

**On Comparing Statistical and Set-Based Methods
in Sensor Data Fusion**

Gregory D. Hager and Sean Engelson

Research Report YALEU/DCS/RR-924
October 1992

On Comparing Statistical and Set-Based Methods in Sensor Data Fusion

Gregory D. Hager and Sean Engelson
Dept. of Computer Science
Yale University
P.O. Box 2158 Yale Station
New Haven, CT 06520-2158

Abstract

We compare the theoretical and practical considerations of two common sensor data fusion methodologies: set-based and statistically based parameter estimation. We first examine their convergence behavior in on a variety of simulated problems. We then describe robot localization systems implemented using both methods and compare their performance. Our conclusion is that set-based methods have performance that sometimes exceeds that of statistical methods, though this result is highly problem dependent. We then characterize these problem dependencies.

1 Introduction

In recent years, it has become common to express sensor data fusion problems in terms of parameter estimation or hypothesis testing, and to solve these problems using statistical estimation methods [6]. Practically without exception, solutions apply variations of classical mean-square estimation techniques [9]. However, the efficacy of these techniques depends to a great extent on the character and fidelity of mathematical sensor models [12, 14, 15]. Specifically, accurate estimates and accurate computation of estimation error depend heavily on the statistical properties of sensing errors and the structure of the observation system.

We have been experimenting with a complementary set of techniques that eschews an explicit statistical representation of error. Instead, these *set-based* estimation techniques combine data under the assumption that sensor error is *bounded*, and focus on establishing bounds on maximal parameter error. There is a great deal of precedent for the set-based approach. A variety of publications examine the use of set-based techniques for linear systems [4, 17, 18, 20]. Taylor [19] examined the use of set-based methods as a means of propagating uncertainty through kinematic chains. Brooks [5] used symbolic techniques to solve nonlinear inequalities and index into a model database. More recently, a variety of recognition techniques have been implemented using bounded error models [10, 11]. Ellis [7] has examined the problem of improving polyhedral object recognition by improving the propagation and combination of set-based uncertainties. Atiya [1] describes the solution to robot location recognition and computation using set-based methods. Similarly, set-based computations are used by Engelson and McDermott [8] to build maps of a robot environment. Hager [13] describes set-based techniques as a basis for computing qualitative decisions from sensor data.

We are now investigating the strengths and weaknesses of both statistical and set-based approaches through theoretical analysis, simulation, and experimentation. We focus on three issues:

- What type of comparative statements can be made between set-based and statistical methods?
- What are the practical differences in applying both methods?
- How can information be combined in cases where both models apply?

This article focuses largely on the first two issues. In the next section, we discuss the general characteristics of statistical and set-based methods. In the following section, we present simulation results that illustrate the behavior of the two methods on specifically chosen problems. In Section 4, we describe a statistically based localization system that is structurally identical to the set-based system described in [2] and compare our experiences with both versions. We close with a short summary of our results and present some ideas for combining set-based and statistical methods.

2 Theoretical Basics

Structurally, geometric sensor data fusion problems involve modelling a variety of data sources relative to an underlying physical or geometrical model, and “solving” the equations for model parameters given sensor inputs. This problem becomes difficult when any or all of the following factors enter the problem:

- Sensor information is in error.
- The relationship between sensor outputs and model parameters is nonlinear.
- Parameter estimates must be propagated through a dynamic system.
- Correct, complete sensor models are not available.
- Sensor models are occasionally violated (the *outlier* problem).

We investigate these problems in a graduated manner by studying observation systems of the general form $z = h(x) + v$. We consider the cases where h is linear and nonlinear, and x is scalar and vector. We evaluate estimation methods both in the context of obtaining point estimates and obtaining confidence intervals. All of these problems are discussed within the framework of time-series estimation. That is, data arrives sequentially, and past observations are available only if explicitly stored for later retrieval.

2.1 Estimation Methods

Mean Square Estimation: Given a prior distribution on model parameters, x , and a sampling (conditional) distribution on observations, z , the optimal mean-square estimator is the posterior mean (assuming the required moments exist). When the prior parameter information and sensor noise are both Gaussian and the observation system is linear, the posterior mean is a linear function of observations. If the unknown parameters follow a linear dynamic law, then the propagation of parameter information is again linear. This estimation system is commonly referred to as the Kalman filter [9].

When some of the assumptions above are violated, the posterior mean usually becomes a non-linear function of observations. However, if the prior and sampling distributions have known first and second moments and the system is linear, then the optimal *linear* estimator is still the Kalman filter. Once the measurement system becomes nonlinear, there are no generally known optimality results, although the extended Kalman filter (EKF) and iterated extended Kalman filter (IEKF), modified Kalman filters applied to a linearized version of the observation system, are often used.

Set-Based Estimation: Assume that the errors in sensor observations are bounded: $v \in \mathcal{V}$. Let $a + \mathcal{V}$ denote the set \mathcal{V} translated by a and define $h^{-1}(z) = \{x \mid z \in h(x) + \mathcal{V}\}$. For any series of sensor observations z_1, z_2, \dots, z_n ,

$$x \in \mathcal{X}_n = \bigcap_{i=1}^n h^{-1}(z_i).$$

Assuming the error bounds are tight, \mathcal{X}_n is the smallest set guaranteed to satisfy these conditions. Furthermore, if h is invertible and the probability of v taking values in any small disk near the border of \mathcal{V} is nonzero, then it can be shown that the sequence of sets \mathcal{X}_i , $i = 1, 2, \dots, n$ converges in probability to a point estimate.

The primary difficulty in implementing set-based methods is representing the sets $h^{-1}(z_i)$, and their intersections. The following general comments can be made:

1. In the scalar case, the linear problem is trivial and exists in closed form. The nonlinear case can usually be solved by computing the inverse function h^{-1} , using interval methods or other forms of extremal analysis.
2. In the multivariate linear case, if v is a convex set, then so is \mathcal{X}_i for all i . If v is polyhedral then the intersection problem can be solved with linear programming methods. Unfortunately the representation and computation of set intersections will grow in complexity over time. Also, the complexity of intersection computation grows quickly with dimensionality. Removing the convexity restriction increases the complexity. In practice, it is common to use some type of regular bounding set that consumes fixed memory and computational resources. Examples are the use of bounding ellipses [4], bounding linear convex sets [3], and bounding intervals [1].
3. In the nonlinear multivariate case, the structure of a given set \mathcal{X}_i is sufficiently complex that there are no general, exact computational results. If h is invertible, the usual approach is to bound \mathcal{X}_i using methods described above. If h cannot be inverted, there are a class of iterative refinement techniques that can be applied [13], although this adds considerably to the memory used and the computational effort required.

When the unknown parameters follow a dynamic law, information is combined over time by projecting the uncertainty set containing the parameters at each time step. In the linear case, this is usually straightforward. In the nonlinear case, it is common to use interval arithmetic [16] for this purpose.

Comparisons and Comments: The fundamental difference between set-based and statistical methods is the character of the sensing systems for which convergence can be demonstrated. The

Kalman filter is convergent for linear systems with very specific statistical properties. Set-based methods are convergent for linear and nonlinear systems when tight bounds on the errors are known and tight approximations to the uncertainty set can be calculated. Practically speaking, Kalman filtering becomes more difficult when nonlinearities are introduced, and set-based methods are more difficult to apply to high-dimensional systems.

3 A Comparison of Efficiency

In the previous section, we outlined the conditions that differentiate between set-based and statistical methods. Practically speaking, approximate versions of the Kalman filter can be applied to many problems that could be attacked with set-based methods. Specifically, if sensor errors can be modeled using a distribution on a bounded sample space, it is possible to apply either set-based or statistical methods. Intuitively, if the distribution on errors is strongly centralized and observation errors are independent, we expect a “tuned” linear estimator to perform quite well. Conversely, if errors are difficult to model statistically and the distribution of errors is not strongly centralized, we expect fairly rapid set-based convergence.

Consider noise models from the class of truncated Gaussian distributions. This class can be described by the following four parameter model:

$$\pi_t(x, \mu, \sigma^2, l, u) = \begin{cases} \frac{1}{k} \pi(x, \mu, \sigma^2) & \text{if } l \leq x - \mu \leq u; \\ 0 & \text{otherwise.} \end{cases}$$

where π is the classical Gaussian pdf, and k is a normalizing coefficient equal to the probability mass of a Gaussian in the interval $[l, u]$.

When computing Bayes’ theorem on truncated Gaussian random variables, the characteristics of the posterior are affected by two independent factors: the mean and covariance of the distributions involved, and the bounds on the distributions. Simple calculations show that the posterior distribution is exactly the result of applying Bayes rule to the untruncated distributions and then truncating to the intersection of the sets of support of the original distributions. Hence, the result is again a truncated Gaussian.¹ The same comments hold for multivariate distributions. Consequently, if the set of support for a truncated Gaussian is large relative to its variance, the posterior mean will be very close to that which the Kalman filter computes—that is, nearly a linear function of observation. Conversely, as the variance becomes large relative to the set of support, the nonlinearity introduced by the truncation begins to dominate and linear estimation rules will not perform as well.

We pursue this observation in the context of two problems: point estimation and bounds estimation for hypothesis testing. The rationale for this two-pronged approach is that the best technique to use often depends on the ultimate goal of the estimation process. Point estimation rewards accuracy and lack of estimation bias. Hypothesis testing rewards accurate bounds on the range of a parameter, and is largely independent of how well the exact parameter values can be located within those bounds.

Estimator efficiency can be measured in simulation as the number of iterations taken to achieve a stopping criterion. When dealing with point estimation, the criterion is to stop when the estimate variance drops below a threshold. For set-based methods, we assume a uniform distribution over

¹This result actually holds for any truncated distributions; it does not depend on Gaussianity.

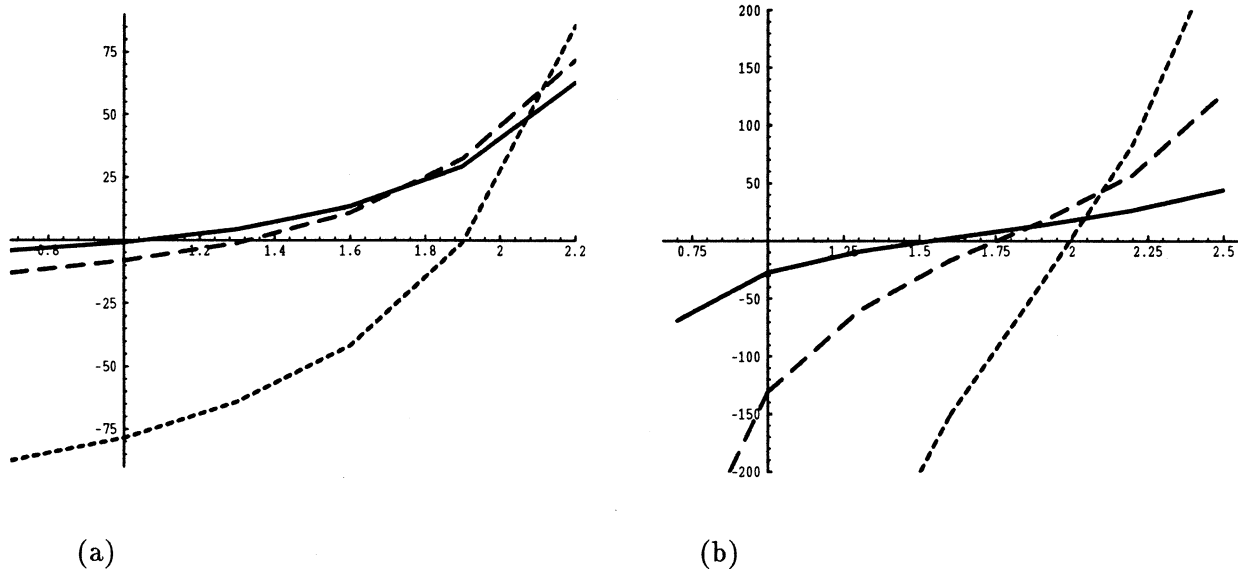


Figure 1: Difference in number of steps between set-based and statistical estimators for a linear system as a function of truncation half-width. (a) Stopping at variances of 0.1, 0.05, and 0.01 (solid, dashed, and dotted, resp.). (b) Stopping at confidence bounds of 0.4, 0.2, and 0.1 (solid, dashed, and dotted, resp.).

the estimate interval for this purpose, to compute this variance. For hypothesis testing, the criterion is to stop when the estimated 99% confidence interval is smaller than a threshold.

3.1 The Linear Case

Univariate systems: We begin with the univariate linear system $z = x + v$, with v distributed as a truncated Gaussian with variance 1 and truncation half-width b . We simulated this system for a variety of values of b , and computed the efficiency of the set-based and statistical estimators as the number of steps to stopping (averaged over 1000 trials). The differences in efficiency between the interval estimator and the Kalman filter are summarized in Figure 1. As expected, the Kalman filter works better when truncation width is large and the distribution approximates a Gaussian; the set-based estimator works better when truncation width is small. One interesting result of this simulation is that the increased accuracy required by confidence set estimation increases the relative efficiency of the set-based method over the Kalman filter.

Multivariate systems: To illustrate the multivariate case, we consider the system $\vec{z} = H\vec{x} + \vec{v}$ where $\vec{x} \in \mathbb{R}^2$, H is a rotation matrix, and \vec{v} is a vector of identical, independently distributed truncated Gaussian random variables. With no rotation, the system is decoupled and acts like two univariate systems. As rotation increases, overestimation will start to affect the interval results. The optimal set-based result is in the non-rotated case. Figure 2(a) compares estimator efficiency with no rotation against a rotation of $\pi/8$. The graph clearly shows that performance of the




Figure 2: Multivariate linear system results. (a) Relative efficiency of set-based and statistical estimators as a function of truncation half-width, under no rotation (solid) and rotation of $\pi/8$ (dashed). (b) Mean squared error of final Kalman filter mean (dotted) and final interval estimate midpoint (solid).

interval-based estimator degrades far more rapidly as error bounds increase, due to parameter overestimation. However, as can be seen in Figure 2(b), the accuracy of the Kalman filter degrades with increasing bounds, as observations further from true become possible. There is thus a tradeoff between efficiency and accuracy, which will depend on the particular application.

3.2 The Nonlinear Case

Set-based methods use essentially the same set of techniques for both linear and nonlinear systems; linearization is not needed. However, as noted above, linear estimation theory is commonly applied to nonlinear systems by using linear approximations. The difficulty with this approach is that it is difficult to demonstrate convergence except in very weak cases where the theory of stochastic approximation can be applied [9]. Except in relatively rare cases, the gains used in the Kalman filter must be adjusted to allow the system to accommodate linearization errors. In [12] we analyzed this problem and noted that in many cases, a lower bound on the derivatives of the system led to convergent behavior. This choice of gain modification is also in accord with stochastic approximation theory. Thus, we adopt the methodology of using the smallest derivative value within some confidence interval about the mean when implementing EKF's.

Figure 3(a) shows the results for a simple univariate polynomial system. Figure 3(b) shows the results for an exponential observation system. The relative efficiencies of the estimators are qualitatively the same as in the linear case. In the exponential case however, the curves show a greater advantage for the interval method. This is due to using the minimum derivative in a confidence in-

terval for linearization in the EKF, which gives stability at the cost of slower convergence. Another feature of that graph is that the crossover points seem to converge for different stopping points; it is unclear whether this is just a coincidence or reflects some unknown structural property.

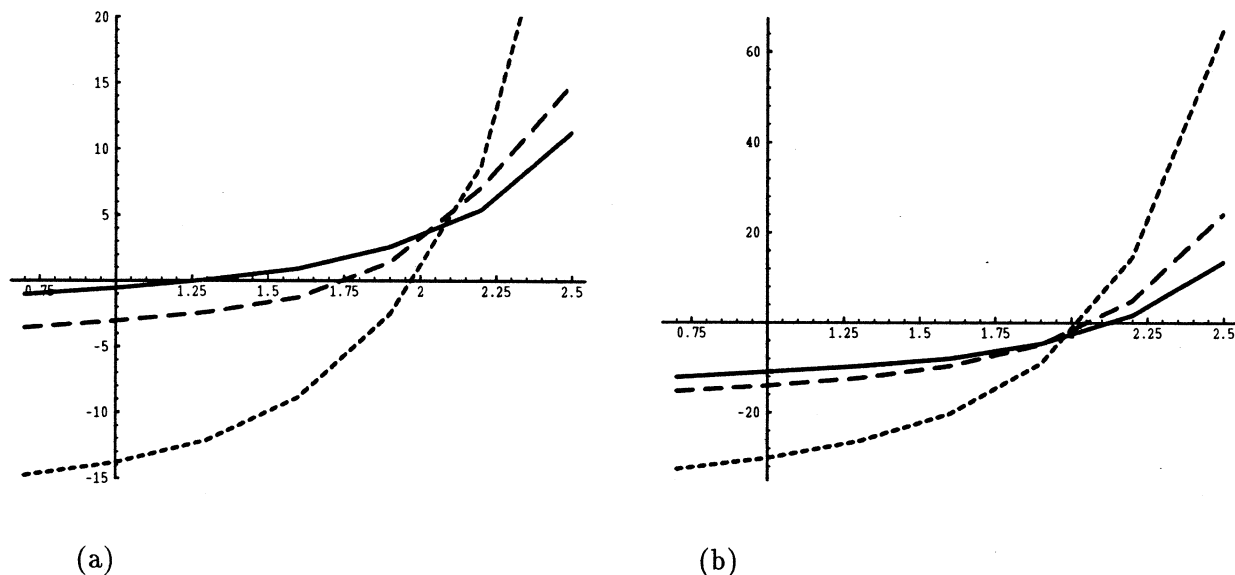


Figure 3: Difference in number of steps between set-based and statistical estimators for the non-linear systems. (a) The system $z = (x^3 + x) + v$ at $x = 1$, stopping at variances of 0.1, 0.05, and 0.01 (solid, dashed, and dotted, resp.). (b) The system $z = e^x + v$ at $x = 1$, stopping at variances of 0.1, 0.05, and 0.01 (solid, dashed, and dotted, resp.).

This may be a case where keeping *both* set-based and statistical information around would be useful. Set-based estimates can be used to get better derivative bounds for statistical estimation. This can produce more efficient accurate estimation.

4 Experimental Results

In [2], Atiya and Hager described a set-based landmark matching and robot localization system. We recently constructed a corresponding estimation-based system using the linear estimation techniques described above. This section compares the performance of both methods for stereo tracking and feature matching for two demonstrative cases.

4.1 System Description

A detailed description of the system can be found in [1]. Briefly, the localization system hardware consists of two cameras mounted on a computer controlled precision translation table with their optical axes perpendicular to the line of translation. The right camera is fixed and the left translates. The localization software determines the robot location (two translations and one rotation) by

matching naturally and artificially occurring vertical “stripes” to a “map” of global locations. Both matching and localization occur in a two-dimensional plane parallel to the direction of slider motion and containing the optical axis of the single fixed camera. The imaging model for the stereo camera system includes extrinsic geometric components relating the cameras frame to a fixed coordinate frame on the slider, and intrinsic parameters describing a perspective and lens distortion model for both cameras.

Stereo data is taken at intervals T as the left camera is in motion on the slider. At each time point $t_k = kT, k = 0, 1, 2, \dots$, the following measurements are gathered: image coordinates $o_i^r(k), i = 1, \dots, m^r(k)$ from the right camera; image coordinates $o_i^l(k), i = 1, \dots, m^l(k)$ from the left camera; and the slider position $o^{sl}(k)$ read from the slider encoders. Examining time series of camera data led us to conclude that the maximum error of our line detector is 0.5 pixels. In order to account for modeling error, we adopted an error bound of 0.55 pixels to be used in all set-based algorithms. Statistically, we assume that coordinate observations from the cameras are independent and that errors are zero mean. The observation standard deviation, σ , was varied from 0.3 pixels (the equivalent of a uniform distribution in the interval $[-.55, .55]$) to 1.0 pixel as described below. The slider positioning accuracy is modeled as zero mean noise with variance $w = 10^{-4} \text{ cm}^2$ based on manufacturer’s specifications.

4.2 Stereo Triangulation

Problem: Computing the coordinates of a point $i, (x_i, y_i)$, given a series of image observations.

At time $t_k = 0$, the cameras are less than a centimeter apart, so initial stereo correspondences are simple to compute. Correspondences in subsequent images are computed by validating new measurements based on previous stereo estimates. Data from validated correspondences is combined sequentially as it becomes available. The state vector $p(k)$ describing the system at step k consists of the coordinates of point $i, (x_i(k), y_i(k))$ and the slider position $s(k)$. The measurement model can be written as:

$$\begin{bmatrix} o_i^r(k) \\ o_i^l(k) \\ o^{sl}(k) \end{bmatrix} = \begin{bmatrix} h_1(p_i(k)) \\ h_2(p_i(k)) \\ s(k) \end{bmatrix} + \begin{bmatrix} v_{or} \\ v_{ol} \\ v_{osl} \end{bmatrix} \quad (1)$$

The components h_1 and h_2 correspond to the composition of the camera transformation equations. During image acquisition, the state vector for point i has the following dynamic model:

$$p_i(k+1) = \begin{bmatrix} x_i(k+1) \\ y_i(k+1) \\ s(k+1) \end{bmatrix} = \begin{bmatrix} x_i(k) \\ y_i(k) \\ s(k) \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ w \end{bmatrix} \quad (2)$$

IEKF Formulation: The prior mean for the state vector is taken to be $\hat{p}_i^-(0) = (0, 300, 0) \text{ cm}$ and the covariance is described by diagonal matrix with 10^6 cm^2 for the first two entries and 10^{-4} cm^2 for the final entry. The first two values are somewhat large, but smaller values introduce too much bias into the initial IEFK estimates and lead to tracking and correspondence failure. Correspondence validation is accomplished by projecting the current estimate into each camera and testing consistency using Mahalanobis distance with threshold ϵ_{track} .²

²We note that we do not consider the Mahalanobis distance to be the best test; instead it was chosen based on its wide use in the literature [6].

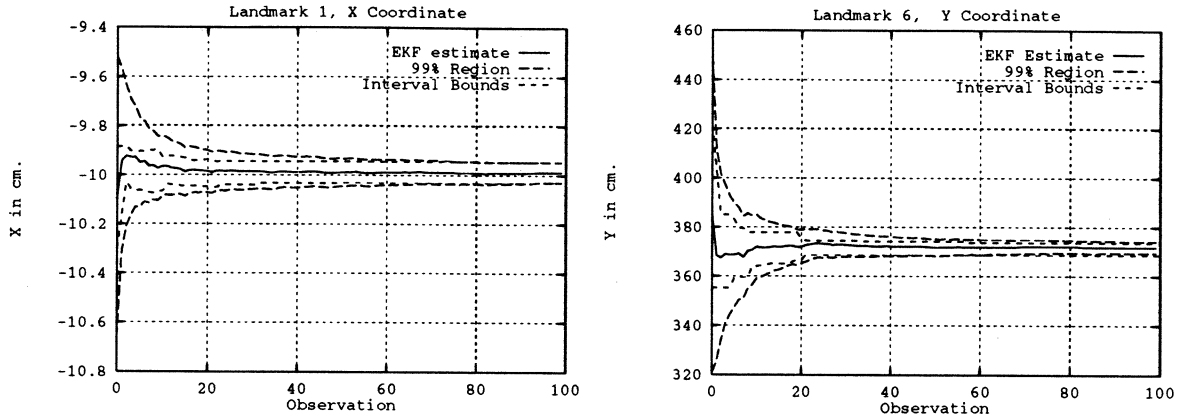


Figure 4: Tracking results for the EKF and set-based estimator.

Set-Based Formulation: Let $\mathcal{P}_i(k)$ denote all stereo locations consistent with the i th camera observations in a series. The stereo measurement equations are invertible, so that $\mathcal{P}_i(k)$ can be computed exactly and forms a diamond-shaped region represented by four points [2]. This region is simple to manipulate, so no set approximations are used for most calculations. Tracking validation takes place by projecting the solution set onto the image plane, and checking if the resulting intervals intersect the error interval about the observed data.

Comparison of Results: Figure 4 compares the results of estimating the x component of a landmark, and the y component of a second landmark using $\sigma = 0.5$ pixels. The center curve is the IEKF estimate. The outer curves are a 99% confidence interval from the IEKF and the interval bounds from the set-based method. In both cases, we see that the initial rate of convergence of the set-based method is more rapid than that of the IEKF, with the two bounds approaching each other after several 10's of observations. We have empirically observed that observation errors are often not zero mean. This introduces some estimation bias in the IEKF. This is seen in the plot on the right.

Tracking performance depends somewhat on correct tuning of the IEKF parameters and tracking tolerances. The parameter tested for data validity is a χ^2 variable with 2 degrees of freedom, so we parameterize ϵ_{track} by the probability, Q , of rejecting valid data. Successful tracking at the variance level described above occurred only when the probability of rejection is less than 0.25. This is a fairly large value considering that 100 observations were validated with *no* failures, leading us to believe that we have chosen a generous data variance. We note that the tracking problem is generally unambiguous, so correct tracking is not highly dependent on specific values of Q and σ^2 . We have found cases, however, where tracking is nearly ambiguous; in such cases, correctness is quite sensitive to changes in Q .

4.3 Correspondence

Problem: Match landmark positions in a global coordinate system to m estimated landmark positions.

Correspondence computation is based on the observation that three points form a triangle, and the lengths and angles of this triangle are invariant with respect to change of coordinates. Given

triangle parameters and some notion of confidence for triples of observed points and triples of points from the map, it is possible to compare invariant parameters to compute potential matches. These matches are further partitioned into consistent categories and the “best” category of those meeting certain minimal size criteria is chosen as the correspondence between stored and observed landmarks.

Statistical Method: Given the mean and covariance estimates of three points and the mapping $f(\cdot)$ from points to invariant parameters, the mean and approximate covariance of invariant parameters is computed using standard linearization techniques. The consistency of triples is checked using a Mahalanobis distance test parameterized by ϵ_{corr} . The variable tested is a χ^2 random variable with 3 degrees of freedom. As before, we parameterize ϵ_{corr} by Q , the probability of rejecting valid data.

Set-Based: The set of all invariant parameter values for a triple of points is computed using straightforward extremal analysis. Two vectors of parameters are consistent if their corresponding uncertainty sets overlap.

Comparison of Results: Figure 5 shows the results of the statistical matching algorithm at the observation positions as a function of observation variance and match confidence. The three numbers in each column are: the total number of matched triangles, the number of false negatives (missing matches) and the number of false positives (superfluous matches). Optimal results for both trials are 20, 0, and 0, respectively. Columns marked by (*) are where not all landmarks were correctly matched. This table shows that matching is correct for nearly all trials shown. Setting $Q = 0.1$ yields the best performance on data set 1, but does not work well on set 2. A compromise, setting σ between 0.3 and 0.5 pixels and Q smaller than 0.05, yields acceptable performance on both data sets. However, performance is quite sensitive to small changes in σ and Q , and the wide acceptance region makes the matching algorithm very sensitive to false positives. This may lead to failure in ambiguous situations. This data leads us to conclude: *there is no single set of parameters that works well over the entire range of operation*. By comparison, Figure 6 shows the results of the set-based matching method. We found that the single tolerance parameter set chosen from our time series analysis gives performance that usually equals or exceeds the performance of the tuned statistical method. Moreover, correct matching results were computed for all values from 0.3 to 0.7 on both data sets (with values of less than 0.3, stereo tracking ceases to function). A value of 0.4 yielded no false negatives or false positives in either data set. Hence, the algorithm appears to be robust to tolerance changes.

5 Conclusions

Clearly, the best estimation technique to use depends crucially on the characteristics of the problem. No single technique is “best” in all cases. The localization problem of the previous section is essentially a nonlinear hypothesis testing problem with bounded, statistically ill-behaved error. For this type of problem, we have found set-based methods to be more reliable and to require less heuristic tuning than comparable statistical techniques. More generally, our experience and results can be summarized as follows:

- The choice between statistical and set-based methods should be based on the character of sensor errors:

	Data Set 1				Data Set 2			
$Q \setminus \sigma$	0.3	0.5	.75	1.0	0.3	0.5	.75	1.0
0.1	20	20	20	20	5*	8	14	22
	0	0	0	0	15	12	6	1
	0	0	0	0	0	0	0	3
0.05	21	22	22	23	12	18	21	27
	0	0	0	0	8	3	1	0
	1	2	2	3	0	1	2	7
0.0005	23	24	26	26	22	28	30	35
	0	0	0	0	0	0	0	0
	3	4	6	6	2	8	10	15

Figure 5: Correspondence results for the statistical method.

Set \ Tol	0.3	0.4	0.5	0.55	0.6	0.7	0.8
Set 1	16	20	20	20	20	20	20
	4	0	0	0	0	0	0
	0	0	0	0	0	0	0
Set 2	8	20	22	25	27	30	81*
	12	0	0	0	0	0	0
	0	0	2	5	7	10	61

Figure 6: Correspondence results for the set-based method.

- When the observation system is invertible and sensing errors are bounded and not strongly centralized, set-based methods typically perform quite well and can be applied with practically no tuning.
- When the observation system is not highly nonlinear, errors can be characterized statistically, and the error distribution is strongly centralized and/or unbounded, statistical techniques will most likely perform well with little or no tuning.
- The precise tradeoffs between the two methods also depends on the ultimate application. An application requiring point estimates of moderate accuracy (where performance drops as the square of accuracy) will tend to prefer mean square estimation. An application interested in reliable bounds on parameters (performances drops as a 0-1 law) will tend to prefer set-based methods.
- Robustness issues are remarkably similar between the two methods. On problems where a method performs well, it can usually be made robust to a variety of deviations.

There are some cases where neither technique is desirable. For example, a strongly correlated, unbounded series of measurements with unmodeled statistical behavior will be difficult to treat with any method.

In many cases, set-based and statistical techniques can be easily and profitably combined. For low-dimensional problems, set-based techniques provide bounds on unknown parameters that can be used to accelerate statistical estimation performance. Such bounds are particularly useful in statistical problems where lower bounds on derivatives are useful in ensuring convergence of EKF-style estimators. In cases such as the robot localization problem where confidence sets are needed and nonlinear transformations are involved, set-based method may provide reliable set computation, while statistical methods can be applied to data that passes the test to get accurate point estimates.

We are currently investigating these ideas from both a theoretical and practical perspective in the framework of the system described above. We expect to have completed more empirical tests on the robot localization problem, soon. We also plan to investigate the theoretical properties of set-based hypothesis testing vs. statistical methods.

Acknowledgements: We wish to acknowledge Sami Atiya of the Fraunhofer Institute in Karlsruhe Germany for generously lending us the data used in Section 4 of this article. This research was supported by DARPA grant N00014-91-J-1577, by National Science Foundation grants IRI-9109116 and DDM-9112458, and by funds provided by Yale University. Sean Engelson is supported by a fellowship from the Fannie and John Hertz Foundation.

References

- [1] S. Atiya. *Zur Lokalisierung von mobilen Robotern mit Hilfe bildgebender Sensoren: Ein mengentheoretischer Ansatz*. draft, Fakultät für Informatik der Universität Karlsruhe, 1992. In German.
- [2] S. Atiya and G. D. Hager. Real-time vision-based robot localization. In *Proc. 1991 IEEE Int'l Conf. on Robotics and Automation*, pages 639–643. IEEE Computer Society Press, April 1991.
- [3] B. R. Barmish and J. Sankaran. The propagation of parametric uncertainty via polytopes. *IEEE Journal on Automatic Control*, AC-24(2):346–349, April 1979.
- [4] D. P. Bertsekas and I. B. Rhodes. Recursive state estimation for a set-membership description of uncertainty. *IEEE Journal on Automatic Control*, AC-16(2):117–128, April 1971.
- [5] R. Brooks. Symbolic reasoning among 3-D models and 2-D images. *Artificial Intelligence*, 17:285–348, 1981.
- [6] H. Durrant-Whyte. *Integration, Coordination, and Control of Multi-Sensor Systems*. Kluwer, Boston, MA, 1988.
- [7] R. Ellis. Geometric uncertainties in polyhedral object recognition. *IEEE Journal on Robotics and Automation*, 7(3):361–371, June 1991.
- [8] S. P. Engelson and D. McDermott. Error correction in mobile robot map learning. In *Proc. IEEE Int'l Conf. on Robotics and Automation*. Nice, France, May 1992.

- [9] A. Gelb, editor. *Applied Optimal Estimation*. MIT Press, Cambridge, MA, 1974.
- [10] W. Grimson. On the recognition of curved objects. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 11(6):632–643, June 1989.
- [11] W. Grimson, D. Huttenlocher, and D. Jacobs. A study of affine matching with bounded sensor error. In G. Sandini, editor, *Proc. European Conf. on Computer Vision*, pages 291–306. Springer Verlag, 1992.
- [12] G. D. Hager. *Task-Directed Sensor Fusion and Planning*. Kluwer, Boston, MA, 1990.
- [13] G. D. Hager. Constraint solving methods and sensor-based decision making. In *Proc. IEEE Int'l Conf. on Robotics and Automation*, pages 1662–1667. IEEE Computer Society Press, May 1992.
- [14] G. D. Hager and M. Mintz. Sensor modeling and robust sensor fusion. In *Proc. Fifth Int'l Symp. on Robotics Research*. MIT Press, Cambridge, MA, August 1989.
- [15] S. Maybank. Filter based estimates of depth. In *Proc. British Machine Vision Conf.*, pages 349–354. Oxford, UK, Sept. 24-27, 1990.
- [16] R. E. Moore. *Interval Analysis*. Prentice-Hall, Englewood Cliffs, N.J., 1966.
- [17] A. Sabater and F. Thomas. Set membership approach to the propagation of uncertain geometric information. In *Proc. 1991 IEEE Int'l Conf. on Robotics and Automation*, pages 2718–2723. IEEE Computer Society Press, Washington, D.C., 1991.
- [18] F. C. Schweppe. Recursive state estimation: Unknown but bounded errors and system inputs. *IEEE Journal on Automatic Control*, AC-13(1):22–28, February 1968.
- [19] R. H. Taylor and V. T. Rajan. The efficient computation of uncertainty spaces for sensor-based robot programming. In *Proc. IEEE Workshop on Intelligent Robots and Systems*. IEEE Press, November 1988.
- [20] H. S. Witsenhausen. Sets of possible states of linear systems given perturbed observations. *IEEE Journal on Automatic Control*, AC-13(5):556–558, October 1968.