
Randall C. Smith*

SRI International
Menlo Park, California 94025

Peter Cheeseman

NASA Ames
Moffett Field, California 94025

On the Representation and Estimation of Spatial Uncertainty

Abstract

This paper describes a general method for estimating the nominal relationship and expected error (covariance) between coordinate frames representing the relative locations of objects. The frames may be known only indirectly through a series of spatial relationships, each with its associated error, arising from diverse causes, including positioning errors, measurement errors, or tolerances in part dimensions. This estimation method can be used to answer such questions as whether a camera attached to a robot is likely to have a particular reference object in its field of view. The calculated estimates agree well with those from an independent Monte Carlo simulation. The method makes it possible to decide in advance whether an uncertain relationship is known accurately enough for some task and, if not, how much of an improvement in locational knowledge a proposed sensor will provide. The method presented can be generalized to six degrees of freedom and provides a practical means of estimating the relationships (position and orientation) among objects, as well as estimating the uncertainty associated with the relationships.

1. Introduction

In many applications it is necessary to reason on the basis of inaccurate information about spatial relationships among objects. For example, a mobile robot needs to represent and reason about the approximate relationships between itself and other objects. In addi-

tion, the robot must be able to use sensor information to reduce locational uncertainty (in both position and orientation) to a degree sufficient for achieving particular tasks. This problem is complicated in practice because the location of one object relative to another may be known only indirectly through a sequence of relative frames of reference with uncertainty. In this paper, we present a method for explicitly representing and manipulating the uncertainty associated with these transformations. We also show how sensors can be used to reduce this uncertainty. This formalism makes it possible to estimate, *in advance*, the probability of a robot going through a door given the current uncertainty of the robot and door location, the probability of a camera having an object in its field of view, whether a particular sensor will have sufficient accuracy to accomplish a particular task, and so on.

Brooks (1985) argues that it is not appropriate for mobile robots to use a global reference frame. He feels that a set of local reference frames linked via uncertain transformations is better. We show how the uncertainty of a frame relative to another can be estimated and how the reduction in uncertainty due to sensing can be mapped into any frame, regardless of where the sensing was performed. Because of this flexibility, no particular frame is necessary as an absolute reference.

In the following sections, we show how to estimate the uncertainty in three degrees of freedom (x , y , θ) using a mobile robot as the example. As the robot moves from one place to another, its uncertainty about its location with respect to its initial location grows. We present a method for making quantitative estimates of the resulting error, provided the moves are discrete. If the robot uses sensors (e.g., acoustic range or vision sensors) to make *relative* measurements between it and other objects, this new information can be used to improve knowledge about the *global* locations of the robot and objects with respect to any coordinate

The work reported in this paper was supported by the National Science Foundation under Grant ECS-8200615 and the Air Force Office of Scientific Research under Contract F49620-84-K-0007.

* Currently at General Motors Research Laboratories, Warren, Michigan.

The International Journal of Robotics Research,
Vol. 5, No. 4, Winter 1986,
© 1986 Massachusetts Institute of Technology.

Fig. 1. A sequence of approximate transformations.


frame connecting them in a network of relationships (a relational map). Because the estimation and update formulas and associated procedures require only simple matrix computations, the proposed method is computationally simple and fast. The theory assumes that the contributing errors are "small," allowing a good approximation with a first-order model of errors, and that the sensor errors are independent of the locational error. Although the example presented is for a mobile robot, the formulas for estimating locational uncertainty are applicable to many other domains. In fact, this work was originally motivated by the need to represent and reason about uncertainty in off-line programming applications for industrial manipulators.

2. The Basic Estimation Operations

An uncertain, or *approximate transformation* (AT), consists of an estimated mean relation of one coordinate frame relative to another and a covariance matrix that expresses the uncertainty of the estimate. ATs typically arise from relative motions or from sensing operations. The term AT A describes the uncertain relative location (A) of an object with respect to some defining coordinate frame.

In Figs. 1–4, ATs are represented symbolically by an arrow going from the reference frame to the relative location. An ellipse, centered about the nominal estimate, depicts the uncertainty in the estimated relationship. The ellipse is the contour of constant probability (in two dimensions) for a multivariate Gaussian distribution, and can be used to bound a high-confidence region in which the actual location should be found. Though the procedures described in this paper for estimating the mean and covariance of ATs do not assume Gaussian distributions of the errors, some distribution must be assumed when the explicit calculation of probabilities is required. Section 6.3 describes the rationale for assuming an underlying Gaussian distribution in such cases, and Appendix A describes how to extract the ellipse parameters from the estimated covariance matrix of an AT.

In this paper, we introduce two basic operations that allow the estimation of the relationship between any two coordinate frames, given the uncertain relative transformations linking them. The first operation,



called *compounding*, allows a chain of ATs to be collapsed (recursively) into a single AT. The final compounded transformation has greater uncertainty than its components. The second operation, called *merging*, combines information from parallel ATs to produce a single resultant AT with uncertainty less than any of its components.

2.1. COMPOUNDING OF APPROXIMATE TRANSFORMATIONS

A simple robot example of compounding is similar to that discussed in Brooks (1985) (see Fig. 1).

In this example, the robot makes a number of moves and ends up near its initial position W. (All coordinate frames besides the initial one are omitted in Fig. 1 for clarity.) The uncertainty of the robot's final location with respect to W is large, as indicated by the large error ellipse. The solid error ellipses express the relative uncertainty of the robot with respect to its last position, while the dashed ellipses express the uncertainty of the robot with respect to W. Note that the uncertainty of the robot's location with respect to the world frame grows with each move. In Section 3, we show how to calculate compounded ATs, such as E and G.

2.2. MERGING OF APPROXIMATE TRANSFORMATIONS

A different situation arises when there is more than one independent AT relating the frame of interest to a given reference frame. For example, if the robot in Fig. 1 observes its starting point (the frame **W**) from **L₄**, then we have two ATs from **W** to **L₄**: (1) AT **G** calculated by compounding, and (2) the inverse of AT **S** given by the sensor. These two transformations are combined by weighted averaging (using the Kalman filter equations described in Section 4) to give a more accurate estimate of the location of **L₄** with respect to **W**. The error (covariance) of AT **S** comes from the intrinsic error of the sensor, while the covariance of AT **G** comes from the compounded covariances of the individual robot motions **A**, **B**, **C**, **D**.

The method for calculating compound ATs is given in Section 3, and the method for merging ATs is given in Section 4.

3. Compound Approximate Transformations

In this section, we describe the compounding procedure for calculating the nominal location and associated error (expressed as a covariance matrix) of any object relative to any other object linked through a chain of approximate transformations (e.g., Fig. 1). This approach differs from that presented by Brooks (1982, 1985), who used a max/min representation of the error. The max/min approach assumes the worst case when errors are compounded and so can badly overestimate the error when the results are propagated through several transformations. Also, the interpretation of the max/min bounds are unclear. Do the bounds mean that it is impossible for an observation to fall outside them, or that it is just very unlikely to do so (and if so, how unlikely)? The approach described by Chatila and Laumond (1985) for the HILARE robot is more similar to the one described here; however, they used a scalar error estimate of position and were not concerned with angular error.

3.1. FORMULAS FOR COMPOUNDING

In Fig. 1, we wish to describe the coordinates of **L₂** (X_3, Y_3, θ_3) with respect to reference frame **W**. We are

compounding transformations **A** (X_1, Y_1, θ_1) and **B** (X_2, Y_2, θ_2). The explicit transformation is given in Eq. (1) and is derived from the formulas for transforming one frame into another, as shown in Paul (1981). For example.

$$\begin{aligned} X_3 &= f(X_1, Y_1, \theta_1, X_2, Y_2, \theta_2) \\ &= X_2 \cos \theta_1 - Y_2 \sin \theta_1 + X_1, \\ Y_3 &= g(X_1, Y_1, \theta_1, X_2, Y_2, \theta_2) \\ &= X_2 \sin \theta_1 + Y_2 \cos \theta_1 + Y_1, \\ \theta_3 &= h(X_1, Y_1, \theta_1, X_2, Y_2, \theta_2) = \theta_1 + \theta_2. \end{aligned} \quad (1)$$

We wish to estimate the means and covariances of these three functions. The variables are now assumed to be *random variables*. The functions are approximated by a first-order Taylor series expansion about the means of the variables. The mean values of the functions (to first order) are the functions applied to the variable means: e.g., $\hat{X}_3 \cong f(\hat{X}_1, \hat{Y}_1, \hat{\theta}_1, \hat{X}_2, \hat{Y}_2, \hat{\theta}_2)$. In addition to estimating the mean transformation, an AT includes the associated covariance matrix of this transformation. To estimate the covariance matrix for this case, we express the previous Taylor series expansion in matrix form, resulting in the following (deviate) matrix:

$$\begin{pmatrix} \Delta X_3 \\ \Delta Y_3 \\ \Delta \theta_3 \end{pmatrix} \cong \mathbf{J}(\Delta X_1, \Delta Y_1, \Delta \theta_1, \Delta X_2, \Delta Y_2, \Delta \theta_2)^T, \quad (2)$$

where **J** is the (3×6) Jacobian of the transformation evaluated at the mean values of the variables:

$$\begin{aligned} \mathbf{J} &= \begin{pmatrix} \frac{\partial f}{\partial X_1} & \frac{\partial f}{\partial Y_1} & \frac{\partial f}{\partial \theta_1} & \frac{\partial f}{\partial X_2} & \frac{\partial f}{\partial Y_2} & \frac{\partial f}{\partial \theta_2} \\ \frac{\partial g}{\partial X_1} & \frac{\partial g}{\partial Y_1} & \frac{\partial g}{\partial \theta_1} & \frac{\partial g}{\partial X_2} & \frac{\partial g}{\partial Y_2} & \frac{\partial g}{\partial \theta_2} \\ \frac{\partial h}{\partial X_1} & \frac{\partial h}{\partial Y_1} & \frac{\partial h}{\partial \theta_1} & \frac{\partial h}{\partial X_2} & \frac{\partial h}{\partial Y_2} & \frac{\partial h}{\partial \theta_2} \end{pmatrix} \\ &= \begin{pmatrix} 1 & 0 & -(Y_3 - Y_1) & \cos \theta_1 & -\sin \theta_1 & 0 \\ 0 & 1 & (X_3 - X_1) & \sin \theta_1 & \cos \theta_1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix} \\ &= [\mathbf{H}|\mathbf{K}] \end{aligned} \quad (3)$$

and X_3, Y_3 are given in Eq. (1). *Covariance* is defined as the expectation of the squared deviates. We “square” both sides of Eq. (2) by multiplying both sides of the equation by their respective transposes. Taking the expectation of the result, we get the covariance matrix C_3 :

$$C_3 \cong \mathbf{J} \begin{pmatrix} C_1 & \mathbf{0} \\ \mathbf{0} & C_2 \end{pmatrix} \mathbf{J}^T = \mathbf{H}C_1\mathbf{H}^T + \mathbf{K}C_2\mathbf{K}^T. \quad (4)$$

The 3×3 matrix C_3 expresses the covariances of the coordinates of L_2 with respect to \mathbf{W} (i.e., the AT \mathbf{E}), computed from the given covariance matrices C_1 and C_2 (which express the error on the variables of \mathbf{A} and \mathbf{B}). Because an element of the covariance matrix is by definition $C_{ij} = E(\Delta x_i \Delta x_j)$ and the standard deviation of variable x is $\sigma_x = \sqrt{E(\Delta x^2)}$, an element of the covariance matrix can be expressed as:

$$C_{ij} = \rho_{ij} \sigma_i \sigma_j,$$

$$\rho_{ij} = \frac{E(\Delta x_i \Delta x_j)}{\sqrt{E(\Delta x_i^2)E(\Delta x_j^2)},}$$

where ρ_{ij} is the correlation coefficient. On the diagonal of the covariance matrix, $i = j$; thus ρ_{ij} is 1, and C_{ii} is just the variance. Note that if there is no angular error associated with variables θ_1 and θ_2 , then we can reduce the dimensions of \mathbf{H} and \mathbf{K} and the corresponding covariance matrices. This is achieved by removing rows or columns associated with θ_1 and θ_2 . When no angular errors are involved, the compounding equations are linear in the random variables. The results in this case are the theoretically correct values for the transformed first and second moments and are *not* approximations.

Equations (1) and (4) are needed to estimate the compound AT for more than two component ATs. The method consists of computing the compound of two adjacent ATs and replacing the pair by this result. The AT between any pair of connected frames can be computed by a number of such reductions. That is, the result of a series of relative transformations can be computed by compounding the first two transformations to form a new composite transformation, then compounding this new transformation with the third transformation to form a new composite, and so on. In

Fig. 1, for example, \mathbf{A} and \mathbf{B} are compounded to give AT \mathbf{E} , then \mathbf{E} and \mathbf{C} are compounded to give \mathbf{F} , and finally \mathbf{F} and \mathbf{D} are compounded to give \mathbf{G} . By a simple chain-rule argument, it can be shown that equivalent first-order estimates of \mathbf{G} are obtained by using the recursive approach above, or by defining the final relationship in terms of all the variables at once, and finding the first-order estimate of the mean and covariance. Clearly, the recursive approach is simpler.

A given directed graph of uncertain relationships may not be in a form needed by the above formulas. For example, in Fig. 2A, \mathbf{B} is pointing in the opposite direction to the form assumed above. To create the correct form, it is necessary to first reverse \mathbf{B} , as in Fig. 2B. This reversal is easily accomplished using the following formulas (the inverse transformation):

$$\begin{aligned} X' &= -X \cos \theta - Y \sin \theta, \\ Y' &= X \sin \theta - Y \cos \theta, \\ \theta' &= -\theta. \end{aligned} \quad (5)$$


As before, the first-order estimate of the mean values of the dashed coordinates is simply the given functions of the mean variables. The covariance matrix for the reversed transformation is estimated from the given covariance matrix of \mathbf{B} , (e.g., the matrix \mathbf{C}):

$$\mathbf{C}' \cong \mathbf{R} * \mathbf{C} * \mathbf{R}^T, \quad (6)$$

where \mathbf{R} is the Jacobian of the above transformation equations; that is, \mathbf{R} is given by:

$$\begin{aligned} \mathbf{R} &= \begin{pmatrix} \frac{\partial X'}{\partial X} & \frac{\partial X'}{\partial Y} & \frac{\partial X'}{\partial \theta} \\ \frac{\partial Y'}{\partial X} & \frac{\partial Y'}{\partial Y} & \frac{\partial Y'}{\partial \theta} \\ \frac{\partial \theta'}{\partial X} & \frac{\partial \theta'}{\partial Y} & \frac{\partial \theta'}{\partial \theta} \end{pmatrix} \\ &= \begin{pmatrix} -\cos \theta & -\sin \theta & Y' \\ \sin \theta & -\cos \theta & -X' \\ 0 & 0 & -1 \end{pmatrix} \end{aligned} \quad (7)$$

Fig. 2. Examples of compounding situations.



and \mathbf{R} is evaluated at \hat{X} , \hat{Y} , and $\hat{\theta}$. Again, if there is no angular error, we can omit the row and column of \mathbf{R} associated with θ , and the resulting covariance calculation will be exact.

With this inverse operation and the pairwise compounding operation given above, it is possible to estimate the compound AT between any two frames of reference that are linked by a chain of ATs, as in Fig. 1.

3.2. THE ASSUMPTIONS FOR COMPOUNDING

The assumptions behind this compounding operation are:

The first order approximation expressed by Eq. (2) is sufficiently accurate.

The errors of ATs to be compounded are independent.

The errors of ATs to be compounded have zero mean.

The first-order approximation is reasonable provided that the standard deviations of the variables (e.g., σ_{x_1}) are small, because we are neglecting terms of $(\sigma_x^2)^2$ and higher orders in the Taylor series expansions. More accurately, the function should be "smooth" about the estimated point over an interval roughly the magnitude of a standard deviation of the variable. This approximation can underestimate or overestimate the error covariance, and the accuracy of the estimate depends on the significance of the second-order (and higher) terms.

The assumption that successive ATs have independent error implies the use of a compound covariance matrix with off-diagonal submatrices which are zero, as in Eq. (4). If these errors are known to be dependent, a compound AT with nonzero covariance submatrices can be given directly; that is, the dependent

case can be handled by treating the dependent pair of ATs as a unit, with a given compound covariance matrix containing nonzero submatrices.

Systematic, or nonzero mean errors are not modeled. It is assumed that systematic errors will be removed by suitable calibration.

4. Merging of Approximate Transformations

The second basic operation is the merging of two or more parallel ATs to obtain a more accurate estimate of the relationship. For example, in Fig. 1 we can estimate \mathbf{G} , which expresses the uncertain relationship of L_4 with respect to the frame \mathbf{W} , and we are given AT \mathbf{S} describing \mathbf{W} with respect to L_4 . The problem we are considering is how to estimate the relation of L_4 with respect to \mathbf{W} given both parallel transformations. The procedure we use for merging is based on the use of the Kalman filter equations (for static-state estimation).

4.1. FORMULAS FOR MERGING

The first step is to find all the ATs linking the frames of interest. For each independent chain of transformations between these frames, estimate the compound AT. For the example in Fig. 1, \mathbf{G} must be computed using the methods described in Section 3 to estimate the AT of L_4 with respect to \mathbf{W} . The next step is to ensure that all the parallel ATs to be merged are pointing in the desired direction, and to reverse those that are not. In Fig. 1, \mathbf{S} should be reversed. The reversal of an AT is given by Eqs. (5) and (6).

Once all the parallel ATs to be merged have been estimated (including any necessary reversals), the merging procedure combines them in pairs, using the result of the last merge as an input to the next merge. The merging begins with any pair and proceeds by merging each additional AT with the preceding result until all ATs have been merged. Consequently, it is only necessary to describe the pairwise merge procedure. In the following, let \mathbf{C}_1 and \mathbf{C}_2 be the covariance matrices of the ATs to be merged, and \mathbf{C}_3 be the covariance matrix of the resulting merged pair. Similarly, \hat{X}_1 and \hat{X}_2 are the estimated mean values (expressed as

a vector) of the transformations to be merged, and \hat{X}_3 is the resulting estimated mean. The first step in the merging operation is to compute the Kalman gain factor defined by:

$$\mathbf{K} = \mathbf{C}_1 * [\mathbf{C}_1 + \mathbf{C}_2]^{-1}. \quad (8)$$

This factor is used to compute the required merged covariance matrix

$$\mathbf{C}_3 = \mathbf{C}_1 - \mathbf{K} * \mathbf{C}_1 \quad (9)$$

and the merged mean

$$\hat{X}_3 = \hat{X}_1 + \mathbf{K} * (\hat{X}_2 - \hat{X}_1). \quad (10)$$

These formulas are sufficient for merging any number of parallel ATs. The apparent asymmetry of Eqs. (8), (9), and (10) is not real, because the same result is obtained regardless of which transformation is labeled \mathbf{C}_1 or \mathbf{C}_2 , and so on. The order in which the mergings are performed is irrelevant. In one dimension, these formulas reduce to the following simple forms (where V s are variances):

$$\begin{aligned} V_3 &= \frac{V_1 V_2}{V_1 + V_2}, \\ \hat{X}_3 &= \frac{V_2}{V_1 + V_2} \hat{X}_1 + \frac{V_1}{V_1 + V_2} \hat{X}_2. \end{aligned} \quad (11)$$

The covariance matrices reduce to simple variances combined as shown, and the merged mean is just a weighted average of the contributing means. The scalar formula for the weighted average was also used by Chatila and Laumond (1985). If the compounding operation described in Section 3 is also reduced to one dimension, we find that variances in series transformations simply add to give the compound variance. If this result and Eq. (11) are combined, we find a similarity to electric circuit theory where variances are analogous to resistances. That is, in a network of resistors, series resistances are combined by addition, and parallel resistances combine according to Eq. (11) to give a combined resistance (variance) between any two points. In Section 5, we will exploit the similarity between the ways uncertainties and resistances are combined.

4.2. THE ASSUMPTIONS FOR MERGING

Kalman filter theory, or state estimation, is described in many texts on estimation theory; the derivation of the basic formulas is lengthy and is not presented here (see Nahi 1976). The basic assumptions of the theory in this context are:

The errors in ATs to be merged are independent.

Extensions can be made to handle merging of nonindependent ATs, when the dependence is given.

The formula for merging two parallel ATs is appropriately defined as a weighted, linear combination of the two estimates. If both ATs to be merged have Gaussian distributions, it is known that the estimator that combines them and that minimizes the mean square estimation error has a linear form (Nahi 1976). Furthermore, we only estimate the first two moments of each AT's distribution (for practical reasons). When the mean and variance of an unknown probability distribution are the only information available, a simple maximum entropy derivation gives the Gaussian distribution as the distribution that assumes the least information. Finally, the linear form is commonly used in estimators and has been found in practice to provide good results in a number of applications, including navigation (Gelb 1984).

The errors are unbiased; that is, the mean value of the errors is zero. This assumption excludes systematic errors, which are usually eliminated by suitable calibration.

The ATs to be merged are expressed in the same coordinate system (e.g., Cartesian). If this is not true, suitable mappings must first be performed. We have described a simplified form of the full Kalman update equations applied to merging, which are suitable assuming the mapping (if necessary) is linear. If the mapping is nonlinear the extended Kalman filter equation for updating the estimate should be used (see Gelb 1984, for example).

When these assumptions are satisfied, the merging formula computes an unbiased estimate for the updated mean of the relationship and the covariance of the error. The estimate is optimal for Gaussian vari-


Fig. 3. A network of approximate transformations.

ables and linear mappings between coordinate systems and optimal-linear for non-Gaussian variables and linear mappings. If the mappings are nonlinear, merging (by the extended Kalman filter update equations) is a suboptimal, nonlinear estimator that has been found to work quite well in practice in many problems, including navigation. Note that the equations do not require Gaussian variables, but because we retain only estimates of the first two moments, the resulting distribution must be assumed to be Gaussian in order to compute a probability. Fortunately, in many cases the result of combining a number of different densities quickly tends toward an approximately Gaussian distribution (by the central limit theorem).

5. Combined Estimation Procedure

The previous sections define two basic operations for combining ATs: compounding (denoted by \oplus), which estimates the combined AT for two serially linked ATs; and merging (denoted by \otimes), which combines two parallel ATs. Both \oplus and \otimes are associative operations, but only \otimes is commutative. These two operations can be used to estimate the AT between any two locations in a network of ATs. The procedure is as follows:

1. Mark the two locations (nodes) whose AT is to be estimated.
2. Find any chain of three nodes in the current network of ATs that form a sequence of ATs (as in Fig. 2). Its middle node should not be a marked node and it should not be connected to any other node. Apply the compounding operation to the two ATs that occur in any such chain. A compounding operation reduces the network by replacing the two ATs by a single equivalent AT.
3. Find any pair of nodes that are linked by two parallel ATs and apply the merging operation to them. Such merging operations reduce the network by replacing each pair of parallel ATs by a single merged AT.
4. Repeat Steps 2 and 3 until the network is reduced to a single AT connecting the marked nodes (if possible).




A difficulty with this procedure is that Steps 2 and 3 may not reduce the network to a single AT. For example, consider the network in Fig. 3. The analogous ways in which resistances and variances are combined, previously mentioned in Section 4.1, can be used to illustrate the problem. If the ATs of the network were replaced by resistors, then this network becomes a Wheatstone bridge circuit. It is well known that some electric circuits, such as the Wheatstone bridge, cannot be reduced to a single equivalent resistance between nodes w and u by simple combinations of series and parallel resistance calculations. Likewise, the merging and compounding operations cannot be used to reduce the network of Fig. 3 to an equivalent uncertain relation between that pair of nodes.

The procedure defined in this section *can* be used to estimate the AT from t to v by the formula: $(A^R \oplus B) \otimes C \otimes (D \oplus E)$, where A^R denotes the reversal of A (section 3.1). However, there is no combination of compounding and merging operations that reduces the network to a single AT from w to u .

In such cases, an approximate estimate can be found by deleting loop-forming ATs. In this example, deleting C (or B and D) will produce a reducible network. The best deletion to choose is the one that minimizes the desired components of the final covariance matrix. Clearly, this deletion procedure is nonoptimal in that it does not use all the available information.

The analogy to electric circuit theory suggests a method for using all the information, in cases such as Fig. 3, to derive a better estimate. The method is to replace a triangle (Delta) of ATs (such as A , B , C) by an equivalent "Y" of ATs, as shown in Fig. 4 (known as a Delta-Y transformation in circuit theory).

Fig. 4. Transformation of Wheatstone bridge example.



The problem is to calculate the ATs for the equivalent “Y” (X, Y, Z), given the Delta composed of ATs (A, B, C). First, compute the best estimate for each relation describing a side of the Delta (ignoring D and E). The new estimates are denoted by (A', B', C') in Fig. 4. They are $A' = A \otimes (B \oplus C^R)$, $B' = B \otimes (A \oplus C)$, and $C' = C \otimes (A^R \oplus B)$. Arbitrarily assign the mean value of X to be the identity transformation. With this mean value for X , the mean values for Y and Z are the same as for B' and A' respectively. Next, calculate the covariance matrices for the ATs of the equivalent “Y” using the following formulas:

$$\begin{aligned} C_X &= \frac{1}{2}(C_{A'} + C_{B'} - \mathbf{K}C_C\mathbf{K}^T) \\ &= C_A[C_A + C_B + \mathbf{K}C_C\mathbf{K}^T]^{-1}C_B, \\ C_Y &= \frac{1}{2}(-C_{A'} + C_{B'} + \mathbf{K}C_C\mathbf{K}^T) \\ &= C_B[C_A + C_B + \mathbf{K}C_C\mathbf{K}^T]^{-1}\mathbf{K}C_C\mathbf{K}^T, \\ C_Z &= \frac{1}{2}(C_{A'} - C_{B'} + \mathbf{K}C_C\mathbf{K}^T) \\ &= C_A[C_A + C_B + \mathbf{K}C_C\mathbf{K}^T]^{-1}\mathbf{K}C_C\mathbf{K}^T, \end{aligned}$$

where \mathbf{K} is defined in Eq. (3).

If the system is linear in the random variables (i.e., there is no significant angular error), then the mean locations found in calculating A', B', C' are consistent—that is, any one of the mean relations can be computed from the other two. For example, the mean of $A' \oplus C'$ is equal to the mean of B' . This will not be exactly true if there are significant nonlinearities in the compounding and reversal operations. However, it is possible to develop an iterative method for finding a consistent set of means. This only becomes necessary if the angular errors are very large (e.g., with a standard deviation greater than 5 degrees).

Once the network has been transformed, as in Fig. 4, the formula $X \oplus ((Y \oplus E^R) \otimes (Z \oplus D))$ will give the desired AT between w and u . Unfortunately, the Delta-Y transformation method outlined above cannot be used to reduce any network to a single AT. A general method for estimating the AT between any pair of nodes given an arbitrary network of ATs and other constraints is being investigated. This method, based on recursive state estimation, defines all ATs with respect to a common reference frame. However, the AT between any pair of frames can be extracted from this representation.

6. A Mobile Robot Example

A mobile robot needs to use its sensors to build and update a world map to navigate in both known and unknown environments. At any stage, this map can be viewed as a network of uncertain relations that can be used to decide important navigational questions. Initially, the robot may take its starting position as the “world” frame. As it moves from its initial position, the uncertainty of its location with respect to the world grows with each successive move, as shown in Fig. 1. Each relative move and its error determined from the robot model are represented by an uncertain transformation. After a number of such moves, the robot’s location with respect to the world frame becomes so uncertain that the robot is unlikely to succeed in actions (e.g., going through a doorway) based purely on its current information. The procedures described above allow the robot to estimate the uncertainty of its location relative to any other coordinate frame and decide whether the relationship is known accurately enough to perform a particular task. Note that because this estimate can be made ahead of time, the robot can decide that proposed motions will create too much uncertainty *before* they are performed, or that sensing will be necessary to reduce the uncertainty.

6.1. ROBOT SENSING

A mobile robot is usually equipped with sensors that allow it to determine the location of objects to an accuracy determined by the sensor resolution. Such

sensed relationships are also represented as ATs in the AT network, along with those due to motion. The sensed information allows the robot to reduce the uncertainty of its location with respect to other objects or the world. Because the reduction of uncertainty through sensing can be calculated ahead of time, it is possible to decide if proposed sensing steps will reduce uncertainties enough for the accomplishment of a task *before performing the sensing*. If a proposed sensing step is not sufficiently accurate, alternative sensing strategies can be evaluated. This means that the AT estimation procedure can be used for off-line programming of a robot, as well as for real-time applications, such as mobile robot navigation.

It is important that the robot sense as accurately as possible the location of many reference objects while it is still in its initial position (“world”). Once the robot moves, these objects allow the robot to estimate its position with respect to the world with the greatest possible accuracy. Even if the robot is unable to sense the original reference objects, accurate world locations can still be found if the original reference objects are related to other observable objects through accurately sensed relationships. That is, a richly connected network of ATs allows the robot to locate itself with respect to any object (in the network) with an accuracy largely determined by the accuracy of the sensors. Note that the method advocated here is for the robot to maintain a network of the original ATs resulting from individual motions or sensing operations and to estimate composite ATs as required.

6.2. THE SENSING PROCEDURE

Given that the robot has decided to perform a sensor step, the procedure it should adopt in adding a sensor-derived AT to the current network is as follows:

1. *Determine whether sensing is possible.* After estimating the AT between the object and the sensor by the procedure described above, decide whether the sensor is capable of making the desired observations. For example, decide whether a wall is likely to be close enough to be detected with an acoustic range sensor, or whether the object to be sensed is in the field of view of a camera (see Appendix A).


2. *Make the actual observation and decide if it is reasonable.* Given the uncertain prior location of the object to be sensed (estimated from the current ATs) and the error associated with the sensor, decide whether the probability of the actual observation is below a predetermined threshold. Such errors would occur, for example, if the camera viewed the wrong object.
3. *Combine multiple sensings.* Merge independent measurements by the sensor into a single AT using the merging Eqs. (8), (9), and (10). If the sensing information is in a different coordinate system (e.g., polar rather than Cartesian), then the merging operation is performed in the sensor system and the final AT is mapped back into the Cartesian form (the specific formulas for performing this mapping in the case of polar to Cartesian are given in Appendix B).

6.3. MONTE CARLO SIMULATIONS

To test the assumptions in the theory, we ran an independent Monte Carlo simulation, where the robot's position is calculated many times using Gaussian distributions for the errors in the given relations. Figure 5 shows the resulting robot positions as points with the estimated 90% confidence contour superimposed on them. This contour is formed by assuming that the resulting probability distribution is Gaussian. The relative error in any component of the estimated means and covariances (compared to the simulated values) is typically less than 1%, unless the angular errors are large (e.g., with a standard deviation of greater than 6 degrees).

The assumption that the resulting probability distribution is Gaussian may not be correct, however. For example, the crescent-shaped distribution of points in Fig. 5 does not match the superimposed error ellipse. Note that the mean and covariance estimates may still be accurate—it is the Gaussian assumption that is breaking down. The divergence occurs because of nonlinearities in the Jacobian (i.e., its dependence on the trigonometric functions of θ) when the given angular errors are large. Even in this case, the final probability distribution formed from two crescent distributions is very close to Gaussian, as seen in Fig. 5. If the

ig. 5. A Monte Carlo simulation.



probability distributions of the variables are not Gaussian, the central limit theorem assures us that the final compound distribution will be approximately Gaussian provided that there are a “large” number of ATs.

7. Discussion

The previous sections describe a procedure for making accurate quantitative estimates of the mean and covariance of the location (position and orientation) of any frame relative to any other frame given a network of ATs. If the distribution associated with these estimates is assumed to be Gaussian, then probabilities can be estimated. Examples of the application of this theory are presented and the estimates are compared with Monte Carlo simulations for a three-degree-of-freedom

mobile robot. The only significant divergence between the estimates and the simulation occurs when the angular error is large compared with a subsequent displacement error. Under these circumstances, the resultant probability distribution is significantly non-Gaussian (i.e., noticeably nonelliptical in Cartesian space). In all other cases, the mean and variance estimates agreed extremely well with the simulated values. Even when the angular errors are large, the effect of the nonlinearities on the resulting distribution is reduced when several ATs are combined, as seen in Fig. 5.

The formulas given for three degrees of freedom (X, Y, θ) can be easily extended to include the Z coordinate in the obvious way. Unfortunately, the extension to the full six-degrees-of-freedom case $(X, Y, Z, \theta, \phi, \psi)$ is not so simple. The main problem is that the additional angular terms can introduce singularities

(“poles”) into the Jacobian. The estimated covariance shows extreme sensitivity to the error of particular input angles in the vicinity of a pole. These singularities destroy the value of the estimates provided by the estimation procedure unless corrective measures are taken. One approach we are investigating is to pre-rotate the world frame so that none of the ATs have nominal angles in the vicinity of a pole.

There are many uses for the estimation procedures described in this paper. The major use motivating this work is the ability to estimate *ahead of time* when sensing steps are necessary and to determine whether a particular sensor is capable of making a particular measurement or of supplying the required accuracy. The procedure also supplies a quantitative method for judging when a sensor has “glitched” (i.e., its measured value is too unlikely, given the prior expectations). Another possibility being investigated is whether the results of running this procedure can be used to revise the error models used for particular mobile robots or sensors, giving a form of autocalibration.

The limitation of the estimation procedure described in this paper is that only two moments (the mean and covariance) are estimated, hence the knowledge of the underlying probability distribution is limited. By assuming the (multivariate) distribution is Gaussian, we can still make probabilistic estimates that agree with simulations very well, except when the angular errors are large. A way around this difficulty that we are investigating is to estimate a third moment (a “skewness” measure).

Appendix A. Derivation of Ellipse Parameters

The determination of whether the probability of observing the object is greater than a given threshold assumes that the probability distribution of our knowledge of the object’s location is a multivariate (x, y, θ) Gaussian distribution. The general form is given by:

$$P(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^n \det \mathbf{C}}} e^{-\frac{1}{2}(\mathbf{x}-\hat{\mathbf{x}})^T \mathbf{C}^{-1}(\mathbf{x}-\hat{\mathbf{x}})},$$

where n is the number of dimensions, \mathbf{C} is the covariance matrix, $\hat{\mathbf{x}}$ is the nominal mean vector, and \mathbf{x} is a

vector denoting a particular point. The contours of equal probability of this distribution form ellipsoids in n dimensional space that are centered at the mean location $\hat{\mathbf{x}}$, and whose axes are only aligned with the Cartesian frame if the covariance matrix \mathbf{C} is diagonal. The formulas for extracting the principal axes of a two-dimensional ellipsoid are given below. In the case where we are only interested in the positional error ellipse, \mathbf{C} is the reduced (2×2) covariance matrix formed from the (3×3) matrix by extracting only the X, Y terms. In this case, the resulting marginal probability distribution is:

$$P(x, y) = \frac{1}{2\pi \sqrt{\sigma_x^2 \sigma_y^2 (1 - \rho^2)}} e^{-\frac{1}{2(1-\rho^2)} \left[\frac{x^2}{\sigma_x^2} + \frac{2\rho xy}{\sigma_x \sigma_y} + \frac{y^2}{\sigma_y^2} \right]},$$

$$\mathbf{C} = \begin{pmatrix} \sigma_x^2 & \rho \sigma_x \sigma_y \\ \rho \sigma_x \sigma_y & \sigma_y^2 \end{pmatrix},$$

where ρ is the correlation coefficient for x and y . For example, we might want to check the probability of a camera seeing a particular object given our current locational uncertainty. We decide this question by determining whether the ellipse corresponding to a given confidence limit is completely contained within the field of view.

For decision making purposes (e.g. the field-of-view case), it is necessary to determine the explicit equiprobable contours (ellipses or ellipsoids) of the multivariate Gaussian distribution specified by given mean $\hat{\mathbf{x}}$ vector and covariance \mathbf{C}_x matrix. These ellipses can be used to determine the probability that a given vector will lie within, say, the 90% confidence ellipse. The ellipsoid formula is:

$$(\mathbf{x} - \hat{\mathbf{x}})^T \mathbf{C}_x^{-1} (\mathbf{x} - \hat{\mathbf{x}}) = k^2, \quad (\text{A1})$$

where k is a constant chosen for a particular confidence threshold, and \mathbf{x} is a point on the ellipsoid boundary. The relationship between k and the probability of a point lying within the ellipsoid specified by k is:

$$N = 1; \quad P = -\frac{1}{\sqrt{2\pi}} + 2 \operatorname{erf}(k)$$

$$N = 2; \quad P = 1 - e^{-\frac{k^2}{2}}$$

$$N = 3; \quad P = -\frac{1}{\sqrt{2\pi}} + 2 \operatorname{erf}(k) - \sqrt{\frac{2}{\pi}} k \exp^{-\frac{k^2}{2}},$$

where N is the number of dimensions and erf is the error function.

In the particular case where the two-dimensional (x, y) error ellipse is required given the 3-D (x, y, θ) covariance matrix, the procedure is as follows. Firstly, produce the two-dimensional marginal covariance matrix from the 3-D covariance matrix by striking out the row and column of the unwanted variable. That is:

$$\mathbf{C}_3(3 \times 3) \Rightarrow \mathbf{C}_2(2 \times 2).$$

The corresponding family of two-dimensional ellipses is given by Eq. (A1), and in this case reduces to

$$Ax^2 + 2Bxy + Cy^2 - k^2 = 0,$$

where A, B, C are found from the two-dimensional covariance matrix and Eq. (A1). The angle θ that the major axis of this ellipse makes with the positive x -axis is given by:

$$\theta = \frac{1}{2} \arctan \left(\frac{2B}{A-C} \right) \quad \theta \left[\frac{-\pi}{2}, \frac{\pi}{2} \right].$$

If we define

$$T = \sqrt{A^2 + C^2 - 2AC + 4B^2},$$

then we find the following lengths:

$$\begin{aligned} \text{half major axis} &= \sqrt{\frac{2k^2}{A+C-T}}, \\ \text{half minor axis} &= \sqrt{\frac{2k^2}{A+C+T}}. \end{aligned}$$

As given above, the probability of a point being located inside an ellipse defined by a particular value of k is given by:

$$P(x, y \in \text{ellipse}) = 1 - e^{-\frac{k^2}{2}}, \\ k^2 = -2 \log(1 - Pr),$$

with the following confidence ellipses for different k :

$$\begin{aligned} 50\% &\Rightarrow k^2 = 1.386, \\ 90\% &\Rightarrow k^2 = 4.605. \end{aligned}$$

Appendix B. Coordinate Frame Mappings

This appendix gives the explicit formulas for mapping the mean and covariance matrix from polar (r, ϕ, θ) representation to an equivalent Cartesian form, where θ is the rotation of the coordinate frame in both the polar and Cartesian case. This mapping is necessary, for example, if a camera expresses the location of an object in polar coordinates and the camera error is also given in polar coordinates, but the result of sensing (or the merged results of many sensings) is required in a Cartesian representation. The nominal (mean) transformation is given by

$$\begin{aligned} \hat{x} &\cong f_1(\hat{r}, \hat{\phi}) = \hat{r} \cos(\hat{\phi}), \\ \hat{y} &\cong f_2(\hat{r}, \hat{\phi}) = \hat{r} \sin(\hat{\phi}), \\ \hat{\theta} &= f_3(\hat{\theta}) = \hat{\theta}. \end{aligned}$$

\mathbf{R} , the Jacobian of the transformation, is given by:

$$\mathbf{R} = \begin{pmatrix} \frac{\partial f_1}{\partial r} & \frac{\partial f_1}{\partial \phi} & \frac{\partial f_1}{\partial \theta} \\ \frac{\partial f_2}{\partial r} & \frac{\partial f_2}{\partial \phi} & \frac{\partial f_2}{\partial \theta} \\ \frac{\partial f_3}{\partial r} & \frac{\partial f_3}{\partial \phi} & \frac{\partial f_3}{\partial \theta} \end{pmatrix} = \begin{pmatrix} \frac{x}{r} & -y & 0 \\ \frac{y}{r} & x & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Using a first-order approximation similar to that in Section 3, we get the following Cartesian covariance matrix:

$$\mathbf{C}(x, y, \theta) \cong \mathbf{R} * \mathbf{C}(r, \phi, \theta) * \mathbf{R}^T.$$

A more accurate approach is to use the extended Kalman filter (Gelb 1984), which includes the (nonlinear) coordinate mapping in the Kalman gain matrix \mathbf{K} .

REFERENCES

- Brooks, R. A. 1982. Symbolic error analysis and robot planning. *Int. J. Robotics Res.* 1(4):29–68.
- Brooks, R. A. 1985. Visual map making for a mobile robot. *Proc. IEEE Int. Conf. Robotics and Automation*. St. Louis: IEEE, pp. 824–829.
- Chatila, R., and Laumond, J-P. 1985. Position referencing and consistent world modeling for mobile robots. *Proc. IEEE Int. Conf. Robotics and Automation*. St. Louis: IEEE, pp. 138–145.
- Gelb, A. 1984. *Applied optimal estimation*. Cambridge: MIT Press.
- Nahi, N. E. 1976. *Estimation theory and applications*. New York: R. E. Krieger.
- Paul, R. P. 1981. *Robot manipulators: mathematics, programming and control*. Cambridge: MIT Press.

