REGISTRATION TECHNIQUES FOR MULTIPLE SENSOR SURVEILLANCE

Martin P. Dana

Command and Control Systems Division
P.O. Box 3310
Fullerton, CA 91634

In order to integrate multiple sensor data into a single air picture, the individual sensor data must be expressed in a common coordinate system, free from errors due to site uncertainties, antenna orientation, and improper calibration of range and time. The process of ensuring the requisite "error free" coordinate conversion of sensor data is called registration. This paper develops a Generalized Least–Squares Estimation technique for sensor registration and compares quantitatively this technique with of some the standard methods in use today.

1. Background

Modern Command and Control systems depend on a surveillance subsystem to provide an air situation picture on which decisions must be based. In order to maintain an accurate, complete and current air picture, the surveillance subsystem will, in turn, depend on combinations of netted sensors to provide the raw data from which the air situation picture is developed. To date, unfortunately, attempts to net multiple sensors into a single surveillance system have met with limited success, due in large part to the failure to register adequately the individual sensor (see Ref. [1]).

Why the registration of multiple sensor systems has been, in general, inadequate is not easily explained. The problem does not seem to be understood or even recognized beyond a small circle of systems engineers at a few Government laboratories and aerospace companies. Certainly it has not received the attention which, for example, the problem of tracking or state estimation has received. Literally thousands of papers have been published on Kalman filtering; many excellent (as well as mediocre) texts have been written on the subject of optimal estimation. Publications on the registration problem are limited to a few technical reports funded by various Department of Defense agencies. Registration, it seems, has been an afterthought in most system efforts.

The purpose of this paper is two–fold: first, to define the registration problem in terms of the sources of registration error and their implication on multi–sensor target tracking; and, second, to provide a solution of the registration problem. The solution of the problem discussed below is based on the techniques of multivariate statistical analysis. Thus, there is an obvious parallel between this solution to the registration problem and the Kalman filter to the extent that both can be derived from the theory of Generalized Least–Squares Estimation. More importantly, however, the solution discussed below treats the problem with a similar level of detail and sophistication as has been applied to the tracking function.

Since radars are still the primary sensors in use today, and since the problem of radar registration has not yet been resolved adequately, this paper will address the problem of radar registration only. The same principles can be applied to sensor networks which include other kinds of sensors.

2. The Registration Problem

The fundamental problem in sensor netting to determine whether data reports from two or more remotely located sensors represent a common aircraft or a distinct aircraft. Before this can be accomplished successfully, however, the individual sensor data must be expressed in a common coordinate system, free from errors due to site uncertainties, antenna orientation, and improper calibration of range and time. The process of ensuring the requisite "error free" coordinate conversion of sensor data is called registration. Thus, registration is an ABSOLUTE prerequisite for sensor netting.

The major sources of registration error for radars are listed below in the left–hand column of Table 1, together with some possible corrective actions in the right–hand column.

TABLE 1. Registration Error Sources

Error Source	Corrective Measure

Range:
    Offset
    Scale
    Atmospheric Refraction
	Test Target
	
	Tabular Corrections

Azimuth
    Offset
 	
    Antenna Tilt
	Solar Alignment; North Finders

Elevation (3–D Radars):
    Offset
    Antenna Tilt	

Time:
    Offset
    Scale	

Radar Location:	JITDS, PLRS, GPS, Satellite Survey

Coordinate Conversations:
    Radar Plane
    System Plane	

Source: Fischer, Muehe, Cameron: Registration Errors in a Netted Air Surveillance System (Ref. [1]).

Of the sources of registration error listed in Table 1, there are three sources which have proved to be major problems in air defense systems; they are: (1) position of the sensor with respect to a "system" coordinate origin; (2) alignment of the antennas with respect to a common north reference (i.e., the azimuth offset error); and (3) range offset errors. The other errors may exist in the current radar systems; however, they have not been significant problems in the past. As the radar technology improves, some of the other error sources may become significant factors.

As suggested in Table 1, electronic position-location systems such as GPS or commercial satellite survey systems are available which can locate a position on the earth to within a maximum error of 100 feet (or better). This accuracy is certainly adequate for radar systems in which the standard deviation of the range measurement error is no better than 0.125 nmi. The problem is now to deal with range and azimuth offset errors.

The potential effects of range and azimuth offset errors are illustrated in Figure 1. Registration errors are systematic measurement errors rather than random errors. The figure illustrates the expected or average reports of a common target from two radars, each of which consistently reports (1) a range less than the true range by a fixed amount (the offset) and (2) an azimuth (measured clockwise from north) less than the true azimuth by a fixed offset. For any specific set of measurements, random measurement errors will be superimposed on the bias errors.


Figure 1. Range and Azimuth Registration Errors. Registration errors introduce measurement biases into the system; this will result in degraded tracking performance or even in the initiation of multiple tracks for a single target.

The effect of systematic errors is to introduce biases into the estimation process. Therefore, failure to register a multiple radar system adequately can result in varying degrees of performance degradation, depending on the magnitude of the error or bias with respect to the random measurement errors and the tracking gates. The level of degradation ranges from the formation of multiple, redundant tracks for a single aircraft to reduced track accuracy and stability, or simply the loss of the benefits of multiple radar tracking by reducing the system, in effect, to a single radar tracking system.

3. Registration Procedures

System registration may be considered as a two phase process: sensor initialization and relative alignment. The objective of the initial registration procedure is to register, with respect to absolute coordinates, each sensor independently. Once the position of the sensor has been estimated, the range measurements have been calibrated, and an initial alignment with respect to true north has been completed, the procedures for relative alignment of the system sensors can be initiated. The initialization procedures generally are straightforward; the REAL registration problem is the relative alignment of the system sensors.

Techniques for relative registration depend on common targets, preferably targets of convenience rather than controlled flights. Generally, data is collected until a sufficient number of paired reports have been obtained, and then a set of bias corrections are computed. The usual techniques for obtaining the solutions is either to formulate the problem as an ordinary (or unweighted) least-squares estimation (LSE) problem or to rely on simple averaging to remove the random error components. The major limitation of either approach is that each radar report is treated equally when, in fact, the measurement (i.e., observation) errors are a function of both the individual radar parameters and target range.

The least-squares approach is commonly employed in the NATO air defense systems. This approach obtains a relative solution for a subordinate radar with respect to a master radar, which is assumed to be perfect. Since there is no particular reason to believe that the master radar is "perfect", this approach can only verify that the initial alignment is adequate; estimates of non-zero biases merely indicate that there is a registration error. Range offset errors, in particular, are not relative; a bias at the master site cannot be transferred to the subordinate site.

The alternative approach is the simple averaging process which is employed in the US-Canadian Joint Surveillance System (JSS) for North America and in the FAA National Air Space System (NAS) (for enroute air traffic control). The derivation of this technique assumes a symmetric distribution of points about the line joining the two radars. Consequently, the solutions are very sensitive to the actual target distribution. For radars along many political borders, it may not be possible to obtain any data at all from one side of the line joining the two sites.

Given this situation, it is obvious that a new approach to system registration is needed. The basic objective of this research is to develop a technique for registration with the following characteristics:

- Insensitive to target distribution,

- Applicable to fixed site, mobile and airborne sensor systems,

- Provide alternative solution sets depending on the need,

- Provide a quality estimate for the solution set, and

- Be based on a recognized principle of optimality.

4. Bias Estimation

Fisher, et. al, suggest (Ref. [1], p. 17) three alternative approaches; specifically, the generalized linear least-squares estimation (GLSE) technique and two numerical optimization methods, one based on a grid search technique and the other on Powell's method for steepest descent. The GLSE is dismissed for computational reasons, and the grid search approach is dismissed in favor of Powell's method because of slow convergence.

Commercial array processors or special purpose co-processors are now available which are capable of performing the large scale matrix operations required by the GLSE approach (see Ref. [2]). Consequently, the GLSE approach is reconsidered in this paper. The technique developed by Wax in Ref. [3] can be applied to formulate the generalized Gauss-Markov problem.

The approach suggested by Wax is to formulate the difference dP in the reported positions as a function of the set of measured variables Z (i.e., observations) and the set of biases B (i.e., parameters) to be estimated: $dP = F(Z,B)$. Following the usual linearization technique, but with the roles of the actual values and estimators reversed, the vector equation for position difference can be transformed in the classical Gauss-Markov GLSE model (see Ref. [2]): $X*B + E = Y$, where X is a matrix of known parameters, E is the vector of measurement errors, and Y is the measurement vector.

The solution of the GLSE Problem above is simply

$$\tilde{B} = (Cov) * X^T * V^{-1} * Y$$

where $(Cov) = (X^T * V^{-1} * X)^{-1}$ is the covariance matrix for the estimate $\tilde{B}$ of B.

The difficult part of the formulation is to develop a representation for the covariance matrix V of the measurement error. However, Fischer, et. al., provide a framework for the derivation in Appendix D of Ref. [1]. The details of the derivation of the solution for two range and two azimuth biases are provided in the appendix of this paper.

In general the GLSE approach requires a capability to perform arithmetic with large matrices. For this application, however, the problem may be greatly simplified using the independence of the measurements, both between radars and over the set of targets. As shown in the appendix, the covariance matrix V for the error term E is a block-diagonal matrix; the dimension of the individual blocks is the same as the cardinality of the set Z for the individual samples. For the registration problem, the measurement set Z contains four (4) independent measurements. Therefore, the covariance matrix is the inverse of a sum of 4x4 matrices, which can be computed easily. (See the appendix.)

5. Numerical Evaluation

During the past year the GLSE approach has been formulated and evaluated at Hughes Aircraft Company. The evaluations have considered both theoretical covariance analyses and simulation analyses for comparison of the GLSE technique with the JSS, NATO and the ordinary LSE techniques. Some of the major results of these evaluations are presented below.

5.1 Covariance Analyses.

The GLSE approach was developed for three distinct solution sets; these were the following:

- Two azimuth offset biases,
- Two range offset biases, and
- Two range and two azimuth offset biases.

In the case of the "two azimuth bias" solution, it is assumed that there is a potential azimuth bias at each of the two radars, which are called the master and subordinate for convenience; it is assumed further that there are no range biases at either of the two radars. For the "two range bias" solution, the analogous assumptions were used. The "two range/two azimuth bias" solution is that derived in detail in the appendix.

The covariance matrices for the three alternative solutions sets were analyzed with respect to the number of samples (that is, targets) used in the solution and the distribution of the targets in the (x,y)-plane. The results of the analysis with respect to the sample size are shown in Figures 3, 4 and 5 for the sensor/target geometries illustrated in Figure 2. For these analyses, the standard deviations of the random range and azimuth measurement errors at both radars were assumed to be 0.125 nmi. in range and 0.18 degree (approximately 3.0 milli-radians) in azimuth. These statistics are typical of modern air defense and air traffic control radars.


Figure 2. Sensor/Target Geometry. The covariance analyses were conducted for targets distributed symmetrically along the perpendicular bisector of the line segment joining the two radars and for targets distributed along a line parallel to the line segment joining the two radars.

The ratio of the standard deviation of the azimuth bias estimate to the standard deviation of the (random) azimuth measurement error is plotted in Figure 3. From the graph, approximately 125 samples are required in order to obtain a bias-to-measurement ratio of 0.10, which will ensure that any system track inaccuracies are due to the random errors rather than the systematic or bias errors.

by the positive slopes for each curve on the graph. If there are no range biases, however, these (azimuth-only) solutions will be somewhat more accurate than the JSS and the GLSE-2 solutions.


Figure 6. RMS Azimuth Bias Estimation Error versus Range Bias. The JSS and the GSLE-2 algorithms solve for range biases and, therefore, are insensitive to range biased data. The GLSE-2 algorithm can achieve a 50% reduction in estimation error with respect to the JSS approach.

The performance of the range bias estimation techniques is shown in Figure 7 versus an actual azimuth bias. As was the case with azimuth-only solutions, the presence of a bias which is assumed to be non-existent, will degrade performance severely. Also, as before, the GLSE-2 algorithms can achieve a 2-to-1 reduction of the standard deviation of the estimation error with respect to the JSS method.


Figure 7. RMS Range Bias Estimation Error Versus Azimuth Bias. The range-only GSLE technique is very sensitive to azimuth biases but is still preferred the JSS method when the magnitude of the azimuth bias is less than 0.20 degree.

Based on the data presented above and extensive analyses which were not included here because of the limited space, the following statements summarize the conclusions of the IR&D work conducted by Hughes Aircraft Company over the past two years. In general, the GLSE approach exhibits modest CPU requirements; the algorithm is certainly practical as an off-line or background capability. More importantly, the GLSE approach is less sensitive to sensor/target geometry than any of the other registration approaches, particularly the JSS approach. In most cases, the GLSE algorithms can achieve satisfactory registration accuracies with 50 to 100 point-pairs rather than the 200 often required by the JSS or NATO algorithms.

References

1. W. L. Fischer, C. E. Muehe and A.G. Cameron, Registration Errors in a Netted Air Surveillance System; MIT Lincoln Laboratory Technical Note 1980-40, 2 Sept. 1980.

2. T. W. Anderson: An Introduction of Multivariate Statistical Analysis; John Wiley & Sons, Inc., 1958.

3. Mati Wax: "Position Location from Sensors with Position Uncertainty," IEEE Trans. on Aerospace and Electronic Systems, Vol. AES-19, No. 5 (Sept., 1983).

Appendix: Mathematical Development

In the following derivation, assume that a master radar R(1) is located at the origin of the coordinate system and that a subordinate radar R(2) is located at coordinates (u,v). For this derivation, it is immaterial which radar is the master and which is the subordinate. Also assume that there are N targets in the intersection of the respective fields of view, denoted by T(1), T(2), ... T(N). (See Figure A.)

The basic problem is to determine the range and azimuth biases at each radar from the measurements of the common targets T(1), T(2), ... T(N). That is, it


Figure A. Sensor Measurement Geometry. The registration algorithm must determine the system biases for the measurement set $X(k) = \{r(1, k), \theta(1, k), r(2, k), \theta(2, k)\}$.

Figure 3. Azimuth Bias Estimation Performance Versus Sample Size. The variance of the azimuth bias estimation error decreases as a function of 1/N, where N is the number of samples.

Similarly, the ratio of the standard deviations of the range bias estimate to the random range measurement error is plotted in Figure 4. From the graph, 100 to 150 samples are required in order to obtain a 5-to-1 improvement ratio. Although this is not as dramatic as the 10-to-1 ratio obtained for the


Figure 4. Range Bias Estimation Performance Versus Sample Size. The variance of the range bias estimation error decreases as a function of 1/N, where N is the number of samples.

azimuth case, it will be adequate for target tracking. As is the case for the azimuth bias estimate, the variance of the range bias estimate is approximately inversely proportional to the number N of samples.

Finally, the results of the same analyses are shown in Figure 5 for the "two range and two azimuth biases" solution. As was the case for the range-only or the azimuth-only cases, the solutions are relatively insensitive to target distribution. In this case, however, a symmetric distribution along the perpendicular bisector yields the best performance by a factor of at least 2-to-1 (with respect to the parallel target case). If a symmetrically distributed sample cannot be obtained, then 200 or more points may be required in order to obtain satisfactory performance for tracking.


Figure 5. Range and Azimuth Bias Estimation Performance Versus Sample Size. The performance for this case is moderately sensitive to target distribution; however a sample size of 200 will be adequate of tracking in any case.

5.2 Simulation Results

In addition to the covariance analyses described above, simulation analyses were conducted, primarily in order to compare the GLSE algorithms with other techniques (for which covariance estimates are not available). A secondary goal of the analyses were to determine the sensitivity of the current registration assumptions of the techniques. For example, the NATO registration algorithm assumes that there are no range biases at the radars; if there are range biases, then the algorithm will translate them into azimuth biases.

For the simulation data presented in Figures 6 and 7 below, the sensor/target geometry is similar to the "perpendicular bisector case" in the preceding section except that the targets are not symetrically distributed along the bisector. For this analysis, all of the targets are in the first quadrant of the coordinate system.

In Figure 6, six algorithms for estimation of two azimuth biases are compared in terms of the RMS error between the estimated bias and the true bias, which was 0.50 degree. The algorithms included the two relevant versions of the GSLE (or Gauss-Markov) approach discussed in Section 5.1; a trilateration technique which uses only the range measurements from each radar; the JSS method; the NATO (or 407L) method; and an ordinary least-squares algorithm. The NATO algorithm differs from the ordinary LSE algorithm only to the extent that the NATO algorithm solves for an azimuth bias and a position bias of the subordinate radar; the position bias is then translated into an azimuth bias at the master radar.

The results of the analysis are shown in Figure 6, in which the RMS error in the azimuth bias estimates are shown on the ordinate versus an actual range bias on the abscissa. As one would expect, the JSS Algorithm and the GLSE/Gauss-Markov algorithm for the "two azimuth and two range" solution are insensitive to the presence of range biases since both approaches solve for range biases in addition to azimuth biases. The other approaches are derived from the assumption that there are no range biases and, therefore, produce inaccurate azimuth bias estimates when there are range biases, as indicated

is necessary to estimate the azimuth biases a(1) and a(2) at R(1) and R(2), respectively, and the range biases b(1) and b(2) at R(1) and R(2). Denote the set of biases by

$$\beta = \{ a(1), a(2), b(1), b(2) \} \qquad (1)$$

For each target T(k), define the set of radar measurements

$$\psi(k) = \{ r(1,k), \theta(1,k), r(2,k) \, \theta(2,k)\} \qquad (2)$$

where $r(1,k)$, $\theta(1,k)$ and $r(2,k)$, $\theta(2,k)$ denote the range and azimuth measurements from radar R(1) and radar R(2), respectively.

For each set of measurements, $\psi(k)$, the observations are the separations in the (x,y)-plane of the reported target positions. These are:

$$dx(k) = [r(1,k) + b(1)] \cos [\theta(1,k) + a(1)] - u$$
$$- [r(2,k) + b(2)] \cos [\theta(2,k) + a(2)] \qquad (3)$$

$$dy(k) = [r(1,k) + b(1)] \sin [\theta(1,k) + a(1)] - v$$
$$- [r(2,k) + b(2)] \sin [\theta(2,k) + a(2)] \qquad (4)$$

Equations (3) and (4) above relate the set β of parameters to be estimated to the set of measurements $\psi(k)$ and the vector of observations $[dx(k), dy(k)]$. However, these functional relationships are non-linear.

In order to apply the Gauss-Markov theory of Generalized Least Squares Estimation (GLSE), it will be necessary to represent the observations as a linear function of the parameters to be estimated, namely β. This can be accomplished by defining a function f as follows:

$$f(\psi(k), \beta) = [dx(k), dy(k)]^T$$

where the superscript T denotes the transposition of the vector (or, later, the matrix). Further, let $\psi'(k)$ and β' denote the actual measurement sets and an initial estimate of β, respectively. Now, Taylor's Theorem can be used to approximate the function f at the true values of $\psi(k)$ and β in terms of the measurements $\psi'(k)$ and the initial estimate β':

$$f(\psi(k), \beta) = f(\psi'(k), \beta')$$
$$+ \nabla_\beta f(\psi'(k), \beta') (\beta - \beta') \qquad (5)$$
$$+ \nabla_\psi f(\psi'(k), \beta') [\psi(k) - \psi'(k)]$$

where

$$F(k) = \nabla_\psi f[\psi'(k), \beta'] \qquad (6)$$

$$= \begin{bmatrix} \dfrac{\delta [dx(k)]}{\delta r(1,k)} & \dfrac{\delta [dx(k)]}{\delta \theta(1,k)} & \dfrac{\delta [dx(k)]}{\delta r(2,k)} & \dfrac{\delta [dx(k)]}{\delta \theta(2,k)} \\[2mm] \dfrac{\delta [dy(k)]}{\delta r(1,k)} & \dfrac{\delta [dy(k)]}{\delta \theta(1,k)} & \dfrac{\delta [dy(k)]}{\delta r(2,k)} & \dfrac{\delta [dy(k)]}{\delta \theta(2,k)} \end{bmatrix}$$

and

$$G(k) = \nabla_\beta f(\psi'(k), \beta') \qquad (7)$$

$$= \begin{bmatrix} \dfrac{\delta[dx(k)]}{\delta a(1)} & \dfrac{\delta[dx(k)]}{\delta a(2)} & \dfrac{\delta[dx(k)]}{\delta b(1)} & \dfrac{\delta[dx(k)]}{\delta b(2)} \\[2mm] \dfrac{\delta[dy(k)]}{\delta a(1)} & \dfrac{\delta[dy(k)]}{\delta a(2)} & \dfrac{\delta[dy(k)]}{\delta b(1)} & \dfrac{\delta[dy(k)]}{\delta b(2)} \end{bmatrix}$$

If the errors $[\psi(k) - \psi'(k)]$ and $(\beta - \beta')$ are sufficiently small that the higher order terms can be neglected, then the approximation in (5) may be regarded as an equality. Also, note that

$$f(\psi(k), \beta) = 0 \qquad (8)$$

by definition; therefore:

$$G(k)\beta + F(k) \, d\psi'(k) = G(k) \, \beta' - f(\psi'(k), \beta') \qquad (9)$$

where $d\psi'(k) = \psi'(k) - \psi(k)$. Note that the matrix G(k) is a matrix of known parameters, $F(k) \, d\psi'(k)$ is the error due to the measurement noise, and that the terms on the right-hand side of equation (9) now represent the observations.

With all of this notation and the approximation of equation (5), equations (3) and (4) may now be reformulated as the classical Gauss-Markov model of GLSE theory:

$$X \beta + \varepsilon = Y \qquad (10)$$

by setting

$$X = [G(1), G(2), \ldots G(N)]^T \qquad (11)$$

$$\varepsilon = [F(1) \, d\psi'(1), F(2) \, d\psi'(2), \ldots F(N) d\psi'(N)]^T \qquad (12)$$

$$Y = [G(1) \, \beta' - f(\psi'(1), \beta'), G(2) \, \beta' -$$
$$f(c'(2), \beta'), \ldots G(N) \ldots]^T \qquad (13)$$

Note that X is a 2Nx4 matrix, ε is 2N vector, and that the observation vector is also of dimension 2N.

The last step in this application of the Gauss-Markov model is to develop the covariance Σ_ε matrix for the error vector ε. To this end, define:

$$\Sigma_\varepsilon = E[\varepsilon\varepsilon^T] = \text{diag} \{F(k) \, E[d\psi'](d\psi')^T] \, F^T(k)\} \qquad (14)$$

where the notation "diag" indicates a diagonal matrix with the non-zero terms enclosed in the brackets. Note that

$$\Sigma_\psi = E\left[(d\psi')(d\psi')^T\right] = \text{diag}\left[\sigma^2_{R(1)}, \sigma^2_{\theta(1)}, \sigma^2_{R(2)}, \sigma^2_{\theta(2)}\right] \qquad (15)$$

Further, note that F(k) is a 2 x 4 matrix and that Σ_ψ is a 4 x 4 matrix; therefore

$$\Sigma_k = F(k) \, \Sigma_\psi \, F^T(k) \qquad (16)$$

is a 2 x 2 matrix. This implies that Σ_ε is a block-diagonal matrix of the form

$$\Sigma_\varepsilon = \text{diag}\left[\Sigma_1, \Sigma_2, \ldots \Sigma_N\right]. \qquad (17)$$

The solution of the Gauss-Makrov equation (10) is simply

$$\hat{\beta} = (X^T \Sigma_\varepsilon^{-1} X)^{-1} X^T \Sigma_\varepsilon^{-1} Y \qquad (18)$$

where

$$\text{Cov} (\hat{\beta}) = (X^T \Sigma_\varepsilon^{-1} X)^{-1}. \qquad (19)$$

Since Σ_ε is a 2NX2N block-diagonal matrix, it follows that

$$X^T \Sigma_\varepsilon^{-1} X = \sum_{k=1}^{N} G^T(k) \Sigma_k^{-1} G(k) \qquad (20)$$

where the individual terms of the sum are 4X4 matrices. Similarly

$$X \Sigma_\varepsilon^{-1} Y = \sum_{k=1}^{N} G^T(k) \Sigma_k^{-1} [G(k)\beta' - f(\psi'(k),\beta')] \qquad (21)$$

If the individual radar measurement errors are normally distributed, then $d\psi'(k)$ is a normally distributed vector; and $F(k)\, d\psi'(k)$ is a linear combination of normal variables and is, therefore,

normally distributed. Thus ψ is $N(0, \Sigma_\psi)$.

Equation (18) is the minimum variance solution under any error distribution. For the normal distribution (i.e., $\varepsilon \sim N(0, \Sigma_\varepsilon)$), $\hat{\beta}$ is also the maximum likelihood solution. By these criteria, β in (18) is the "best" solution to the problem as defined by equation (10):

$$Y = X\beta + \varepsilon$$

where the error term ε is distributed as $N(0, \Sigma_\varepsilon)$, X is a matrix of known parameters, and Y is the vector of observations.