A General Method for Approximating Nonlinear Transformations of Probability Distributions

Simon Julier and **Jeffrey K. Uhlmann**
Robotics Research Group
Department of Engineering Science
University of Oxford
Oxford, OX1 3PJ United Kingdom

1st November, 1996

Abstract

In this paper we describe a new approach for generalised nonlinear filtering. We show that the technique is more accurate, more stable, and far easier to implement than an extended Kalman filter. Several examples are provided, including the application of the new filter to problems involving discontinuous functions.

1 Introduction

Possibly the most important problem arising in tracking and control applications is the representation and maintenance of uncertainty. The state of a system, whether measured or estimated, is rarely known perfectly because (a) measuring instruments and processes have limited precision, and/or (b) estimates of evolving systems are based on process models that fail to include all governing parameters. The uncertainty associated with a state estimate can be represented most generally by a probability distribution incorporating all knowledge about the state. Because the amount of knowledge about the state is inherently finite, a complete parameterisation of the state probability distribution will also be finite. Unfortunately, measurements of an evolving system generally implies that the number of parameters necessary to specify the state probability distribution will increase without bound[8].

In order to permit tractable algorithms for tracking and control applications, an approximate state estimate must be generated. The most common approach is to maintain a fixed number of moments of the state distribution so as to limit the computational demands of the algorithm to the constraints of available resources. The most successful approach for fixed-moment estimation is the Kalman filter[7]. For linear systems the Kalman filter provides the optimal solution for maintaining a consistent estimate of the first two moments of the state distribution: the mean and variance. The Kalman filter (and its many variants) is the most widely used tracking and control algorithm both because of its mathematical rigour, and because most measurement and process models are inadequate to provide reliable information about higher order moments[9, 10].

The Kalman filter exploits the fact that (a) given only the mean and variance (or covariance in multiple dimensions) of a distribution, the most conservative assumption that can be made about the distribution is that it is a Gaussian having the given mean and variance, and (b) the

fact that the application of a linear operator to a Gaussian distribution always yields another Gaussian distribution. Given the assumptions of (a) and (b) it is straightforward to show that the Kalman filter will yield the best possible estimate of the mean and variance of the state. The requirement that the mean and variance of the state is measurable represents little practical difficulty, but the requirement that all observation and process models be linear is rarely satisfied in nontrivial applications.

In order to apply the mechanics of the Kalman filter to nonlinear problems, the *extended Kalman filter* (EKF) was developed[3]. The EKF is not so much an extension of the Kalman filter, but a crude approach for approximating nonlinear systems with linear ones. More specifically, the EKF simply calls for the replacement of every nonlinear transformation with a linearised approximation. In this article we examine an alternative generalisation of the Kalman filter that accommodates nonlinear transformations through the use of a new representation of the mean and variance information about the current state [15]. We will argue that the new approach is superior to the EKF in every important respect.

In the next section we establish the mechanics and notation associated with the Kalman filter, and we show the steps where the EKF deviates from the ordinary Kalman filter. In the subsequent sections we describe the new approach and its applications to filtering. The discussion references a number of simple examples and a detailed application will be given.

2 The Linear Filtering Paradigm

We seek the unbiased, minimum-mean squared error estimator of the state vector of the system of interest, $\mathbf{x}(k)$. The system evolves according to the discrete-time nonlinear state transition equation

$$\mathbf{x}(k+1) = \mathbf{f}[\mathbf{x}(k), \mathbf{u}(k+1), \mathbf{v}(k+1), k+1] \tag{1}$$

where $\mathbf{f}[\cdot,\cdot,\cdot,\cdot]$ is the process model, $\mathbf{x}(k)$ is the state of the system at timestep k, $\mathbf{u}(k+1)$ is the input vector and $\mathbf{v}(k+1)$ is a q-dimensional noise process.

The only information available about this system are its sequence of control inputs and a set of observations, which are acquired at discrete times. These observations are related to the state vector by the equation

$$\mathbf{z}(k+1) = \mathbf{h}[\mathbf{x}(k+1), \mathbf{u}(k+1), k+1] + \mathbf{w}(k+1), \tag{2}$$

where $\mathbf{z}(k+1)$ is the observation vector, $\mathbf{h}[\cdot,\cdot,\cdot]$ is the observation model which transforms the plant state space into observation space and $\mathbf{w}(k)$ is additive measurement noise.

We assume that the additive noise vectors, $\mathbf{v}(k)$ and $\mathbf{w}(k)$, are Gaussian and form uncorrelated white sequences: $\mathrm{E}[\mathbf{v}(k)] = \mathrm{E}[\mathbf{w}(k)] = \mathbf{0}$, for all k, and

$$\mathrm{E}\left[\mathbf{v}(i)\mathbf{v}^T(j)\right] = \delta_{ij}\mathbf{Q}(i), \tag{3}$$

$$\mathrm{E}\left[\mathbf{w}(i)\mathbf{w}^T(j)\right] = \delta_{ij}\mathbf{R}(i), \tag{4}$$

$$\mathrm{E}\left[\mathbf{v}(i)\mathbf{w}^T(j)\right] = \mathbf{0}, \quad \forall i, j. \tag{5}$$

The true system state vector is not known and must be estimated as it evolves through time. The representation of this uncertainty arises in that the estimate is a probability distribution conditioned on all prior observations and control inputs. However, determining the minimum

mean squared error is equivalent to finding the conditional mean. We use the notation $\hat{\mathbf{x}}(i \mid j)$ to be the state estimate at time i conditioned on all observations up to time j,

$$\hat{\mathbf{x}}(i \mid j) = \mathrm{E}\left[\mathbf{x}(i)|\mathbf{Z}^j\right] \tag{6}$$

where $\mathbf{Z}^j = [\mathbf{z}(1), \mathbf{z}(2), \ldots, \mathbf{z}(j)]^T$. The conditional covariance of this estimate is

$$\mathbf{P}(i \mid j) = \mathrm{E}\left[\{\mathbf{x}(i) - \hat{\mathbf{x}}(i \mid j)\}\{\mathbf{x}(i) - \hat{\mathbf{x}}(i \mid j)\}^T|\mathbf{Z}^j\right]. \tag{7}$$

In general it is extremely difficult to determine the values of these quantities and a linear estimator is often employed. This type of estimator incorporates observation information linearly but does *not* require that the process and observation models are linear themselves. As shown in [2], the minimum mean squared estimator has a "predictor-corrector" structure [2, 9]. First, the current state estimate and covariance are transformed through the state transition and observation equations. These quantities are also known as the one step ahead predictions since they represent an estimate of the state of the system at time $k+1$ given all observations up to time k. The system is observed at time $k+1$ and the observation information is used to "update" these predictions to produce the estimates $\hat{\mathbf{x}}(k+1 \mid k+1)$ and $\mathbf{P}(k+1 \mid k+1)$. The update equations are

$$\hat{\mathbf{x}}(k+1 \mid k+1) = \hat{\mathbf{x}}(k+1 \mid k) + \mathbf{W}(k+1)\nu(k+1) \tag{8}$$

and

$$\mathbf{P}(k+1 \mid k+1) = \mathbf{P}(k+1 \mid k) - \mathbf{W}(k+1)\mathbf{P}_{\nu\nu}(k+1 \mid k)\mathbf{W}^T(k+1). \tag{9}$$

The vector $\nu(k+1)$ is the innovation, which is equal to the difference between the actual observation at k, $\mathbf{z}(k+1)$, and the prior mean observation, $\hat{\mathbf{z}}(k+1 \mid k)$:

$$\nu(k+1) = \mathbf{z}(k+1) - \hat{\mathbf{z}}(k+1 \mid k). \tag{10}$$

The covariance of this quantity is $\mathbf{P}_{\nu\nu}(k+1 \mid k)$. $\mathbf{W}(k+1)$ is the Kalman gain and its value is given by

$$\mathbf{W}(k+1) = \mathbf{P}_{xz}(k+1 \mid k)\mathbf{P}_{\nu\nu}^{-1}(k+1 \mid k). \tag{11}$$

Within this framework of the estimator problem the accuracy of the prior means and covariances are the determining factors in estimator performance. The extended Kalman filter and new filter both use the same structure but employ different assumptions in determining the predicted means and covariances.

3 The extended Kalman filter

The EKF assumes that

$$\begin{aligned} \hat{\mathbf{x}}(k+1 \mid k)_{EKF} &= \mathbf{f}[\mathrm{E}\left[\mathbf{x}(k)|\mathbf{Z}^k\right], \mathrm{E}\left[\mathbf{v}(k+1)\right], \mathbf{u}(k+1), k+1] \\ &\approx \mathrm{E}\left[\mathbf{f}[\mathbf{x}(k), \mathbf{u}(k+1), \mathbf{q}_{q\times 1}, k+1]|\mathbf{Z}^k\right]. \end{aligned} \tag{12}$$

This is equivalent to the assumption that the estimated mean at the previous time step, $\hat{\mathbf{x}}(k \mid k)$, is approximately equal to the true state of the system at that time. Therefore, assuming that the size of the errors are small, the state dynamic model is expanded as a Taylor series about the estimate $\hat{\mathbf{x}}(k \mid k)$. By neglecting second and higher order terms, the state prediction propagates

through the nonlinear equations whilst the state errors propagate through a separate linear system. The predicted state and covariances are

$$\hat{\mathbf{x}}(k+1 \mid k)_{EKF} = \mathbf{f}[\hat{\mathbf{x}}(k \mid k), \mathbf{u}(k+1), k+1], \tag{13}$$

$$\mathbf{P}(k+1 \mid k)_{EKF} = \mathcal{J}_f \mathbf{P}(k \mid k) \mathcal{J}_f^T + \mathbf{Q}(k+1), \tag{14}$$

where $\mathcal{J}_f$ is the Jacobian matrix of the state transition equation $\mathbf{f}[\cdot,\cdot,\cdot]$ evaluated around $\hat{\mathbf{x}}(k \mid k)$ and $\mathbf{Q}(k+1)$ is the covariance of the dynamic driving noise injected during the transition from k to $k+1$.

By a similar process, the Taylor series for the observation equation is expanded about $\hat{\mathbf{x}}(k+1 \mid k)$ and is truncated at the first order. The predicted observation and observation covariances are

$$\hat{\mathbf{z}}(k+1 \mid k)_{EKF} = \mathbf{h}[\hat{\mathbf{x}}(k+1 \mid k), \mathbf{u}(k+1), k+1], \tag{15}$$

$$\mathbf{P}_{\nu\nu}(k+1 \mid k)_{EKF} = \mathcal{J}_h \mathbf{P}(k+1 \mid k) \mathcal{J}_h^T + \mathbf{R}(k+1). \tag{16}$$

The cross-covariance is

$$\mathbf{P}_{xz}(k+1 \mid k)_{EKF} = \mathbf{P}(k+1 \mid k) \mathcal{J}_h^T. \tag{17}$$

The form of the EKF is chosen to resemble that of the linear Kalman filter which is optimal for linear systems. However, the extended Kalman filter is generally suboptimal for nonlinear systems. The estimates of $\hat{\mathbf{x}}(k+1 \mid k)_{EKF}$, $\mathbf{P}(k+1 \mid k)_{EKF}$, $\hat{\mathbf{z}}(k+1 \mid k)_{EKF}$, $\mathbf{P}_{zz}(k+1 \mid k)_{EKF}$ and $\mathbf{P}_{xz}(k+1 \mid k)_{EKF}$ are all made on the assumption that the errors in truncating the Taylor series to the first order are small.

The impact of this assumption can be illustrated by considering the motion of a vehicle moving along a circular arc. In the process model we shall describe, the state space is the position and orientation of the vehicle, $\mathbf{x}(k) = [x(k), y(k), \psi(k)]^T$. The velocity of the vehicle is $V(k)$ and the radius of curvature is $R(k)$. The discrete-time process model for this system is

$$x(k+1) = x(k) + R(k)\left(\sin\left(\psi(k) + \frac{V(k)\Delta T}{R(k)}\right) - \sin\psi(k)\right) \tag{18}$$

$$y(k+1) = y(k) + R(k)\left(\cos\left(\psi(k) + \frac{V(k)\Delta T}{R(k)}\right) - \cos\psi(k)\right) \tag{19}$$

$$\psi(k+1) = \psi(k) + \frac{V(k)\Delta T}{R(k)}, \tag{20}$$

where these equations yield the exact solution for circular motion on an arc of constant radius.

It is assumed that the vehicle's speed $V(k)$ is disturbed by a zero-mean uncorrelated process. The result, as shown in Figure 1 which represents the condition at time k, is that the covariance ellipse in position uncertainty is orientated in the direction of travel.

Now consider the motion of the vehicle after it has turned through 90^o. Figure 2 shows the true position of the vehicle and its covariance ellipse at time $k+1$, one step ahead from the time at k. As can be seen, the covariance ellipse has been expanded and rotated: the initial uncertainty in position is augmented by additional uncertainties of vehicle motion across the arc and the change in rotation of the vehicle.

The EKF predicts the covariance forwards using Equation 14. The Jacobian matrix for the state transition equations is

$$\mathcal{J}_f = \begin{bmatrix} 1 & 0 & \Delta T V(k) \cos\psi(k) \\ 0 & 1 & -\Delta T V(k) \sin\psi(k) \\ 1 & 0 & 0 \end{bmatrix} \tag{21}$$

which is a constant velocity model tangential to the circle. The effect of this on covariance prediction is shown in Figure 3. While the mean position estimate is predicted forward using the circular motion model, the covariance is linearly projected in the initial direction of travel. This leads to a failure of the filter to maintain the critical information that the largest component of the uncertainty in the vehicle's position is in the direction in which it is travelling. In other words, the predicted covariance reflects the uncertainty about the vehicle in its previous state rather than its current state. This error must be compensated for by injecting additional dynamic noise via $\mathbf{Q}(k)$ which expands the predicted covariance ellipse so that the true covariance ellipse always lies within it. This is illustrated in Figure 4 which shows the true ellipse and one such adjusted EKF-predicted ellipse.

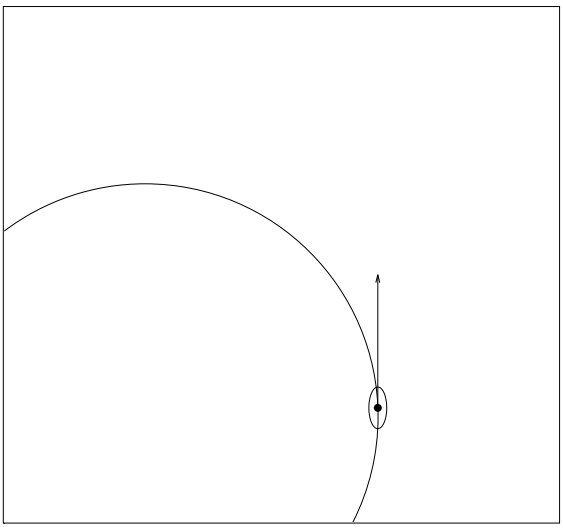

Fig. 1: Mean and covariance of a vehicle at time $t = m$

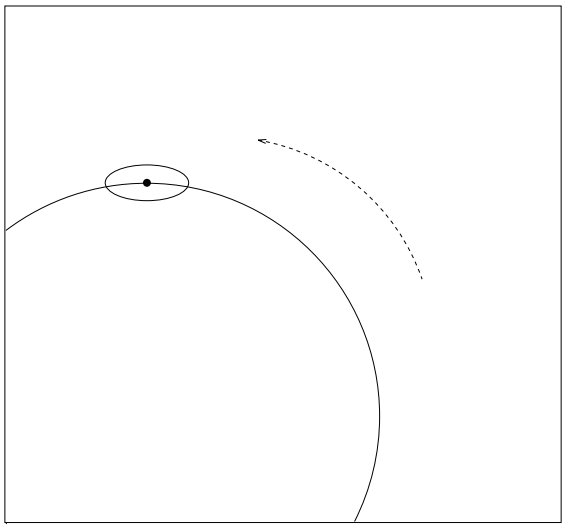

Fig. 2: True mean and covariance prediction to time $t = m + \Delta t$

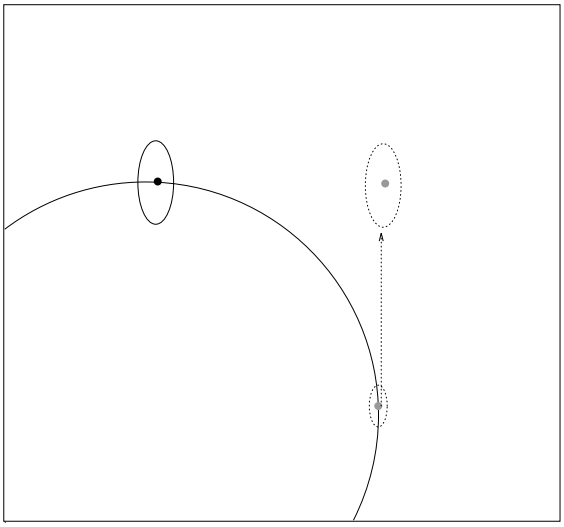

Fig. 3: EKF prediction of mean with linear covariance propagation

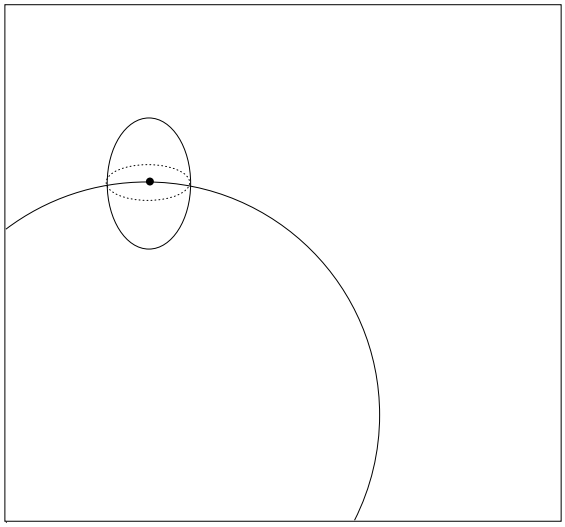

Fig. 4: EKF prediction "adjusted" to compensate for linearisation error

In summary, the failing of the EKF is its inability to make predictions of the system state, the observations and the associated covariance matrices when the system and/or observation models are non-linear. At the heart of this is the need to determine the mean and covariance of a posterior distribution which is acted upon by a non-linear transformation when the prior has a known mean and covariance and is assumed to be Gaussian. The EKF uses linearisations

- the non-linear transformation is approximated using a linear transformation which leads to appreciable errors for realistic observation models, process models, and their associated error covariances. We seek an alternative, tractable, general method for calculating these statistics without the need to use these linearising assumptions.

4 A General Method for Predicting Mean and Covariance

To state the general problem, we have an n-dimensional vector random variable $\mathbf{x}$ with mean $\bar{\mathbf{x}}$ and covariance $\mathbf{P}_{xx}$ and would like to predict the mean $\bar{\mathbf{y}}$ and covariance $\mathbf{P}_{yy}$ of a m-dimensional vector random variable $\mathbf{y}$ where $\mathbf{y}$ is related to $\mathbf{x}$ by the non-linear transformation

$$\mathbf{y} = \mathbf{g}\left[\mathbf{x}\right]. \tag{22}$$

In filtering there are two such transformations - $\mathbf{x}$ could be $\hat{\mathbf{x}}(k \mid k)$ and $\mathbf{y}$ is $\hat{\mathbf{x}}(k+1 \mid k)$ (for predicting the state) and $\mathbf{x}$ is $\hat{\mathbf{x}}(k \mid k)$ and $\mathbf{y}$ is $\hat{\mathbf{z}}(k+1 \mid k)$ for predicting the observation.

We begin with the following intuition: *With a fixed number of parameters it should be easier to approximate a Gaussian distribution than it is to approximate an arbitrary nonlinear function/transformation.* Following this intuition we wish to find a parameterisation which captures the mean and covariance information while at the same time permitting the direct propagation of the information through an arbitrary set of nonlinear equations. This can be accomplished by generating a discrete distribution having the same first and second (and possibly higher) moments, where each point in the discrete approximation can be directly transformed. The mean and covariance of the transformed ensemble can then be computed as the estimate of the nonlinear transformation of the original distribution.

Given an n-dimensional Gaussian distribution having covariance $\mathbf{P}$, we can generate a set of $O(n)$ points having the same sample covariance from the columns (or rows) of the matrices $\pm\sqrt{n\mathbf{P}}$ (the positive and negative roots). This set of points is zero mean, but if the original distribution has mean $\bar{\mathbf{x}}$, then simply adding $\bar{\mathbf{x}}$ to each of the points yields a symmetric set of $2n$ points having the desired mean and covariance. Because the set is symmetric its odd central moments are zero, so its first three moments are the same as the original Gaussian distribution. This is the minimal number of points capable of encoding this information. A random sampling of points from the distribution, on the other hand, will generally introduce spurious modes in the transformed distribution even if the set of sample points has the correct mean and covariance. In a filtering application these modes will take the form of high frequency noise that may completely obscure the signal.

We summarise the basic method as follows:

1. Compute the set $\boldsymbol{\sigma}$ of $2n$ points from the rows or columns of the matrices $\pm\sqrt{n\mathbf{P}}$. This set is zero mean with covariance $\mathbf{P}$. Compute a set of points with the same covariance, but with mean $\bar{\mathbf{x}}$, by translating each of the points as $\mathcal{X} = \boldsymbol{\sigma} + \bar{\mathbf{x}}$.

2. Transform each point as $\mathcal{Y}_i = \mathbf{g}\left[\mathcal{X}_i\right]$.

3. Compute $\bar{\mathbf{y}}$ and $\mathbf{P}_{yy}$ by computing the mean and covariance of the $2n$ points in the set $\{\mathcal{Y}_i\}$.

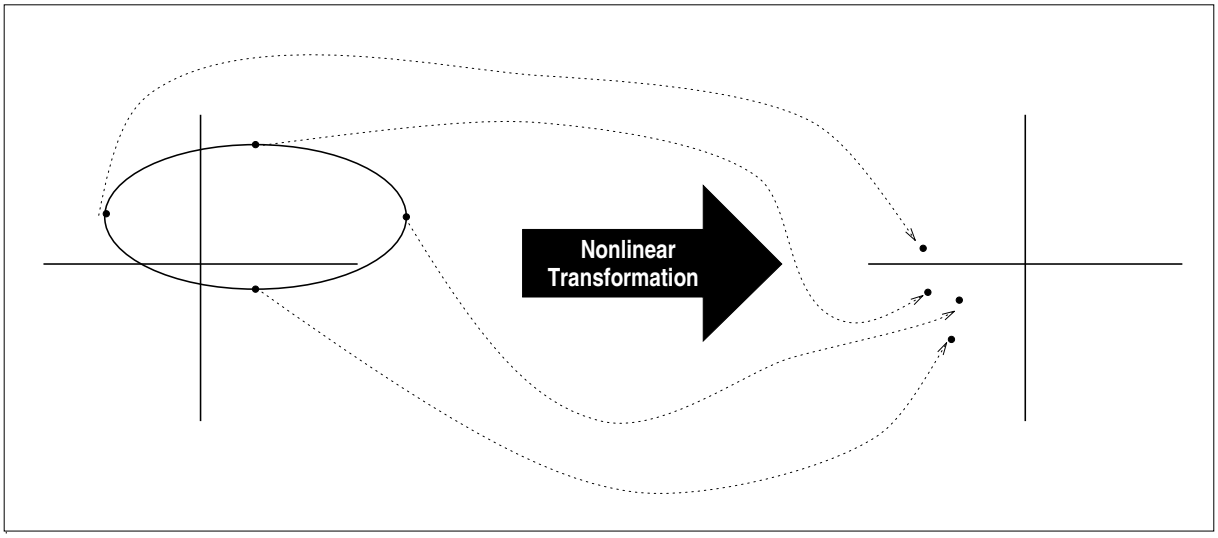

Fig.5: Sigma points capturing the mean and covariance of the distribution are transformed independently. The mean and covariance of the transformed sigma points define the predicted state.

This formulation can be generalised by exploiting our freedom to choose which of the infinite number of possible square roots we use[1] (including non-square matrix roots) and by our freedom to include any multiple of the mean, $\bar{\mathbf{x}}$, in the set $\boldsymbol{\sigma}_i$. Our freedom to choose an arbitrary matrix square root comes from the fact that any square root can be found from any other root by applying an orthonormal transformation [14]. If our original matrix is $\mathbf{A}$ then a matrix square root $\sqrt{\mathbf{A}}_1$ has the property that,

$$\sqrt{\mathbf{A}}_1 \sqrt{\mathbf{A}}_1^T = \mathbf{A}.$$

Hoewever, if we define a second matrix square root $\sqrt{\mathbf{A}}_2 = \sqrt{\mathbf{A}}_1 \mathbf{U}$ where $\mathbf{U}$ is an orthonormal transformation, then

$$\begin{aligned} \sqrt{\mathbf{A}}_2 \sqrt{\mathbf{A}}_2^T &= \left(\sqrt{\mathbf{A}}_1 \mathbf{U}\right)\left(\sqrt{\mathbf{A}}_1 \mathbf{U}\right)^T \\ &= \sqrt{\mathbf{A}}_1 \mathbf{U}\mathbf{U}^T \sqrt{\mathbf{A}}_1^T \\ &= \sqrt{\mathbf{A}}_1 \sqrt{\mathbf{A}}_1^T. \end{aligned}$$

We are not restricted to using orthogonal or symmetric matrix square roots which are numerically sensitive and computationally expensive to find. Rather, efficient and stable methods such as the Cholesky decomposition can be used, a vital consideration for real-time application.

Using multiple copies of the mean obviously will not affect the mean of the set, and will only affect the scaling factor for the calculation of the other points. The implications of this are discussed below and described in more detail in Appendix A.

To illustrate the difference between the new method and linearisation, consider again the example presented in the last section. The motion of the vehicle is a non-linear transform carried out over time. Figures 6 and 7 show how the new method projects the position of the rotating vehicle from the earlier example:

[1] If the matrix square root $\mathbf{A}$ is of the form $\mathbf{A}^T\mathbf{A}$, then the sigma points are formed from the *columns* of $\mathbf{A}$. However, for a root of the form $\mathbf{A}\mathbf{A}^T$ the *rows* of $\mathbf{A}$ are used.

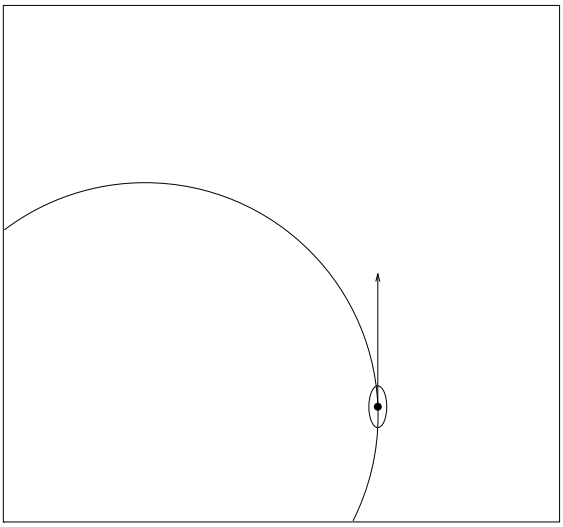

Fig.6: Vehicle travelling on a circular path with uncertainty in speed

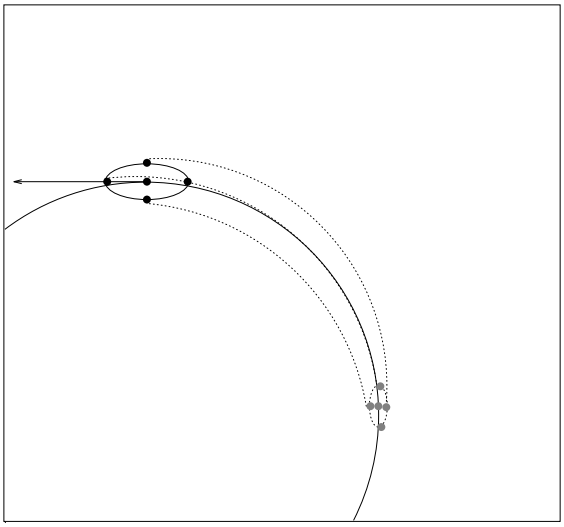

Fig.7: New filter correctly predicts effects of rotation on the covariance

Rather than project the mean and covariance through separate equations, the covariance ellipse is approximated by a discrete set of points as shown in Figure 6. As shown in Figure 7, each one of these points is separately projected along a circular path, and the final covariance ellipse is rotated and scaled.

We summarise the general method:

1. The set of translated sigma points is computed from the $n \times n$ matrix $\mathbf{P}_{xx}$ as

$$\begin{aligned} \boldsymbol{\sigma} &\longleftarrow \quad 2n \text{ rows or columns from } \pm\sqrt{(n+\kappa)\mathbf{P}_{xx}} \\ \mathcal{X}_0 &= \bar{\mathbf{x}} \\ \mathcal{X}_i &= \boldsymbol{\sigma}_i + \bar{\mathbf{x}}, \end{aligned}$$

   which assures that

$$\mathbf{P}_{xx} = \frac{1}{2(n+\kappa)} \sum_{i=1}^{2n} [\mathcal{X}_i - \bar{\mathbf{x}}][\mathcal{X}_i - \bar{\mathbf{x}}]^T.$$

2. The transformed set of sigma points are evaluated for each of the $0 - 2n$ points by

$$\mathcal{Y}_i = \mathbf{g}\left[\mathcal{X}_i\right] \tag{23}$$

3. The predicted mean is computed as

$$\bar{\mathbf{y}} = \frac{1}{n+\kappa} \left\{ \kappa \mathcal{Y}_0 + \frac{1}{2} \sum_{i=1}^{2n} \mathcal{X}_i \right\}. \tag{24}$$

4. And the predicted covariance is computed as

$$\mathbf{P}_{yy} = \frac{1}{n+\kappa} \left\{ \kappa [\mathcal{Y}_0 - \bar{\mathbf{y}}][\mathcal{Y}_0 - \bar{\mathbf{y}}]^T + \frac{1}{2} \sum_{i=1}^{2n} [\mathcal{Y}_i - \bar{\mathbf{y}}][\mathcal{Y}_i - \bar{\mathbf{y}}]^T \right\} \tag{25}$$

Comparing this algorithm with the linearisation algorithm we see a number of significant advantages:

- It is not necessary to calculate the Jacobian or make any other approximations of $\mathbf{g}[\cdot]$.
- The prediction stage only consists of standard linear algebra operations (matrix square roots, outer products, matrix and vector summation)
- The number of computations (including an efficient matrix square root algorithm) scales with dimensions at the same rate as linearisation.
- Constraints can be readily incorporated by applying the constraint to each of the projected set $\mathcal{Y}_i$.

In Appendix A we analyse of the performance of the new transformation algorithm against linearisation in detail. It is shown that the most natural framework to use to compare the two algorithms is a Taylor Series expansion of $\mathbf{g}[\cdot]$ evaluated about $\bar{\mathbf{x}}$. Linearisation introduces errors in the mean calculation at second order and in the covariance at the fourth order. The new method, however, yields errors in the mean and covariance which are both of fourth order. Further, on an absolute term-by-term basis, the errors at all higher orders can be made smaller than those introduced by linearisation. In many applications we expect the effects of the lower order terms to be significant, and so the reduction in errors can lead to significant improvements in estimation accuracy. This is demonstrated in a number of examples given below. When the function is discontinuous linearised estimates are almost incapable of capturing this information. If the discontinuity does not lie at the point of linearisation then the estimate does not contain the information. If the linearisation point lies on the discontinuity then the Jacobian matrix may not exist and hence the covariance cannot be predicted. The new filter uses a distribution of points and captures the effects of the discontinuity if it influences a significant proportion of the distribution.

The analysis also reveals the role which is played by κ: it affects the scaling of the fourth and higher moments of the distribution. Thus κ is a convenient parameter for exploiting knowledge (if available) about the higher moments of the given distribution [2]. It is shown that choosing $\kappa = 2$ for a scalar system leads to errors in the mean and covariance which are sixth order. For multi-dimensional systems choosing $\kappa = 3 - n$ minimises the mean squared error up to the fourth order. There is no restriction on the sign of κ but, if κ is negative, then we cannot interpret the distribution of the sigma points as a probability distribution. Further, when κ is negative there is the possibility, as with all approximation algorithms, that the predicted covariance will be non-positive semi-definite. In this case, we can use a modified form of the algorithm for calculating the covariance:

$$\mathbf{P}_{yy} = \frac{1}{2(n+\kappa)} \sum_{i=1}^{2n} [\mathcal{Y}_i - \bar{\mathbf{y}}][\mathcal{Y}_i - \bar{\mathbf{y}}]^T. \tag{26}$$

Both the original and modified algorithms have the property that, as $n + \kappa$ tends to zero, the mean tends to the value obtained by a truncated second order prediction algorithm. Further, the modified algorithm has the useful property that, in this limit, the covariance tends towards that calculated by linearisation[3].

[2] This leads us to the conclusion that the method is applicable for any prior, symmetric, unimodal distribution and not just Gaussian.

[3] This can be contrasted with an alternative approach of the initial intuition which was explored in [13]. Under that scheme no copies of the previously estimated mean are included in the sample set and the sigma points are scaled using a parameter α. In the limit as α tends to infinity this algorithm predicts the same mean and covariance as the EKF. However when $\alpha = 1$, this method estimates the same mean and covariance as that of the new filter but with $\kappa = 0$.

Any choice of κ yields MMSE results for linear $\mathbf{g}[\cdot]$.

To show the significant improvements which can be made using the new method instead of linearisation, we consider the two following simple examples.

4.1 Example 1: a simple continuous transform

Suppose there is a one-dimensional random y which is related to the random variable x by the nonlinear transformation

$$y = \mathbf{g}[x] = x^2. \tag{27}$$

Given that x is a normally distributed variable with mean $\bar{x}$ and covariance σ_x^2, what is the mean and covariance of y, $\bar{y}$ and σ_y^2?

First consider the "true" situation (denoted by the subscript T). The true random variable x_T is written as

$$x_T = \bar{x} + \delta x, \tag{28}$$

which is the sum of the mean $\bar{x}$ and δx, a zero mean, normal random variable with covariance σ_x^2. The series expansion for one realisation of δx is

$$\mathbf{g}[\bar{x} + \delta x] = \bar{x}^2 + 2\bar{x}\delta x + (\delta x)^2. \tag{29}$$

Taking expectations,

$$\bar{y}_T = \bar{x}^2 + \sigma_x^2. \tag{30}$$

The mean squared error in this realisation is

$$\begin{aligned} (y - \bar{y})_T^2 &= (\bar{x}^2 + 2\bar{x}\delta x + (\delta x)^2 - \bar{x}^2 - \sigma_x^2)^2 \\ &= (2\bar{x}\delta x + \delta x^2 - \sigma_x^2)^2 \\ &= (\delta x)^4 + 4\bar{x}(\delta x)^3 + (4\bar{x}^2 - 2\sigma_x^2)(\delta x)^2 - 4\sigma_x^2\bar{x}\delta x + \sigma_x^4. \end{aligned} \tag{31}$$

Taking expectations of this gives the true covariance,

$$(\sigma_y^2)_T = \mathrm{E}\left[(\delta x)^4\right] + 4\bar{x}^2\sigma_x^2 - \sigma_x^4 \tag{32}$$

where the first term is the kurtosis of the series. From moment generating functions it can be shown to have the value $3\sigma_x^4$ [11]. Therefore, the true covariance is

$$(\sigma_y^2)_T = 2\sigma_x^4 + 4\bar{x}^2\sigma_x^2. \tag{33}$$

The linearisation algorithm predicts its mean from Equation 14

$$\bar{y}_{LIN} = \mathbf{G}[\bar{x}] = \bar{x}^2 \tag{34}$$

and covariance from Equation 14

$$(\sigma_y^2)_{LIN} = 4\bar{x}^2\sigma_x^2. \tag{35}$$

Comparing these equations with those for the true system, it can be seen that the linearisation assumptions eliminate a number of significant terms in the mean and covariance. This leads to a biased mean and an under prediction in the value of the covariance.

Now consider the new algorithm. Using the steps given earlier, the position of the sigma points are determined first. Since the system is one dimensional there are only three points of interest: the two sigma points and the mean. Denoting the ith point by $\mathcal{X}_i$, these points lie at

$$\{\mathcal{X}_0, \mathcal{X}_1, \mathcal{X}_2\} = \{\bar{x}, \bar{x} - \sigma, \bar{x} + \sigma\}, \tag{36}$$

where $\sigma = \sqrt{(n+\kappa)\sigma_x^2}$ and $n = 1$. Projecting these points through $\mathbf{g}[\cdot]$ according to Equation 23 yields a new set of points $\mathcal{X}_i\prime$ that lie at

$$\{\mathcal{X}_0\prime, \mathcal{X}_1\prime, \mathcal{X}_2\prime\} = \left\{\bar{x}^2, \bar{x}^2 - 2\bar{x}\sigma + \sigma^2, \bar{x}^2 + 2\bar{x}\sigma + \sigma^2\right\}. \tag{37}$$

Applying Equation 24, the mean is calculated as

$$\begin{aligned} \bar{y} &= \frac{1}{2(1+\kappa)}\left(2\kappa\bar{x}^2 + 2\bar{x}^2 + 2(1+\kappa)\sigma_x^2\right) \\ &= \bar{x}^2 + \sigma_x^2. \end{aligned} \tag{38}$$

As can be seen, the predicted mean is equal to the true mean and is independent of κ. This result arises because of the properties of the new filter's projection equations. As explained earlier, the mean and covariance are correctly predicted up to and including the second order terms. Both errors and the scaling effects of κ begin to act on the fourth order. Since the mean is only a function of the first two orders, none of these properties of higher order moments affect the results. The true covariance, however, is a function of the kurtosis of the distribution and we expect the estimated value to contain a term related to κ. Since the state space is one dimensional, we use the original form of the covariance prediction equation (Equation 25):

$$\begin{aligned} \sigma_y^2 &= \frac{1}{2(1+\kappa)}\left(\sum_{i=1}^{2n}\{\mathcal{X}_i\prime - \bar{y}\}^2 + 2\kappa\sigma_x^4\right) \\ &= \kappa\sigma_x^4 + 4\bar{x}^2\sigma_x^2. \end{aligned} \tag{39}$$

To find a solution the value of κ must be specified. The kurtosis of the true distribution is $3\sigma_x^4$ and that of the sigma points is σ_x^2. Since the kurtosis of the sigma points is scaled by an amount $(1+\kappa)$, the kurtosis of both distributions only agree when $\kappa = 2$. By substituting in the above equation, this result is confirmed.

4.2 Example 2: a simple discontinuous system

We now consider an example of a discontinuous process model in which we wish to estimate the mean position of a projectile, $[x(k), y(k)]^T$. The projectiles are initially released at time 1 from a random position $[x(1), y(1)]^T$ and travel at a constant and known speed v_x in the x direction. The path of the projectiles is obscured by a wall. If a projectile hits the wall there is an elastic, instantaneous impact, and the projectile is reflected back on itself at the same velocity as it travels forwards. This situation is illustrated in the figure which also shows the covariance ellipse of the initial distribution. We wish to estimate the mean position and covariance of the position at time 2, $[x(2), y(2)]^T$, where $\Delta T \triangleq t_2 - t_1$.

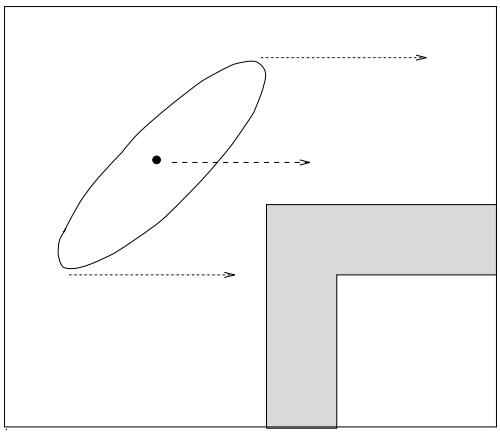

Fig.9: Part of the distribution reflects from the wall while the remainder passes unimpeded.

The process model for this system is

$$x(2) = \begin{cases} x(1) + \Delta T v_x & y(1) \geq 0, \\ \Delta T v_x - x(1) & y(1) < 0 \end{cases} \quad (40)$$

$$y(2) = y(1). \quad (41)$$

At time 1 the estimated is $[x(1), y(1)]^T$. The error in this estimate is Gaussian, zero mean and has covariance $\mathbf{P}(1 \mid 1)$. The true conditional mean and covariance was determined using Monte Carlo simulation for different choices of the initial mean of y. The figure shows that the new filter estimates the mean very closely, suffering only small spikes as the translated sigma points successively pass the wall. Further analysis shows that the covariance for the filter is only slightly larger than the true covariance, but conservative enough to account for the deviation of its estimated mean from the true mean. The linearised method, however, bases its entire estimate of the conditional mean on the projection of the prior mean, so its estimates bear no resemblance to the true mean except when most of the distribution either hits or misses the wall.

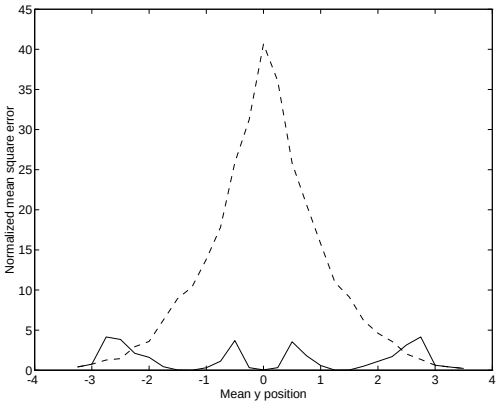

Fig.10: Normalised mean square error of the EKF (dashed line) and the new filter (solid line).

As can be seen, in each case the new method gives significant improvements over linearisation. In the first example, the method is exact. In the second, the errors are much smaller.

5 The New Filter

In the previous section we presented a method for determining the mean and covariance of a distribution which is superior to linearisation in many important respects. In this section we return to our original motivation and use this methodology to derive a filtering algorithm. From the discussion in Section 2, the following steps must be carried out:

- predict the new state of the system $\hat{\mathbf{x}}(k+1 \mid k)$ and covariance $\mathbf{P}(k+1 \mid k)$. This prediction must include the effects of the process noise vector $\mathbf{v}(k+1)$,
- predict the expected observation $\hat{\mathbf{z}}(k+1 \mid k)$ and the innovation covariance $\mathbf{P}_{\nu\nu}(k+1 \mid k)$, and
- predict the cross-correlation matrix $\mathbf{P}_{xz}(k+1 \mid k)$.

As will be shown, carrying out the latter two steps is trivial once we have established the first step. In the general formulation, $\mathbf{v}(k+1)$ can enter in a very flexible fashion. Components of it can, for example, reflect additive noises on $\hat{\mathbf{x}}(k \mid k)$, $\mathbf{u}(k)$ or even on the time index k. Further, the noise can be injected in a non-linear fashion as, for example, multiplicative terms. In the face of such generality, it is not sufficient to treat the noise as separate, additional terms. Further, if we follow the route of the EKF and simply add a process noise covariance term then we do not account for the effect of the process noise on the mean. However, with the new method the

effects of process noise can be incorporated in a very simple fashion. We define an augmented $(n+q)$ dimensional state vector $\hat{\mathbf{x}}^a(k \mid k)$ where

$$\hat{\mathbf{x}}^a(k \mid k) \triangleq \left(\frac{\mathbf{x}(k)}{\mathbf{v}(k)} \right) \tag{42}$$

This augmented vector has mean

$$\mathrm{E}\left[\hat{\mathbf{x}}^a(k \mid k)\right] = \left(\frac{\hat{\mathbf{x}}(k \mid k)}{\mathbf{0}_{q\times 1}} \right) \tag{43}$$

and covariance

$$\mathbf{P}^a(k|k) = \begin{bmatrix} \mathbf{P}(k \mid k) & \mathbf{P}_{xw}(k \mid k) \\ \mathbf{P}_{wx}(k \mid k) & \mathbf{Q}(k) \end{bmatrix}, \tag{44}$$

where, in general, $\mathbf{P}_{xw}(k \mid k)$ is the correlation between the injected noise and the error in the current state estimate.

The non-linear transformation Equations 23 to 25 are now used on the $2(n+q)$ sigma points from $\hat{\mathbf{x}}^a(k \mid k)$. Repeating the equations here for clarity,

1. The set of translated sigma points is computed from the $(n+q)\times(n+q)$ matrix $\mathbf{P}^a(k|k)$ as

$$\begin{aligned} \boldsymbol{\sigma}^a(k|k) &\longleftarrow \quad 2(n+q) \text{ rows or columns from } \pm\sqrt{(n+q+\kappa)\mathbf{P}^a(k|k)} \\ \mathcal{X}_0(k \mid k) &= \hat{\mathbf{x}}^a(k \mid k), \\ \mathcal{X}_i(k \mid k) &= \boldsymbol{\sigma}^a_i(k|k) + \hat{\mathbf{x}}^a(k \mid k), \end{aligned}$$

   which assures that

$$\mathbf{P}^a(k|k) = \frac{1}{2(n+q+\kappa)} \sum_{i=1}^{2(n+q)} [\mathcal{X}_i(k \mid k) - \hat{\mathbf{x}}(k \mid k)][\mathcal{X}_i(k \mid k) - \hat{\mathbf{x}}(k \mid k)]^T.$$

2. The transformed set of sigma points are evaluated for each of the $0 - 2(n+q)$ points by

$$\mathcal{X}_i(k+1 \mid k) = \mathbf{f}\left[\mathcal{X}_i(k \mid k), \mathbf{u}(k+1), k\right] \tag{45}$$

3. The predicted mean is computed as

$$\hat{\mathbf{x}}(k+1 \mid k) = \frac{1}{n+q+\kappa} \left\{ \kappa \mathcal{X}_0(k+1 \mid k) + \frac{1}{2} \sum_{i=1}^{2(n+q)} \mathcal{X}_i(k+1 \mid k) \right\}. \tag{46}$$

4. And the predicted covariance is computed as

$$\begin{aligned} \mathbf{P}(k+1|k) &= \frac{1}{n+q+\kappa} \Big\{ \kappa [\mathcal{X}_0(k+1 \mid k) - \hat{\mathbf{x}}(k+1 \mid k)][\mathcal{X}_0(k+1 \mid k) - \hat{\mathbf{x}}(k+1 \mid k)]^T \\ &\quad + \frac{1}{2} \sum_{i=1}^{2(n+q)} [\mathcal{X}_i(k+1 \mid k) - \hat{\mathbf{x}}(k+1 \mid k)][\mathcal{X}_i(k+1 \mid k) - \hat{\mathbf{x}}(k+1 \mid k)]^T \Big\} \end{aligned} \tag{47}$$

Using these equations ensures that the prediction with uncertainty in the state and process noise yields estimation errors in fourth order and above.

To complete the description of the new filter, the equivalent statistics for the innovation sequence and the cross correlation must be determined. Instantiating each point through the observation model to yield $\mathcal{Z}_i(k+1 \mid k) = \mathbf{h}[\mathcal{X}_i(k+1 \mid k), \mathbf{u}(k+1), k+1]$, the mean observation is found from

$$\hat{\mathbf{z}}(k+1 \mid k) = \frac{1}{n+q+\kappa} \left\{ \kappa \mathcal{Z}_0(k+1 \mid k) + \frac{1}{2} \sum_{i=1}^{2(n+q)} \mathcal{Z}_i(k+1 \mid k) \right\}. \tag{48}$$

and the covariance is determined from

$$\begin{aligned} \mathbf{P}_{zz}(k+1 \mid k) &= \frac{1}{n+q+\kappa} \Big\{ \kappa [\mathcal{Z}_0(k+1 \mid k) - \hat{\mathbf{z}}(k+1 \mid k)][\mathcal{Z}_0(k+1 \mid k) - \hat{\mathbf{z}}(k+1 \mid k)]^T \\ &\quad + \frac{1}{2} \sum_{i=1}^{2(n+q)} [\mathcal{Z}_i(k+1 \mid k) - \hat{\mathbf{z}}(k+1 \mid k)][\mathcal{Z}_i(k+1 \mid k) - \hat{\mathbf{z}}(k+1 \mid k)]^T \Big\}. \end{aligned} \tag{49}$$

The innovation covariance is equal to the sum of $\mathbf{P}_{zz}(k+1 \mid k)$ and the observation noise covariance matrix,

$$\mathbf{P}_{\nu\nu}(k+1 \mid k) = \mathbf{P}_{zz}(k+1 \mid k) + \mathbf{R}(k+1). \tag{50}$$

Finally, noting that the additive disturbances $\mathbf{w}(k)$ and $\mathbf{v}(k)$ are uncorrelated, the cross correlation matrix is

$$\begin{aligned} \mathbf{P}_{xz}(k+1 \mid k) &= \frac{1}{n+q+\kappa} \Big\{ \kappa \left[\mathcal{X}_0(k+1 \mid k) - \hat{\mathbf{x}}(k+1 \mid k)\right] \left[\mathcal{Z}_0(k+1 \mid k) - \hat{\mathbf{z}}(k+1 \mid k)\right]^T \\ &\quad + \frac{1}{2} \sum_{i=1}^{2(n+q)} \left[\mathcal{X}_i(k+1 \mid k) - \hat{\mathbf{x}}(k+1 \mid k)\right] \left[\mathcal{Z}_i(k+1 \mid k) - \hat{\mathbf{z}}(k+1 \mid k)\right]^T \Big\} \end{aligned} \tag{51}$$

When κ is negative then $\mathbf{P}(k+1 \mid k)$ may not be positive semidefinite. In that case a modified form of the covariance equation can be used. Further, in the limit as $n+q+\kappa$ tends to zero the mean tends to that calculated by the truncated second order filter. The modified covariance equation leads to the same covariance as that calculated by the EKF.

To summarize, we have presented a new filtering algorithm based on a new method for determining the mean and covariance of a probability distribution. The resulting filter has many advantages over the EKF. In particular, linearisation is not required, process noise can be readily and consistently incorporated, and the estimates are accurate up to fourth order.

6 Application

In this section we compare the performance of the new filter against that of the EKF for a problem which was initially presented in [1]. We choose this example because it has significant nonlinearities in the process and observation models and has been analysed extensively in the literature.

We wish to estimate the position, velocity and constant ballistic coefficient of a body as it re-enters the atmosphere at a very high altitude at a very high velocity. Acceleration due to gravity is negligible compared to the altitude and velocity dependent drag terms. The body is constrained so that it falls vertically and at discrete points in time the range of the body is measured using a radar in the presence of white, uncorrelated noise. The geometry of the situation is shown in Figure 8: the radar is at an altitude of H (100,000ft) and the horizontal range between the body and the radar, M, is (100,000ft). This system has three states which are defined as follows:

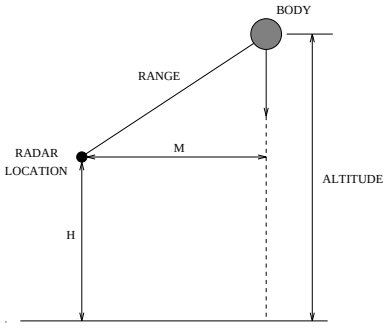

Fig.8: The geometry of the example

- $x_1(t)$ altitude (in feet)
- $x_2(t)$ velocity (positive downwards, in feet per second)
- $x_3(t)$ (constant) ballistic coefficient (in per feet)

The continuous time dynamics of this system are

$$\dot{x}_1(t) = -x_2(t) + w_1(t) \tag{52}$$
$$\dot{x}_2(t) = -e^{-\gamma x_1(t)} x_2(t)^2 x_3(t) + w_2(t) \tag{53}$$
$$\dot{x}_3(t) = w_3(t) \tag{54}$$

where $w_1(t), w_2(t)$ and $w_3(t)$ are zero mean, uncorrelated noises with covariances given by $\mathbf{Q}(t)$ and γ is a constant (5×10^{-5}) which relates the air density with altitude. The range at time t, $z(t)$, is

$$z(t) = \sqrt{\left(M^2 + [x_1(t) - H]^2\right)} + r(t) \tag{55}$$

where $r(t)$ is the uncorrelated observation noise with covariance $\mathbf{R}(t) = 10^4 \text{ft}^2$.

Both filters were implemented in discrete time and observations were taken once per second. However, considering the nonlinearities of the process model and the high velocity (initially 20,000fts^{-1}), numerical integration of Equations 52 to 54 produced reasonable predictions of the state variables only for extremely small time steps. In accordance with [1], a fourth order Runge-Kutta scheme was employed with 64 steps between each observation. This additional complication did not greatly affect the implementation of the new filter. The numerical scheme was applied to predict the position of each sigma point individually and the mean and covariance were calculated just before an observation was made. Since $n = 3$ we chose $\kappa = 0$ to minimise the maximum error up to the fourth order. With this choice of κ there is no possibility of predicting a non-positive semidefinite covariance matrix and Equation 47 was used. However, the implementation of the EKF was a rather more involved process. The mean was calculated using the numerical scheme. However, since the values of the state variables changed significantly across the time interval, $\mathcal{J}_f$ exhibited strong time variations as well. This could be compensated for by repeatedly linearising the prediction equations and using Equation 14. Let δ be the duration of each time step,

$$\delta \triangleq \frac{t_{k+1} - t_k}{64}. \tag{56}$$

Then the covariance was propagated from the nth to $n+1$th step using

$$\begin{aligned} \mathbf{P}(t_k + (n+1)\delta \mid t_k) &= \boldsymbol{\phi}(t_k + (n+1)\delta, t_k + n\delta)\mathbf{P}(t_k + n\delta \mid t_k) \\ &\quad \times \boldsymbol{\phi}^T(t_k + (n+1)\delta, t_k + n\delta) \end{aligned} \tag{57}$$

where

$$\boldsymbol{\phi}(t_k + (n+1)\delta, t_k + n\delta) \approx \mathbf{I} + \delta\mathcal{J}_f + \frac{\delta^2}{2}(\mathcal{J}_f)^2 \tag{58}$$

and $\mathcal{J}_f$ was evaluated about $t_k + n\delta$. The final predicted covariance, $\mathbf{P}(k+1 \mid k)$, was given by $\mathbf{P}(t_k + 64\delta \mid t_k)$.

The initial true state of the system is

$$\left.\begin{array}{rcl} x_1(0) & = & 300,000\text{ft}, \\ x_2(0) & = & 20,000\text{fts}^{-1}, \\ x_3(0) & = & 10^{-3}\text{ft}^{-1}. \end{array}\right\} \tag{59}$$

and the initial estimates of these states are

$$\left.\begin{array}{rcl} \hat{x}_1(0 \mid 0) & = & 300,000\text{ft}, \\ \hat{x}_2(0 \mid 0) & = & 20,000\text{fts}^{-1}, \\ \hat{x}_3(0 \mid 0) & = & 3 \times 10^{-5}\text{ft}^{-1} \end{array}\right\} \tag{60}$$

with covariance

$$\mathbf{P}(0 \mid 0) = \begin{bmatrix} 10^6 & 0 & 0 \\ 0 & 4 \times 10^6 & 0 \\ 0 & 0 & 10^{-4} \end{bmatrix} \tag{61}$$

Thus the initial estimate of altitude and velocity is correct, while the initial estimate of the ballistic parameter $\hat{x}_3(0 \mid 0)$ is very bad. Physically, this corresponds to assuming that the body is "heavy" whereas in reality the body is "light." We wish to see how quickly the filters converge to the true state of the system and, to this end, no process noises are injected into the system. Further, we set $\mathbf{Q}(k) = \mathbf{0}$ for each filter. These conditions are in accordance with those for the original test [1]. Referring to the earlier analysis we expect that there will be significant differences in the predictions made by the EKF and the new filter which will, in turn, lead to differences in the magnitudes of the state errors. Note that if both filters are tuned with non-zero $\mathbf{Q}(k)$s, better performance will be achieved.

In Figure 9 we show the average magnitude of the state errors committed by each filter across a Monte Carlo simulation consisting of 50 runs. As can be seen, there is initially very little difference between the two filters since, at high altitude, air density is low and $x_3(k)$ has little effect on body motion. However, after about 10s the body has fallen sufficiently far and the effects of drag become significant. Here the performance for the two filters differs quite dramatically. Figure 10 shows the performance for the velocity estimates. As can be seen, there are large error spikes for both filters. These are related to the fact that, at 10s, the altitude of the body is the same as that of the radar. Range information provides less data about body motion and so leads to increasing errors. Even so it can be seen that the EKF has a larger error spike in this region and only slowly converges at lower altitudes. Finally Figure 11 shows the errors in predicting $x_3(\cdot)$. As can be seen, the error in the EKF estimate is converging to be an order of magnitude larger than that for the new filter.

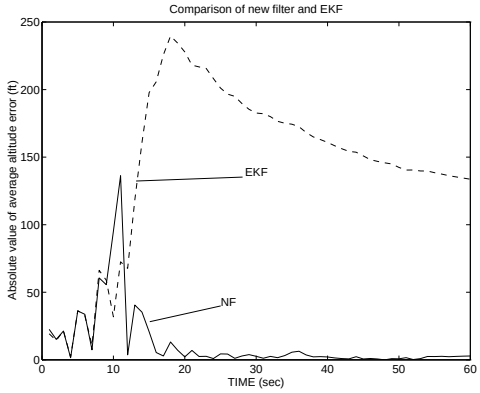

Fig. 9: Absolute mean position error.

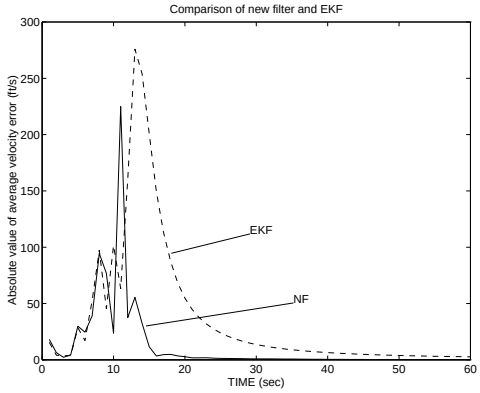

Fig. 10: Absolute mean velocity error.

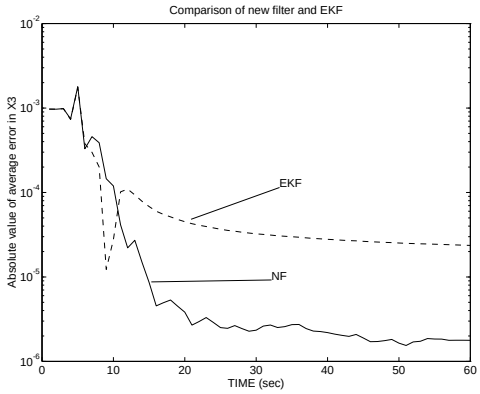

Fig. 11: Absolute mean error in $x_3(\cdot)$.

In Figures 12 and 13 we show the errors in the position estimates made by the EKF and the new filter and the associated estimates of the 2 standard deviation bounds. As can be seen, the error in the EKF falls outside the deviation bounds after about 30s indicating that the filter is consistently over predicting the accuracy of its estimates. However, the errors in the new filter estimate always lie in the estimated bounds.

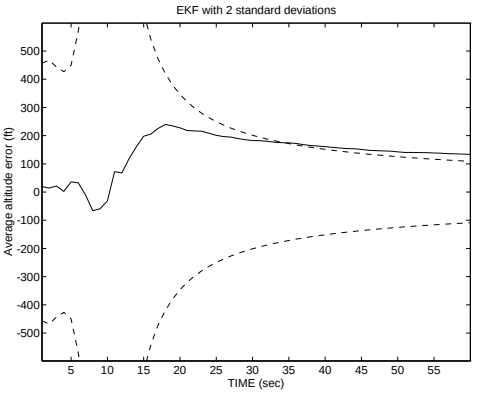

Fig. 12: EKF mean position error.

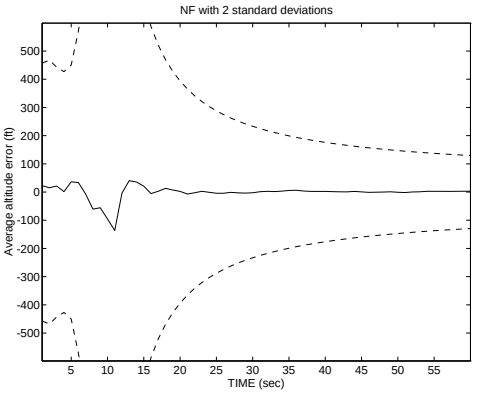

Fig. 13: NF mean position error.

Therefore we conclude that in this example the new filter has substantial advantages both in implementation and performance. The value of the filter has been demonstrated in other applications including high-speed road vehicle navigation [4–6], map based localisation [15] and satellite navigation [12].

7 Conclusions

In this article we have shown that a recently proposed alternative to the extended Kalman filter is preferable in terms of performance and ease of implementation. More specifically, we have shown:

1. The new filter is provably superior to the EKF in terms of expected error for all absolutely continuous nonlinear transformations. The new filter can be applied with non-differentiable functions in which the EKF is undefined.

2. The new filter avoids the derivation of Jacobian matrices for linearising nonlinear kinematic and observation models. This makes the new filter conducive to the creation of efficient, general purpose "black box" code libraries.

3. Our exposition of the filter clearly demonstrates that the new filter is a more direct generalisation of the linear Kalman filter than is the EKF. In particular, we suggest that the new approach is a true extension of the Kalman paradigm, whereas the so-called EKF is more of a recipe for pounding nonlinear models into linear holes.

4. Empirical results for a highly nonlinear problem arising in vehicle control[4], satellite navigation and fault detection[12] suggest that the new filter yields results which are at least as good as those obtained from a well-tuned EKF. Moreover, the performance of the new filter is more robust than the EKF when stressed.

The fact that the new filter is provably superior to the EKF in every important performance respect is sufficient evidence to conclude that the EKF is an obsolete methodology. The implications of the new filter, however, go far beyond simply improving performance. The fact that the new filter does not require the derivation of Jacobians eliminates the major obstacle to the development of high fidelity kinematic models in practical applications. Real world engineering experience with the EKF has led most implementors to conclude that the modelling of subtle dynamic effects usually entails a large effort (deriving Jacobians, tuning the many fudge factors necessitated by the EKF, etc.) that is usually defeated by linearisation errors. The new filter, on the other hand, permits highly complex models to be implemented and tested quickly.

In summary we should emphasize that the important result in this paper is a new method for applying nonlinear transformations to multivariate Gaussian distributions. This result is therefore not limited to nonlinear extensions of the Kalman filter. It can be applied to almost any filter that employs linearisation approximations to nonlinear functions (e.g., information filters, H_∞ controllers, etc.).

A Theoretical Analysis of Performance

In this section we compare the performance of the new prediction algorithm against that of linearisation in terms of predicting the mean and covariance of the posterior random variable. In the next section we lay out the basic framework of the analysis, which is the Taylor Series. In the subsequent sections we examine the prediction of the mean and the covariance. Finally we conclude with some loose comments about discontinuous functions.

A.1 The Multi-Dimensional Taylor Series

For the purpose of this analysis we assume that all nonlinear transformations are analytic across the domain of all possible values of $\mathbf{x}$. This condition means that the process model can be expressed as a multidimensional Taylor series consisting of an arbitrary number of terms. As the number of terms tends to infinity, the residual in the series always tends to zero and so the series always converges to the true value of the function. Note that these assumptions are more restrictive than those required for both of the algorithms. To apply linearisation it must be necessary to differentiate the function to form the Jacobian matrix. The new algorithm does not even place this restriction.

If we now consider the prior variable $\mathbf{x}$ as a mean $\bar{\mathbf{x}}$ plus a zero-mean disturbance $\boldsymbol{\Delta}\mathbf{x}$ with covariance $\mathbf{P}$, then the Taylor series expansion for the non-linear transformation $\mathbf{g}[\mathbf{x}]$ about $\bar{\mathbf{x}}$ is

$$\mathbf{g}[\bar{\mathbf{x}} + \boldsymbol{\Delta}\mathbf{x}] = \mathbf{g}[\bar{\mathbf{x}}] + \mathbf{D}_{\Delta x}\mathbf{g} + \frac{\mathbf{D}^2_{\Delta x}\mathbf{g}}{2!} + \frac{\mathbf{D}^3_{\Delta x}\mathbf{g}}{3!} + \frac{\mathbf{D}^4_{\Delta x}\mathbf{g}}{4!} + \cdots \tag{62}$$

The $\mathbf{D}_{\Delta x}\mathbf{g}$ operator evaluates the total differential of $\mathbf{g}[\cdot]$ when perturbed around a nominal value $\bar{\mathbf{x}}$ by $\boldsymbol{\Delta}\mathbf{x}$,

$$\mathbf{D}_{\Delta x}\mathbf{g} \stackrel{\triangle}{=} \left(\left(\boldsymbol{\Delta}\mathbf{x}^T \nabla \right) \mathbf{g}[\mathbf{x}]^T \right)^T \Bigg|_{\mathbf{x}=\bar{\mathbf{x}}} \tag{63}$$

This operator can be arranged in two ways. First, the $\mathbf{D}$ operator can be expressed as $\mathcal{J}_g \boldsymbol{\Delta}\mathbf{x}$, where $\mathcal{J}_g$ is the Jacobian matrix of $\mathbf{g}[\cdot]$ evaluated about $\bar{\mathbf{x}}$. Second, the operator can be written as the scalar operator

$$\mathbf{D}_{\Delta x} = \sum_{i=1}^{n} \Delta x_i \frac{\partial}{\partial x_i} \tag{64}$$

which acts on $\mathbf{g}[\cdot]$ on a component-by-component basis. The ith term in the Taylor series for $\mathbf{g}[\cdot]$ is given by

$$\frac{\mathbf{D}^i_{\Delta x}\mathbf{g}}{i!} = \frac{1}{i!} \left(\sum_{j=1}^{n} \Delta x_j \frac{\partial}{\partial x_j} \right)^i \mathbf{g}[\mathbf{x}] \Bigg|_{\mathbf{x}=\bar{\mathbf{x}}} \tag{65}$$

This can be expressed as a sum of components, each of which is an ith order differential of $\mathbf{g}[\cdot]$ with respect to $\mathbf{x}$ and an ith order product of $\boldsymbol{\Delta}\mathbf{x}$. If $\boldsymbol{\Delta}\mathbf{x}$ is a random variable then the expected value of the ith term in the Taylor series for $\mathbf{g}[\bar{\mathbf{x}} + \boldsymbol{\Delta}\mathbf{x}]$ is

$$\begin{aligned} \mathrm{E}\left[\frac{\mathbf{D}^i_{\Delta x}\mathbf{g}}{i!} \right] &= \frac{1}{i!} \mathrm{E}\left[\left(\sum_{j=1}^{n} \Delta x_j \frac{\partial}{\partial x_j} \right)^i \mathbf{g}[\mathbf{x}] \right] \Bigg|_{\mathbf{x}=\bar{\mathbf{x}}} \\ &= \frac{1}{i!} \mathrm{E}\left[\left(\sum_{j=1}^{n} \Delta x_j \frac{\partial}{\partial x_j} \right)^i \right] \mathbf{g}[\mathbf{x}] \Bigg|_{\mathbf{x}=\bar{\mathbf{x}}} \end{aligned} \tag{66}$$

It is expressed in terms of the ith moments of the distribution of $\boldsymbol{\Delta}\mathbf{x}$ and the ith order partial derivatives of $\mathbf{g}[\cdot]$. To illustrate this, we consider a simple case where $\mathbf{g}[\cdot]$ is the mapping from a two dimensional space to another two-dimensional space. Letting $\mathbf{x} = [x_1, x_2]^T$, $\mathbf{y} = [y_1, y_2]^T$ and $\mathbf{g}[\mathbf{x}] = [g_1(\mathbf{x}), g_2(\mathbf{x})]^T$ then

$$\mathbf{D}_{\Delta x}\mathbf{g} = \Delta x_1 \frac{\partial \mathbf{g}}{\partial x_1} + \Delta x_2 \frac{\partial \mathbf{g}}{\partial x_2} \tag{67}$$

$$\frac{\mathbf{D}^2_{\Delta x}\mathbf{g}}{2!} = \Delta {x_1}^2 \frac{\partial^2 \mathbf{g}}{\partial {x_1}^2} + \Delta x_1 \Delta x_2 \frac{\partial^2 \mathbf{g}}{\partial x_1 \partial x_2} + \Delta {x_2}^2 \frac{\partial^2 \mathbf{g}}{\partial {x_2}^2} \tag{68}$$

When we compare the accuracy of the different filters, we shall employ the same nonlinear equation $\mathbf{g}[\cdot]$ and mean $\bar{\mathbf{x}}$ but with different distributions (which are not necessarily probability distributions) of $\boldsymbol{\Delta}\mathbf{x}$. Using the above form, we can compare the performance of each filter simply by comparing the moments of $\boldsymbol{\Delta}\mathbf{x}$ without knowledge of the value or behaviour of the state transition equations.

A.2 Performance in predicting the mean of a continuous function

The predicted state estimate is found by taking expectations of Equation 62,

$$\bar{\mathbf{y}} = \mathrm{E}\left[\mathbf{g}[\bar{\mathbf{x}} + \boldsymbol{\Delta}\mathbf{x}]\right] = \mathbf{g}[\bar{\mathbf{x}}] + \mathrm{E}\left[\mathbf{D}_{\Delta x}\mathbf{g} + \frac{\mathbf{D}^2_{\Delta x}\mathbf{g}}{2!} + \frac{\mathbf{D}^3_{\Delta x}\mathbf{g}}{3!} + \frac{\mathbf{D}^4_{\Delta x}\mathbf{g}}{4!} + \cdots\right] \tag{69}$$

For the true mean, denoted by the subscript T, $\boldsymbol{\Delta}\mathbf{x}$ is a zero-mean, Gaussian process with covariance $\mathbf{P}$. By symmetry, all odd ordered moments in this distribution are zero. Therefore the expected value of all odd terms in this series are zero and

$$\bar{\mathbf{y}}_T = \mathbf{g}[\bar{\mathbf{x}}] + \mathrm{E}\left[\frac{\mathbf{D}^2_{\Delta x}\mathbf{g}}{2!} + \frac{\mathbf{D}^4_{\Delta x}\mathbf{g}}{4!} + \cdots\right] \tag{70}$$

The second order even terms can be written as

$$\frac{\mathbf{D}^2_{\Delta x}\mathbf{g}}{2!} = \frac{\mathbf{D}_{\Delta x}\left(\mathbf{D}_{\Delta x}\mathbf{g}\right)}{2!} = \left(\frac{\boldsymbol{\Delta}\mathbf{x}^T\nabla\boldsymbol{\Delta}\mathbf{x}^T\nabla}{2!}\right)\mathbf{g} = \left(\frac{\nabla^T\boldsymbol{\Delta}\mathbf{x}\boldsymbol{\Delta}\mathbf{x}^T\nabla}{2!}\right)\mathbf{g}. \tag{71}$$

Using the second interpretation of the $\mathbf{D}$ operator and noting that $\mathrm{E}\left[\boldsymbol{\Delta}\mathbf{x}\boldsymbol{\Delta}\mathbf{x}^T\right] = \mathbf{P}_{xx}$, the second term in the Taylor series can be written as

$$\mathrm{E}\left[\frac{\mathbf{D}^2_{\Delta x}\mathbf{g}}{2!}\right] = \left(\frac{\nabla^T\mathbf{P}_{xx}\nabla}{2!}\right)\mathbf{g}. \tag{72}$$

Equation 70 can now be written as

$$\bar{\mathbf{y}}_T = \mathbf{g}[\bar{\mathbf{x}}] + \left(\frac{\nabla^T\mathbf{P}_{xx}\nabla}{2!}\right)\mathbf{g} + \mathrm{E}\left[\frac{\mathbf{D}^4_{\Delta x}\mathbf{g}}{4!} + \cdots\right]. \tag{73}$$

Linearisation however, truncates this series at the first order and predicts the conditional mean as

$$\bar{\mathbf{y}}_{LIN} = \mathbf{g}[\bar{\mathbf{x}}]. \tag{74}$$

This estimate is independent of the covariance and higher moments of the distribution of $\bar{\mathbf{x}}$. However, comparing this with Equation 73 reveals that this is accurate only if the expected value of the second and higher order terms in the series are zero. This is always true for a linear system since the second and higher derivatives of the transformation are zero. However, for a general nonlinear system these terms are non-zero and this condition does not hold. Therefore errors are introduced at the second order.

The new filter predicts the mean from the projected set of points using Equation 24. Consider the Taylor series for the transition of each point $\mathcal{X}_i$. This can be expressed as the Taylor series about $\bar{\mathbf{x}}$,

$$\mathcal{Y}_i = \mathbf{g}[\mathcal{X}_i] = \mathbf{g}[\bar{\mathbf{x}}] + \mathbf{D}_{\sigma_i}\mathbf{g} + \frac{\mathbf{D}^2_{\sigma_i}\mathbf{g}}{2!} + \frac{\mathbf{D}^3_{\sigma_i}\mathbf{g}}{3!} + \frac{\mathbf{D}^4_{\sigma_g}\mathbf{f}}{4!} + \cdots \tag{75}$$

where $\boldsymbol{\sigma}_i(k) = \mathcal{X}_i kk - \bar{\mathbf{x}}$. Applying Equation 24, the predicted estimate is

$$\begin{aligned}\bar{\mathbf{y}} &= \frac{1}{n+\kappa}\left\{\kappa\mathbf{g}[\bar{\mathbf{x}}] + \frac{1}{2}\sum_{i=1}^{2n}\left(\mathbf{g}[\bar{\mathbf{x}}] + \mathbf{D}_{\sigma_i}\mathbf{g} + \frac{\mathbf{D}^2_{\sigma_i}\mathbf{g}}{2!} + \frac{\mathbf{D}^3_{\sigma_i}\mathbf{g}}{3!} + \frac{\mathbf{D}^4_{\sigma_i}\mathbf{g}}{4!} + \cdots\right)\right\} \\ &= \mathbf{g}[\bar{\mathbf{x}}] + \frac{1}{2(n+\kappa)}\sum_{i=1}^{2n}\left(\mathbf{D}_{\sigma_i}\mathbf{g} + \frac{\mathbf{D}^2_{\sigma_i}\mathbf{g}}{2!} + \frac{\mathbf{D}^3_{\sigma_i}\mathbf{g}}{3!} + \frac{\mathbf{D}^4_{\sigma_i}\mathbf{g}}{4!} + \cdots\right). \qquad (76)\end{aligned}$$

Comparing this series with the true series, we see that different values for the predicted mean occur only if the moments of $\boldsymbol{\Delta}\mathbf{x}$ and $\boldsymbol{\sigma}_i(k)$ are different[4]. The distribution of $\boldsymbol{\sigma}_i(k)$ is symmetric. All odd moments are zero and hence all odd terms sum to zero. Recalling that the sigma points are found from the column (or row) vectors of the matrix square root of $\sqrt{(n+\kappa)\bar{\mathbf{x}}}$, the second order even term is

$$\frac{\mathbf{D}^2_{\sigma_i}}{2!} = \left(\frac{\nabla^T\boldsymbol{\sigma}_i(k)\boldsymbol{\sigma}_i^T(k)\nabla}{2!}\right)\mathbf{g}. \qquad (77)$$

Therefore the predicted mean is

$$\bar{\mathbf{y}} = \mathbf{f}[\bar{\mathbf{x}}] + \left(\frac{\nabla^T\mathbf{P}_{xx}\nabla}{2!}\right)\mathbf{g} + \frac{1}{2(n+\kappa)}\sum_{i=1}^{2n}\left(\frac{\mathbf{D}^4_{\sigma_i}\mathbf{g}}{4!} + \cdots\right). \qquad (78)$$

Comparing this series with Equation 73 we see that the mean predicted by the new filter agrees with the true mean up to the third order and that errors are introduced in the fourth and higher order terms. This does not necessarily guarantee that the estimate is more accurate since we have not examined the behaviour of the higher order terms in the series. We now consider this problem.

To examine the higher order errors we observe that the random vector $\boldsymbol{\Delta}\mathbf{x}$ with covariance $\mathbf{P}_{xx}$ can be *stochastically decoupled* - it can be expressed in terms of an uncorrelated random vector $\boldsymbol{\Delta}\mathbf{x}'$ with covariance $\mathbf{I}$ (where $\mathbf{I}$ is the identity matrix). This decoupling is achieved by means of the linear transformation

$$\boldsymbol{\Delta}\mathbf{x} = \mathbf{A}(k)\boldsymbol{\Delta}\mathbf{x}', \qquad (79)$$

where $\mathbf{A}(k)$ is a matrix square root of $\mathbf{P}_{xx}$. If a_{ij} is the ith element in the jth column vector of $\mathbf{A}(k)$ then the $\mathbf{D}$ operator can be expressed as

$$\mathbf{D}_{\Delta x} = \sum_{i=1}^{n}\Delta x'_i\left(\sum_{j=1}^{n}a_{ij}\frac{\partial}{\partial x_j}\right). \qquad (80)$$

Similarly, we can identify a set of sigma points $\boldsymbol{\sigma'}_i(k)$ which capture the mean and covariance of a normal distribution with covariance identity. This set of points is related to $\boldsymbol{\sigma}_i(k)$ by

$$\boldsymbol{\sigma}_i(k) = \mathbf{A}(k)\boldsymbol{\sigma'}_i(k), \qquad (81)$$

where $\mathbf{A}(k)$ is, again, any matrix square root of $\mathbf{P}_{xx}$[5] Rather than handle a correlated random vector we now only need to consider the uncorrelated random vectors $\boldsymbol{\Delta}\mathbf{x}'$ and $\boldsymbol{\sigma'}_i(k)$. For the Gaussian case it can be shown that the fourth order moments (or *kurtosis*) are [11]

$$\begin{aligned}\mathrm{E}\left[\Delta x_i'^4\right] &= 3 \quad \forall i, \\ \mathrm{E}\left[\Delta x_i'^2\Delta x_j'^2\right] &= 1 \quad \forall i{\neq}j.\end{aligned}$$

[4]Our terminology here is somewhat lax. We can only talk about moments of the sigma points if they are a probability distribution - that is, when $\kappa \geq 0$. When $\kappa < 0$ then we are referring to a weighted average of components raised to a particular power.

[5]This result leads to an alternative insight into the sigma points - namely that the sigma points capture the mean and covariance of a distribution with mean zero and covariance the identity. This distribution is scaled by the square root of the covariance matrix and translated by the mean.

All other fourth order moments are zero. For the sigma points, the kurtosis of the jth components are

$$\frac{1}{2(n+\kappa)}\sum_{i=1}^{2n}{\sigma'_{ij}}^4(k) = n+\kappa \quad \forall j \tag{82}$$

and all other fourth order products are zero.

This analysis shows the effect of κ: Although the first three moments are unaltered, it affects the scaling of the fourth and higher order moments of the distribution of $\boldsymbol{\sigma}_i(k)$. If information is known about the predicted distribution then this can be incorporated into the choice of κ so that the error in the predicted mean is minimised. However if there is *no* information about the higher order terms of $\mathbf{g}[\cdot]$, the best choice of κ is motivated by ensuring that the errors committed by the new algorithm are smaller than those committed using linearisation.

Comparing the kurtosis of the true distribution against that for the sigma points, we observe two differences. First, the kurtosis of a single state has a value of 3 for the Gaussian distribution but $n+\kappa$ for the sigma distribution. Therefore there is a difference in the scaling of the moments. Second, the sigma point distribution has a zero cross kurtosis (and indeed they are zero for all higher order moments as well) but the Gaussian distribution has nonzero cross kurtoses. Therefore, except for the one dimensional case (when cross-kurtoses do not exist), the "shape" of the moments are different. If κ is chosen such that $n+\kappa = 3$ then the kurtosis of the single states for both distributions are the same. For a one dimensional state space, the errors are introduced in the sixth and higher order moments. However, when the space is multidimensional fourth order errors are introduced through the cross-kurtoses terms. Linearisation, by comparison, assumes that all fourth order moments are zero. Therefore since the error in the kurtosis is smaller for the new algorithm than that assumed by linearisation, the absolute errors in the fourth order errors in predicting the conditional mean using the new algorithm are smaller than those using linearisation.

To consider the sixth and higher order moments we consider how the values of the higher order moments grow. For a Gaussian distribution these moments grow factorially. However, the moments for the sigma point distribution grow geometrically with common factor $n+\kappa$. Therefore for any choice of κ it is possible to select a sufficiently large order such that the moments of the true series exceeds those for the sigma points. When $n+\kappa = 3$ the moments coincide at the fourth moment. For all higher terms, the Gaussian moments are larger in magnitude than the sigma point moments. However the linearisation enforces the condition that all of these higher order terms are zero. Therefore, on a term-by-term basis the errors in the terms of the new algorithm are smaller than those for linearisation.

We observe that as $n+\kappa$ tends to zero, the kurtosis and higher order moments for the sigma points tend to zero. The predicted mean converges to

$$\lim_{(n+\kappa)\to 0} \bar{\mathbf{y}} = \mathbf{g}[\bar{\mathbf{x}}] + \left(\frac{\nabla^T \mathbf{P}_{xx} \nabla}{2!}\right)\mathbf{g}, \tag{83}$$

which is accurate to the second order. This prediction approximation is equivalent to that employed in the well-known [10] *truncated second order filter*: the Taylor series expansion is truncated after the second term in the series. Note, however, that this result is achieved without the need to computer Jacobians and Hessians.

We conclude that the new algorithm can predict the mean more accurately than linearisation for all continuous nonlinear transformations. Performance is determined by the choice of κ since this factor scales the fourth and higher order moments of the distribution. If information about

the true conditional mean (from, for example Monte Carlo simulations) then κ can be adjusted to minimise the error. If $0 < n + \kappa \leq 3$ then the absolute error in the predicted mean is smaller than that with linearisation. We note that in many filtering applications the first and second terms are dominant and κ has a minimal effect on estimation performance.

A.3 Performance in predicting the covariance of a continuous function

The true $(\mathbf{P}_{yy})_T$ is given by

$$(\mathbf{P}_{yy})_T = \mathrm{E}\left[[\mathbf{y} - \bar{\mathbf{y}}_T][\mathbf{y} - \bar{\mathbf{y}}_T]^T\right] \tag{84}$$

where the expectation is taken over the distribution of $\mathbf{y}$. The realisation of the state error is

$$\begin{aligned} \mathbf{y} - \bar{\mathbf{y}}_T &= \mathbf{g}[\bar{\mathbf{x}} + \boldsymbol{\Delta}\mathbf{x}] - \bar{\mathbf{y}}_T \\ &= \mathbf{D}_{\Delta x}\mathbf{g} + \frac{\mathbf{D}^2_{\Delta x}\mathbf{g}}{2!} + \frac{\mathbf{D}^3_{\Delta x}\mathbf{g}}{3!} + \frac{\mathbf{D}^4_{\Delta x}\mathbf{g}}{4!} - \mathrm{E}\left[\frac{\mathbf{D}^2_{\Delta x}\mathbf{f}}{2!} + \frac{\mathbf{D}^4_{\Delta x}\mathbf{g}}{4!} + \cdots\right] \end{aligned} \tag{85}$$

with substitutions from Equations 62 and 70. The true covariance is found by post multiplying the state error by the transpose of itself and taking expectations. Recalling the symmetry of $\boldsymbol{\Delta}\mathbf{x}$, the expected value of all odd order terms of $\boldsymbol{\Delta}\mathbf{x}$ evaluate to zero and the true covariance is

$$\begin{aligned} (\mathbf{P}_{yy})_T &= \mathrm{E}\left[\mathbf{D}_{\Delta x}\mathbf{g}(\mathbf{D}_{\Delta x}\mathbf{g})^T + \frac{\mathbf{D}_{\Delta x}\mathbf{g}(\mathbf{D}^3_{\Delta x}\mathbf{g})^T}{3!} + \frac{\mathbf{D}^2_{\Delta x}\mathbf{g}(\mathbf{D}^2_{\Delta x}\mathbf{g})^T}{2 \times 2!} + \frac{\mathbf{D}^3_{\Delta x}\mathbf{g}(\mathbf{D}_{\Delta x}\mathbf{g})^T}{3!}\right] \\ &- \mathrm{E}\left[\frac{\mathbf{D}^2_{\Delta x}\mathbf{g}}{2!}\right]\mathrm{E}\left[\frac{\mathbf{D}^2_{\Delta x}\mathbf{g}}{2!}\right]^T + \cdots. \end{aligned} \tag{86}$$

Recalling the identity

$$\mathbf{D}_{\Delta x}\mathbf{g} = \mathcal{J}_g \boldsymbol{\Delta}\mathbf{x} \tag{87}$$

and using the expected given in Equation 72, we rewrite the above equation as

$$\begin{aligned} (\mathbf{P}_{yy})_T &= \mathcal{J}_g \mathbf{P}_{xx} \mathcal{J}_g^T + \mathrm{E}\left[\frac{\mathbf{D}_{\Delta x}\mathbf{g}(\mathbf{D}^3_{\Delta x}\mathbf{g})^T}{3!} + \frac{\mathbf{D}^2_{\Delta x}\mathbf{g}(\mathbf{D}^2_{\Delta x}\mathbf{g})^T}{2 \times 2!} + \frac{\mathbf{D}^3_{\Delta x}\mathbf{g}(\mathbf{D}_{\Delta x}\mathbf{g})^T}{3!}\right] \\ &- \left[\left(\frac{\nabla^T \mathbf{P}_{xx} \nabla}{2!}\right)\mathbf{g}\right]\left[\left(\frac{\nabla^T \mathbf{P}_{xx} \nabla}{2!}\right)\mathbf{g}\right]^T + \cdots \end{aligned} \tag{88}$$

The linearisation algorithm predicts the covariance using

$$(\mathbf{P}_{yy})_{LIN} = \mathcal{J}_g \mathbf{P}_{xx} \mathcal{J}_g^T, \tag{89}$$

which is the true series truncated after the first term. Therefore the errors in the predicted covariance are in the fourth and higher orders. In general it is not possible to determine whether the predicted covariance is conservative since this depends upon the higher order differentials of $\mathbf{g}[\cdot]$.

The new algorithm predicts the covariance using Equation 25 which requires the values of $\mathcal{Y}_i - \bar{\mathbf{y}}$ and $\mathcal{Y}_0 - \bar{\mathbf{y}}$. These values are given by

$$\mathcal{Y}_i - \bar{\mathbf{y}} = \mathbf{D}_{\Delta x}\mathbf{g} + \frac{\mathbf{D}^2_{\Delta x}\mathbf{g}}{2!} + \frac{\mathbf{D}^3_{\Delta x}\mathbf{g}}{3!} + \frac{\mathbf{D}^4_{\Delta x}\mathbf{g}}{4!} + \cdots$$

$$-\frac{1}{2(n+\kappa)}\sum_{i=1}^{2n}\left(\frac{\mathbf{D}^2_{\sigma_i}\mathbf{g}}{2!}+\frac{\mathbf{D}^4_{\sigma_i}\mathbf{g}}{4!}+\cdots\right). \tag{90}$$

$$\mathcal{Y}_0-\bar{\mathbf{y}} = -\frac{1}{2(n+\kappa)}\sum_{i=1}^{2n}\left(\frac{\mathbf{D}^2_{\sigma_i}\mathbf{g}}{2!}+\frac{\mathbf{D}^4_{\sigma_i}\mathbf{g}}{4!}+\cdots\right). \tag{91}$$

Noting that

$$\begin{aligned}\frac{1}{2(n+\kappa)}\sum_{i=1}^{2n}\mathbf{D}_{\sigma_i}\mathbf{g}(\mathbf{D}_{\sigma_i}\mathbf{g})^T &= \frac{1}{2(n+\kappa)}\sum_{i=1}^{2n}\mathcal{J}_g\boldsymbol{\sigma}_i(k)\boldsymbol{\sigma}_i^T(k)\mathcal{J}_g^T\\ &= \mathcal{J}_g\mathbf{P}_{xx}\mathcal{J}_g^T,\end{aligned} \tag{92}$$

the predicted covariance is

$$\begin{aligned}\mathbf{P}_{yy} &= \mathcal{J}_g\mathbf{P}_{xx}\mathcal{J}_g^T \\ &+ \frac{1}{2(n+\kappa)}\sum_{i=1}^{2n}\left(\frac{\mathbf{D}_{\sigma_i}\mathbf{g}(\mathbf{D}^3_{\sigma_i}\mathbf{g})^T}{3!}+\frac{\mathbf{D}^2_{\sigma_i}\mathbf{g}(\mathbf{D}^2_{\sigma_i}\mathbf{g})^T}{2\times 2!}+\frac{\mathbf{D}^3_{\sigma_i}\mathbf{g}(\mathbf{D}_{\sigma_i}\mathbf{g})^T}{3!}\right)\\ &- \left[\left(\frac{\nabla^T\mathbf{P}_{xx}\nabla}{2!}\right)\mathbf{g}\right]\left[\left(\frac{\nabla^T\mathbf{P}_{xx}\nabla}{2!}\right)\mathbf{g}\right]^T+\cdots\end{aligned} \tag{93}$$

Comparing this with the true series we see that the predicted covariance agrees with the true covariance up to the second order terms in the series. Since the kurtoses of the true and sigma point distributions are different, errors are introduced at the fourth and higher orders. By employing similar arguments as used with the conditional mean, we argue that the absolute errors in the prediction of the covariance are smaller using this formulation than that used by the EKF. However because we are attempting to faithfully approximate the covariance matrix, we do not ensure that this approximation is positive semidefinite if κ is negative. Similar problems are experienced with other sophisticated schemes which approximate higher order moments or probability density distributions [3, 10]. The situation in which this arises can be illustrated by considering the limit as $n+\kappa$ tends to zero,

$$\lim_{(n+\kappa)\to 0}\mathbf{P}_{yy}=\mathcal{J}_g\mathbf{P}_{xx}\mathcal{J}_g^T-\left[\left(\frac{\nabla^T\mathbf{P}_{xx}\nabla}{2!}\right)\mathbf{g}\right]\left[\left(\frac{\nabla^T\mathbf{P}_{xx}\nabla}{2!}\right)\mathbf{g}\right]^T. \tag{94}$$

The last term is of the order of covariance squared (and hence of the kurtosis of $\mathbf{\Delta x}$) but does not scale with κ. We can ensure positive semi-definiteness at the cost of a more conservative covariance prediction by calculating the "covariance" about $\bar{\mathbf{y}}$.

$$(\mathbf{P}_{yy})_{MOD}=\frac{1}{2(n+\kappa)}\left(\sum_{i=1}^{2n}[\mathcal{Y}_i-\bar{\mathbf{y}}][\mathcal{Y}_i-\bar{\mathbf{y}}]^T\right). \tag{95}$$

Positive semi-definiteness is guaranteed by the fact that the covariance matrix is evaluated as the sum of outer products of vectors.

To quantify the errors committed by this new method of predicting the covariance, we examine the Taylor expansion of $\mathcal{Y}_i-\bar{\mathbf{y}}$ which is

$$\mathcal{Y}_i-\bar{\mathbf{y}}=\mathbf{D}_{\Delta x}\mathbf{g}+\frac{\mathbf{D}^2_{\Delta x}\mathbf{g}}{2!}+\frac{\mathbf{D}^3_{\Delta x}\mathbf{g}}{3!}+\frac{\mathbf{D}^4_{\Delta x}\mathbf{g}}{4!}+\cdots \tag{96}$$

Predicting the covariance using Equation 95 yields

$$
\begin{aligned}
(\mathbf{P}_{yy})_{MOD} &= \mathcal{J}_g \mathbf{P}_{xx} \mathcal{J}_g^T \\
&+ \frac{1}{2}\sum_{i=1}^{2n}\left(\frac{\mathbf{D}_{\sigma_i}\mathbf{g}(\mathbf{D}^3_{\sigma_i}\mathbf{g})^T}{3!} + \frac{\mathbf{D}^2_{\sigma_i}\mathbf{g}(\mathbf{D}^2_{\sigma_i}\mathbf{g})^T}{2\times 2!} + \frac{\mathbf{D}^3_{\sigma_i}\mathbf{g}(\mathbf{D}_{\sigma_i}\mathbf{g})^T}{3!}\right) + \cdots
\end{aligned}
\tag{97}
$$

which does not include the subtractive terms. Therefore, in general, the covariance will be larger using this form than that initially guided by our intuition. This form also has the useful property that, in the limit as $n + \kappa$ tends to zero, the predicted covariance is

$$
\lim_{(n+\kappa)\to 0} (\mathbf{P}_{yy})_{MOD} = \mathcal{J}_g \mathbf{P}_{xx} \mathcal{J}_g^T, \tag{98}
$$

which is the same as that estimated through linearisation, but without the use of Jacobians

In conclusion this analysis shows that the covariance predicted by the new filter is at least as accurate as that predicted by the EKF. Although both approaches predict the covariance correctly up to the second order, the absolute errors in the fourth and higher order terms for the new filter are smaller. However, the analysis also shows that the original method of calculating the covariance may lead to a matrix which is not positive semidefinite. An alternative method for calculating the covariance has been presented. This method ensures positive semi-definiteness and is still more accurate than the EKF. The performance of all three methods of covariance prediction only converge with a linear system.

A.4 Predicting the mean and covariance of discontinuous functions

The preceding analysis is valid only if the state transition equations are continuous across the range of all possible values of the state estimates and predictions. However, in many practical situations it is possible to conceive of discontinuous process models, which we now briefly discuss.

Since the state transition equation has a finite expected value, all discontinuities must involve *finite* discontinuities in behaviour of the function. Further, suppose that $\mathbf{f}[\cdot]$ is piecewise approximated by a number of continuous functions, as illustrated in Figure 8. Each continuous function has its own Taylor series and when evaluating the function about a point the appropriate Taylor series must be used. The EKF can exhibit two types of behaviour. First, if the state estimate $\hat{\mathbf{x}}(k \mid k)$ does not lie at a discontinuity then the projected state estimate and covariance will not reflect the presence of the discontinuity at all. Second, if the discontinuity lies at the state estimate, there are difficulties in applying the EKF. Specifically, if the function is non-differentiable at that point then it is not possible to predict the covariance using the conventional EKF covariance prediction equations.

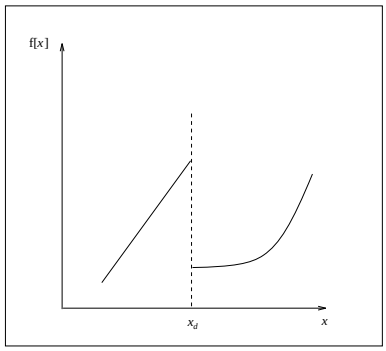

Fig.8: A discontinuous state transition equation.

The new filter does not require that the state transition equation be differentiable and can be used with *any* valid process model. However if the function is discontinuous then performance will be worse than that for a continuous function. If a discontinuity does not lie within the covariance ellipse formed by the sigma points then the state and covariance predictions do not acknowledge the existence of it. However, if the discontinuity lies outside the sigma points then it is unlikely that it affects a significant proportion of the distribution. If the discontinuity

lies at a sigma point it is not possible, in general, to make any comments about performance. However, if a discontinuity lies within the sigma points then first order errors are introduced into the mean and covariance predictions. When the mean and covariance are calculated for the new filter, the relevant Taylor series for each sigma point must be employed. Odd terms do not, in general, cancel out in these summations and an error is introduced into the first term in the series. Although the first moment is correctly represented using the sigma points, it is scaled by $1/\sqrt{n+\kappa}$. The size of this error can be reduced by reducing the value of κ at the cost of distorting the higher order terms in the series. Further, as $n+\kappa$ tends to zero the sigma points converge towards one another and there is the possibility that the sigma points will miss the discontinuity. This error is significant only if a substantial proportion of the distribution is affected by the discontinuity.

A.5 Summary

We have considered the performance of the EKF and the new filter prediction equations for both continuous and discontinuous state transition equations. In general, when the process model is continuous the errors in the EKF prediction are second order. The new filter, however, is accurate as far as the third order, and errors are only introduced at the fourth order. In many practical applications the lower terms are significant and the new filter can be significantly more accurate. When the function is discontinuous the EKF only incorporates this fact in its estimates if the discontinuity lies on the current state estimate. If the function is not differentiable at that point the covariance cannot be calculated. The new filter uses a distribution of points and captures the effects of the discontinuity if it influences a significant proportion of the distribution.

References

[1] M. Athans, R. P. Wishner and A. Bertolini. Suboptimal state estimation for continuous-time nonlinear systems from discrete noisy measurements. *IEEE Transactions on Automatic Control*, TAC-13(6):504–518, October 1968.

[2] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association.* The Academic Press, 1988.

[3] A. H. Jazwinski. *Stochastic Processes and Filtering Theory.* Academic Press, 1970.

[4] S. J. Julier. Process models for the high-speed navigation of conventional road vehicles. Technical report, University of Oxford, 1994. Transfer Thesis.

[5] S. J. Julier and H. F. Durrant-Whyte. Navigation and Parameter Estimation of High Speed Road Vehicles. In *Robotics and Automation Conference, Nagoya, Japan.*, pages 101–105, 1995.

[6] S. J. Julier, J. K. Uhlmann and H. F. Durrant-Whyte. A New Approach for Filtering Nonlinear Systems. In *American Control Conference, Seattle, Washington.*, pages 1628–1632, 1995.

[7] R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME, Journal of Basic Engineering*, 82:34–45, March 1960.

[8] H. J. Kushner. Dynamical equations for optimum non-linear filtering. *Journal of Differential Equations*, 3:179–190, 1967.

[9] P. S. Maybeck. *Stochastic Models, Estimation, and Control*, volume 1. Academic Press, 1979.

[10] P. S. Maybeck. *Stochastic Models, Estimation, and Control*, volume 2. Academic Press, 1982.

[11] A. Papoulis. *Probability, Random Variables, and Stochastic Processes.* McGraw-Hill, Inc., 1965.

[12] B. M. Quine. Fault detection in spacecraft guidance systems. Technical report, University of Oxford, August 1994. Transfer Thesis.

[13] B. M. Quine, J. K. Uhlmann and H. F. Durrant-Whyte. Implicit jacobians for linearized state estimation in nonlinear systems. In *The 1995 American Automatic Control Conference, Seattle U.S.A.*, pages 1645–1646, 1995.

[14] G. Strang. *Linear Algebra and its Applications.* Harcourt Brace Jovanich, 3 edition, 1988.

[15] J. K. Uhlmann. Simultaneous map building and localization for real time applications. Technical report, University of Oxford, 1994. Transfer thesis.