Extensive Representations and Algorithms for Nonlinear Filtering and Estimation

Ethan Stump[1], Ben Grocholsky[2], and Vijay Kumar[1]

[1] GRASP Laboratory, University of Pennsylvania
   {estump,kumar}@seas.upenn.edu
[2] Robotics Institute, Carnegie Mellon University
   grocholsky@ri.cmu.edu

Summary. Most estimation problems in robotics are difficult because of (a) the non-linearity in observation models; and (b) the lack of suitable probabilistic models for the process and observation noise. In this paper we develop a set-valued approach to estimation that overcomes both these limitations and illustrates the application to localization of multiple, mobile sensor platforms with range sensors.

1 Introduction

Practical estimation tasks require us to deal with nonlinearities that are inherent in process dynamics and observation models. Common solutions to deal with such nonlinear state transition and measurement models require linearization of at least some portion of the problem. The concern is that linearization can lead to inconsistent error handling, removes the ability to directly represent ambiguous confidence sets, and can not be applied when the underlying state is unobservable.

Many techniques have been proposed to either avoid or delay linearization. Some of the most popular in recent years are sampling-based approaches such as Monte Carlo Localization, introduced by Fox *et al.* [9]. They rely on estimating a probability distribution for the system state, but instead of maintaining a simple parameterized distribution, which may require linearization, the distribution is discretely sampled to allow for arbitrary densities. Further refinements have allowed the fusion of sampled representations with standard parameterized ones to solve more challenging problems such as simultaneous localization and mapping tasks [10].

The opposite approach is to delay linearization by simply storing all measurement and state updates until a later time when the larger data base allows for the use of consistency to improve the estimates. The GraphSLAM algorithm presented by Thrun, Burgard, and Fox [4] is an example of such a technique; after several measurements have been taken, an estimate is formed by iteratively linearizing the state propagation and measurement equations and solving a least squares problem to maximize agreement with measurements and problem dynamics. Several different variations on this theme have been used by others such as Folkesson and Christensen [11] and Konolige [12].

The disadvantage of such delayed measurement integration is that best estimates are not available during data acquisition, making it impossible to knowledgeably improve the collection process. As a compromise, Thrun *et al.* [4] propose the Sparse Extended Information Filter to proactively incorporate each measurement as it is taken while being able to explicitly manage the information links between different entities formed through measurements and motion. The final picture allows for intuitive identification of how features are related but still requires linearization.

For many applications, such linearization may prove to be acceptable, but not for our current application of localization using range-only measurements. In the most general form of this problem, there is no sense of direction and so any attempt to linearize a range measurement will likely result in crippling inconsistency after further measurements and motion. Compelling sampling-based estimation approaches to this problem have been demonstrated by Djugash, Singh and Corke [5], but such implementations may require large numbers of samples making them computationally unattractive.

An important evolution in the methodology of Information-type filters was presented by Hanebeck [6]. The central idea is a nonlinear embedding, sometimes referred to as an over-parameterization, that maps the system states into a extended state space in a way such that the measurement equations become linear. The resulting framework lends itself to the application of the methodology introduced by Schweppe in the field of dynamic estimation under bounded noise [3]. Successful application of such set-based estimation techniques to static localization tasks involving range-only measurements and relative bearing measurements are demonstrated in [7]. While set-based techniques have been investigated by other researchers [14,15,16,17], Hanebeck's representation allows exact representations in the extended state space.

We build on Hanebeck's work and address the range-only localization problem that is frequently encountered in robotics. First, we present techniques that allow estimates of the actual state from the extended state representation. In addition, we present the first incorporation of dynamics into the framework, bringing the system closer to useful implementation on mobile robotics. The resulting filter is simple, robust, recursive, and avoids linearization. The forms are Information-like and so the filter behaves much like the SEIF [4] when it comes to identifying information links between different entities and watching how these links change through motion.

In Section 2, we present notation and equations for a mobile robot system using range-only measurements in Section 2. This work makes extensive use of ellipsoidal calculus, which we introduce in Section 3. The nonlinear transformation framework is presented and demonstrated in Section 4 and we present new techniques for approximate inversion of the filtered sets in Section 5. Finally, we introduce motion in Section 6 and discuss in Section 7.

2   Problem Formulation

We consider a mobile sensor network equipped with relative-range measurement capabilities with some capability of local sensing such as odometry or inertial

measurements. This network consists of n standard nodes that have either unknown or only partially known positions and m anchor nodes that have fully known positions with respect to some global reference frame. For convenience, we will assume these m nodes are stationary.

Expressed in the global frame, the position of the ith standard node is a variable, $\bar{x}_i = [\,x_i\ y_i\,]^T$, and the position of the lth anchor node is a constant, $\bar{a}_l$. The total state of the network is $\tilde{x} = [\,\bar{x}_1^T\ \dots\ \bar{x}_n^T\,]^T$, and belongs to the space $\mathcal{S} = \Re^{2n}$. A measurement between standard node i and standard node j has the form:

$$z_{ij} = h_{ij}(\tilde{x}) + e = \|\bar{x}_i - \bar{x}_j\| + e \tag{1}$$

while a measurement between standard node i and anchor node l has the form:

$$z_i^l = h_i^l(\tilde{x}) + e = \|\bar{x}_i - \bar{a}_l\| + e \tag{2}$$

These measurements have noise e; for the purposes of this paper we assume that this noise is bounded with constant bound ϵ. Thus $e \in [-\epsilon, \epsilon]$. These assumptions could be relaxed to include other models of bounded noise.

In a mobile sensor network, the n standard nodes can be considered to be attached to mobile robots. We adopt, for simplicity, a point model:

$$\bar{x}_{i,k+1} = \bar{x}_{i,k} + \bar{u}_{i,k}, \quad i = 1,\, 2,\, \dots\, n \tag{3}$$

where $\bar{u}_i$ is the control input for the ith mobile node at time k. The state of the system evolves discretely, with the dynamic transition from step k to $k+1$ given by (3), and a set of inter-node measurements taken at each step k:

$$\tilde{z}_k = \tilde{h}_k(\tilde{x}_k) + \tilde{e}_k$$

where $\tilde{h}_k$ is a combination of the measurement types expressed in (1) and (2).

3   Ellipsoids

Because we rely extensively on the results in [2], we now summarize their notation and definitions. We will use x, x_0 to denote the state and z to denote observations without worrying about the notation $(\bar{\cdot})$ or $(\tilde{\cdot})$ in this section.

An ellipsoid can be defined by two quantities: a vector specifying the position of its center, and a symmetric positive semi-definite matrix that encodes the directions and lengths of its semi-axes as the eigenvectors and eigenvalues respectively. Given $x_0 \in \Re^n$ and $E \in \mathcal{S}_+^n$, an n-dimensional ellipsoid is defined by the set:

$$\varepsilon_n(x_0, E) = \left\{\ x\ |\ (x - x_0)^T E (x - x_0) \le 1\ \right\} \tag{4}$$

If E is singular, then the resulting ellipsoid is degenerate and possesses directions, corresponding to the eigenvectors of the zero eigenvalues, where x is unconstrained. The center in this case is actually only a single representative point of the affine set at the center of the ellipsoid.

3.1 Fusion

Analogous to the fusion operation in sensor fusion, we define fusion for set valued estimates to be the operation that takes two ellipsoids and finds an ellipsoid that tightly bounds their intersection. The minimum-volume bounding ellipsoid can be found using iterative algorithms such as that of [13], but, as noted there, the complexity of this procedure is an open problem. However, the suboptimal approach taken in [2], repeated here, involves the minimization of a convex function over a bounded interval and so is simple and fast. Given two n-dimensional ellipsoids, $B_1 = \varepsilon_n(x_1, E_1)$ and $B_2 = \varepsilon_n(x_2, E_2)$, a one-parameter family of fusing ellipsoids is $\varepsilon_n^\lambda(x_0, E)$, $\lambda \in [0,1]$, defined by:

$$
\begin{aligned}
X &= \lambda E_1 + (1-\lambda)E_2 \\
k &= 1 - \lambda(1-\lambda)(x_2 - x_1)^T E_2 X^{-1} E_1 (x_2 - x_1) \\
x_0 &= X^{-1}(\lambda E_1 x_1 + (1-\lambda)E_2 x_2) \\
E &= \frac{1}{k}X
\end{aligned}
\tag{5}
$$

and the fused ellipsoid is taken as the ε_n^λ that has minimum volume. This amounts to either solving a bounded minimization problem over λ using the above, or finding the zero of the derivative of the volume as in Theorem 3 of [2]. This approximate intersection is denoted by $\tilde{\cap}$:

$$
\varepsilon_n(x_0, E) \leftarrow \varepsilon_n(x_1, E_1) \,\tilde{\cap}\, \varepsilon_n(x_2, E_2)
$$

3.2 Propagations

We have interest in two different ellipsoid propagations, both paralleling the state operations of our system given by (1), (2), and (3). Theorem 1 of [2] provides the general operation that is specialized to these two special cases.

We consider first a 1-dimensional ellipsoid associated with a single measurement (also seen as an interval) $\varepsilon_1(z, 1/\epsilon^2)$. If this measurement is obtained with a linear observation model, $z = H_{1 \times n} x$, its pre-image is the n-dimensional degenerate ellipsoid: $\varepsilon_n(H^\dagger z, (1/\epsilon^2)H^T H)$. $H^\dagger z$ can be any solution of the linear map, but we will use $H^\dagger$ as the pseudoinverse of H; note that the ellipsoid is rank 1 due to the form of its matrix. Thus there is an $n-1$ dimensional affine set of points that are consistent with this one dimensional observation.

Given an n-dimensional ellipsoid $\varepsilon_n(x_0, E)$, its image under the linear map $y = A_{n \times n} x + b_{n \times 1}$ is the n-dimensional ellipsoid: $\varepsilon_n(Ax_0 + b, AEA^T)$. This is the same expression encountered in propagation of Gaussian distributions in linear systems theory.

3.3 Slicing

It can be shown that the intersection of an n-dimensional ellipsoid $\varepsilon_n(x_0, E)$ with an m-dimensional ($m \leq n$) affine set in $\Re^n$, $\mathcal{A} = \{\, y_0 + Yc \mid c \in \Re^m \,\}$, produces an ellipsoid $\varepsilon_m(\eta, F')$, where:

$$F = Y^T E Y$$
$$\eta = F^\dagger Y^T E(x_0 - y_0)$$
$$k = 1 - (x_0 - y_0)^T E(x_0 - y_0) + \eta^T F \eta$$
$$F' = \frac{1}{k} F$$

(6)

If $k < 0$ then the affine set and the ellipsoid do not intersect.

3.4   Projection

An ellipsoid projection finds the "shadow" of an ellipsoid in some of its components. It can be shown that the projection of the n-dimensional ellipsoid $B = \varepsilon_n(x_0, E)$ onto its first m components is given by:

$$\mathcal{P}(B) = \varepsilon_m(x_{0,m}, E_{11} - E_{12} E_{22}^{-1} E_{12}^T)$$

(7)

where $x_{0,m}$ are the first m components of x_0, and $E = \begin{bmatrix} E_{11} & E_{12} \\ E_{12}^T & E_{22} \end{bmatrix}$, with E_{11} as an m-dimensional block.

Furthermore, if Y is a basis for any subspace of dimension m and Z is a basis for its null space of dimension $n - m$, then the ellipsoid $B = \varepsilon_n(x_0, E)$ projected onto this subspace is given by $B_Y = \varepsilon_m(Y^T x_0, E_Y)$, with E_Y given by:

$$E_Y = Y^T E Y - Y^T E Z (Z^T E Z)^{-1} Z^T E Y$$

(8)

The projection operation for ellipsoids is analogous to marginalization of multivariate Gaussians (see [4]).

4   Nonlinear Embedding

Hanebeck [6] introduced a novel framework involving a nonlinear embedding that maps the system states into an extended state space in such a way that the measurement equations become linear. The basic idea is shown in Figure 1.

Fig. 1. The state space $\mathcal{S}$ is extended by adding functionally dependent coordinates of $\mathcal{S}$ to create the extended state space $\mathcal{S}^\star = \mathcal{S} \oplus \tilde{\mathcal{S}}$

The *extended space* $\mathcal{S}^\star$ is formed by augmenting the *base space* $\mathcal{S}$ with additional dimensions. Define a smooth map $f : \mathcal{S} \to \tilde{\mathcal{S}}$. Recall that $\tilde{x}$ denotes elements of $\mathcal{S}$ and let $\tilde{x}^\star$ denote elements of $\mathcal{S}^\star$. The map f defines a $2n$-dimensional smooth sub-manifold in $\mathcal{S}^\star$ with coordinates $\tilde{x}$:

$$\tilde{x}^\star = g(\tilde{x}) = \begin{bmatrix} \tilde{x} \\ f(\tilde{x}) \end{bmatrix}$$

The Jacobian of g,

$$\frac{\partial g}{\partial \tilde{x}} = \begin{bmatrix} I \\ \frac{\partial f}{\partial \tilde{x}} \end{bmatrix}$$

is always full rank and the resulting manifold M is diffeomorphic to $\mathcal{S}$. Our goal is to choose f and transform the system equations in such a way as to make them linear in the p-dimensional extended space $\mathcal{S}^\star$ ($p > 2n$). Note that we will ultimately be interested only in those points that lie on the manifold, *i.e.*, $\tilde{x}^\star \in M$. The procedure is illustrated with an example next.

It should be noted that we have no automatic procedure for choosing f and instead rely on inspection of the system equations. However, Hanebeck [6] suggests that general polynomial bases such as Bernstein polynomials could be useful in this regard.

4.1   Application to Range Measurements

As an illustration, begin with the range measurement equation between two standard nodes:

$$z_{ij} - e = \|\bar{x}_i - \bar{x}_j\|$$

and square both sides. The left hand side represents the interval $[z_{ij} - \epsilon, z_{ij} + \epsilon]$ which, when squared using interval arithmetic, becomes $[z_{ij}^2 - 2z_{ij}\epsilon + \epsilon^2, z_{ij}^2 + 2z_{ij}\epsilon + \epsilon^2]$. By letting $z_{ij}^\star = z_{ij}^2 + \epsilon^2$ and $w \in [-2z_{ij}\epsilon, 2z_{ij}\epsilon]$, the transformed measurement equation is:

$$z_{ij}^\star = \bar{x}_i \cdot \bar{x}_i + \bar{x}_j \cdot \bar{x}_j - 2\bar{x}_i \cdot \bar{x}_j + w$$

If a different bounded-noise model is used for e, such a transformation is still possible as long as care is taken to ensure that the modified estimate and bounds are conservative.

Notice that this equation is nonlinear in the system variables $\bar{x}_i$ and $\bar{x}_j$ but is linear in the variables $\bar{x}_i \cdot \bar{x}_i$, $\bar{x}_j \cdot \bar{x}_j$, and $\bar{x}_i \cdot \bar{x}_j$:

$$z_{ij}^\star = [\,1\ 1\ -2\,] \begin{bmatrix} \bar{x}_i \cdot \bar{x}_i \\ \bar{x}_j \cdot \bar{x}_j \\ \bar{x}_i \cdot \bar{x}_j \end{bmatrix} + w$$

Applying this process to the range measurement equation between a standard node and an anchor node leads to:

$$z_i^{l,\star} - \bar{a}_l \cdot \bar{a}_l = [\,-2\bar{a}_l^T\ 1\,] \begin{bmatrix} \bar{x}_i \\ \bar{x}_i \cdot \bar{x}_i \end{bmatrix} + w$$

Accordingly, we define f so that:

$$f(\tilde{x}) = [\ldots, \ \bar{x}_i \cdot \bar{x}_i, \ \bar{x}_j \cdot \bar{x}_j, \ \bar{x}_i \cdot \bar{x}_j, \ \ldots]^T \, .$$

Thus $\mathcal{S}^\star$ is constructed by adding at most $n + {}^nC_2$ dimensions to $\mathcal{S}$, corresponding to all dot product combinations of the positions of the nodes that appear in the measurement equations. The measurement equations, after suitable modifications to the additive noise, are now linear in $\mathcal{S}^\star$ while having bounded noise.

We are not limited to range-only sensors. Indeed, the measurement equations for bearing-only sensors can also be made linear with a similar embedding (see [7]).

4.2   Recursive Filtering

We saw that an appropriate definition of the map f allows us to write each measurement equation at time step k in the form:

$$z_k^\star - w_k^\star = H_k^\star \tilde{x}_k^\star$$

where $\tilde{x}_k^\star \in \mathcal{S}^\star$. By viewing the interval quantity on the left-hand side as a 1-dimensional ellipsoid, we apply the results of Section 3.2 to define $\mathcal{Z}_k = \varepsilon_P((H_k^\star)^\dagger z_k^\star, (1/(w_k^\star)^2)(H_k^\star)^T H_k^\star)$ as the feasibility ellipsoid in $\mathcal{S}^\star$ consistent with this measurement.

Fig. 2. A measurement $\mathcal{Z}$, transformed into a set bounded by two hyperplanes, defines an interesting set in the base space $\mathcal{S}$ when intersected with M

Fig. 3. The feasibility set $\mathcal{X} \subset \mathcal{S}$ is found by intersecting an ellipsoid $\mathcal{E}^\star \subset \mathcal{S}^\star$ with the manifold M

Each $\mathcal{Z}_k$ can be seen as a pair of bounding hyperplanes constraining the possible embedded states. However, since only the values of $\mathcal{S}^\star$ lying in M have meaning, the actual set described by $\mathcal{Z}_k$ can be interesting, as portrayed in Figure 2. As more measurements are included, more bounds need to be incorporated. However, rather than tracking an increasing number of hyperplanes, each

new $\mathcal{Z}_k$ is incorporated into an aggregate state estimate ellipsoid $\mathcal{E}_k^\star$ using the fusion introduced in Section 3.1:

$$\mathcal{E}_k^\star \leftarrow \mathcal{E}_k^\star \tilde{\cap} \mathcal{Z}_k$$

When the filtering is started, $\mathcal{E}_k^\star$ can be initialized with $\varepsilon_p(0,0)$ to reflect the fact that nothing is known about the state. This is analogous to the initialization of the Information form of the Kalman filter (see [4]). Unlike the case of an extended Kalman filter, problematic estimate initialization procedures are not necessary.

4.3   Set Inversion

After performing filtering steps, the feasibility ellipsoid $\mathcal{E}_k^\star$ contains all $\tilde{x}_k^\star$ consistent with measurements up to step k, but not all of these elements have physical meaning. Only the $\tilde{x}_k^\star \in M$ actually represent the images of states in $\mathcal{S}$. This feasible set, $\mathcal{X}_k \subset \mathcal{S}$ is found by:

$$\mathcal{X}_k = \{\ \tilde{x} \in \mathcal{S}\ |\ g(\tilde{x}) \in \mathcal{E}_k^\star\ \}$$

The basic idea is shown in Figure 3.

This inversion can be carried out exactly using the implicit form (4) of $\mathcal{E}_k^\star = \varepsilon_p(x_0, E)$ together with g:

$$(g(\tilde{x}_k) - x_0)^T E(g(\tilde{x}_k) - x_0) \leq 1 \tag{9}$$

Any $\tilde{x}_k$ satisfying this implicit nonlinear inequality belong to the true feasibility set.

4.4   Single Robot Application

In order to demonstrate the operation of this framework, we present results for the simulated localization of one robot using range measurements to known anchors in Figure 4. The state space is given by a single pair $\tilde{x} = [\,x\ y\,]^T \in \mathcal{S} = \Re^2$ and is embedded into $\mathcal{S}^\star$ by
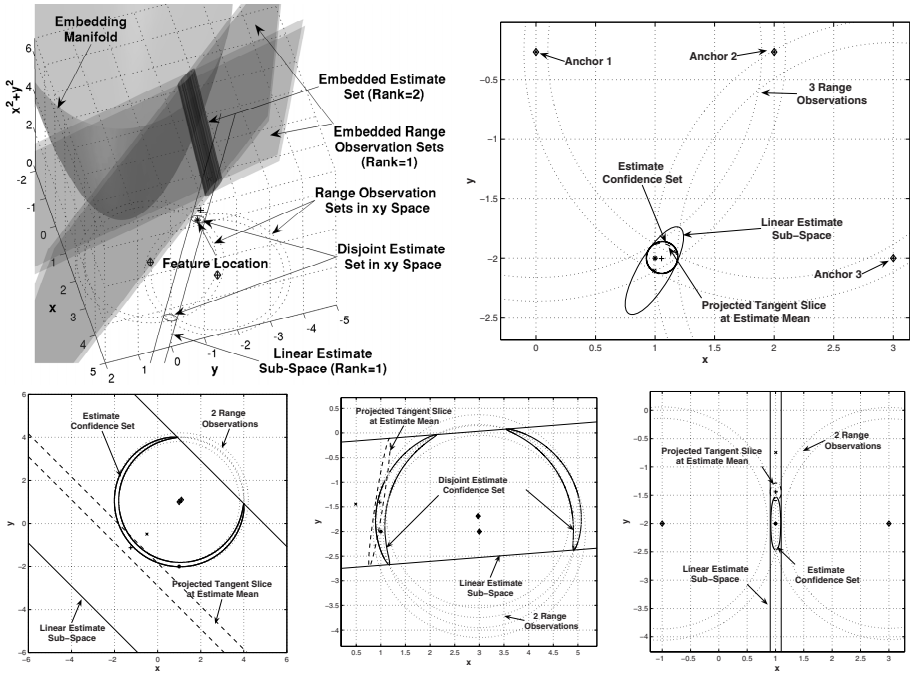
$$\tilde{x}^\star = [\,x,\ y,\ x^2 + y^2\,]^T. \tag{10}$$

The only measurements are from the single standard node to one of n anchor nodes at positions $\bar{a}_i$, $i \in \{1, \ldots, n\}$. The transformed measurement equation to anchor i is given by:

$$((z^i)^2 + \epsilon^2 - \bar{a}_i \cdot \bar{a}_i) + w^i = [-2\bar{a}_i^T\ 1]\,\tilde{x}^\star$$

where z^i is the measured range from the robot to the anchor, ϵ is the symmetric noise bound on this measurement, and $w^i \in [-2\epsilon z^i, 2\epsilon z^i]$.

Fig. 4. Visualization of the proposed embedding for a single robot range-only localization problem along with examples of possible estimate uncertainties represented by this formulation. In this case a closed form expression is available for the position estimate set. This set can take the form of an annulus or either a single or two disjoint transformed ellipses. The linear estimate sub-space and projected tangent slice results are discussed in Section 5.

5   Approximate Inversion

Since (9) describes the feasible set in $\mathcal{S}$ as a single nonlinear implicit inequality in $\dim(\mathcal{S})$ variables, it may not be suitable for finding feasible sets of large systems. We propose two methods for approximating this inversion: the first makes use of the projection idea introduced in Section 3.4 and is provably conservative, and the second makes use of the recognition of the embedded manifold M in an intuitive yet not provably conservative result. In order to demonstrate the effectiveness of these approximations, the simulation begun in Section 4.4 is scaled up to include more standard nodes, making it impossible to meaningfully use the exact inversion (9).

5.1   Base Projection

Working within an extended space, an ellipsoid $\varepsilon_p(x_0, E)$ can have its parameters split according to terms belong to the base space and the augmentations:

$$x_0 = \begin{bmatrix} x_{0,L} \\ x_{0,N} \end{bmatrix} \quad E = \begin{bmatrix} E_L & E_C \\ E_C{}^T & E_N \end{bmatrix}$$

with L, N, and C denoting linear, nonlinear, and coupling terms, respectively.

The exact inversion would be, using (9):

$$\mathcal{X}_k = \left\{ \tilde{x} \in \mathcal{S} \ \middle| \ \begin{bmatrix} \tilde{x} - x_{0,L} \\ f(\tilde{x}) - x_{0,N} \end{bmatrix}^T \begin{bmatrix} E_L & E_C \\ E_C{}^T & E_N \end{bmatrix} \begin{bmatrix} \tilde{x} - x_{0,L} \\ f(\tilde{x}) - x_{0,N} \end{bmatrix} \leq 1 \right\}$$

but using ellipsoid projection gives a conservative bound on $\mathcal{X}_k$. We also refer to this base projection process as the linear estimate sub-space in the figures. The validity of the process follows straight from the operation of projection, but a proof of the result serves to identify potentially important conditions:

Theorem 1 (Base Projection). *If E_N is invertible, then $\mathcal{X}_k \subset \varepsilon_{2n}(x_{0,L}, E_L - E_C E_N^{-1} E_C^T)$.*

Proof. First define $y \equiv \tilde{x} - x_{0,L}$ and $q(y) \equiv f(y + x_{0,L}) - x_{0,N}$. Then, points in $\mathcal{X}_k$ must satisfy:

$$y^T E_L y \leq 1 - 2 y^T E_C q(y) - q(y)^T E_N q(y)$$

The ellipsoid $\varepsilon_N(x_{0,L}, E_L - E_C E_N^{-1} E_C^T)$, once shifted, is defined by the inequality:

$$y^T (E_L - E_C E_N^{-1} E_C^T) y \leq 1$$

and so all points of $\mathcal{X}_k$ belong to this ellipsoid if:

$$1 - 2 y^T E_C q(y) - q(y)^T E_N q(y) - y^T E_C E_N^{-1} E_C^T y \leq 1$$

The maximum of the left hand side is found by the program:

$$\max_y \quad 1 - \begin{bmatrix} y \\ q(y) \end{bmatrix}^T \begin{bmatrix} E_C E_N^{-1} E_C^T & E_C \\ E_C{}^T & E_N \end{bmatrix} \begin{bmatrix} y \\ q(y) \end{bmatrix} \tag{11}$$

Call this matrix A; its Schur complement is $S = E_C E_N^{-1} E_C^T - E_C E_N^{-1} E_C^T = 0$. Using results from [1], $A \succeq 0$ since $S \succeq 0$, and so the program (11) has a finite maximum at 1. Thus, $\mathcal{X}_k \subset \varepsilon_n(x_{0,L}, E_L - E_C E_N^{-1} E_C^T)$. $\diamond$

The first question to ask is whether it can be expected that E_N will be full rank. Intuition suggests that this will be true after many measurements have been incorporated. Consider the fact that the nonlinear terms from f have only been introduced to correspond with terms in the measurement equations. After a full set of measurements have been taken, each nonlinear term will have shown up at least once in the measurement equations and so information will be known about it.

However, this intuition only makes sense in well-conditioned cases where there are several interconnected measurement that can serve to isolate the contributions of each nonlinear term to the final estimate. If this is not the case, then

further reference to [1] tells us that the claim $S \succeq 0 \implies A \succeq 0$ made above will be true if $(I - E_N E_N^\dagger)E_C = 0$.

These conditions only come about because we would like to completely ignore the details of the nonlinear transformation when trying to approximate the true set. If the simple tests fail, then we must take f into account, and the base projection is only conservative if it can be shown that:

$$\inf_y \quad \begin{bmatrix} y \\ q(y) \end{bmatrix}^T \begin{bmatrix} E_C E_N^\dagger E_C^T & E_C \\ E_C^T & E_N \end{bmatrix} \begin{bmatrix} y \\ q(y) \end{bmatrix} \geq 0$$

over the domain in question.

If the base projection is indeed conservative, then the coordinates of the true state must lie within it.

5.2   Tangent Slices

We now present an alternate technique that does take the contributions of the nonlinear transformation into account when approximating the true set inversion. By recognizing M as a manifold embedded into $\mathcal{S}^\star$ by g, the tangent space of M at a point $\tilde{x}^\star \in M$ can be found using the Jacobian of g.

Now $\frac{\partial g}{\partial \tilde{x}}$ and $\tilde{x}^\star$ define an affine set in $\mathcal{S}^\star$ with the same dimension as $\mathcal{S}$. By restricting $\mathcal{E}_k^\star$ to this set using the slicing operation of Section 3.3, an approximate representation is found in $\mathcal{S}$.

Use of this operation requires two steps: first finding a single point of M to use, and then calculating the Jacobian at this point and slicing $\mathcal{E}_k^\star$. The point could be chosen by an optimization procedure that sought to find the closest suitable point to the mean of $\mathcal{E}_k^\star = \varepsilon_p(x_0, E)$:

$$\min_{\tilde{x}} \quad \|g(\tilde{x}) - x_0\|_2$$

where x_0 could possibly be an affine set, $x_0 + Null(E)\lambda$, if $\mathcal{E}_k^\star$ has degenerate directions.

This optimization problem has the potential to be very nasty and the effects of choosing a non-optimal point are not known. As a simple heuristic, we have taken $\tilde{x}^\star = g(x_{0,L})$ as the point to linearize about.

The over- or under-estimation of the true set would seem to be intimately related with the curvature of M, but we have no proofs regarding the quality of this approximation, only demonstrations of its use in the following examples. Accordingly, the true state coordinates need not lie within the tangent slice estimate.

5.3   Multiple Robot Application

As a comparison of the set approximation techniques, we present simulation results of an experiment with two standard nodes and three anchor nodes in Figure 5. Both the base projection and tangent slicing approximations are compared against a brute force calculation found by griding the x, y coordinates of

both standard nodes and then checking for inclusion in $\mathcal{E}_k^\star$ by using the implicit form (9).

The system state consists of a two position vectors $\tilde{x} = [\bar{x}_1^T \ \bar{x}_2^T]^T = [x_1 \ y_1 \ x_2 \ y_2]^T \in \mathcal{S} = \Re^4$ and the state is transformed into $\mathcal{S}^\star$ according to:

$$\tilde{x}^\star = \left[\, \bar{x}_1, \bar{x}_2, \bar{x}_1 \cdot \bar{x}_1, \bar{x}_2 \cdot \bar{x}_2, \bar{x}_1 \cdot \bar{x}_2 \,\right]^T \tag{12}$$

As a representative example, the measurement between standard node 1 and anchor i would take the form:

$$((z_1^i)^2 + \epsilon^2 - \bar{a}_i \cdot \bar{a}_i) + w_1^i = \begin{bmatrix} -2\bar{a}_i^T \ 0_{1\times 2} \ 1 \ 0 \ 0 \end{bmatrix} \tilde{x}^\star$$

where z_1^i is the measured range from robot 1 to the anchor, ϵ is the symmetric noise bound on this measurement, and $w_1^i \in [-2\epsilon z_1^i, 2\epsilon z_1^i]$. Measurements to robot 2 take a similar form.

The inter-robot measurement is:

$$((z_{1,2})^2 + \epsilon^2) + w_{1,2} = \begin{bmatrix} 0_{1\times 2} \ 0_{1\times 2} \ 1 \ 1 \ -2 \end{bmatrix} \tilde{x}^\star$$

where $z_{1,2}$ is the measured range from robot 1 to robot 2, ϵ is the symmetric noise bound on this measurement, and $w_{1,2} \in [-2\epsilon z_{1,2}, 2\epsilon z_{1,2}]$.

When calculating the tangent slice approximation, the Jacobian of this system is:

Fig. 5. Two robot localization examples that illustrate the ability of the proposed representation to capture complex structured estimate uncertainty. Note that the base projection of the higher-dimensional uncertainty bounds a smaller estimate recoverable from the intersection of $\mathcal{E}_k^\star$ with M.

$$
\frac{\partial T}{\partial \tilde{x}} =
\begin{bmatrix}
1 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 \\
0 & 0 & 1 & 0 \\
0 & 0 & 0 & 1 \\
2x_1 & 2y_1 & 0 & 0 \\
0 & 0 & 2x_2 & 2y_2 \\
x_2 & y_2 & x_1 & y_1
\end{bmatrix}
$$

This example is taken further by including more standard nodes in Figure 6. The brute force representation of the feasible set is computationally daunting, so we present only the approximate representations.

6   Incorporating Motion

The static examples seen so far have shown the power of the representation, but have failed to take advantage of the dynamics of the problem. In applications, the incorporation of motion often transforms a poorly posed problem and makes it possible to estimate the state effectively. In [6], Hanebeck makes no use of dynamics, and further work in [7] and [8] also deals only with static systems. The main problem with incorporating dynamics into this framework is due to the fact that the chosen embedding renders the measurement equations linear but does not necessarily do so for the state update and, even worse, can make linear dynamics become nonlinear.

However, for the problem at hand, we will show that dynamic updates can in fact be incorporated using the fairly restrictive assumptions of a point model and perfect input. Despite these limitations, this inclusion represents a step forward for the theory and relaxations may be possible.

We consider a single robot once more and apply the transformation (10) to the dynamics given by (3). By inserting (3) into (10), we derive an expression relating the state at a given time step to the state at a previous time step:

$$
D
\begin{bmatrix}
x \\
y \\
x^2 + y^2
\end{bmatrix}
=
\begin{bmatrix}
x + u_x \\
y + u_y \\
(x + u_x)^2 + (y + u_y)^2
\end{bmatrix}
$$

$$
=
\begin{bmatrix}
1 & 0 & 0 \\
0 & 1 & 0 \\
2u_x & 2u_y & 1
\end{bmatrix}
\tilde{x}^\star +
\begin{bmatrix}
u_x \\
u_y \\
u_x^2 + u_y^2
\end{bmatrix}
$$

A similar derivation is possible for the two robot case as well, using (12):

$$
D
\begin{bmatrix}
\bar{x}_1 \\
\bar{x}_2 \\
\bar{x}_1 \cdot \bar{x}_1 \\
\bar{x}_2 \cdot \bar{x}_2 \\
\bar{x}_1 \cdot \bar{x}_2
\end{bmatrix}
=
\begin{bmatrix}
1 & 0 & 0 & 0 & 0\,0\,0 \\
0 & 1 & 0 & 0 & 0\,0\,0 \\
0 & 0 & 1 & 0 & 0\,0\,0 \\
0 & 0 & 0 & 1 & 0\,0\,0 \\
2u_{1,x} & 2u_{1,y} & 0 & 0 & 1\,0\,0 \\
0 & 0 & 2u_{2,x} & 2u_{2,y} & 0\,1\,0 \\
u_{2,x} & u_{2,y} & u_{1,x} & u_{1,y} & 0\,0\,1
\end{bmatrix}
\tilde{x}^\star +
\begin{bmatrix}
u_{1,x} \\
u_{1,y} \\
u_{2,x} \\
u_{2,y} \\
u_{1,x}^2 + u_{1,y}^2 \\
u_{2,x}^2 + u_{2,y}^2 \\
0
\end{bmatrix}
$$

Fig. 5. (*Continued*) Two additional range observations provide a unique estimate solution captured by this representation. In this case the projected linear subspace is unbounded.
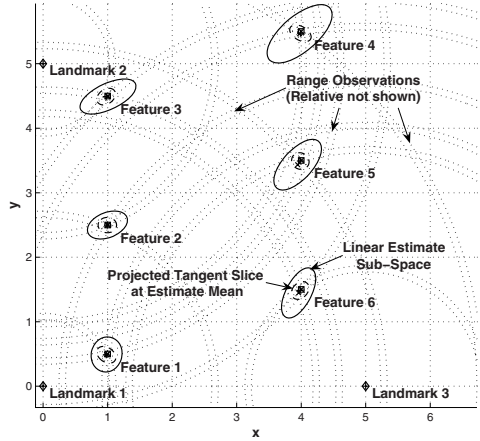
Fig. 6. Three landmarks are sufficient to yield a unique bounded estimate from the linear sub-space in this network localization example. The tangent slice yields much tighter estimates.
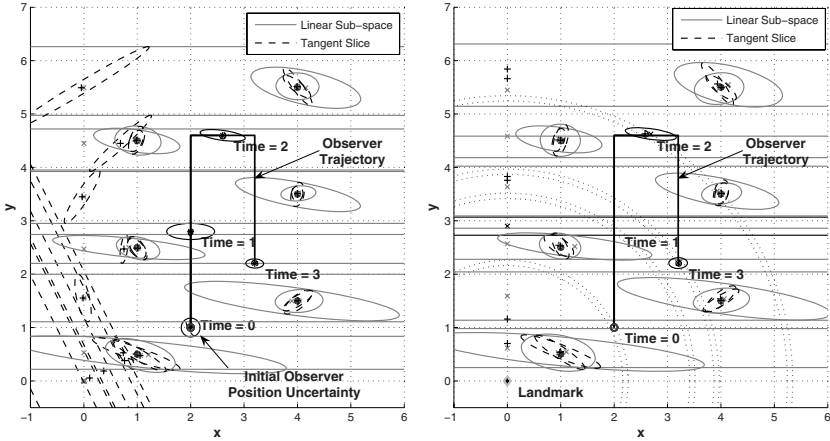

Fig. 7. Two solutions illustrating unique and consistent simultaneous localization and landmark mapping through robot motion: Given an initial robot location estimate *(left)*, and the aid of a known landmark given no prior estimate information *(right)*. A 27 state sparse linear filter fully captures the complex uncertainty structure. The linear sub-space approximation is unbounded when the true uncertainty set is unbounded or non-unique. This approximation becomes bounded in both cases once the robot has performed motion and measurements in two distinct directions.

This demonstrates that the process can extend to any number of point robots.

The resulting update equation now has a nonlinear coupling between the state and input appearing in the transition matrix, making it impossible to allow input disturbances under the current framework. If we allow for perfect input knowledge, then each dynamic update step is performed using the propagation results of Section 3.2 to find $\mathcal{E}_{k+1}^{\star}$. Simulations show that this approach works well for small time steps but may break down over larger intervals – the reasons for this have not been fully explored.

6.1   Application to Landmark Mapping

To demonstrate the effectiveness of incorporating motion into the filtering procedure, we present simulated results of an experiment with a standard node mounted on a mobile robot moving through a system of unlocalized static nodes. Two cases are presented in Figure 7. Global reference is provided by either giving the robot an initial position fix or observing a known landmark. Perfect inputs are applied in accordance with the dynamic update assumptions stated above. After the robot has completed multiple turns, the static nodes have been mapped.

7   Conclusion

We present a novel application of set-based estimation theory that lends itself to simultaneous localization and mapping with multiple mobile sensor platforms. While this paper focussed on range sensors, it is easy to include other types of sensors. The main advantage is our ability to incorporate sensors with nonlinear observation models without any knowledge of the noise.

While our approach shows robustness to modeling uncertainties and to initialization, there are two main limitations that we are currently addressing. First, the approach, as presented, is limited to Euclidean dynamic models. We are exploring alternative representations that will allow us to treat the dynamics and the observation as linear processes with additive noise. Second, while our extensive representation incorporates *all* relevant information in $\mathcal{S}^{\star}$ and makes checking data for consistency and correspondence very easy, it is computationally difficult to translate this information to the base space $\mathcal{S}$. This is because of the complexity of computing the intersection of the feasibility set with the manifold $\mathcal{M}$. However, because the underlying representation is algebraic, it is possible to use symbolic computation software for polynomial algebra to delineate this set and this is an area of ongoing investigation.

Finally, we are also addressing the control of vehicles to actively reduce the volume of the uncertainty with the set-valued representation discussed here. In [18] we present an experimental study with multiple robots localizing static features.

Acknowledgments

References

1. Boyd, S., Vandenberghe, L.: Convex optimization. University Press, Cambridge (2004)
2. Ros, L., Sabater, A., Thomas, F.: An ellipsoidal calculus based on propagation and fusion. IEEE Transactions on Systems, Man., and Cybernetics 32(4) (August 2002)
3. Schweppe, F.C.: Recursive state estimation: unknown but bounded errors and system inputs. IEEE Transactions on Automatic Control AC-13(1), 22–28 (1968)
4. Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics. The MIT Press, Cambridge (2005)
5. Djugash, J., Singh, S., Corke, P.: Further results with localization and mapping using range from radio. In: Proceedings, Fifth Int'l Conf. on Field and Service Robotics, Pt. Douglas, Australia (July 2005)
6. Hanebeck, U.D.: Recursive nonlinear set- theoretic estimation based on pseudo-ellipsoids. In: Int'l Conf. on Multisensor Fusion and Integration for Intelligent Systems, pp. 159–164 (2001)
7. Briechle, K., Hanebeck, U.D.: Localization of a mobile robot using relative bearing measurements. IEEE Transactions on Robotics and Automation 20(1), 36–44 (2004)
8. Horn, J., Hanebeck, U.D., Riegel, K., Heesche, K., Hauptmann, W.: Nonlinear set-theoretic position estimation of cellular phones. In: Proceedings of SPIE. AeroSense Symposium, vol. 5084, pp. 51–58 (2003)
9. Fox, D., Burgard, W., Dellaert, F., Thrun, S.: Monte Carlo localization: Efficient position estimation for mobile robots. In: Proceedings of the National Conference on Artificial Intelligence (AAAI), Orlando (1999)
10. Murphy, K.: Bayesian map learning in dynamic environments. In: Advances in Neural Information Processing Systems (NIPS), MIT Press, Cambridge (2000)
11. Folkesson, J., Christensen, H.I.: Graphical SLAM: A self-correcting map. In: Proceedings of the International Symposium on Autonomous Vehicles, Lisboa, PT (2004)
12. Konolige, K.: Large-scale map-making. In: Proceedings of the AAAI National Conference on Artificial Intelligence, San Jose, CA, pp. 457–463 (2004)
13. Yildirim, E.A.: On the Minimum Volume Covering Ellipsoid of Ellipsoids. Technical Report, Dept. of Applied Mathematics and Statistics, Stony Brook University (2005)
14. Atiya, S., Hager, G.D.: Real-time vision-based robot localization. IEEE Trans. on Robotics and Automation 9(6), 785–800 (1993)
15. Hanebeck, U.D., Schmidt, G.: Set theoretical localization of fast mobile robots using an angle measurement technique. In: Proc. IEEE Int. Conf. on Robotics and Automation, pp. 1387–1394 (1996)
16. Di Marco, M., Garulli, A., Giannitrapani, A., Vicino, A.: A set theoretic aproach to dynamic robot localization and mapping. Autonomous Robots 16, 23–47 (2004)
17. Spletzer, J., Taylor, C.: A bounded uncertainty approach to multi-robot localization. In: Proc. IROS, pp. 1258–1265 (2003)
18. Grocholsky, B., Stump, E., Shiroma, P., Kumar, V.: Control for Localization of Targets Using Range-Only Sensors. In: Proceedings of the International Symposium on Experimental Robotics, Rio de Janeiro, Brazil (2006)