

Decision Based Uncertainty Propagation Using Adaptive Gaussian Mixtures

Gabriel Terejanu^a
terejanu@buffalo.edu

Puneet Singla^b
psingla@buffalo.edu

Tarunraj Singh^b
tsingh@buffalo.edu

Peter D. Scott^a
peter@buffalo.edu

^aDepartment of Computer Science & Engineering

^bDepartment of Mechanical & Aerospace Engineering
University at Buffalo, Buffalo, NY-14260.

Abstract – Given a decision process based on the approximate probability density function returned by a data assimilation algorithm, an interaction level between the decision making level and the data assimilation level is designed to incorporate the information held by the decision maker into the data assimilation process. Here the information held by the decision maker is a loss function at a decision time which maps the state space onto real numbers which represent the threat associated with different possible outcomes or states. The new probability density function obtained will address the region of interest, the area in the state space with the highest threat, and will provide overall a better approximation to the true conditional probability density function within it. The approximation used for the probability density function is a Gaussian mixture and a numerical example is presented to illustrate the concept.

Keywords: Adaptive Gaussian Sum, Decision Making, Uncertainty Propagation, Expected Loss.

1 Introduction

Chemical, Biological, Radiological, and Nuclear (CBRN) incidents are rare events but very consequential, which mandates extensive research and operational efforts in mitigating their outcomes. For such critical applications the accuracy in predicting the future evolution of toxic plumes in a timely fashion represents an important part in the Decision Maker (DM) toolbox. Based on these forecasts, decisions can be made on evacuating cities, sheltering or medical gear deployment. Such decisions are taken based on a loss function or region of interest such as the population density in a given area.

Many research projects try to model the atmospheric transport and diffusion of toxic plumes. While inherently stochastic and highly nonlinear, these mathematical models are able to capture just a part of the dynamics of the real phenomenon and the forward integration yields an uncertain prediction. The decision maker takes actions based on expected loss computed using both the predicted uncertainty and the loss function, which here maps a region of interest in

the state space into a threat level, such as the population density in a city. Thus the ability to propagate the uncertainty and errors throughout the dynamic system is of great importance, and the evolution of the probability density function (pdf) has received much attention recently.

Due to the uncertainty accumulation in integrating the model, the forecasts of the system become less and less useful for the decision maker. Data Assimilation (DA) offers a way to reduce the uncertainty by combining measurements provided by sensors with model prediction in a Bayesian way [14]. This gives an improved situation assessment for the hindcast and nowcast. Unfortunately the forecast, used to evaluate the impact assessment, is still affected by the accuracy of probability density function evolution. Even in the hindcast and nowcast cases if the sensors provide ambiguous measurements such a quadratic measurement model, the improvement brought by DA may be marginal.

For nonlinear systems, the exact description of the transition pdf is provided by a linear partial differential equation (pde) known as the Fokker Planck Kolmogorov Equation (FPKE) [12]. Analytical solutions exist only for stationary pdf and are restricted to a limited class of dynamical systems [12]. Thus researchers are looking actively at numerical approximations to solve the FPKE [8, 10], generally using the variational formulation of the problem. For discrete-time dynamical systems, solving for the exact solution, which is given by the Chapman-Kolmogorov Equation (CKE), yields the same problems as in the continuous-time case. Several other techniques exist in the literature to approximate the pdf evolution, the most popular being Monte Carlo (MC) methods [5], Gaussian closure [7] (or higher order moment closure), Equivalent Linearization [13], Stochastic Averaging [9], Gaussian mixture approximations [1, 6, 16]. Furthermore, all these approaches provide only an approximate description of the uncertainty propagation problem by restricting the solution to a small number of parameters.

All these assumptions employed make the problem tractable and computational efficient, which satisfies the requirement of minimizing decision latency. But the approximation given may be of little use when computing the ex-

pected loss, since the method is not sensitive to the region of interest. Such an example may be given using Monte Carlo approximations or Gaussian Sum approximations, when the propagated uncertainty offers no or very little probabilistic support in the region of interest. In other words, no particles or Gaussian components are centered in the region of interest, and even though the probabilistic support may be infinite, the expected loss computed might be underestimated.

Methods to deal with such situations have been developed, from risk sensitive filters [3] to risk sensitive particle filters [17]. The risk sensitive filters minimize the expected exponential of estimation error, controlling this way how much to weigh the higher-order moments versus the lower-order moments. While weighting more the higher-order moments, these methods are not designed to be selectively sensitive to particular regions of interest in the state space. The risk sensitive particle filter is able to generate more samples in the region of interest, but at the expense of biasing the proposal distribution, thus the particles obtained are biased towards the region of interest. While this method is appropriate for fault detection, it provides a limited output for the decision maker who is interested in querying the probability density function for different numerical quantities used in the decision process such as the expected loss, the mean or the mode of the pdf.

The present paper is concerned with providing a better approximation to the probability density function by incorporating contextual loss information held by the decision maker into the DA process. In this work we use a Gaussian mixture approximation to the probability density function. We propose a “non-intrusive” way in computing an approximate pdf that addresses the region of interest and it is closer to the true pdf. Non-intrusive refers here to the fact the we do not require a new DA method when incorporating the loss function into the derivation.

A progressive selection method is designed to add new Gaussian components to the initial Gaussian mixture, in assuring that probabilistic support is reaching the region of interest at the decision time. The initial weights of the new Gaussian components are set to zero and they are modified when propagated throughout the nonlinear dynamical system to minimize the error in the FPKE [16]. Therefore if there is any probability density mass in the region of interest it will be represented by the non-zero weight of the new Gaussian components at the decision time.

The problem is stated in Section 2 and the Gaussian Sum approximation to the conditional pdf is presented in Section 3. The progressive selection of Gaussian components is derived in Section 4. An example to illustrate the concept is given in Section 5 and the conclusions and future work are discussed in Section 6.

2 Problem Statement

Consider a general n -dimensional continuous-time noise driven nonlinear dynamic system with uncertain initial conditions and discrete measurement model, given by the equa-

tions:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(t, \mathbf{x}(t)) + \mathbf{g}(t, \mathbf{x}(t))\Gamma(t) \quad (1)$$

$$\mathbf{z}_k = \mathbf{h}(t_k, \mathbf{x}_k) + \mathbf{v}_k \quad (2)$$

and a set of k observations, $\mathbf{Z}_k = \{\mathbf{z}_i \mid i = 1 \dots k\}$.

We denote, $\mathbf{x}_k = \mathbf{x}(t_k)$, $\Gamma(t)$ represents a Gaussian white noise process with the correlation function $\mathbf{Q}\delta(t_{k+1} - t_k)$, and the initial state uncertainty is captured by the pdf $p(t_0, \mathbf{x}_0)$. The random vector \mathbf{v}_k denotes the measurement noise, which is temporally uncorrelated, zero-mean random sequence with known covariance, \mathbf{R}_k . The process noise and the measurement noise are uncorrelated with each other and with the initial condition.

We are interested in finding the conditional probability density function $p(t, \mathbf{x}(t) \mid \mathbf{Z}_k)$. For $t > t_k$ we get the forecast pdf by integrating only Eq.1 forward, if $t = t_k$ we are interested in the nowcast or filtered pdf and for $t < t_k$ we obtain the hindcast or the smoothed pdf.

Given a state space region of interest at a particular decision time, t_d , which may be represented as a *loss function* by the decision maker, $L(\mathbf{x}_d, a_d)$, the expected loss of an action a_d is given by:

$$L(a_d) = \int L(\mathbf{x}_d, a_d)p(t_d, \mathbf{x}_d \mid \mathbf{Z}_k)d\mathbf{x}_d \quad (3)$$

Here we will consider only the cases where $t_d > t_k$. Given approximate computational methods for the conditional pdf, $\hat{p}(t_d, \mathbf{x}_d \mid \mathbf{Z}_k)$, we are able to obtain an estimate of the expected loss and also find the optimal Bayesian decision, if a set of decisions exists.

$$\hat{L}(a_d) = \int L(\mathbf{x}_d, a_d)\hat{p}(t_d, \mathbf{x}_d \mid \mathbf{Z}_k)d\mathbf{x}_d \quad (4)$$

$$\hat{a}_d = \arg \min_{a_d} \int L(\mathbf{x}_d, a_d)\hat{p}(t_d, \mathbf{x}_d \mid \mathbf{Z}_k)d\mathbf{x}_d \quad (5)$$

The decision making process in the data assimilation framework is presented in Fig.1(left). Obviously if we have a good approximation for the conditional pdf in the region of interest the same can be said for the expected loss. This situation becomes more dramatic when a large deviation exists between the actual and the estimated conditional pdf in the region of interest. In the case of evaluation of a single decision, the algorithm can underestimate the actual expected loss, $\hat{L}(a_d) \ll L(a_d)$, misleading the decision maker with respect to the magnitude of the situation. In the case when an optimal decision has to be chosen, the large difference between conditional pdfs may result in picking not only a sub-optimal decision but a very consequential one.

While one can derive a new method to approximate the conditional pdf by including the loss function in the derivation and reduce the difference in the region of interest to better approximate the expected loss, it will accomplish this at the expense of worsening the approximation of the conditional pdf in the rest of the state space.

This will affect other estimates based on the conditional pdf, but the expected loss, that may be required in guiding

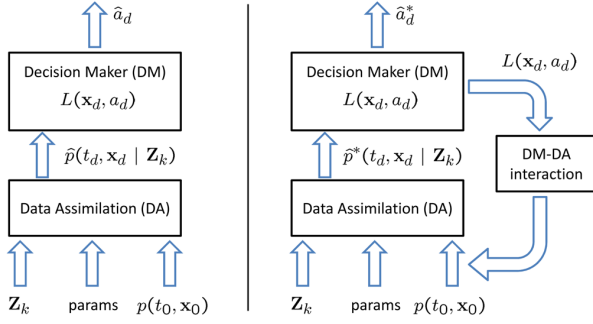


Figure 1: Left figure represents the classic approach to decision making in the data assimilation context. The right figure shows the proposed model.

the decision process, like the mean of the pdf, the modes of the pdf, etc. These will be biased towards the region of interest, thus misleading the decision maker.

In other words, if we call the computation of the expected loss of a given action as *impact assessment* and the computation of the moments and other quantities based on the conditional pdf as *situation assessment*, one will require that both to be as accurate as possible. At the limit, if we can compute exactly the conditional pdf we obtain both *impact assessment* and *situation assessment* since we can quantify exactly the probability of all the outcomes.

Since the decision maker holds important information regarding the use of the conditional pdf obtained from the data assimilation method, we can incorporate this information in the data assimilation process in a non-intrusive manner (do not have to derive a new method), by supplementing the inputs into the data assimilation module. The proposed method is shown in Fig.1(right), where a new interaction level is introduced between the decision maker and the data assimilation, that uses the contextual information provided by the decision maker to supplement the inputs of DA / change the environment in which DA is running.

Therefore we want to find an approximation to the conditional pdf, $\hat{p}^*(t_d, \mathbf{x}_d | \mathbf{Z}_k)$, that addresses the interest held by the decision maker and provides both a better impact and situation assessment than $\hat{p}(t_d, \mathbf{x}_d | \mathbf{Z}_k)$. These objectives can be captured by the following two relations:

$$\int |p(t_d, \mathbf{x}_d | \mathbf{Z}_k) - \hat{p}(t_d, \mathbf{x}_d | \mathbf{Z}_k)|^2 d\mathbf{x}_d \leq \int |p(t_d, \mathbf{x}_d | \mathbf{Z}_k) - \hat{p}(t_d, \mathbf{x}_d | \mathbf{Z}_k)|^2 d\mathbf{x}_d \quad (6)$$

$$\left| \hat{L}^*(a_d) - L(a_d) \right| \leq \left| \hat{L}(a_d) - L(a_d) \right| \quad (7)$$

In the present paper, we will design an interaction level between the decision maker and the data assimilation module that approximates the conditional pdf using a Gaussian mixture. The interaction level is adding new Gaussian components to the initial uncertainty, such that they will be positioned near the region of interest at the decision time. Their

initial weights will be set to zero, thus the initial uncertainty is not changed, but the evolution of the weights is dictated by the error in the Fokker-Planck equation. Thus if any probability density mass is moving naturally towards the region of interest, the weights of the new Gaussian components will become greater than zero. Therefore the method will just make sure that if there is any probability density mass in the region of interest it will be found by the data assimilation method.

In this paper we will consider only the forecast of the conditional pdf when no measurements are available between the current time and the decision time. A suggestion, on how this can be used in the case when we have observations to assimilate between the current time and the decision time, is given in Section 4.

3 Approximation of the Conditional Probability Density Function

The nonlinear filtering problem has been extensively studied and various methods are provided in literature. The Extended Kalman Filter (EKF) is historically the first, and still the most widely adopted approach to solve the nonlinear state estimation problem. It is based on the assumption that the nonlinear system dynamics can be accurately modeled by a first-order Taylor series expansion [4]. Since the EKF provides us only with a rough approximation to the a posteriori pdf and solving for the exact solution of the conditional pdf is very expensive, researchers have been looking for mathematically convenient approximations.

In Ref. [1], a weighted sum of Gaussian density functions has been proposed to approximate the conditional pdf. The probability density function of the initial condition is given by the following Gaussian sum,

$$p(t_0, \mathbf{x}_0) = \sum_{i=1}^N w_0^i \mathcal{N}(\mathbf{x}_0 | \boldsymbol{\mu}_0^i, \mathbf{P}_0^i) \quad (8)$$

$$\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}, \mathbf{P}) = |2\pi\mathbf{P}|^{-1/2} \exp \left[-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \mathbf{P}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right]$$

Let us assume that the underlying conditional pdf can be approximated by a finite sum of Gaussian pdfs

$$\hat{p}(t, \mathbf{x}(t) | \mathbf{Z}_k) = \sum_{i=1}^N w_{t|k}^i \underbrace{\mathcal{N}(\mathbf{x}(t) | \boldsymbol{\mu}_{t|k}^i, \mathbf{P}_{t|k}^i)}_{p_{gi}} \quad (9)$$

where $\boldsymbol{\mu}_{t|k}^i$ and $\mathbf{P}_{t|k}^i$ represent the conditional mean and covariance of the i^{th} component of the Gaussian pdf with respect to the k measurements, and $w_{t|k}^i$ denotes the amplitude of i^{th} Gaussian in the mixture. The positivity and normalization constraint on the mixture pdf, $\hat{p}(t, \mathbf{x} | \mathbf{Z}_k)$, leads to following constraints on the amplitude vector:

$$\sum_{i=1}^N w_{t|k}^i = 1, \quad w_{t|k}^i \geq 0, \quad \forall t \quad (10)$$

A Gaussian Sum Filter [1] may be used to propagate and update the conditional pdf. Since all the components of the mixture pdf (9) are Gaussian and thus, only estimates of their mean and covariance need to be propagated between t_k and t_{k+1} using the conventional Extended Kalman Filter time update equations:

$$\dot{\boldsymbol{\mu}}_{t|k}^i = \mathbf{f}(t, \boldsymbol{\mu}_{t|k}^i) \quad (11)$$

$$\dot{\mathbf{P}}_{t|k}^i = \mathbf{A}_{t|k}^i \mathbf{P}_{t|k}^i + \mathbf{P}_{t|k}^i (\mathbf{A}_{t|k}^i)^T + \mathbf{g}(t, \boldsymbol{\mu}_{t|k}^i) \mathbf{Q} \mathbf{g}^T(t, \boldsymbol{\mu}_{t|k}^i) \quad (12)$$

$$\mathbf{A}_{t|k}^i = \left. \frac{\partial \mathbf{f}(t, \mathbf{x}(t))}{\partial \mathbf{x}(t)} \right|_{\mathbf{x}(t) = \boldsymbol{\mu}_{t|k}^i} \quad (13)$$

In Ref. [16] an update method of adapting the weights of different Gaussian components during propagation is given based on minimizing the error in the Fokker-Planck equation for continuous dynamical systems, and on minimizing the integral square difference between the true forecast pdf, given by the Chapman-Kolmogorov equations and its Gaussian sum approximation in the discrete time dynamical systems.

The new weights are given by the solution of the following quadratic programming problem:

$$\begin{aligned} \mathbf{w}_{t|k} &= \arg \min_{\mathbf{w}_{t|k}} \frac{1}{2} \mathbf{w}_{t|k}^T (\mathbf{L} + \mathbf{I}) \mathbf{w}_{t|k} - \mathbf{w}_{t|k}^T \mathbf{w}_{k|k} \quad (14) \\ &\text{subject to } \mathbf{1}_{N \times 1}^T \mathbf{w}_{t|k} = 1 \\ &\quad \mathbf{w}_{t|k} \geq \mathbf{0}_{N \times 1} \end{aligned}$$

where $\mathbf{w}_{t|k} \in \mathbb{R}^{N \times 1}$ is a vector of Gaussian weights, and the elements of $\mathbf{L} \in \mathbb{R}^{N \times N}$ are given by:

$$\begin{aligned} L_{ij} &= \int_V \mathcal{L}_i \mathcal{L}_j d\mathbf{x} \quad (15) \\ \mathcal{L}_i(t, \mathbf{x}) &= \left[\frac{\partial p_{g_i}}{\partial \boldsymbol{\mu}_{t|k}^i}{}^T \mathbf{f}(t, \boldsymbol{\mu}_{t|k}^i) + \sum_{j=1}^n \sum_{k=1}^n \frac{\partial p_{g_i}}{\partial P_{t|k}^{i,jk}} \dot{P}_{t|k}^{i,jk} \right. \\ &\quad + \sum_{j=1}^n \left(f_j(t, \mathbf{x}) \frac{\partial p_{g_i}}{\partial x_j} + p_{g_i} \frac{\partial f_j(t, \mathbf{x})}{\partial x_j} \right. \\ &\quad \left. \left. + \frac{1}{2} \frac{\partial d_j^{(1)}(t, \mathbf{x}) p_{g_i}}{\partial x_j} - \frac{1}{2} \sum_{k=1}^n \frac{\partial^2 d_{jk}^{(2)}(t, \mathbf{x}) p_{g_i}}{x_j x_k} \right) \right] \\ d^{(1)}(t, \mathbf{x}) &= \frac{1}{2} \frac{\partial \mathbf{g}(t, \mathbf{x})}{\partial \mathbf{x}} \mathbf{Q} \mathbf{g}(t, \mathbf{x}) \\ d^{(2)}(t, \mathbf{x}) &= \frac{1}{2} \mathbf{g}(t, \mathbf{x}) \mathbf{Q} \mathbf{g}^T(t, \mathbf{x}) \end{aligned}$$

Notice, to carry out this minimization, we need to evaluate integrals involving Gaussian pdfs over volume V which can be computed exactly for polynomial nonlinearity and in general can be approximated by the Gaussian quadrature method.

The measurement update is done using Bayes rule, where the state and the covariance matrix are updated using the Extended Kalman Filter measurement update equations, and the weights are updated as it is shown in Ref. [2]. The equations can be found in Ref. [15].

By updating the forecast weights, not only can we obtain a more accurate estimate but also a better approximation to the conditional probability density function [15]. This weight update method during uncertainty propagation is very useful when the measurement model offer limited or no information in updating the states of the system.

The estimated conditional pdf is used to compute the expected loss. We require that the loss function provided is positive, finite everywhere and it is able to distinguish the important states from the unimportant ones. For simplicity the loss function used in this work has the following form:

$$L(\mathbf{x}_d, a_d) = \mathcal{N}(\mathbf{x}_d | \boldsymbol{\mu}_L, \boldsymbol{\Sigma}_L) \quad (16)$$

Observe that even with a better approximation of the weights of different Gaussian components, these components may drift away from the loss function due to first order Taylor series uncertainty propagation and limited information in the measurement update, situation which may be avoided if the conditional pdf can be found in an exact way.

Due to the approximations used in propagating the conditional pdf it may happen that no or very little probability density mass exists in the region of interest at the decision time, depicted here by the loss function. Thus the expected loss will be significantly underestimated, misleading this way the decision maker.

4 Decision Maker - Data Assimilation Interaction Level

The iterative method proposed here, is adding a set of Gaussian components to the initial pdf that are sensitive to the loss function at the decision time. After propagation, these Gaussian components will be located near the center of support of the loss function at the decision time. Initially the weights of these components are set to zero, and they will be updated in the propagation step if any probability density mass is moving in their direction. The weights at the decision time will give their relative contributions in computing the expected loss with respect to the entire conditional pdf.

An algorithm that bears similarity to the simulated annealing and the progressive correction used in particle filters [11], is proposed in selecting the initial Gaussian components.

The main idea is to select a set of Gaussian components initially, propagate each one of them using the time update equations in the Extended Kalman Filter until the decision time is reached and based on their contributions to the expected loss, select their means and variances in the initial distribution such that after propagation the expected loss is maximized.

Let the initial pdf be given by $p(t_0, \mathbf{x}_0)$ as a Gaussian sum,

Eq.8. Compute the mean and the variance of this mixture.

$$\boldsymbol{\mu}_0 = \sum_{i=1}^N w_0^i \boldsymbol{\mu}_0^i \quad (17)$$

$$\mathbf{P}_0 = \sum_{i=1}^N w_0^i [\mathbf{P}_0^i + (\boldsymbol{\mu}_0^i - \boldsymbol{\mu}_0)(\boldsymbol{\mu}_0^i - \boldsymbol{\mu}_0)^T] \quad (18)$$

Assume that we want to add another M new Gaussian components to the initial pdf with zero weights and sensitive to the loss function. We sample the means of these Gaussian components from the initial distribution such that their equally weighted sum gives the mean in Eq.17.

$$\boldsymbol{\mu}_i \sim p(t_0, \mathbf{x}_0) \quad \text{for } i = 1 \dots M-1 \quad (19)$$

$$\boldsymbol{\mu}_M = M\boldsymbol{\mu}_0 - \sum_{i=1}^{M-1} \boldsymbol{\mu}_i \quad (20)$$

The default covariance of the Gaussian components is \mathbf{D} . We want to find the new covariance \mathbf{D}^* such that the covariance of the new Gaussian components matches the covariance of the initial pdf. Let $\mathbf{D}^* = \gamma\mathbf{D}$. Thus we want to find γ such that we minimize the following expression:

$$J_\gamma = \text{Tr} \left[\mathbf{P}_0 - \frac{1}{M} \sum_{i=1}^M \left(\gamma\mathbf{D} + (\boldsymbol{\mu}_i - \boldsymbol{\mu}_0)(\boldsymbol{\mu}_i - \boldsymbol{\mu}_0)^T \right) \right] \quad (21)$$

$$\gamma = \frac{1}{\text{Tr}[\mathbf{D}]} \text{Tr} \left[\mathbf{P}_0 - \frac{1}{M} \sum_{i=1}^M (\boldsymbol{\mu}_i - \boldsymbol{\mu}_0)(\boldsymbol{\mu}_i - \boldsymbol{\mu}_0)^T \right] \quad (22)$$

Only solutions $\gamma > 0$ are accepted. Otherwise we repeat the sampling of the means, Eq. 19. Once we have the initial Gaussian sum components we propagate them using the time update equations in the Extended Kalman Filter until we reach the decision time. Let $\boldsymbol{\mu}_{t_d}^i$ and $\mathbf{P}_{t_d}^i$ be their means and covariances. The Gaussian components will then be weighted based on their contribution to the expected loss. A larger contribution means a more sensitive component to the loss function, thus a larger weight.

To be able to compute the weights of the Gaussian components, make sure that all of them are fairly weighted and we are not running into numerical problems and also create an indicator to mark the end of the algorithm, we compute an inflation coefficient for the loss function. Let $\boldsymbol{\Sigma}_L^* = \alpha\boldsymbol{\Sigma}_L$ be the inflated covariance of the loss function.

The inflation coefficient α is found such that the expected loss computed using the most distant Gaussian component from the loss function is maximized. Let the mean and the covariance of the most distant component be denoted by $\boldsymbol{\mu}_{t_d}^{max}$ and $\mathbf{P}_{t_d}^{max}$ respectively.

$$\begin{aligned} J_{max} &= \int \mathcal{N}(\mathbf{x}_d | \boldsymbol{\mu}_L, \alpha\boldsymbol{\Sigma}_L) \mathcal{N}(\mathbf{x}_d | \boldsymbol{\mu}_{t_d}^{max}, \mathbf{P}_{t_d}^{max}) d\mathbf{x}_d \\ &= \mathcal{N}(\boldsymbol{\mu}_L | \boldsymbol{\mu}_{t_d}^{max}, \alpha\boldsymbol{\Sigma}_L + \mathbf{P}_{t_d}^{max}) \end{aligned} \quad (23)$$

An equivalent way to seek α is by minimizing the negative log of the above expectation.

$$\begin{aligned} J_{min} &= \log[\det(\alpha\boldsymbol{\Sigma}_L + \mathbf{P}_{t_d}^{max})] + \\ &(\boldsymbol{\mu}_L - \boldsymbol{\mu}_{t_d}^{max})^T (\alpha\boldsymbol{\Sigma}_L + \mathbf{P}_{t_d}^{max})^{-1} (\boldsymbol{\mu}_L - \boldsymbol{\mu}_{t_d}^{max}) \end{aligned} \quad (24)$$

Let us denote $\mathbf{K} = \alpha\boldsymbol{\Sigma}_L + \mathbf{P}_{t_d}^{max}$ and $\mathbf{U} = (\boldsymbol{\mu}_L - \boldsymbol{\mu}_{t_d}^{max})(\boldsymbol{\mu}_L - \boldsymbol{\mu}_{t_d}^{max})^T$. We seek $\alpha > 0$ such that

$$\frac{\partial J_{min}}{\partial \alpha} = 0 \quad (25)$$

$$\text{Tr} \left[\mathbf{K}^{-1} \boldsymbol{\Sigma}_L - \mathbf{K}^{-1} \mathbf{U} \mathbf{K}^{-1} \boldsymbol{\Sigma}_L \right] = 0 \quad (26)$$

After a few mathematical manipulations, Eq.26 can be written in the following format:

$$\text{Tr} \left[\mathbf{K}^{-1} \boldsymbol{\Sigma}_L (\alpha\mathbf{I} + \mathbf{P}_{t_d}^{max} \boldsymbol{\Sigma}_L^{-1} - \mathbf{U} \boldsymbol{\Sigma}_L^{-1}) \mathbf{K}^{-1} \boldsymbol{\Sigma}_L \right] = 0 \quad (27)$$

Denote $\mathbf{A} = \mathbf{K}^{-1} \boldsymbol{\Sigma}_L$ and $\mathbf{B} = \alpha\mathbf{I} + \mathbf{P}_{t_d}^{max} \boldsymbol{\Sigma}_L^{-1} - \mathbf{U} \boldsymbol{\Sigma}_L^{-1}$. Observe that for $\alpha > 0$ the matrix \mathbf{A} is symmetric and positive definite.

Lemma: If $\text{Tr}[\mathbf{A}\mathbf{B}\mathbf{A}] = 0$ and \mathbf{A} is symmetric and positive definite then $\text{Tr}[\mathbf{B}] = 0$.

Proof: Let $\mathbf{A} = \mathbf{V}\mathbf{S}\mathbf{V}^T$ be a singular value decomposition of matrix \mathbf{A} , where \mathbf{V} is a unitary matrix and \mathbf{S} is a diagonal matrix. Our trace can now be written as $\text{Tr}[\mathbf{A}\mathbf{B}\mathbf{A}] = \text{Tr}[\mathbf{V}\mathbf{S}\mathbf{V}^T \mathbf{B} \mathbf{V}\mathbf{S}\mathbf{V}^T] = \text{Tr}[\mathbf{S}^2 \mathbf{B}]$.

If $\text{Tr}[\mathbf{S}^2 \mathbf{B}] = 0$ then $\mathbf{S}^2 \mathbf{B}$ is a commutator. Thus there is \mathbf{X} and \mathbf{Y} such that $\mathbf{S}^2 \mathbf{B} = \mathbf{X}\mathbf{Y} - \mathbf{Y}\mathbf{X}$. But $\mathbf{B} = \mathbf{S}^{-2} \mathbf{X}\mathbf{Y} - \mathbf{S}^{-2} \mathbf{Y}\mathbf{X} = \mathbf{X}^* \mathbf{Y} - \mathbf{Y}\mathbf{X}^*$, where $\mathbf{X}^* = \mathbf{S}^{-2} \mathbf{X}$. Therefore \mathbf{B} is also a commutator, hence $\text{Tr}[\mathbf{B}] = 0$.

Applying the previous lemma to Eq.27 we get

$$\text{Tr} \left[\alpha\mathbf{I} + \mathbf{P}_{t_d}^{max} \boldsymbol{\Sigma}_L^{-1} - \mathbf{U} \boldsymbol{\Sigma}_L^{-1} \right] = 0 \quad (28)$$

Therefore we accept solutions $\alpha > 1$ that satisfy the following relation

$$\alpha = \frac{1}{n} \text{Tr} \left[(\mathbf{U} - \mathbf{P}_{t_d}^{max}) \boldsymbol{\Sigma}_L^{-1} \right] \quad (29)$$

For $\alpha \leq 1$ we stop the algorithm. Otherwise we continue getting the weights of the Gaussian components by solving the following optimization problem:

$$\mathbf{w} = \arg \min_{\mathbf{w}} \frac{1}{2} \mathbf{w}^T \mathbf{M} \mathbf{w} - \mathbf{w}^T \mathbf{N} \quad (30)$$

$$\text{subject to } \mathbf{1}_{M \times 1}^T \mathbf{w} = 1 \quad (31)$$

$$\mathbf{w} \geq \mathbf{0}_{M \times 1} \quad (32)$$

where $\mathbf{w} \in \mathbb{R}^{M \times 1}$, $\mathbf{M} \in \mathbb{R}^{M \times M}$ and $\mathbf{N} \in \mathbb{R}^{M \times 1}$ and their entries are given by:

$$m_{ij} = \mathcal{N} \left\{ \boldsymbol{\mu}_{t_d|0}^j \mid \boldsymbol{\mu}_{t_d|0}^i, \mathbf{P}_{t_d|0}^i + \mathbf{P}_{t_d|0}^j \right\} \quad (33)$$

$$n_i = \mathcal{N} \left\{ \boldsymbol{\mu}_L \mid \boldsymbol{\mu}_{t_d|0}^i, \mathbf{P}_{t_d|0}^i + \boldsymbol{\Sigma}_L^* \right\} \quad (34)$$

The new pdf used to sample the new means is given by:

$$p_{new}(t_0, \mathbf{x}_0) = \sum_{i=1}^M \mathcal{N}(\mathbf{x}_0 \mid \boldsymbol{\mu}_i, \beta \mathbf{D}^*) \quad (35)$$

Algorithm 1 Progressive Selection of Gaussian Components

Require: t_d - decision time

$p(t_0, \mathbf{x}_0)$ - initial probability density function

M - number of extra Gaussian components

\mathbf{D} - Gaussian component covariance

w_{tol} - add only Gaussian components with weights greater than this threshold

$L(\mathbf{x})$ - loss function $\mathcal{N}\{\mathbf{x}|\boldsymbol{\mu}_L, \boldsymbol{\Sigma}_L\}$

- 1: $\hat{p}_0 = p(t_0, \mathbf{x}_0)$, $\alpha = \infty$, $\gamma = -1$
 - 2: **while** ($\alpha > 1$) & maxiter **do**
 - 3: The mean and the covariance of the initial pdf
 $\boldsymbol{\mu}_0 = \sum_{i=1}^N w_0^i \boldsymbol{\mu}_0^i$
 $\mathbf{P}_0 = \sum_{i=1}^N w_0^i [\mathbf{P}_0^i + (\boldsymbol{\mu}_0^i - \boldsymbol{\mu}_0)(\boldsymbol{\mu}_0^i - \boldsymbol{\mu}_0)^T]$
 - 4: **while** ($\gamma < 0$) **do**
 - 5: Get the means of the Gaussian components
 Draw $\boldsymbol{\mu}_i \sim p(t_0, \mathbf{x}_0)$ for $i = 1 \dots M - 1$
 Set $\boldsymbol{\mu}_M = M\boldsymbol{\mu}_0 - \sum_{i=1}^{M-1} \boldsymbol{\mu}_i$
 - 6: $\gamma = \frac{1}{\text{Tr}[\mathbf{D}]} \text{Tr} \left[\mathbf{P}_0 - \frac{1}{M} \sum_{i=1}^M (\boldsymbol{\mu}_i - \boldsymbol{\mu}_0)(\boldsymbol{\mu}_i - \boldsymbol{\mu}_0)^T \right]$
 - 7: **end while**
 - 8: Get the covariance of the Gaussian components
 $\mathbf{P}_0^i = \gamma \mathbf{D}$
 - 9: Propagate the moments from $t = 0$ to $t = t_d$
 $\dot{\boldsymbol{\mu}}_{t|0}^i = \mathbf{f}(t, \boldsymbol{\mu}_{t|0}^i)$
 $\dot{\mathbf{P}}_{t|0}^i = \mathbf{A}_{t|0}^i \mathbf{P}_{t|0}^i + \mathbf{P}_{t|0}^i (\mathbf{A}_{t|0}^i)^T + \mathbf{g} \mathbf{Q} \mathbf{g}^T$
 - 10: Get the most distant component
 by computing the Mahalanobis distance
 $d_i = (\boldsymbol{\mu}_L - \boldsymbol{\mu}_{t_d|0}^i)^T (\mathbf{P}_{t_d|0}^i + \boldsymbol{\Sigma}_L)^{-1} (\boldsymbol{\mu}_L - \boldsymbol{\mu}_{t_d|0}^i)$
 $\boldsymbol{\mu}_{t_d|0}^{\text{max}}, \mathbf{P}_{t_d|0}^{\text{max}} = \arg \max(d_i)$
 - 11: Compute optimal α and the inflated matrix $\boldsymbol{\Sigma}_L^*$
 $\alpha = \frac{1}{n} \text{Tr} \left[\left((\boldsymbol{\mu}_{t_d|0}^{\text{max}} - \boldsymbol{\mu}_L)(\boldsymbol{\mu}_{t_d|0}^{\text{max}} - \boldsymbol{\mu}_L)^T - \mathbf{P}_{t_d|0}^{\text{max}} \right) \boldsymbol{\Sigma}_L^{-1} \right]$
 - 12: **if** $\alpha < 1$ **then** $\alpha = 1$ **end if**
 $\boldsymbol{\Sigma}_L^* = \alpha \boldsymbol{\Sigma}_L$
 - 13: Elements of $\mathbf{M} \in \mathbb{R}^{M \times M}$ and $\mathbf{N} \in \mathbb{R}^{M \times 1}$
 $m_{ij} = \mathcal{N} \left\{ \boldsymbol{\mu}_{t_d|0}^j \mid \boldsymbol{\mu}_{t_d|0}^i, \mathbf{P}_{t_d|0}^i + \mathbf{P}_{t_d|0}^j \right\}$
 $n_i = \mathcal{N} \left\{ \boldsymbol{\mu}_L \mid \boldsymbol{\mu}_{t_d|0}^i, \mathbf{P}_{t_d|0}^i + \boldsymbol{\Sigma}_L^* \right\}$
 - 14: Compute the weights
 $\mathbf{w} = \arg \min_{\mathbf{w}} \frac{1}{2} \mathbf{w}^T \mathbf{M} \mathbf{w} - \mathbf{w}^T \mathbf{N}$
 subject to $\mathbf{1}_{M \times 1}^T \mathbf{w} = 1$
 $\mathbf{w} \geq \mathbf{0}_{M \times 1}$
 - 15: Set $\hat{p}_0 = \sum_{j=1}^M w_j \mathcal{N}\{\mathbf{x}|\boldsymbol{\mu}_0^j, \beta \mathbf{P}_0^j\}$
 - 16: **end while**
 - 17: Set $p_{\text{NEW}}(t_0, \mathbf{x}_0) = p(t_0, \mathbf{x}_0) + \sum_{j=1}^M w_j \mathcal{N}\{\mathbf{x}|\boldsymbol{\mu}_0^j, \beta \mathbf{P}_0^j\}$
 - 18: **return** $p_{\text{NEW}}(t_0, \mathbf{x}_0)$
-

Where $\beta \leq 1$ is a coefficient that controls the decrease of the initial variance. If α has decreased from the previous iteration this means that the Gaussian components are getting closer to the loss function and therefore we can decrease the

variance of the initial distribution to finely tune the position of the Gaussian components, otherwise $\beta = 1$. We continue to sample new means from the new pdf until $\alpha < 1$ or the maximum number of time steps has been reached. The entire algorithm is presented in Table 1. While not the scope of this paper, the above method can also be applied when measurements are available between the current time and the decision time. The progressive selection algorithm will be applied every time a measurement has been assimilated and the a posteriori pdf has been found. The drawback of this procedure is that the number of Gaussian components will increase linearly with the number of measurements. Better ways to deal with the measurement updates are set as future work.

In the case of multiple loss functions, the algorithm is run once for each one of the loss functions, creating sets of initial Gaussian components sensitive to their loss function.

5 Numerical Results

To illustrate the concept of incorporating contextual information into the uncertainty propagation algorithm, we consider the following continuous-time dynamical system with uncertain initial condition given by (36):

$$\begin{aligned} \dot{x} &= \sin(x) + \Gamma(t) \quad \text{where} \quad Q = 1 \quad (36) \\ x_0 &\sim \mathcal{N}(-0.3, 0.3^2) \end{aligned}$$

The state space region of interest is depicted by the following loss function, and the time of decision is at $t_d = 8$ sec.

$$L(x) = \mathcal{N}(x \mid \frac{\pi}{2}, 0.1^2) \quad (37)$$

First we compute an accurate numerical solution based on the FPKE, and this will stand as the true probability density function. The performance measures for this method will be labeled as TRUTH. The evolution of the pdf using this method can be seen in Fig.2a.

Three other approximations for the pdf are provided including the method presented in this paper. The first approximation propagates the initial uncertainty using the first order Taylor expansion, Eq.(11) and Eq.(12), also known as the Extended Kalman Filter time update equations, labeled later as EKF. The evolution of the pdf for this method is presented in Fig.2b.

For the next approximation method, we add another 5 Gaussian components to the initial one, creating this way a Gaussian mixture with 6 components. The means of the new components are just the result of back propagation (from $t_d = 8$ sec to $t_c = 0$ sec) of 5 equidistant samples taken in the 3 sigma bound of the loss function support. The variance of the new components is set to 10^{-10} and their initial weights are set to zero. The label used for this method is GS_BCK and the evolution of the pdf is shown in Fig.2c. While all the means of the new Gaussian components are positioned in the loss function support region, their variances are very being difficult to see the probability density mass in that region.

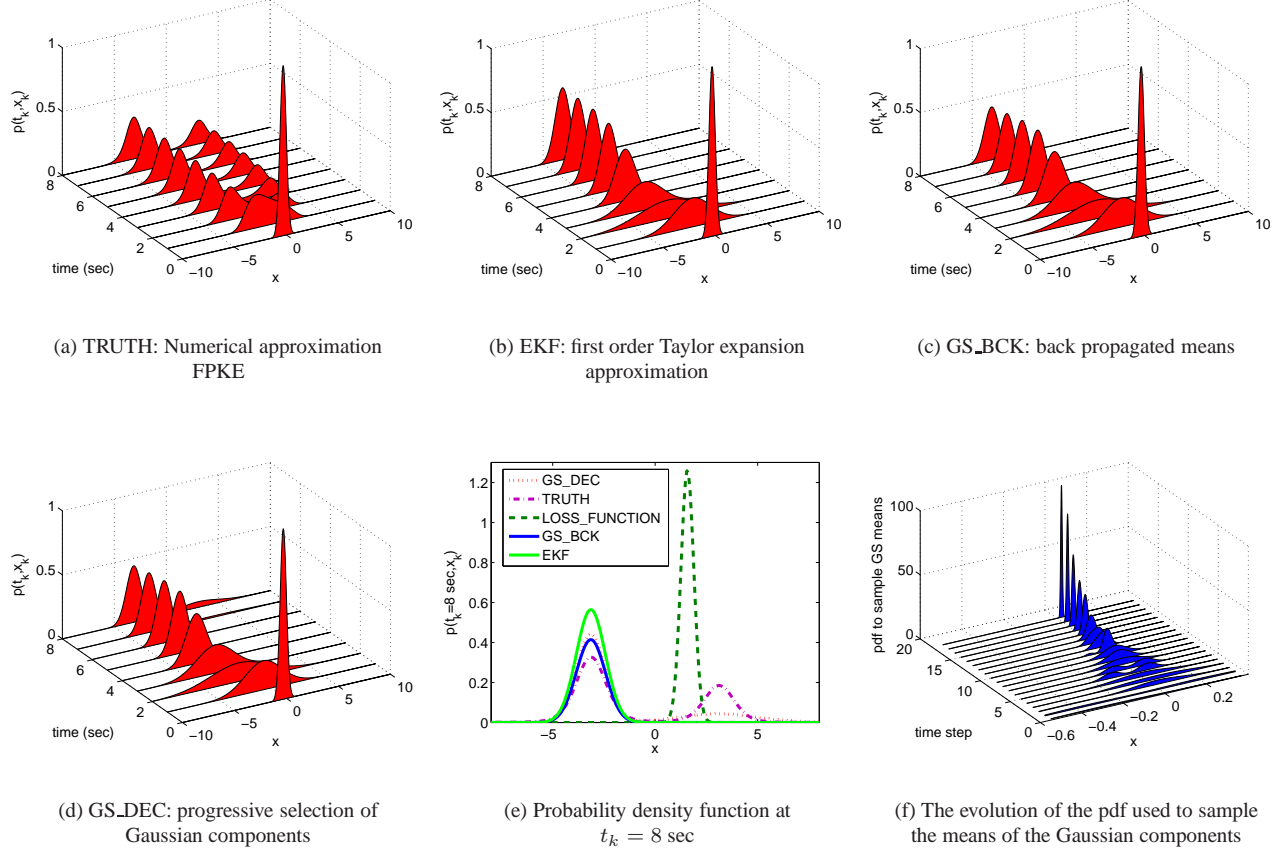


Figure 2: The evolution of the Forecast pdf and the Sampling pdf

We apply the method presented in this paper to generate at most 5 new Gaussian components to be added to the initial condition. Their means and variances are returned by the progressive selection algorithm, Alg.1. The initial weights of the new Gaussian components have also been set to zero. The default value for the β coefficient is 0.9 and Gaussian components are included only if their weights are greater than $w_{tol} = 10^{-3}$. The label used for this method is GS_DEC and its corresponding pdf is presented in Fig.2d.

The evolution the Gaussian components for the last two methods is also achieved using the first-order Taylor expansion, but it is interrupted every $\Delta t = 0.5$ sec to adjust the weights of different Gaussian components using the optimization in Eq.(14).

The following performance measures have been computed for the methods used in the experiment:

$$\hat{L}_d = \int L(x) \hat{p}(t_d, x_d) dx_d \quad (38)$$

$$\hat{R}_{err} = \frac{1}{L_d} |L_d - \hat{L}_d| \quad (39)$$

$$ISD = \int |p(t_d, x_d) - \hat{p}(t_d, x_d)|^2 dx_d \quad (40)$$

$$WISD = \int L(x) |p(t_d, x_d) - \hat{p}(t_d, x_d)|^2 dx_d \quad (41)$$

In Fig.2e it is plotted the forecast pdf at time t_d for all the methods. Our method, GS_DEC, is able to better estimate the probability density mass in the region of interest.

In Table I, we present the performance measures after 500 Monte Carlo runs. The expected loss given by the GS_DEC method is consistently better approximated over all the Monte Carlo runs than the EKF and the GS_BCK method. We also are able to consistently give a better approximation to the pdf overall and also in the region of interest than the EKF method, which justify the use of this method. Compared with the GS_BCK we do a better job in average in approximating the pdf which suggests that there is a trade off in selecting the Gaussian components regarding their means and variances.

In Fig. 2f it is plotted the evolution of the pdf used to sample the means of the new Gaussian components for one particular run. The pdf used in the first iteration is our initial uncertainty and we see how it converges, as the number of iterations increases, to a particular region in the state space that is sensitive to the loss function at the decision time.

6 Conclusion

An interaction level between the decision maker and the data assimilation module has been designed, such that we

Table 1: Performance measures - 500 Monte Carlo runs

	\hat{L}_d	\hat{R}_{err}	ISD	$WISD$
TRUTH	0.0332	N/A	N/A	N/A
EKF	4.93E-09	1.0000	0.1840	0.0015
GS_BCK	0.0001	0.9968	0.0536	0.0015
GS_DEC (mean)	0.0256	0.2300	0.0470	0.0004
GS_DEC: Percentile Table - 500 Observations				
Percent	\hat{L}_d	\hat{R}_{err}	ISD	$WISD$
0.0%	0.0010	0.0151	0.0368	0.0002
5.0%	0.0142	0.0230	0.0378	0.0003
10.0%	0.0177	0.0271	0.0380	0.0003
25.0%	0.0229	0.0566	0.0387	0.0003
50.0%	0.0257	0.2270	0.0491	0.0003
75.0%	0.0313	0.3090	0.0514	0.0004
90.0%	0.0323	0.4670	0.0574	0.0006
95.0%	0.0324	0.5710	0.0601	0.0007
100.0%	0.0327	0.9700	0.0705	0.0014

can incorporate contextual information held by the decision maker into the data assimilation process to better approximate the conditional probability function.

The progressive selection algorithm is run once at the beginning of the simulation to supplement the initial uncertainty with new Gaussian components that are sensitive to the loss function at the decision time. The weights of all the Gaussian components are then updated during the propagation based on the Fokker Planck Equation. This way we obtain not only a better approximation of the probability density function in the region of interest but also a better approximation overall.

The cost of this overall improvement is an increase in the number of Gaussian components. The principal benefit is not the modest increase in accuracy overall, but the significantly enhanced accuracy within the decision maker's region of interest.

Acknowledgment: *This work was supported under Contract No. HMI582-08-1-0012 from ONR.*

References

- [1] H. Alspach, D.; Sorenson. Nonlinear bayesian estimation using gaussian sum approximations. *Automatic Control, IEEE Transactions on*, 17(4):439–448, 1972.
- [2] Brian D.O. Anderson and John B. Moore. *Optimal Filtering*. Prentice-Hall, 1979.
- [3] R.N. Banavar and J.L. Speyer. Properties of risk-sensitive filters/estimators. *IEE Proc.Control Theory and Applications*, 145:106–112, 1998.
- [4] J. Crassidis and J. Junkins. *Optimal Estimation of Dynamic Systems*. CRC Press, 2004.
- [5] A. Doucet, N. de Freitas, and N. Gordon. *Sequential Monte-Carlo Methods in Practice*. Springer-Verlag, April 2001.
- [6] K. Ito, K.; Xiong. Gaussian filters for nonlinear filtering problems. *Automatic Control, IEEE Transactions on*, 45(5):910–927, 2000.
- [7] R. N. Iyengar and P. K. Dash. Study of the random vibration of nonlinear systems by the gaussian closure technique. *Journal of Applied Mechanics*, 45:393–399, 1978.
- [8] M. Kumar, P. Singla, S. Chakravorty, and J. L. Junkins. The partition of unity finite element approach to the stationary fokker-planck equation. In *2006 AIAA/AAS Astrodynamics Specialist Conference and Exhibit, Keystone, CO*, Aug. 21-24, 2006.
- [9] T. Lefebvre, H. Bruyninckx, and J. De Schutter. Kalman filters of non-linear systems: A comparison of performance. *International journal of Control*, 77(7):639–653, 2004.
- [10] G. Muscolino, G. Ricciardi, and M. Vasta. Stationary and non-stationary probability density function for non-linear oscillators. *International Journal of Non-Linear Mechanics*, 32:1051–1064, 1997.
- [11] C. Musso, N. Oudjane, and F. Le Gland. *Sequential Monte Carlo Methods in Practice, Statistics for Engineering and Information Science*, chapter 12: Improving regularized particle filters, pages 247–271. Springer-Verlag, New York, 2001.
- [12] H. Risken. *The Fokker-Planck Equation: Methods of Solution and Applications*. Springer, 1989.
- [13] J. B. Roberts and P. D. Spanos. *Random Vibration and Statistical Linearization*. Wiley, 1990.
- [14] Gabriel Terejanu, Tarunraj Singh, and Peter D. Scott. Unscented Kalman Filter/Smother for a CBRN Puff-Based Dispersion Model. In *The 10th Intl. Conference on Information Fusion, Quebec City, Canada*, 2007.
- [15] Gabriel Terejanu, Puneet Singla, Tarunraj Singh, and Peter D. Scott. A novel gaussian sum filter method for accurate solution to nonlinear filtering problem. In *The 11th International Conference on Information Fusion, Cologne, Germany*, 2008.
- [16] Gabriel Terejanu, Puneet Singla, Tarunraj Singh, and Peter D. Scott. Uncertainty propagation for nonlinear dynamical systems using gaussian mixture models. *Journal of Guidance, Control, and Dynamics*, 31:1623–1633, 2008.
- [17] S. Thrun, Langford. J., and V. Verma. Risk sensitive particle filters. In *Advances in Neural Information Processing Systems 14*. MIT Press, 2002.