

Proof. The matrix $uu^T - vv^T$ has the span of $\{u, v\}$ as its range. It thus has rank 2 if and only if u and v are linearly independent. The vectors x and w are linearly independent if and only if $x - w$ and w are linearly independent. These observations and (4) prove the first assertion.

It is straightforward to show that when u and v are linearly independent, the matrix $A = uu^T - vv^T$ has one positive and one negative eigenvalue (and the rest are zero). Indeed, if z is in the span of u and v and is orthogonal to v then $z^T Az = (z^T u)^2 > 0$, and if z is in the span of u and v and is orthogonal to u then $z^T Az = -(z^T v)^2 < 0$, which implies that A is indefinite and hence has nonzero eigenvalues of both signs. Since B is of this form, and has rank 2, we know that its eigenvalues are $\beta_1 > \beta_2 = 0 > \beta_3$. Since $HH^T = I + B$, the singular values of H and the eigenvalues of B are related by $\sigma_i^2 = 1 + \beta_i$. This allows us to conclude that the ordering relations in (6) hold.

Again, if u and v are linearly independent and $A = uu^T - vv^T$ then $z = u + v$ is not an eigenvector of A. For if it were, say $Az = \lambda z$, then equating coefficients of u and v (valid due to their independence) we find that $\lambda = u^T u + u^T v = -(v^T v + u^T v)$ which leads to $-2u^T v = u^T u + v^T v$, or $0 = (u + v)^T (u + v)$ which implies that $u = -v$, which is impossible because they were assumed to be linearly independent. Now note that x is the sum of $x - w$ and w, and that (4) holds, whence we can conclude that x is not an eigenvector of B. Similar reasoning applied to y shows that it is not an eigenvector of C. **QED**

Lemma 5 has as an immediate consequence the following:

Theorem 3 *The second singular value of a ROPR is equal to one.*

Proof. If $\mathrm{rank}(B) = 0$ then $\Sigma = I$; all of the singular values are equal to one. If $\mathrm{rank}(B) = 1$ then either the largest singular value is greater than one or the smallest is less than one, but the other two, and always the second, are equal to one. If $\mathrm{rank}(B) = 2$ then Lemma 5 applies. Finally, as noted above, $\mathrm{rank}(B) > 2$ is impossible. **QED**

Let $A(:, k)$ denote the k^{th} column of the matrix A.

Lemma 6 *Let H be a ROPR that is not orthogonal. Then $U(:, 2)$ is orthogonal to both x and Ry, and $V(:, 2)$ is orthogonal to both y and $R^T x$.*

Proof. If x and Ry are linearly dependent then so are $R^T x$ and y. In that case, and since $B \neq 0$, we have by Lemma 2 that both x and Ry are multiples of either $U(: .1)$ or of $U(:, 3)$, and similarly that both $R^T x$ and y are multiples of either $V(: .1)$ or of $V(:, 3)$. In case $\mathrm{rank}(B) = 2$, Lemma 5 shows that the eigenvalues of B are distinct, and the second of them is zero. By (4), the one-dimensional null space of B is the set of vectors orthogonal to both x and Ry. And $U(:, 2)$ is the normalized null vector of B, so it is orthogonal to x and Ry. Similar arguments based on $H^T H$ yield the corresponding conclusion concerning $V(:, 2)$. **QED**

Lemmas 5 and 6 give us the following guide to where to look for the vectors x and y.

Corollary 1 *When the ROPR H is not orthogonal, x is a linear combination of $U(:, 1)$ and $U(:, 3)$ and y is a linear combination of $V(:, 1)$ and $V(:, 3)$. Moreover, when the singular values of H are distinct, then $x = aU(:, 1) + cU(:, 3)$ and $y = bV(:, 1) + dV(:, 3)$ and none of the four scalars a, b, c, d is zero.*

Suppose that H is a ROPR, not orthogonal. Then

$$U^T x = \begin{pmatrix} a \\ 0 \\ c \end{pmatrix}$$

and

$$V^T y = \begin{pmatrix} b \\ 0 \\ d \end{pmatrix}$$

We know (Theorem 3) that $\sigma_2 = 1$. Thus,

$$\begin{aligned} Q &= \Sigma + U^T x (V^T y)^T \\ &= \begin{pmatrix} \sigma_1 + ab & 0 & ad \\ 0 & 1 & 0 \\ cb & 0 & \sigma_3 + cd \end{pmatrix} \end{aligned} \tag{7}$$

Clearly, the 2×2 matrix

$$Q_2 = Q_2(a, b, c, d) = \begin{pmatrix} \sigma_1 + ab & ad \\ cb & \sigma_3 + cd \end{pmatrix} \tag{8}$$

is orthogonal along with Q, and it has the same determinant.

Thus, the problem of computing the motion parameters reduces to the diagonal, 2×2 case. Given the computed first and third singular values

$$\sigma_1 > 1 > \sigma_3 \tag{9}$$

of H, (suitably adjusted as above to enforce these relations) determine the scalars a, b, c, and d so that $Q_2(a, b, c, d)$ is orthogonal and has the desired determinant, namely Δ.

An orthogonal matrix of order 2 having positive determinant is a plane-rotation matrix of the form

$$\begin{pmatrix} C & S \\ -S & C \end{pmatrix} \tag{10}$$

where $C = \cos\theta$ and $S = \sin\theta$. Any orthogonal matrix of order 2 whose determinant is -1 is of the form

$$\begin{pmatrix} C & S \\ S & -C \end{pmatrix} \tag{11}$$

for a sine, cosine pair.

While we already showed that none of the four scalars can be zero, we now have another simple proof of the fact. We claim that the off-diagonal elements of Q_2 are nonzero, which implies this. For if they, the sines, are zero, then the diagonal elements, the cosines, are ± 1. In order for this to occur and in view of (9), we would need that both ab and cd be nonzero, whence the off diagonal entries would also be nonzero, a contradiction. And if the sines are nonzero then by (8) none of the four parameters can vanish.

Suppose that $Q_2(a, b, c, d)$ has the desired properties. We claim that $Q_2(-a, -b, c, d)$ has them as well. In the case that Q_2 is a rotation, the sign change alters Q_2 from the plane rotation through the angle θ to the rotation through the angle $-\theta$; it is still a rotation, orthogonal, with determinant one. When $\Delta = -1$, the change in the signs of a and b again changes only the signs of the off-diagonal elements of Q_2, leaving it orthogonal with determinant -1. As none of the four scalars is zero, this sign change represents an actual change to the motion parameters R, x, and y. Thus, when $\mathrm{rank}(B) = 2$, the solutions to the basic problem come in pairs. We need to prove now that there is just one pair that satisfy (1), always up to mutual rescaling of x and y.

To do so we simply note that the requirements on Q_2 amount to a set of quadratic equations, and that these admit exactly two real solutions. We have two equations corresponding to the form of Q_2, either (10) or (11), and one equation that amounts to $C^2 + S^2 = 1$. Because we know that $d \neq 0$, for the moment we take $d = 1$, and we renormalize later. Let $Y \equiv 1 + b^2$ and let $D \equiv \sigma_1 - \sigma_3$. From the assumption that Q_2 is a plane rotation we derive the requirement

$$Y = \frac{\sigma_1^2 - \sigma_3^2}{1 - \sigma_3^2}\,.$$

Thus, with Y uniquely determined, we have only two possiblities for b, namely $b = \pm\sqrt{Y - 1}$. It turns out that $c = D/Y$ is unique and $a = -bc$ changes sign along with b. Similar elementary algebra solves the case $\Delta = -1$ as well.

Theorem 4 *Let H have SVD (2) with distinct singular values satisfying (6). Up to reversal of the signs of x and y, there are ezactly two triples R, x, y with $||y|| = 1$ satisfying (1), corresponding to the two sets, (a, b, c, d) and $(-a, -b, c, d)$ of scalars for which the matrix Q_2 of (8) is orthogonal and has determinant Δ.*

Theorem 5 *A matrix is a ROPR if its second singular value is one.*

Proof. Let H be the given matrix. There are three cases. If all of the singular values of H are one, then Theorem 1 shows that H is a ROPR. If exactly one of the singular values of H differs from one, then Theorem 2 shows that H is a ROPR. And if the second is the only one of H's singular values equal to one, then Theorem 4 shows that H is a ROPR. **QED**

Note that Theorems 3 and 2 are the two halves of a proof that H is a ROPR if and only if its second singular value is equal to one. An analog of this result holds for matrices of arbitrary order. All the singular values except possibly the first and last are one.

3 Finding the closest ROPR

A measured homography might be corrupted by noise and hence lose the ROPR property. How to recover it best? An obvious question is whether one can compute the closest ROPR to the given matrix H. A theorem of Nievergelt gives us a convenient way to do this. Compute the SVD (2) – (3) of H. If $\sigma_2 \neq 1$ then H is not a ROPR. To find the closest ROPR, set σ_2 equal to one, set σ_1 to one of it is not already greater than or equal to one, and set σ_3 to one if it is not already less than or equal to one. After this adjustment of the singular values, the reconstituted $\hat{H} = U\Sigma V^T$ is the ROPR closest to the given H in any unitarily invariant norm (such as the Frobenius norm). This follows from a theorem of Nievergelt [4] When a is an n-vector we write $D(a)$ for the diagonal matrix of order n having the elements of a on the diagonal.

Theorem 6 *Let A be a given matrix, with SVD $A = U_A \Sigma_A V_A^T$. Let a be the ordered vector of singular values, $a = (\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_n)$, for which $\Sigma_A = D(a)$. Let $||\cdot||$ be a unitarily invariant matrix norm. Among all matrices of the same shape as A whose singular values satisfy a given set of linear equations, the closest approximation in the given norm to A is the matrix $B = U_A \Sigma_B V_A^T$ that has the same singular vectors as A and whose singular values $\Sigma_B = D(b)$ where the vector $b = (\beta_1 \geq \beta_2 \geq \cdots \geq \beta_n)$ is nonincreasing, satisfies the given linear equations and is closest to a among all such vectors. It is closest with respect to the induced vector norm $||u - v|| = ||D(u - v)||$, the second quantity measured with the given matrix norm.*

Our procedure for adjusting the singular values to enforce the linear equation $\sigma_2 = 1$ finds a sorted vector of singular values closest to the given computed singular values, whether we are using the matrix Frobenius norm (for which the induced vector norm is the 2-norm, and the closest vector is ours, uniquely) or the matrix spectral norm, for which the induced vector norm is the uniform norm, and our adjusted vector is closest although possibly not uniquely so.

4 Computation of the parameters

Let $H = R - xy^T$ be a given ROPR. How can we compute the motion parameters?

We begin with the computation of the singular value decomposition $H = U\Sigma V^T$. If the computed singular values fail to satisfy (6) we adjust them to make sure that it holds. This process is

1. If $\sigma_1 < 1$ set $\sigma_1 = 1$.
2. Set $\sigma_2 = 1$.
3. If $\sigma_3 > 1$ set $\sigma_3 = 1$.

This ensures that we are working with the singular values of a ROPR, in fact to the ROPR closest to $U\Sigma V^T$.

The solution procedure, naturally, checks the singular values, and considers three distinct cases. In each case, we show how the four scalars a, b, c, and d are computed. The vectors x and y are then obtained as $x = aU(:,1) + cU(:,3)$ and $y = bV(:,1) + dV(:,3)$, followed by rescaling to get the desired normalization of y. Given the four scalar parameters, we may determine Q from (7), and R as $R = UQV^T$.

When all three of the singular values of H are one, there are infinitely many solutions to the basic problem. In addition to indicating that this is the case, we can provide an exemplary solution by taking

$$a = b = 0;$$

and

$$c = d = 0 \quad \text{when} \quad \Delta = 1$$

while

$$c = -2, \quad d = 1 \quad \text{when} \quad \Delta = -1\ .$$

When $\sigma_1 > 1 = \sigma_3$ we compute

$$c = d = 0;$$

and

$$a = 1 - \sigma_1, \quad b = 1 \quad \text{when} \quad \Delta = 1$$

while

$$c = -1 - \sigma_1, \quad b = 1 \quad \text{when} \quad \Delta = -1\ .$$

When $\sigma_1 = 1 > \sigma_3$ we compute

$$a = b = 0;$$

and

$$c = 1 - \sigma_3, \quad d = 1 \quad \text{when} \quad \Delta = 1$$

while

$$c = -1 - \sigma_3, \quad d = 1 \quad \text{when} \quad \Delta = -1\ .$$

In the general case of $\mathrm{rank}(B) = 2$, Tsai offered the following formulas for computation of the elements of Q and the scalars a, b, c, d:

$$\begin{aligned} b &= \pm\left(\frac{\sigma_1^2 - 1}{1 - \sigma_3^2}\right)^{1/2} \\ C &= \frac{\sigma_1 + \Delta\sigma_3 b^2}{1 + b^2} \\ S &= \mp(1 - C^2)^{1/2} \\ a &= -S \\ c &= \sigma_3 - \Delta C \\ d &= 1\ . \end{aligned}$$

The two solutions are obtained with the two choices for the sign of b, the sign of S (and thus of a) is always taken to be the opposite of the sign of b. [1]

There are several things to observe about these formulas, from the numerical analysis viewpoint. First, square roots are involved, and these may be costly. Second and more important, the occurrences of the differences of the squares of computed quantities, such as $1 - C^2$ and $1 - \sigma_3^2$, is problematic. Suppose that x is known to nearly full working precision Suppose, furthermore, that x is close to one: $x = 1 - \delta$ and the difference δ is also quite small. When we square x we get $1 - 2\delta + \delta^2$, and if $|\delta|$ is smaller than the square root of machine precision, then this rounds to $1 - 2\delta$. This roundoff loses important information; the loss is revealed when we compute the difference $1 - x^2 = 2\delta$, which is now known to only half of machine precision. When we take the square root, we therefore get a result for which only half the digits are meaningful. That these things lead to difficulties will be shown by an example below. The problem at hand is in fact one of solving a certain quadratic equation. [2]

The Tsai formulas can produce poor results in floating point for matrices H that are close to orthogonal. Our criterion is backward error; in other words, we want the computed R, x, and y produce small residuals: $||H - (R - xy^T)||$ and $||I - R^T R||$.

All the computations reported here were done in Matlab on a Pentium PC. On that machine, the machine precision is

[1] Note that Tsai work with an arbitrarily scaled H, so that they do not assume that $\sigma_2 = 1$. Their formulas use σ_2 explicitly, but what they compute is unchanged when H and hence its singular values are divided by the original σ_2. So, in effect, we have given the Tsai formulas in their application to a matrix that has been rescaled to make it a ROPR.

[2] The difficulties arising out of the use of textbook formulas in general and for quadratic solvers in particular were described by Forsythe in the 1960s [1, 2].

```
>> eps
    2.220446049250313e-016
```

This is the smallest floating-point number ϵ for which, in floating-point arithmetic, $1 + \epsilon > 1$. It is an upper bound on the relative error due to roundoff of all the individual floating-point operations.

We start with a ROPR whose singular values are all close to one.

```
H =

    8.704900920846258e-001   -1.934310566425376e-001   -4.525830596792723e-001
    2.129569923832601e-001    9.770290004164705e-001   -7.978204075402978e-003
    4.437300068482838e-001   -8.943577959264387e-002    8.916865605979933e-001

(svd(H) - 1) / eps

    5.000000000000000e+000
                         0
   -1.500000000000000e+000
```

so $\sigma_1 = 1 + 5\epsilon$ and $\sigma_3 = 1 - 3\epsilon/2$.

Using the Tsai formulas we compute

```
R =

    8.704900825471498e-001   -1.934310566425377e-001   -4.525830780234804e-001
    2.129569922151317e-001    9.770290004164706e-001   -7.978208563136184e-003
    4.437300256391724e-001   -8.943577959264365e-002    8.916865512470823e-001


x =

    8.704900944689936e-001
    2.129569924252918e-001
    4.437300021505610e-001


y =

   -3.847463276194947e-008
                         0
   -2.107342425544702e-008
```

There is no problem with orthogonality:

```
>> I - R'*R
    1.110223024625157e-016                         0                         0
                         0   -2.220446049250313e-016   -2.220446049250313e-016
                         0   -2.220446049250313e-016   -2.220446049250313e-016
```

but there is a considerable problem with the residual:

```
>> H - (R-x*y'):

   -2.395431064616815e-008    1.665334536937735e-016    9.436895709313831e-016
   -8.025313719128846e-009   -1.110223024625157e-016    1.561251128379126e-016
   -3.586323743531850e-008   -2.220446049250313e-016    4.440892098500626e-016
```

We have evidently lost half of the machine precision in the first column of the residual. The difficulty is with the formula for S. The correct S is of size $O(\epsilon)$, but the roundoff in forming $S^2 = (1 - C^2)$ makes this quantity, which in exact arithmetic is $O(\epsilon^2)$, of size $O(\epsilon)$, whence we compute an approximate S of size $O(\sqrt{\epsilon})$.

How can we correct the problem? We shall give alternative computational formulas, that avoid the roundoff error issues. The key is to work through the problematic roundoff sensitive places in exact arithmetic (algebraic simplification of the formulas). By algebraic simplifications, we avoid taking the square root of the difference of squares of computed quantities that can be close to one another. Instead we compute

$$\begin{aligned} b^2 &= \tfrac{\sigma_1-1}{1-\sigma_3}\tfrac{\sigma_1+1}{\sigma_3+1} \\ b &= \sqrt{b^2} \\ d &= 1\,. \end{aligned}$$

If $\Delta = 1$ then

$$c = \frac{(1-\sigma_3)(1+\sigma_3)}{\sigma_1+\sigma_3}$$

and

$$a = -cb\,,$$

while if $\Delta = -1$ then

$$c = \frac{(1-\sigma_3)(1+\sigma_3)}{\sigma_3-\sigma_1}$$

and

$$a = cb\,.$$

We complete the calculation of R thus:

$$C = \sigma_3 + c$$

and

$$R = UQV^T$$

where

$$Q = \begin{pmatrix} \Delta C & 0 & a \\ 0 & 1 & 0 \\ cb & 0 & C \end{pmatrix}.$$

Note that the off-diagonal elements of Q_2 are now computed as the product cb rather than from the relation $S = \sqrt{1-C^2}$, avoiding the concellation and loss of precision when $S = O(\epsilon)$. It is easy to see that in the difficult case in which both $\sigma_1 - 1$ and $1 - \sigma_3$ are $O(\epsilon)$, the computed $b = O(1)$ and $c = O(\epsilon)$, so that the computed $S = O(\epsilon)$ as it should be.

5 Computation of the parameters: Experiment

We tested the proposed procedures for a set of matrices generated so as to present some challenges to the software. We formed ROPR test matrices by generating a pair of random 3×3 orthogonal matrices U and V, chosen by creating random rotation matrices, and then negating the third column of V in half the cases so as to get some with positive and some with negative determinant. We specified the singular values, keeping $\sigma_2 = 1$ of course, and taking σ_1 and σ_3 either far from, or very close to (a small multiple of machine precision) or exactly equal to one. Although we begin with specified singular values, we don't give these to the software. Rather we form $H = U\Sigma V^T$ and present it. The subsequent computation of the SVD of this computed H will, due to roundoff in forming it and in computing its SVD, yield slightly perturbed singular values. We view this as an advantage, causing additional, and realistic, difficulties to be presented to our code.

In tests on over 100,000 such randomly generated matrices, our code never fails to produce accurate results. We check the determinant of R, check its orthogonality by measuring the largest element of $R^T R - I$, and check the largest element of the residual $H - (R - xy^T)$. We have never found a case in which the error exceeds 16ϵ where ϵ is the machine precision. As above, the tests were done in Matlab running on a Pentium PC.

The Tsai formulas were also tested in this way, and found to suffer from a number of other failures. They fail with a divide-by-zero if $\sigma_2 = \sigma_3$. They can attempt to take the square root of a tiny negative quantity due to roundoff. Most important is the failure illustrated above, in which the use of the computation of a sine from the relation $\sin = \sqrt{1 - \cos^2}$ leads to the loss of half the working precision in cases in which the correct rotation is very close to the identity.

6 Conclusion

We have described several noteworthy properties of the SVD of a rank-one perturbation of a rotation, which is a matrix of the form $R - xy^T$ with R a rotation. Using these facts, we were able to show how to find the rank-one perturbation of a rotation closest to a given matrix. And we were have given numerically robust formulas that allow the efficient and accurate computation of the parameters R, x, and y given $H = R - xy^T$.

References

[1] George E. Forsythe. How do you solve a quadratic equation. Technical Report CS 40, Stanford University, 1966.

[2] George E. Forsythe. What is a satisfactory quadratic equation solver? In Bruno Dejon and Peter Henrici, editors, *Constructive aspects of the fundamental theorem of algebra: Proceedings of a symposium conducted at the IBM Research Laboratory, Zürich-Rüschlikon, Switzerland, June 5–7, 1967*, pages 53–61. Wiley-Interscience, New York, 1969.

[3] Alston Householder. *The theory of matrices in numerical analysis.* Blaisdell, New York, 1964.

[4] Yves Nievergelt. Schmidt-Mirsky matrix approximation with linearly constrained singular values. *Linear Algebra and Its Applications*, pages 207–219, 1997.

[5] Roger Y. Tsai, Thomas S. Huang, and Wei-Li Zhu. Estimating three-dimensional motion parameters of a rigid planar patch, II: Singular value decomposition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, pages 525–534, 1982.