Javad Rezaie

A Bayesian Inversion Approach to Filtering and Decision Making with Applications to Reservoir Characterization

Thesis for the degree of Philosophiae Doctor

Trondheim, October 2013

Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics and
Electrical Engineering
Department of Mathematical Sciences

NTNU – Trondheim
Norwegian University of
Science and Technology

NTNU
Norwegian University of Science and Technology

Thesis for the degree of Philosophiae Doctor

Faculty of Information Technology,
Mathematics and Electrical Engineering
Department of Mathematical Sciences

ISBN 978-82-471-4679-8 (printed version)
ISBN 978-82-471-4680-4 (electronic version)
ISSN 1503-8181

Doctoral theses at NTNU, 2013:275

Printed by NTNU-trykk

تقدیم به پدر و مادرم

To my parents

Til foreldrene mine

ز هر دانشی چون سخن بشنوی   از آموختن یک زمان نغنوی

چو دیدار یابی به شاخ سخن   بدانی که دانش نیاید به بن

Preface

This thesis is submitted in partial fulfilment of the requirements for the degree of philosophiae doctor (PhD) at the Norwegian University of Sciences and Technology (NTNU) in Trondheim. The research is funded by the Uncertainty in Reservoir Evaluation (URE) consortium at NTNU.

Acknowledgements

First and foremost, I would like to thank Professor Jo Eidsvik, my main supervisor, for his invaluable helps throughout these four years. He is a fantastic supervisor and his guidance has been much appreciated and will be missed. He have spent too much time for me. I have learnt a lot from him, both professionally and personally. He have been always available and open to discuss. I have started learning statistics with him as a beginner in this field.

I gratefully acknowledge Professor Henning Omre, my co-supervisor, for his advice, supervision and guidance throughout these four years. The discussions I have had with Henning have been extremely rewarding. He is always full of energy and positive. I also have learnt a lot from him in both scientific and social points of view.

Jo and Henning are very warm-hearted and would like to appreciate them again because of providing me a very good atmosphere for study and live. They have been more than a supervisor for me, and I feel they are my family members.

I'm grateful to Eivind Smørgrav for giving me the opportunity to do some part of my thesis at Statoil. He is a brilliant scientist, humble and gentleman. I learnt a lot from him and I had a very good collaboration with him and Jon Sætrom.

Besides, I would like to thank Professor Tapan Mukerji for giving me the opportunity to visit Stanford University, USA. He was really helpful during that period and after that. I learnt many things from him and the Stanford

Center for Reservoir Forecasting (SCRF) group.

Thanks to my teachers at the Department of Mathematical Sciences, Professor Håkon Tjelmeland, Professor Nikolai Ushakov, Professor Bo Henry Lindqvist, Professor John Tyssedal, Professor Arvid Næss and Professor Mette Langaas.

Moreover, I would like to thank Anne Kajander for all of her helps. She is a very kind lady, and she have helped from before starting my PhD and after finishing it.

I highly appreciate the support and encouragement of my friends in Trondheim, who have been very important to me in this period.

My final thanks are reserved for my parents, brothers and sisters for their moral support and their encouragements throughout my life, and specially these years that I was far a way from them

Javad Rezaie
Trondheim, Norway, June 2013

Abstract

The main objectives of this thesis are the estimation/filtering and decision making problems with a Bayesian inversion point of view, and in geophysical systems. In addition, determining the information content in the measured data is also a challenge in estimation problems, and we use the dimension reduction techniques to deal with this problem. The main applications of the proposed algorithms are for reservoir characterization, and the seismic amplitude versus offset (AVO) data is the most used measurement.

The first part of this thesis tries to address some of the existing problems in the state estimation of high dimensional and complex systems. Our first proposal is a robustified Gaussian mixture filter. Simulations show promising results and the performance of the proposed filter is at least as good as the ensemble Kalman filter (EnKF) and particle filter (PF). In addition, we extend the traditional KF and EnKF for capturing the skewness of the distributions. They automatically converge to the KF or EnKF if there is no skewness in the probability density function (pdf). Simulation results confirm our claim, and they seem to have better performance in the presence of skewness. Furthermore, we investigate the nature of geophysical observations from a filtering point of view by testing several data reduction techniques. We show how to assess the information content in the data, compress the data, and use this compressed data in a reservoir conditioning setting. The methods we present are generic; they apply equally well to all geophysical attributes regardless of representation and can be applied with any filtering algorithm. The last part of this thesis relates to the value of information (VOI) analysis and decision making. We extend the previous method for computing the VOI of seismic AVO data by using a closed skew normal pdf model instead of the Gaussian. The previous method is an special case of the proposed method, and simulation results seems to result in more reliable decisions.

Publication lists

The following four papers construct the main body of this thesis (the last four chapters):

- **Paper I:** Javad Rezaie and Jo Eidsvik (2012). *Shrinked* $(1-\alpha)$ *ensemble Kalman filter and* α *Gaussian mixture filter,* Computational Geoscience, 16(3):837-852, 2012.
- **Paper II:** Javad Rezaie and Jo Eidsvik (2013) *Kalman Filter Variants in the Closed Skew Normal Setting*, In revision.
- **Paper III:** Javad Rezaie, Jon Sætrom and Eivind Smørgrav (2012) *Reducing the Dimensionality of Geophysical Data in Conjunction with Seismic History Matching,* Proceedings of the 74^{th} EAGE Conference & Exhibition incorporating SPE EUROPEC 2012, Copenhagen, Denmark, SPE153924.
- **Paper IV:** Javad Rezaie, Jo Eidsvik and Tapan Mukerji *Value of Information Analysis and Bayesian Inversion for Closed Skew-Normal Distributions: Applications to Seismic Amplitude Versus Offset Data*, In revision.

Furthermore, we presented the following papers and extended abstracts at international conferences and congress:

- **Paper V:** Javad Rezaie and Jo Eidsvik (2010) *A shrinked Gaussian mixture filter for non-linear dynamics*, The 23^{rd} Nordic conference on Mathematical Statistics (NORDSTAT), June 2010, Voss, Norway.
- **Paper VI:** Javad Rezaie and Jo Eidsvik (2011) *The Value of Information for Monitoring Surveys*, The 73^{rd} EAGE Conference & Exhibition incorporating SPE EUROPEC 2011, May 2011, Vienna, Austria.
- **Paper VII:** Javad Rezaie and Jo Eidsvik (2012) *A Gaussian mixture Monte Carlo filter for state estimation in very high dimensional*

nonlinear systems, The Ninth International Geostatistics Congress (Geostats 2012), June 2012, Oslo, Norway.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation and Background

Statistics is a part of mathematical sciences which deals with uncertain quantities. Dealing with uncertain variables is problem specific. The statistical inference may consist of various tasks such as reducing the uncertainty range, estimating the conditional probability density function (pdf) or some statistics such as mean and variance, decision making under uncertainty etc. Among them, estimation of latent variables using noisy and incomplete observations is one of the most important problems in engineering and science.

Assume that the parameters of a system are modeled in the state space form with a set of differential or difference equations. This model contains lots of uncertainty because of un-modeled dynamics, modeling error, simplification errors etc. Moreover, the measurement sensors and tools provide incomplete and noisy observations. The estimation problem is now defined as inferring all system parameters using the available observations. Generally speaking, the estimation procedure increases the accuracy of parameters (by reducing the associated prior uncertainty).

In order to go into more details, it is better to explain one of the most used systems in this thesis. Consider a petroleum reservoir with $n_1 \times n_2 \times n_3$ grid cells, a single injection well and a single production well. The water is pumped into the injection well for replacing and moving oil to the production well. The simulator works on fluid dynamics laws, and it contains some parameters such as porosity, permeability and saturation (see Figure 1.1). Based on fluid dynamics, the flow is faster where the permeability (porosity) is high. The values of these parameters depend on the geological and geophysical characteristics of the reservoir.

Assume the parameters of interest at time t are arranged in a state

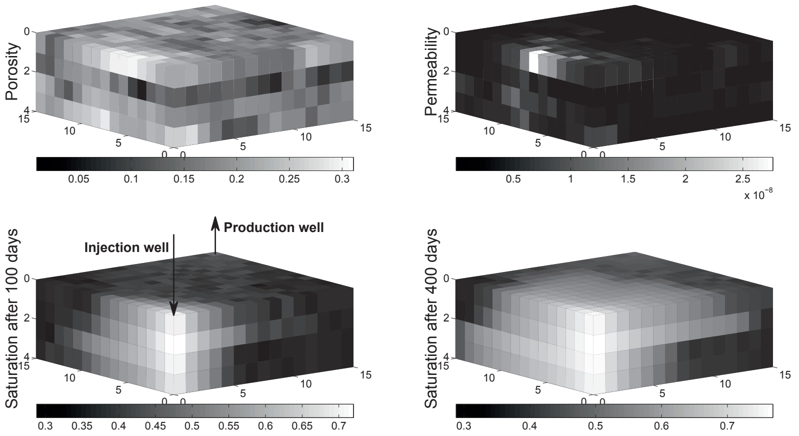

Figure 1.1: The reservoir description: upper plots are the reference porosity and permeability from SPE10 data set; lower left plot shows the position of the injection and production wells; the lower left plot is the saturation after 100 days of running the simulator and the lower right plot is the saturation after 400 days.

vector $\boldsymbol{x}_t \in \Re^{n_x \times 1}$. The fluid laws present the dynamics of the reservoir parameters as partial differential equations (PDE), which may be written in an explicit form as $\boldsymbol{x}_t = \boldsymbol{f}(\boldsymbol{x}_{t-1}; \boldsymbol{\eta}_t)$. This equation is also known as the process model. Here, the nonlinear function $\boldsymbol{f}(\cdot) : \Re^{(n_x+n_\eta)\times 1} \rightarrow \Re^{n_x \times 1}$ is the dynamic equation and $\boldsymbol{\eta}_t \in \Re^{n_\eta \times 1}$ is the process noise with known pdf. The process noise includes the uncertainties in the modeling (i.e. unmodeled dynamics). In this thesis, we assume additive process noise, i.e. $\boldsymbol{x}_t = \boldsymbol{f}(\boldsymbol{x}_{t-1}) + \boldsymbol{\eta}_t$.

The Society of Petroleum Engineers (SPE) organized a series of projects, known as SPE comparative projects, in order to provide benchmark data sets which can be used to compare the performance of different algorithms and methods applied for petroleum reservoir evaluation. The 10th SPE comparative project is the latest one in this series, and known as the SPE10 data set (Christie and Blunt, 2001). The SPE10 data set consists of porosity and permeability for $60 \times 220 \times 85$ Cartesian grid cells. By using this data set as the input to a reservoir flow solver (simulator), the saturation of that reservoir can be used for further evaluations. There are several commercial and non-commercial reservoir simulators (e.g. ECLIPSE by GeoQuest). In this thesis, we will use the MATLAB Reservoir Simulation Toolbox (MRST), see Lie et al. (2012), as flow solver for the saturation (More details about

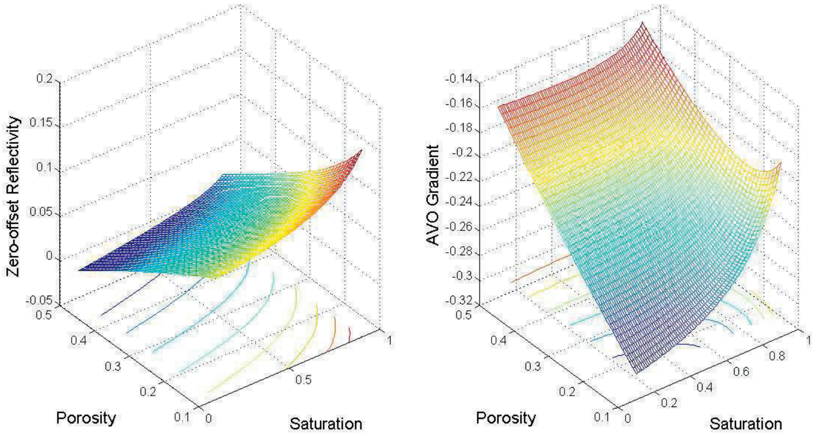

Figure 1.2: Zero-offset reflectivity (left plot) and AVO gradient (right plot) versus saturation and porosity.

the simulator are available at http://www.sintef.no/Projectweb/MRST/).

In addition, assume that at each time t some measurements, $\boldsymbol{d}_t \in \Re^{n_d \times 1}$ are available by the sensors. These measurements relate to the system state by $\boldsymbol{d}_t = \boldsymbol{h}(\boldsymbol{x}_t; \boldsymbol{\epsilon}_t)$, where $\boldsymbol{h}(\cdot) : \Re^{(n_x+n_\epsilon)\times 1} \rightarrow \Re^{n_d \times 1}$ is the measurement equations and $\boldsymbol{\epsilon}_t \in \Re^{n_\epsilon \times 1}$ is the observation noise. Throughout this thesis, the observation noise is assumed to be additive, $\boldsymbol{d}_t = \boldsymbol{h}(\boldsymbol{x}_t) + \boldsymbol{\epsilon}_t$.

Geophysical data are directly or indirectly informative of important subsurface parameters. Seismic measurements, for example, may provide a rich source of information about structures, lithologies and hydrocarbon indicators. The seismic amplitude versus offset (AVO) data is related to the reservoir saturation and porosity through rock physics relations (for more details see Mavko et al. (2003)). In most of this thesis, the data $\boldsymbol{d}_t$ consist of zero-offset reflectivity and AVO gradient at top reservoir (see Figure 1.2). We assume that the observations are made at all lattice cells. We assume conditional independent data: the seismic AVO data at one grid cell only depends on the saturation and porosity at that cell. The expected value of the AVO data is displayed in Figure 1.2, as a function of porosity and saturation. Assimilation of these data reduces the uncertainty, improves prediction of the reservoir and production. For instance small observation of zero-offset reflectivity and large observation of AVO gradient indicate large porosity and saturation. The likelihood is slightly nonlinear, but we linearize the measurement equation using first order Taylor series expansion to get $\boldsymbol{h}(\boldsymbol{x}_t) \approx \boldsymbol{h}_0 + \boldsymbol{H}\boldsymbol{x}_t$, where $\boldsymbol{h}_0 \in \Re^{n_d \times 1}$ and $\boldsymbol{H} \in \Re^{n_d \times n_x}$.

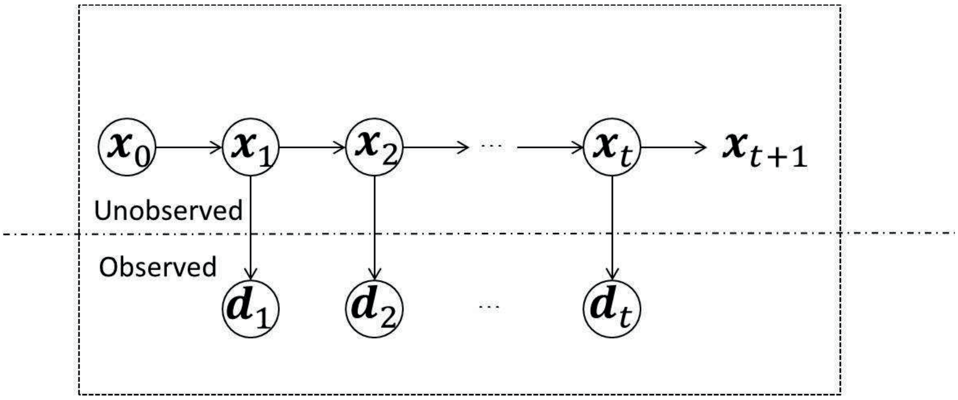

Figure 1.3: Graphical representation of state variables $\boldsymbol{x}_t$, at discrete time points $t = 0, 1, \ldots$ and observations $\boldsymbol{d}_t$, $t = 1, 2, \ldots$. The process model is assumed to follow a Markov structure. The data are assumed to be conditionally independent, given the state variable at the indicated time steps. The filtering problem characterizes the distribution of states over time, given all currently available data.

The reservoir dynamics and observation equations in the state space form then becomes:

$$\begin{aligned} \boldsymbol{x}_t &= \boldsymbol{f}(\boldsymbol{x}_{t-1}) + \boldsymbol{\eta}_t, \\ \boldsymbol{d}_t &= \boldsymbol{h}(\boldsymbol{x}_t) + \boldsymbol{\epsilon}_t. \end{aligned} \tag{1.1}$$

We use the state space formulation with the usual conditional independence assumptions. This means that the conditional distribution of $\boldsymbol{x}_t$, given all previous state and observation variables, only depends on the state at time $t-1$. Moreover, the conditional distribution of observation $\boldsymbol{d}_t$, given the state at that time and all previous states and observations, only depends on the state at time t. Mathematically these two assumptions entail that the conditional distributions are $\pi(\boldsymbol{x}_t|\boldsymbol{X}_{t-1}, \boldsymbol{D}_{t-1}) = \pi(\boldsymbol{x}_t|\boldsymbol{x}_{t-1})$ and $\pi(\boldsymbol{d}_t|\boldsymbol{X}_t, \boldsymbol{D}_{t-1}) = \pi(\boldsymbol{d}_t|\boldsymbol{x}_t)$, respectively (see Figure 1.3). Here, $\boldsymbol{X}_t = [\boldsymbol{x}_0, \boldsymbol{x}_1, \cdots, \boldsymbol{x}_t]$ and $\boldsymbol{D}_t = [\boldsymbol{d}_1, \boldsymbol{d}_2, \cdots, \boldsymbol{d}_t]$ denote the collection of states and data respectively.

The main goal of this thesis is monitoring the state variables (here, the saturation is the most important one) over time. State variable monitoring can be defined in an estimation setting by estimating the posterior distribution $\pi(\boldsymbol{x}_t|\boldsymbol{D}_t)$. In the prediction context, it is also defined as finding the prediction distribution $\pi(\boldsymbol{x}_{t+1}|\boldsymbol{D}_t)$. In this thesis, the Bayesian inversion is the main approach for finding the posterior and prediction distribution.

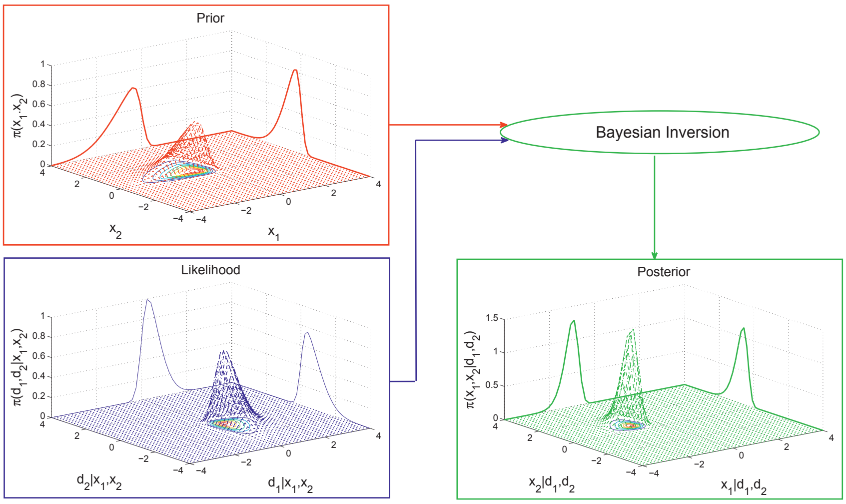

Figure 1.4: Illustration of Bayesian inversion. The upper left plot is the prior distribution, the lower left plot is the likelihood, and the right plot is the resulting posterior.

Bayesian inversion directly relates to the filtering, prediction and decision making problems. There are some difficulties in the Bayesian inversion for very high dimensional and complicated systems. Consequently, similar difficulties exist for Bayesian inversion based algorithms (i.e. the filtering and prediction problems), and in this thesis we try to introduce and address some of them.

1.2 Bayesian Inversion

The scientists are very interested in reducing the uncertainty associated with the prior pdf denoted $\pi(\boldsymbol{x})$. This is achieved by assimilating some measurements which relate to $\boldsymbol{x}$. More generally speaking, one updates the distribution of latent variables given observations $\boldsymbol{d}$. The result is known as the posterior distribution $\pi(\boldsymbol{x}|\boldsymbol{d})$. Bayes' rule is very helpful for dealing with conditional distributions. It says that the posterior distribution of the latent variables given observations is proportional to the likelihood times the prior. In geoscience and petroleum engineering this problem is known as Bayesian inversion. Consider $\pi(\boldsymbol{d}|\boldsymbol{x})$ as the likelihood of the seismic AVO data, given the reservoir variables. Bayes' theorem combines prior information $\pi(\boldsymbol{x})$ and current observations, in the probability domain, in order to get the posterior distribution of the variables of interest:

$$\pi(\boldsymbol{x}|\boldsymbol{d}) \propto \pi(\boldsymbol{d}|\boldsymbol{x})\pi(\boldsymbol{x}). \tag{1.2}$$

Figure 1.4 shows a schematic diagram of a Bayesian inversion problem from a distributional point of view. The prior knowledge (upper left plot) is here a bivariate distribution. Bayes rule adds the information content of the likelihood distribution (lower left plot) to it. The posterior distribution is shown in the right plot, and we see that the uncertainties are reduced.

Bayesian inversion solely or as a part of other problems plays an important role in the data assimilation/state estimation or decision making problems. Thus, we highlight those problems with Bayesian inversion as an ingredient in the following sections.

Before going into more details, we summarize the challenges in this setting:

- The first challenge is the state estimation over time in high dimensional systems. Most of the filtering methods in high dimensional system are ensemble based. In these methods, a limited number of ensembles, say $100 - 200$, is used. The reason is that the forward propagation of ensembles is very time consuming (i.e. couple of days). Thus, for real time implementation we should restrict the algorithms to use this limited number of ensembles. On the other hand, the actual probability distribution of practical systems is complicated (i.e. skewed, multi-modal) and the Gaussian assumption (which is used in many ensemble based filters) may fail. Many ensemble members are required to estimate the complicated distributions. It is a dilemma, and the filtering algorithm proposal(s) should work sufficiently accurate with this limited number of ensembles.

- Another challenging problem in data assimilation is the existence of collinear or highly correlated data. These correlated data sets cause problems related to rank deficiency and overfitting. Hence, reducing the dimensionality of the data, by extracting the useful information and discarding the noise (prior to assimilation), is natural to consider.

- Finally, in decision making problems (i.e. decision on drilling/not drilling new wells for finding hidden pockets of oil) we use for example seismic AVO measurements to make better, informed, decisions. The seismic AVO data requires data processing before use. Different data processing levels directly relates to the price of data. Thus, we should find the optimal data processing level such that it is accurate to an acceptable level, and also it is not so expensive. This relates to the value of information.

Different parts of this thesis propose some novel algorithms or extend the previous algorithms to address these challenges. In all tasks, Bayesian inversion plays an important role. Thus, various flavors of equation (1.2) become useful throughout the thesis.

1.3 Filtering Problem as a Spatio-temporal Bayesian Inversion

Filtering or state estimation (also known as data assimilation) applies Bayesian inversion over time. At each time t, it finds the updated distribution of the state given all observations $\pi(\boldsymbol{x}_t|\boldsymbol{D}_t)$. In constructing this, we will rely on the conditional independence assumptions imposed via equation (1.1). The filtering problem can be defined as a recursive Bayesian estimation problem:

$$\begin{aligned}
\pi(\boldsymbol{x}_t|\boldsymbol{D}_t) &= \frac{\pi(\boldsymbol{d}_t|\boldsymbol{x}_t)\pi(\boldsymbol{x}_t|\boldsymbol{D}_{t-1})}{\pi(\boldsymbol{d}_t|\boldsymbol{D}_{t-1})} \\
\pi(\boldsymbol{x}_t|\boldsymbol{D}_t) &\propto \pi(\boldsymbol{d}_t|\boldsymbol{x}_t)\pi(\boldsymbol{x}_t|\boldsymbol{D}_{t-1}) \propto \pi(\boldsymbol{d}_t|\boldsymbol{x}_t)\int \pi(\boldsymbol{x}_t,\boldsymbol{x}_{t-1}|\boldsymbol{D}_{t-1})d\boldsymbol{x}_{t-1} \\
&\propto \pi(\boldsymbol{d}_t|\boldsymbol{x}_t)\int \pi(\boldsymbol{x}_t|\boldsymbol{x}_{t-1},\boldsymbol{D}_{t-1})\pi(\boldsymbol{x}_{t-1}|\boldsymbol{D}_{t-1})d\boldsymbol{x}_{t-1} \\
&\propto \pi(\boldsymbol{d}_t|\boldsymbol{x}_t)\int \pi(\boldsymbol{x}_t|\boldsymbol{x}_{t-1})\pi(\boldsymbol{x}_{t-1}|\boldsymbol{D}_{t-1})d\boldsymbol{x}_{t-1} \qquad (1.3)
\end{aligned}$$

It consists of two steps, i) prediction, and ii) update. The posterior distribution at the previous time step $\pi\left(\boldsymbol{x}_{t-1}|\boldsymbol{D}_{t-1}\right)$ is considered as the prior for the prediction step. The predictive distribution in equation (1.3) is achieved by $\pi\left(\boldsymbol{x}_t|\boldsymbol{D}_{t-1}\right) = \int \pi\left(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}\right)\pi\left(\boldsymbol{x}_{t-1}|\boldsymbol{D}_{t-1}\right)d\boldsymbol{x}_{t-1}$. Where $\pi\left(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}\right)$ is evaluated according to the process model and process noise distribution. The update step uses $\pi(\boldsymbol{x}_t|\boldsymbol{D}_{t-1})$ as prior and $\pi(\boldsymbol{d}_t|\boldsymbol{x}_t)$ as likelihood in Bayes rule. This step is also known as the correction step because the predicted values are corrected based on measurements $\boldsymbol{d}_t$.

Recursive Bayesian estimation is the exact solution to the filtering problem. Unfortunately, this solution is not applicable to most practical systems because it contains complicated multi-dimensional integrals. However, these integrals can be simplified by putting some limitations on the system assumptions, i.e. linear models and Gaussian distributions.

The Kalman filter (KF) is the optimal solution to the recursive Bayesian estimation for a system with linear dynamics and observation, Gaussian initial conditions, and independent Gaussian process and observation noise (Kalman, 1960). Jazwinsky (1970) extended the KF to handle the state estimation in nonlinear systems. The extended Kalman filter (EKF) uses the KF on the first order Taylor series approximation of the nonlinear equations (Jazwinsky, 1970). Although the EKF is the most used variant of the KF, it has some intrinsic deficiency. For example, if the system nonlinearity is high, the EKF may diverge. Besides, EKF requires the Jacobian. So, the exact equations of the system should be available (white box model). Thus, it is not applicable to black box models, where only the input/output information of the system is available.

Sigma point Kalman filters (SPKF) were proposed to handle the deficiencies of EKF in nonlinear systems (Julier and Uhlmann, 1997; Ito and Xiong, 2000; Nørgaard et al., 2000). These filters are based on nonlinear sampling propagation and a discretized representation of the sample space of the state variables (Merwe and Wan, 2001; Rezaie et al., 2007). The filters work well with black box models because they just need to know the input/output relationships of the system without detail information of equations. Such approaches have shown relevant for low to medium size system dimensions, but for high dimensional systems they become computationally infeasible.

The ensemble Kalman filter (EnKF) was introduced as a sampling representation for very high dimensional systems, see e.g. Evensen (2003), Sakov and Oke (2008) and Evensen (2009). It incorporates the nonlinear process model, whereas a Gaussian approximation is used for the updating with respect to new measurements. This approach has been very useful for practical applications, but the filter solution may be biased or underestimate uncertainty.

The key ingredient of the previous Kalman based filters is the Gaussian assumption. So, they can not handle complicated distributions accurately and may diverge. The Particle filter (PF) was proposed to address complicated distributions (Doucet et al., 2001). The PF is a Monte Carlo (MC) based algorithm which approximates the posterior distribution with weighted samples (Doucet et al., 2001). The PF converges to the optimal filtering distribution under weak regularity conditions, when the sample size goes to infinity. The PF works well for small dimension systems, but for high dimensional systems it suffers from sample degeneracy, i.e. all samples collapse to one sample. In theory one can overcome this problem by increasing the number of MC samples, but this has to increase faster than

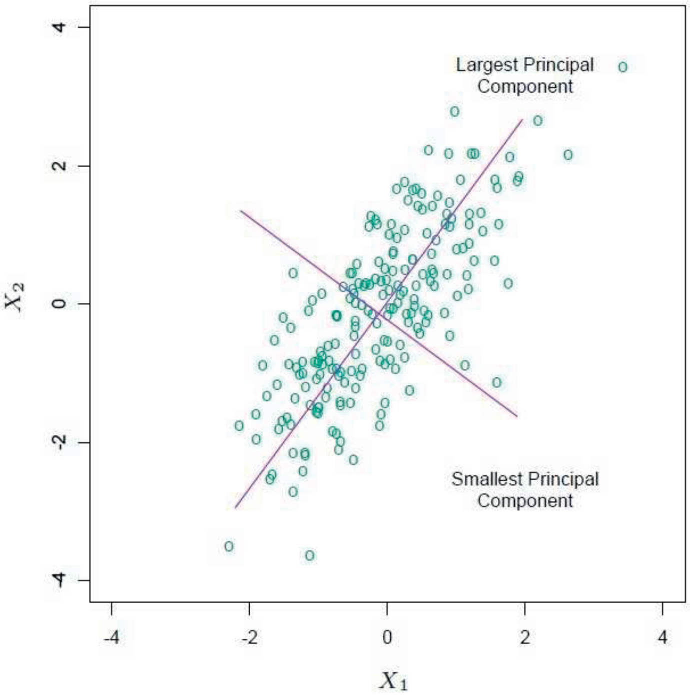

Figure 1.5: The principal components for a two dimensional example (Hastie et al., 2009)

the system dimension, and for most high dimensional practical purposes the computational burden becomes too large (Doucet and Johansen, 2011).

Thus, we can summarize the two major problems as i) real time state estimation in high dimensional systems with limited number of samples, and ii) complicated distributions. The first two papers in this thesis address these challenges to some extent. The simulations show promising results and it seems they have better or equal performance with some of the earlier proposed filters.

1.4 Dimension Reduction in Data assimilation and Integration

In recent years repeated seismic surveys have become more frequent and quantitative methods that can be used to integrate the information found in these surveys into reservoir models are called for. One major challenge in data assimilation is determining the information content in the measured data. Collinear and highly correlated measurements do not provide more information and further causes rank deficiency problems and model overfitting. Thus, finding the most informative data is vital. In the statistical literature, this is a well-known problem and dimension reduction methods have been proposed to deal with the situation (Hastie et al., 2009).

Principle component analysis (PCA) is one of the most frequently used dimension reduction techniques. Its implementation is straight forward by using singular value decomposition (SVD). PCA focuses on finding the structure of data ensemble matrix. By finding the structure of data, we mean that PCA finds the directions which data has maximum variability (for more details see Hastie et al. (2009)). Thus, the first principal component is defined as a vector which represents the first maximum variability direction of data, the second principal component is the direction of the second maximum variability of data etc. By transforming these data from data space to PC space the structure of the data and its actual dimension are found and we can remove the less significant part of the data by removing the last PCs (Figure 1.5 presents the concept of PC direction and data variability).

One challenge when conditioning on seismic data, is the potentially large data dimension, $n_d \gg 1$. This follows because of potentially high computational demands for the assisted filtering algorithms and the possibility of a dramatic model overfitting. When considering MC based techniques such as the EnKF, all the information regarding the prior and posterior distribution can be extracted from the ensemble of realizations. The same is true for the ensemble of forecasted data, found by evaluating the likelihood model (For example the likelihood $\pi(\boldsymbol{d}_t|\boldsymbol{x}_t)$ is Gaussian if $\boldsymbol{\epsilon}_t$ is a Gaussian noise in equation (1.1)). Although the dimension of the data (production and/or seismic) can be high, all of them are not useful. According to linear algebra, the rank of the a matrix is less than or equal to the smaller matrix dimension. So, if the ensemble of forecasted data is arranged in a matrix, the rank is equal to the minimum of the ensemble size and the data dimension. In addition, there is possibly a high spatial and temporal correlation in this forecasted data. This implies that the true rank of the ensemble of forecasted data can be significantly smaller than the data dimension. Paper III in this thesis relates to the dimension reduction of the geophysical data and removing the non-informative data before conditioning.

1.5 Value of Information

Rather than reducing the data size, one may ask if accurate processing and dense data acquisition was really necessary. More data will reduce uncertainty, but will it help the decision making?

Value of information (VOI) analysis relates to making better decisions under uncertainty (Raiffa, 1968; Howard, 1996). It is an old concept in the petroleum industry (Grayson et al., 1962), and it seems to have gained more interest in recent years (Branco et al., 2005; Bickel et al., 2008; Bratvold et al., 2009; Bhattacharjya et al., 2010). The VOI is useful in several petroleum applications where one considers purchasing more data before making a decision. The data comes with a price, and one might ask if it is really worth it, or which data to acquire at the current stage?

The VOI is defined as the maximum cost that we should pay for new information. If the price of data is larger than the VOI, the data are not worth purchasing (Bratvold et al., 2009). The calculation is tied to the underlying decision problem (Say, drill or not), and based on expected monetary units. Consider the expected revenues WITH additional information minus the expected revenues WITHOUT additional information. The former is sometimes called posterior value, while the latter is the prior value. The VOI is defined as:

$$\text{VOI} = \text{Posterior Value} - \text{Prior Value}. \tag{1.4}$$

Here, the prior value is positive only if the expected revenues are larger than the costs, i.e. $\max\{E[R(\boldsymbol{x})] - C, 0\}$. Similarly for the Posterior Value $= \int \max\{E[R(\boldsymbol{x})|\boldsymbol{d}] - C, 0\}\pi(\boldsymbol{d})d\boldsymbol{d}$. Here, $R(\boldsymbol{x})$ is the revenue which is a function of the state, and C is the cost. The state vector here consists of the saturation and porosity as the reservoir parameters. In addition the data is the seismic AVO attributes, which are informative of the saturation and porosity values. The decision may usually be done at multiple locations, and data carry information across space, because of dependencies. For evaluating the expectations in equation (1.4) we need the related distribution. For the posterior value, we are integrating over all possible data, and for each data we compute a conditional expectation. Finding the distribution, and the expected value rely on a Bayesian inversion problem.

The evaluation of geophysical information sources depends on input modeling assumptions. Reservoir parameters and seismic amplitudes are often skewed and by using models which capture the skewness of distributions, the input assumptions are less restrictive and the VOI analysis is more reliable (Figure 1.6). Eidsvik et al. (2008) proposed an algorithm

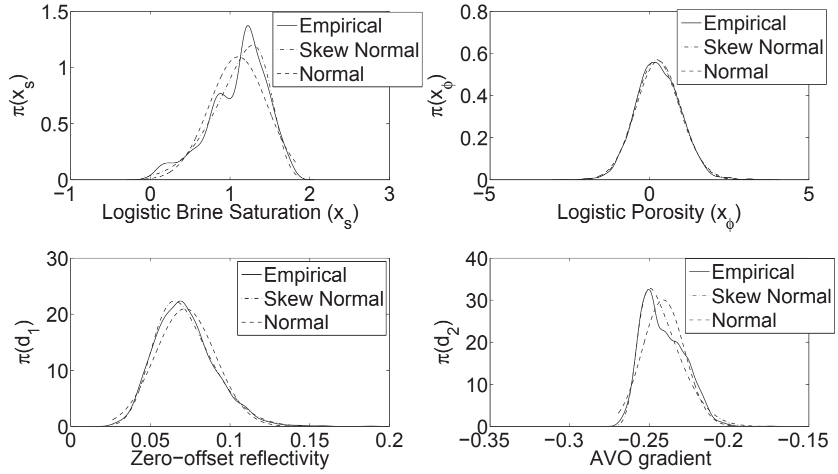

Figure 1.6: A graphical description of distribution fitting based on the SPE10 data, for saturation (upper left plot), porosity (upper right plot), zero-offset reflectivity (lower left plot) and AVO gradient (lower right plot), where the solid curve is the empirical distribution, dash curve is the fitted Gaussian and dash-dot is the fitted skewed.

for the VOI analysis of the seismic AVO and/or controlled source electromagnetic (CSEM) data with Gaussian assumption for reservoir parameters and measurements. But, as Figure 1.6 shows, the reservoir parameters and seismic amplitude measurements are skewed and non-Gaussian. By using models which capture the skewness of distributions, the input assumptions are better-fitting. Bayesian inversion of the seismic AVO data and the VOI analysis are then more reliable. Paper IV in this thesis relates to the VOI analysis of seismic AVO data with more flexible input modeling assumption. This extension results in more reliable decisions if the distributions are skewed.

1.6 Outline of the Papers

In this section, we are going to briefly explain the main contribution of the papers and the implementation results. Note that these papers were presented at several international conferences, URE annual meeting, student seminars in Department of Mathematical Sciences at NTNU, Statoil company and Department of Energy Resource Engineering at Stanford University.

Paper I: Shrinked $(1-\alpha)$ ensemble Kalman filter and α Gaussian mixture filter

Javad Rezaie and Jo Eidsvik

Published in Computational Geoscience

In this paper, we outline a method going between the EnKF and the PF. We are interested in maintaining the robust properties of the EnKF in order to avoid degeneracy. At the same time we encourage useful theoretical properties as the sample size increases. The filtering method we propose here fits a Gaussian mixture distribution to the predictive distribution. With additive noise terms for the dynamical model, it is a special case of the PF, and the algorithm is asymptotically correct when the number of samples, B, goes to infinity (Pitt and Shephard, 1999). Assuming the mean of predicted samples is a robust estimator, our construction shrinks the predicted particles towards this mean. We use a tuning parameter α to control the degree of shrinkage. The extreme cases are the PF ($\alpha = 1$) and the EnKF ($\alpha = 0$). When the number of samples increases, the α goes towards 1. The covariance of the predictive distribution is controlled by scaling the covariance of elements in the mixture. If the main computational cost is the forward propagation, which is often the case in high dimensional applications such as fluid flow simulation, the computation time of our method is at the order of the EnKF.

We test the proposed filter on a simulation study for target tracking, on the Lorenz 40 model, and on a reservoir simulation example. Results indicate that the filter works better or as well as Gaussian mixture PF and the EnKF for systems with different dimensions and complexities.

This work was published in Computational Geoscience, 16(3):837-852, 2012.

Paper II: Kalman Filter Variants in the Closed Skew Normal Setting

Javad Rezaie and Jo Eidsvik

In revision

In this paper, we extend the KF to capture possible skewness of distributions. The proposed filter is based on the closed skew normal (CSN) distributions. The CSN distribution is an extension of the Gaussian distribution (Azzalini and Dalla-Valle, 1996; Gupta et al., 2004). The CSN distribution consists of the Gaussian distribution as a special case.

The general recursive Bayesian filtering method is defined in this setting. Different filtering methods are derived for linear and nonlinear systems. In particular we derive the KF and the EnKF as special cases of the proposed filters. Finally we implement the proposed methods on a synthetic linear system, re-entering a body to the atmosphere at very high altitude and velocity, and a seismic history matching case studies. Simulation results show the good performance of the proposed methods compared with the KF and EnKF.

This work is in revision.

Paper III: Reducing the Dimensionality of Geophysical Data in Conjunction with Seismic History Matching

Javad Rezaie, Jon Sætrom and Eivind Smørgrav

Published in the proceedings of the 74^{th} *EAGE Conference & Exhibition incorporating SPE EUROPEC 2012, Copenhagen, Denmark, SPE153924.*

In this paper, we study the data dimension reduction problem using a Bayesian approach, where the statistical properties of the reservoir can be described through an ensemble of realizations generated from the posterior distribution of the variables of interest. When considering ensemble based methods for high dimensional data, the rank of the ensemble of the forecasted data is at most equal to its size of 20-200 members. However, in practice there are collinearities between the ensemble members, which means that the effective rank can be dramatically smaller than the data dimension. This implies that we can essentially represent high dimensional data, such as time-lapse seismic data, in terms of a few scalars.

We therefore apply dimension reduction techniques to efficiently handle problems related to the data dimensionality. We considered both basis vector and pixel based approaches. The proposed algorithms are applied on a reservoir model inspired by a North Sea oil field as the test benchmark. We find that the basis vector approach based on PCA is particularly promising, resulting from the versatility and robustness of the method. PCA applies equally well for any ensemble based history matching algorithm, while the pixel based methods are designed mainly for the EnKF.

This work was published in the proceedings of the SPE EUROPEC 2012, SPE-153924.

Paper IV: Value of Information Analysis and Bayesian Inversion for Closed Skew-Normal Distributions: Applications to Seismic Amplitude Versus Offset Data

Javad Rezaie, Jo Eidsvik and Tapan Mukerji

In revision

In this paper, we introduce skewed distributions for the modeling of reservoir parameters and geophysical measurements. We fit CSN models to the SPE10 data set and check the sensitivity of the VOI to the model parameters. The Gaussian assumption directs the decision maker to wrong decisions when the data processing cost is in between the evaluated VOIs of the Gaussian and the CSN. Simulation results show that by increasing the variance and/or skewness parameters of data distribution, the VOI decreases, and it is more sensitive to variance than skewness.

One of the computational challenges in the VOI calculation is evaluation of the integrals. We use analytical approximations, and MC approximations for calculation. Besides, we use the power of parallel computing and graphics processing units (GPUs) for computation speed up.

This work is in revision.

1.7 Ideas for Future Work

The four parts of this thesis are connected in the way they use Bayesian updating to assimilate geophysical data. Hence, they focus on different modeling challenges and techniques. These ideas can be used independently or as a part of another algorithm. Thus, combining some of these ideas together or tie-in with other algorithms can be defined as future works. The following ideas are proposed as the future research.

In the VOI part, we work on the VOI analysis of the seismic AVO data. In this part of the thesis, the VOI may help us decide whether to purchase the new data or not. Consider a case where one a priori decided to purchase the data, but the decision making problem is defined as when is the optimal time of conditioning? In the other words, we have the opportunities to condition on the seismic AVO data at different time steps, and we are going to find the time step(s) that has the best conditioning result. Actually, for this problem we must combine the VOI analysis and the filtering problem. Here, we can define different cost functions for optimization. It is obvious that the best results are achieved by conditioning at all time steps. But, we have some more limitations, i.e. the total data price, computation time etc. In preliminary simulation results, we assume that we are allowed to do one conditioning in a time window. By this restriction, we found that the maximum VOI occurs when we have the first maximum correction between the prior and the posterior of the reservoir parameters. This time index is where the saturation goes to the steady state region.

The skewed version of the SPKFs in a CSN setting can also be considered as future research. (Julier, 1998) tried to incorporate the skewness of latent variables into the sigma point selection of the famous unscented Kalman filter (UKF), because the UKF was originally designed to capture the first and second order moments of random variables with a proofed accuracy level Julier and Uhlmann (1997). The problem with this proposal is the number of sigma points (discretization points in the algorithm), which is in the order of $O(n^3)$ for an n dimensional system. Since the number of sigma points then gets huge for medium size systems, it is practically inconvenient. But, by combining the idea of the CSN filter into the UKF the number of necessary sigma points reduces dramatically. According to this idea, we proposed a method to estimate the CSN distribution parameters using some deterministically chosen ensemble members named sigma points. Furthermore, we can combine the skewed UKF and the Bayesian inversion results of Rimstad and Omre (2012b,a), which uses an extended version of the CSN, including multi-modal distributions.

Bibliography

Azzalini, A. and Dalla-Valle, A. (1996). The multivariate skew-normal distribution. *Biometrika*, 83(4):715–726.

Bhattacharjya, D., Eidsvik, J., and Mukerji, T. (2010). The value of information in spatial decision making. *Mathematical Geosciences*, 73(2):141–163.

Bickel, J. E., Gibson, R. L., McVay, D. A., Pickering, S., and Waggoner, J. (2008). Quantifying 3D land seismic reliability and value. *SPE Reservoir Evaluation and Engineering*, 11(5):832–841.

Branco, C. C. M., Pinto, A. C. C., Tinoco, P. M. B., Vieira, P. M. F., Sayd, A. M., Santos, R. L. A., and Prais, F. (2005). The role of the value of information and long horizontal wells in the appraisal and development studies of a brazilian offshore heavy-oil reservoir. In *International Thermal Operations and Heavy Oil Symposium, Calgary, Alberta, Canada*, page SPE97846. Society of Petroleum Engineers.

Bratvold, R. B., Bickel, J. E., and Lohne, H. P. (2009). Value of information in the oil and gas industry: Past, present, and future. *SPE Reservoir Evaluation and Engineering*, 12(4):630–638.

Christie, M. A. and Blunt, M. J. (2001). Tenth SPE comparative solution project: A comparison of upscaling techniques. *SPE Reservoir Engineering and Evaluation*, 4(4):308–317.

Doucet, A., de Freitas, N., and Gordon, N. (2001). *Sequential Monte-Carlo Methods in Practice*. Springer-Verlag.

Doucet, A. and Johansen, A. M. (2011). *A tutorial on particle filtering and smoothing: fifteen years later. In: Oxford Handbook of Nonlinear Filtering, 2011*. Oxford University Press.

Eidsvik, J., Bhattacharjya, D., and Mukerji, T. (2008). Value of information of seismic amplitude and CSEM resistivity. *Geophysics*, 73(4):59–69.

Evensen, G. (2003). The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dynamics*, 53(1):343–367.

Evensen, G. (2009). *Data assimilation, The Ensemble Kalman Filter*. Springer.

Grayson, J. R., Jackson, C., and Tulane, U. (1962). Bayesian analysis-a new approach to statistical decision making. *Journal of Petroleum Technology*, 14(6):603–607.

Gupta, A. K., González-Farías, G., and Domínguez-Molina, J. A. (2004). A multivariate skew normal distribution. *Multivariate Analysis*, 89(1):181–190.

Hastie, T., Tibshirani, R., and Freidman, J. (2009). *The Elements of Statistical Learning; Data Mining, Inference, and Prediction*. Springer.

Howard, R. A. (1996). Information value theory. *Systems Science and Cybernetics*, 2(1):22–26.

Ito, K. and Xiong, K. (2000). Gaussian filters for nonlinear filtering problems. *Automatic Control, IEEE Transactions on*, 45(5):910–927.

Jazwinsky, A. (1970). *Stochastic Processes and Filtering Theory*. Academic Press, New York.

Julier, S. (1998). Skewed approach to filtering. In *Aerospace/Defense Sensing and Controls*, pages 271–282. International Society for Optics and Photonics.

Julier, S. J. and Uhlmann, J. K. (1997). A new extension of the Kalman filter to nonlinear systems. *SPIE - Int. Soc. Opt. Eng.*, 3068(1):182–193.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *ASME Journal of Basic Engineering*, 1(1):35–45.

Lie, K., Krogstad, S., Ligaarden, I. S., Natvig, J. R., Nilsen, H. M., and Skaflestad, B. (2012). Open source MATLAB implementation of consistent discretisations on complex grids. *Computational Geosciences*, 16(2):297–322.

Mavko, G., Mukerji, T., and Dvorkin, J. (2003). *Rock Physics Handbook - Tools for Seismic Analysis in Porous Media.* Cambridge University Press.

Merwe, R. V. D. and Wan, E. (2001). The square-root unscented Kalman filter for state and parameter estimation. In *International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Salt Lake City, UT*, pages 3461–3464. IEEE.

Nørgaard, M., Poulsen, N. K., and Ravn, O. (2000). New developments in state estimation for nonlinear systems. *Automatica*, 36(11):1627 – 1638.

Pitt, M. K. and Shephard, N. (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association*, pages 590–599.

Raiffa, H. (1968). *Decision Analysis: Introductory Lectures on Choices under Uncertainty.* Addison-Wesley.

Rezaie, J., Moshiri, B., Araabi, B. N., and Asadian, A. (2007). GPS/INS integration using nonlinear blending filters. In *SICE Annual Conference*, pages 1674 – 1680. IEEE.

Rimstad, K. and Omre, H. (2012a). Generalized Gaussian random fields using hidden selections. *Technical Report.*

Rimstad, K. and Omre, H. (2012b). Skew Gaussian random fields. *Submitted for publication.*

Sakov, P. and Oke, P. R. (2008). Implications of the form of the ensemble transformation in the ensemble square root filters. *Monthly Weather Review*, 136:1042–1053.

Paper I

Shrinked $(1-\alpha)$ ensemble Kalman filter and α Gaussian mixture filter

J. Rezaie and J. Eidsvik.

Published in Computational Geoscience, 16(3):837-852, 2012.

Chapter 2

Shrinked $(1-\alpha)$ ensemble Kalman filter and α Gaussian mixture filter

---

Abstract. State estimation in high dimensional systems remains a challenging part of real time analysis. The ensemble Kalman filter addresses this challenge by using Gaussian approximations constructed from a number of samples. This method has been a large success in many applications. Unfortunately, for some cases, Gaussian approximations are no longer valid, and the filter does not work so well. In this paper, we use the idea of the ensemble Kalman filter together with the more theoretically valid particle filter. We outline a Gaussian mixture approach based on shrinking the predicted samples to overcome sample degeneracy, while maintaining non-Gaussian nature. A tuning parameter determines the degree of shrinkage. The computational cost is similar to the ensemble Kalman filter. We compare several filtering methods on three different cases: a target tracking model, the Lorenz 40 model, and a reservoir simulation example conditional on seismic and electromagnetic data.

2.1 Introduction

State estimation is an important problem in engineering and science. If we represent the system dynamics (differential or difference equations) in state space form, the measurements are transformed, noisy, and become an incomplete representation of the system state. Filtering methods extract the probability distribution of the state at every time point, given all measurements until that time. For dynamic systems, it is natural to perform the estimation process as soon as new observations arrive. Thus, recursive Bayesian estimation algorithms are powerful for dealing with filtering problems. This consists of sequentially going forward in time according to a two-step routine: (a) a forward propagation step using the system dynamics, and (b) an updating step when the new data gets available. Step a is known as the *prediction problem*, while step b is the *filtering problem*.

The celebrated Kalman filter (KF) is the optimal solution to the recursive estimation challenge under certain model restrictions (Kalman, 1960). These assumptions include linear dynamics and observation model, Gaussian initial conditions, and independent Gaussian process and observation noise. For non-linear systems, we can use linearization and apply the KF update for the resulting system. This is known as the extended Kalman filter (EKF) (Jazwinsky, 1970). If system non-linearity is high, the EKF may diverge. Sigma point Kalman filters were proposed to overcome some of the weak points of linearized filters (Julier and Uhlmann, 1997; Ito and Xiong, 2000; Nørgaard et al., 2000; Merwe and Wan, 2001). They have very good performance for small- to medium-sized systems; but for high dimensional systems, the computational burden becomes too large. The ensemble Kalman filter (EnKF) uses Monte Carlo realizations along with a Gaussian approximation in the updating step. It has shown very good results for high dimensional systems (Evensen, 2003, 2009).

All the algorithms mentioned above have a Gaussian approximation as the key ingredient. Even though the algorithms have been successful for high dimensional non-linear problems, one cannot really justify a Gaussian approximation in these problems, apart from using computational convenience as an argument. Noteably, there are no asymptotic results saying that the EnKF converges to the optimal filtering distribution when the number of samples goes to infinity. The particle filter (PF) is a Monte Carlo-based algorithm which approximates the posterior distribution with weighted samples (Doucet et al., 2001). Under weak regularity conditions, the PF converges to the optimal filtering distribution when the sample size goes to infinity. The conditions are similar to that of importance sampling. One must be able to draw samples from a proposal mechanism, and since the method relies on constructing weighted filtering samples, the proposal distribution cannot have lighter tails than the target filtering distribution (see Doucet and Johansen (2011)). The PF works well for small dimension systems with general non-Gaussian and non-linear models, but for high dimensional systems, it suffers from sample degeneracy, i.e. all samples collapse to one sample. In theory, one can overcome this problem by increasing the number of Monte Carlo samples; but this has to increase faster than the system dimension, and for most high dimensional practical purposes, the computational burden becomes too large (Doucet and Johansen, 2011). There are Gaussian mixture-based filters (Anderson and Moore, 1979) which use EKF for constructing each Gaussian mixture, but this approach is not asymptotically correct because EKF just uses first-order Taylor series expansion, and it looses the sparsity of random variables (Wan et al., 2000).

In this paper, we outline a method going between the EnKF and the PF. We are interested in maintaining the robust properties of the EnKF in order to avoid degeneracy. At the same time, we encourage useful theoretical properties as the sample size increases. The filtering method we propose here fits a Gaussian mixture distribution to the predictive distribution. With additive noise terms for the dynamical model, it is a special case of PF, and the algorithm is asymptotically correct when the number of samples, B, goes to infinity (Pitt and Shephard, 1999). Assuming the mean of predicted samples is a robust estimator, our construction shrinks the predicted particles towards this mean. We use a tuning parameter α to control the degree of shrinkage. The extreme cases are the PF ($\alpha = 1$) and the EnKF ($\alpha = 0$). When the number of samples increases, the shrinkage α goes towards 1. The covariance of the predictive distribution is controlled by scaling the covariance of elements in the mixture. If the main computational cost is the forward propagation, which is often the case in high dimensional applications such as fluid flow simulation, the computation time of our method is at the order of the EnKF.

Some recent publications are similar, but different to the current paper: Sætrom and Omre (2011) change the EnKF updating schemes based on shrinkage methods known from multi-variate linear regression such as partial least square regression and principal component regression. Their method reduces the collinearity between the samples, but the proposed method suffers from the same Gaussian assumptions as standard EnKF. The method proposed in Stordal et al. (2010) shrinks the updating weights, while we shrink the samples and get a natural modification of the weights as a natural consequence. Dovera and Rossa (2011), extend the standard EnKF to a mixture of Gaussians for the predictive distribution. They consider only multi-modal problems. We attempt to handle the more general non-Gaussian cases. In Hoteit et al. (2008) and Hoteit et al. (2011), methods for combining the KF and PF are proposed. They first update the propagated samples using the KF scheme and modify weights according to a PF selection using the updated samples. We apply a different intermediate version of the EnKF and PF. Leeuwen (2010) and Leeuwen (2011) proposes a PF with an intelligent selection of the proposal distribution. The filter is similar to the EnKF updating part with a nudging noise or relaxation term. The predicted samples are guided to high likelihood regions, and sample degeneracy is avoided. In practice, it may be hard to choose the relaxation term, which is case-specific and depends on the number of forward propagation steps for each updating or assimilation step.

The main contribution of our paper is a method of shrinking predictive samples to maintain flexibility in the distribution while avoiding degeneracy. This new approach is different from all of the above because the shrinkage allows another intermediate step between the EnKF and a particular PF. The updating and the weights follow directly from the degree of shrinkage. The degree of shrinkage is the only tuning parameter. Moreover, it is not restricted to bimodal systems, and it is applicable to systems with one or more forward propagation steps before conditioning on the observations. The paper is organized as follows: In Section 2.2, we define the model assumptions used in this paper. Section 2.3 outlines the α shrinked ensemble particle filter. Section 2.4 provides examples from target tracking, the Lorenz 40 model, and a reservoir simulation example.

2.2 Notation and Modeling Assumptions

Denote the state variable at time t by $\boldsymbol{x}_t$, and let $\boldsymbol{X}_t = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_t)$ be the collection of the state variables from time 1 to the current time t. Further, the observations at time t are denoted $\boldsymbol{y}_t$, and $\boldsymbol{Y}_t = (\boldsymbol{y}_1, \ldots, \boldsymbol{y}_t)$ is the collection of observations at this current time step. We assume continuous state and observation variables, i.e. $\boldsymbol{x}_t \in \mathcal{R}^n$ and $\boldsymbol{y}_t \in \mathcal{R}^m$, where the dimensions n and m tend to get large in most modern applications.

We use a state space formulation with the usual conditional independence assumptions. This means that the conditional probability density function of $\boldsymbol{x}_t$, given all previous state and observation variables, only depends on the state at time $t-1$. Moreover, the conditional distribution of observation $\boldsymbol{y}_t$, given the state at that time and all previous states and observations, only depends on the state at time t. Mathematically, these two assumptions entail that the conditional distributions are $\pi(\boldsymbol{x}_t|\boldsymbol{X}_{t-1}, \boldsymbol{Y}_{t-1}) = \pi(\boldsymbol{x}_t|\boldsymbol{x}_{t-1})$ and $\pi(\boldsymbol{y}_t|\boldsymbol{X}_t, \boldsymbol{Y}_{t-1}) = \pi(\boldsymbol{y}_t|\boldsymbol{x}_t)$, respectively. The joint density of $\boldsymbol{Y}_t$ and $\boldsymbol{X}_t$ can then be factorized using the conditional independence assumptions:

$$\pi(\boldsymbol{Y}_t, \boldsymbol{X}_t) = \prod_{i=1}^{t} \pi(\boldsymbol{y}_i|\boldsymbol{x}_i) \prod_{i=2}^{t} \pi(\boldsymbol{x}_i|\boldsymbol{x}_{i-1})\pi(\boldsymbol{x}_1). \tag{2.1}$$

Here, $\pi(\boldsymbol{x}_1)$ is the specified probability density function for the initial state variable. The joint model is defined once we have specified this initial distribution, along with the density for the dynamic propagation model $\pi(\boldsymbol{x}_t|\boldsymbol{x}_{t-1})$ and the likelihood model $\pi(\boldsymbol{y}_t|\boldsymbol{x}_t)$. Depending on the dynamic model and the information content in the data, there will often be a transient phase for small t, revealing the initial conditions.

We next specify our particular assumptions about the forward propagation and the likelihood model. Generically, we let $N(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denote the Gaussian probability density function of random variable $\boldsymbol{x}$, with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. We model the dynamics of the system in the following way:

$$\pi(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}) = N(\boldsymbol{x}_t; \boldsymbol{g}_t(\boldsymbol{x}_{t-1}), \boldsymbol{P}), \tag{2.2}$$

where the expectation term is defined by a non-linear function $\boldsymbol{g}_t(\cdot)$. This function is usually the computationally hard part, involving a forward propagation of complex physical phenomena. For instance, in a reservoir simulation application, this function consists of numerical solutions of the partial differential equations for fluid flow in porous media. The covariance matrix $\boldsymbol{P}$ can incorporate non-modeled physics, and may be a result of using coarse scales in a simulator or of compensating for simplified physics, such as a treating certain physical properties as fixed in the dynamical model. We assume a linear likelihood model, with additive Gaussian noise, i.e.,

$$\pi(\boldsymbol{y}_t|\boldsymbol{x}_t) = N(\boldsymbol{y}_t; \boldsymbol{H}_t\boldsymbol{x}_t, \boldsymbol{R}) \ . \tag{2.3}$$

The matrix $\boldsymbol{H}_t$ is defined by the data acquisition of the problem, while $\boldsymbol{R}$ is the covariance matrix of the measurement noise. Whereas the dynamical model could be very non-linear, the likelihood is assumed linear. In a physical application, this assumption for the likelihood entails that a measurement equation is well known and can be solved analytically. In the simplest case, we have $m = n$ and $\boldsymbol{H}_t = \boldsymbol{I}_n$, indicating that we measure the state directly, with additive noise. Note that we let the function $\boldsymbol{g}_t$ depend on the time variable, and the same holds for the expectation operator $\boldsymbol{H}_t$ in the likelihood. For simplicity, we specify the covariance matrices $\boldsymbol{P}$ and $\boldsymbol{R}$ as fixed over time, but this is easily generalized. If the expectation term in the likelihood is weakly non-linear, one can linearize it at each stage of the filtering scheme. Some asymptotic properties derived below would then not hold in general. Nevertheless, we get a solution as a mixture of approximate Gaussians.

The filtering task consists of sequential propagation and updating as we obtain new observations. At time $t-1$, consider that we have the updated (filtering) distribution of the state given all observations until that time, denoted by the density $\pi(\boldsymbol{x}_{t-1}|\boldsymbol{y}_1,...,\boldsymbol{y}_{t-1}) = \pi(\boldsymbol{x}_{t-1}|\boldsymbol{Y}_{t-1})$. The one-step prediction density is constructed from the dynamical model

$$\pi(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1}) = \int \pi(\boldsymbol{x}_t|\boldsymbol{x}_{t-1})\pi(\boldsymbol{x}_{t-1}|\boldsymbol{Y}_{t-1})d\boldsymbol{x}_{t-1}, \tag{2.4}$$

where the model assumptions simplify the integrand according to $\pi(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}, \boldsymbol{Y}_{t-1}) = \pi(\boldsymbol{x}_t|\boldsymbol{x}_{t-1})$. When the new observation $\boldsymbol{y}_t$ is available, we combine the system dynamics and the likelihood in Bayes rule for the updating:

$$\pi(\boldsymbol{x}_t|\boldsymbol{Y}_t) = \frac{\pi(\boldsymbol{y}_t|\boldsymbol{x}_t)\pi(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1})}{\pi(\boldsymbol{y}_t|\boldsymbol{Y}_{t-1})} \propto \pi(\boldsymbol{y}_t|\boldsymbol{x}_t)\pi(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1}), \tag{2.5}$$

where the conditional independence assumption of the data $\boldsymbol{y}_t$ is used, given the state $\boldsymbol{x}_t$. This recursive Bayesian method gives the exact solution to the general filtering problem; but for practical applications, we cannot implement it for large systems because we must calculate multi-dimensional complicated integrals. Thus, some simplified conditions on the system dynamics and observations inducing some consistent approximations have to be considered.

A sampling approximation of the filtering distribution can be established using Monte Carlo realizations (Fishman, 1996). Suppose we have B independent and identically distributed samples $\boldsymbol{x}^1_{t-1},...,\boldsymbol{x}^B_{t-1}$ from $\pi(\boldsymbol{x}_{t-1}|\boldsymbol{Y}_{t-1})$. The sample approximation to the filtering distribution at time $t-1$ is then

$$\pi(\boldsymbol{x}_{t-1}|\boldsymbol{Y}_{t-1}) = \frac{1}{B}\sum_{b=1}^{B}\delta\big(\boldsymbol{x}_{t-1} - \boldsymbol{x}^b_{t-1}\big), \tag{2.6}$$

with the Dirac function $\delta(\boldsymbol{x}) = 1$ if $\boldsymbol{x} = 0$, and $\delta(\boldsymbol{x}) = 0$ otherwise. For the prediction step, all B samples are run through the dynamical model, i.e., $\boldsymbol{g}_t\big(\boldsymbol{x}^b_{t-1}\big)$, $b = 1,...,B$. In several applications, the $\boldsymbol{g}_t(\cdot)$ evaluation is so time-consuming that this forward propagation can only be run about $B \sim 100$ times. In the next section, we discuss various approaches for recursively updating this size-B sample approximation of the predictive and filtering distributions.

2.3 Shrinked Gaussian Mixture Filters

The algorithms presented below differ in the construction of a predictive density, and the induced filtering density. We first present a particular PF, which is exact when the sample size $B \to \infty$. The PF is represented by a marginalized Gaussian mixture distribution and denoted by the Gaussian mixture filter (GMF). In the practical situation, B cannot get large enough, and this filter will degenerate for high dimensional systems. The (square-root) EnKF algorithm is next presented as a collapsed Gaussian mixture, with all mean values in the mixture being identical. Finally, we outline the shrinkage idea to overcome sample degeneracy, while maintaining asymptotic sample properties when B goes to infinity. We approximate the predictive distribution with a mixture of Gaussians having mixture component mean values between that of the GMF and the EnKF. A tuning parameter, α, is used to adjust the particles between the two extremes given by the GMF and overall mean (EnKF). We denote the resulting filter by the robustified Gaussian mixture filter (RGMF).

2.3.1 Gaussian mixture Monte Carlo filter

The prediction formula is given in Eq. 2.4. In our model formulation, with the sample approximation in Eq. 2.6, this prediction step becomes an integral over a Gaussian weighted with B Dirac functions. The prediction distribution is thus a mixture of B Gaussian densities:

$$\begin{aligned}
&\pi(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1}) \\
&\quad = \int N(\boldsymbol{x}_t; \boldsymbol{g}_t(\boldsymbol{x}_{t-1}), \boldsymbol{P})\pi(\boldsymbol{x}_{t-1}|\boldsymbol{Y}_{t-1})d\boldsymbol{x}_{t-1} \\
&\quad = \frac{1}{B}\sum_{b=1}^{B} \int N\big(\boldsymbol{x}_t; \boldsymbol{g}_t(\boldsymbol{x}_{t-1}), \boldsymbol{P}\big)\delta\big(\boldsymbol{x}_{t-1} - \boldsymbol{x}_{t-1}^b\big)d\boldsymbol{x}_{t-1} \\
&\quad = \frac{1}{B}\sum_{b=1}^{B} N\big(\boldsymbol{x}_t; \boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big), \boldsymbol{P}\big) \\
&\quad = \sum_{b=1}^{B} \pi(b|\boldsymbol{Y}_{t-1})\pi(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1}, b),
\end{aligned} \tag{2.7}$$

where $\pi(b|\boldsymbol{Y}_{t-1}) = \frac{1}{B}$ is used to clarify the identically weighted components $b = 1, 2, ..., B$ in the mixture. Note that we have one mixture component for every sample. Clustering of samples is also possible, giving less mixture components than samples; however, this requires a criteria for clustering of samples (see Frei and Künsch (2013)). In the development below, we do not consider this topic further.

The filtering step becomes

$$\begin{aligned}\pi(\boldsymbol{x}_t|\boldsymbol{Y}_t) &\propto N(\boldsymbol{y}_t;\boldsymbol{H}_t\boldsymbol{x}_t,\boldsymbol{R})\pi(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1})\\ \pi(\boldsymbol{x}_t|\boldsymbol{Y}_t) &= \sum_{b=1}^{B} N\big(\boldsymbol{x}_t;\hat{\boldsymbol{x}}_t^b,\boldsymbol{S}_t\big)\pi(b|\boldsymbol{Y}_t),\end{aligned} \tag{2.8}$$

where $\hat{\boldsymbol{x}}_t^b$ and $\boldsymbol{S}_t$ are the updated mean and covariance matrix, given component b, i.e., $\pi(\boldsymbol{x}_t|\boldsymbol{Y}_t,b)$. This Gaussian density is obtained by the usual KF formula for fixed component b:

$$\begin{aligned}\hat{\boldsymbol{x}}_t^b &= \boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big)+\boldsymbol{P}\boldsymbol{H}_t^{'}\boldsymbol{Q}_t^{-1}\big(\boldsymbol{y}_t-\boldsymbol{H}_t\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big)\big),\\ \boldsymbol{S}_t &= \boldsymbol{P}-\boldsymbol{P}\boldsymbol{H}_t^{'}\boldsymbol{Q}_t^{-1}\boldsymbol{H}_t\boldsymbol{P},\quad \boldsymbol{Q}_t=\boldsymbol{H}_t\boldsymbol{P}\boldsymbol{H}_t^{'}+\boldsymbol{R}.\end{aligned} \tag{2.9}$$

The weights $w_b=\pi(b|\boldsymbol{Y}_t)\propto\pi(\boldsymbol{y}_t|\boldsymbol{Y}_{t-1},b)\pi(b|\boldsymbol{Y}_{t-1})$, where $\pi(\boldsymbol{y}_t|\boldsymbol{Y}_{t-1},b)$ is a Gaussian marginalized over $\boldsymbol{x}_t$. Since expression (2.8) must be a probability density, the weights w_b are required to sum to 1. We get

$$w_b=\frac{N\big(\boldsymbol{y}_t;\boldsymbol{H}_t\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big),\boldsymbol{Q}_t\big)}{\sum_{c=1}^{B}N\big(\boldsymbol{y}_t;\boldsymbol{H}_t\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^c\big),\boldsymbol{Q}_t\big)}. \tag{2.10}$$

The GMF defined by the above formulas now proceeds by repeating the following B times:

1. Sample a component b from length B probability vector $(w_1,...,w_B)$.
2. Sample $\boldsymbol{x}_t^b\sim N\big(\boldsymbol{x}_t;\hat{\boldsymbol{x}}_t^b,\boldsymbol{S}_t\big)$.

A Dirac representation for $\pi(\boldsymbol{x}_t|\boldsymbol{Y}_t)$ is given by the equally weighted B samples. This procedure continues, from time t to $t+1$, according to the dynamical model $\boldsymbol{g}_t(\cdot)$, just like our starting point in Eq. 2.7 at time $t-1$. Step 1 of the algorithm above is done by drawing a uniform $u\sim U(0,1)$ and picking the smallest index b such that $\sum_{c=1}^{b}w_c>u$. Step 2 above can either be done by direct Cholesky factorization. Then one first generates a correlated zero-mean Gaussian using the Cholesky factor (square-root) of $\boldsymbol{S}_t$ and afterwards adds the mean value $\hat{\boldsymbol{x}}_t^b$. Alternatively, one can (a) draw prior variate $\boldsymbol{x}_t^{*b}\sim N\big(\boldsymbol{x}_t^{*b};\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big),\boldsymbol{P}\big)$; (b) draw $\boldsymbol{y}_t^{*b}\sim N\big(\boldsymbol{y}_t^{*b};\boldsymbol{y}_t,\boldsymbol{R}\big)$; and (c) set sample $\boldsymbol{x}_t^b=\boldsymbol{x}_t^{*b}+\boldsymbol{P}\boldsymbol{H}_t^{'}\boldsymbol{Q}_t^{-1}\big(\boldsymbol{y}_t^{*}-\boldsymbol{H}_t\boldsymbol{x}_t^{*b}\big)$. The resulting sample is Gaussian with the correct mean and covariance.

The filtering distributions are Gaussian mixtures, where the Monte Carlo step is used to propagate the particles forward in time. Note that under the modeling assumptions with $\boldsymbol{H}_t$ linear and Gaussian error terms, the updated Gaussian mixture for $\pi(\boldsymbol{x}_t|\boldsymbol{Y}_t)$ is exact, given the B-sample Dirac representation of $\pi(\boldsymbol{x}_{t-1}|\boldsymbol{Y}_{t-1})$. The resulting PF, called GMF here, is much more stable than the standard bootstrap PF (Doucet et al., 2001), since we have marginalized over $\boldsymbol{x}_t$ in the weights w_b. In contrast, the bootstrap PF weighs the forward propagated particles according to the likelihood, and the weights w_b are then less uniform. Still, the Monte Carlo sample approximation applied in the GMF might not be so good for small B. The weights w_b could be very non-uniform in high dimensional systems, and sample degeneracy would occur. Clearly, if one $\boldsymbol{H}_t\boldsymbol{g}_t(\boldsymbol{x}_{t-1}^b)$ is much closer to the data $\boldsymbol{y}_t$ than the others, the associated w_b in Eq. 2.10 is close to 1, while all other weights are near 0. The mixture is then focused on one component, and the approach underestimates the variability of the distribution. These effects, of course, depend on the number of particles B, the non-linearity $\boldsymbol{g}_t(\cdot)$, and the covariance matrix $\boldsymbol{Q}_t$. Heuristically, in high dimension, say with a diagonal $\boldsymbol{Q}_t$, none of the particles $\boldsymbol{H}_t\boldsymbol{g}_t(\boldsymbol{x}_{t-1}^b)$ are close to the data $\boldsymbol{y}_t$, but the closest is much closer than the second closest. The squared mismatch distances, summed over m data dimensions,are given by $\sum_{j=1}^m Q_{t,jj}^{-1}\big(y_{jt} - H_j g_{t,j}\,(x_{t-1}^b)\big)^2$, which will blow up linearly with dimension.

One alternative is, of course, to increase the sample size B. However, this sample size B must typically increase faster than the dimension, and for some non-linear problems, the B cannot be very large because of the computation time of $\boldsymbol{g}_t(\cdot)$. Other tricks that slow down degeneracy include tapering or localization of the covariance matrix, which effectively reduces the dimension.

2.3.2 Ensemble Kalman filter as a collapsed Gaussian mixture

The EnKF is based on a Gaussian approximation to the predictive distribution $\pi(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1})$. For our model, this entails matching the mean and covariance matrix of the predictive Gaussian mixture distribution in Eq. 2.7. Assuming the additive noise (with covariance matrix $\boldsymbol{P}$) in the dynamical model, the forecasted or predicted mean and covariance in the EnKF is given by the following:

$$\begin{aligned}
\bar{\boldsymbol{x}}_t &= \frac{1}{B}\sum_{b=1}^{B}\boldsymbol{g}_t(\boldsymbol{x}_{t-1}^b)\\
\bar{\boldsymbol{P}}_t &= \boldsymbol{P} + \frac{1}{B}\sum_{b=1}^{B}\big(\boldsymbol{g}_t(\boldsymbol{x}_{t-1}^b) - \bar{\boldsymbol{x}}_t\big)\big(\boldsymbol{g}_t(\boldsymbol{x}_{t-1}^b) - \bar{\boldsymbol{x}}_t\big)', \qquad (2.11)
\end{aligned}$$

where we have used the formula for double covariance $Var(X) = E(Var(X|b))+Var(E(X|b))$, conditioning on the mixture components. The predictive distribution used in an EnKF update becomes

$$\begin{aligned}
\hat{\pi}(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1}) &= N(\boldsymbol{x}_t;\bar{\boldsymbol{x}}_t,\bar{\boldsymbol{P}}_t) \\
&= \frac{1}{B}\sum_{b=1}^{B} N(\boldsymbol{x}_t;\bar{\boldsymbol{x}}_t,\bar{\boldsymbol{P}}_t), \qquad (2.12)
\end{aligned}$$

where we have regarded it as a Gaussian mixture of B terms, all with the same mean and covariance. Thus, this mixture density has collapsed to a single Gaussian.

The updated distribution for the EnKF is

$$\begin{aligned}
\hat{\pi}(\boldsymbol{x}_t|\boldsymbol{Y}_t) &= \frac{\pi(\boldsymbol{y}_t|\boldsymbol{x}_t)\hat{\pi}(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1})}{\pi(\boldsymbol{y}_t|\boldsymbol{Y}_{t-1})} \\
\hat{\pi}(\boldsymbol{x}_t|\boldsymbol{Y}_t) &\propto N(\boldsymbol{y}_t;\boldsymbol{H}_t\boldsymbol{x}_t,\boldsymbol{R})N(\boldsymbol{x}_t;\bar{\boldsymbol{x}}_t,\bar{\boldsymbol{P}}_t) \\
\hat{\pi}(\boldsymbol{x}_t|\boldsymbol{Y}_t) &= N\left(\boldsymbol{x}_t;\bar{\boldsymbol{x}}_t+\bar{\boldsymbol{P}}_t\boldsymbol{H}_t^{'}\bar{\boldsymbol{Q}}_t^{-1}(\boldsymbol{y}_t-\boldsymbol{H}_t\bar{\boldsymbol{x}}_t),\bar{\boldsymbol{S}}_t\right) \\
\bar{\boldsymbol{S}}_t &= \bar{\boldsymbol{P}}_t-\bar{\boldsymbol{P}}_t\boldsymbol{H}_t^{'}\bar{\boldsymbol{Q}}_t^{-1}\boldsymbol{H}_t\bar{\boldsymbol{P}}_t, \\
\bar{\boldsymbol{Q}}_t &= \boldsymbol{H}_t\bar{\boldsymbol{P}}_t\boldsymbol{H}_t^{'}+\boldsymbol{R}. \qquad (2.13)
\end{aligned}$$

Sampling from this updated distribution is achieved by drawing B i.i.d Gaussian variables from $\hat{\pi}(\boldsymbol{x}_t|\boldsymbol{Y}_t)$. We use Cholesky factorization to draw the updated samples, as described in the previous subsection. For the alternative sampling method, the prior samples in stage a are now directly available. Stage b of perturbing the data variate and stage c of forming the linear combination give a statistically equivalent method for generating updated samples.

The EnKF is robust in the sense that no degeneracy occurs. Viewed as a collapsed mixture of Gaussians,all the weights are constant, equal to $\frac{1}{B}$. All the particles in the predictive distribution have collapsed to the mean, and the components have the same covariance matrix. However, one cannot find theoretical justification for the Gaussian approximation of the prediction density. If this assumption is reasonable, the filter will perform very well. If the approximation is too far from the true non-Gaussian distribution, the filter will introduce bias and possibly diverge. In high dimensional systems, this does not seem to happen so often, and the practical aspects of the EnKF have shown very useful.

2.3.3 Robustified Gaussian mixture Monte Carlo filter

We now present the robustified Gaussian mixture filter, that aims to stabilize the GMF from Section 2.3.1, using the EnKF in Section 2.3.2 as a guide. Definition for $b = 1, ..., B$

$$\boldsymbol{z}_t^b = \alpha \boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big) + (1-\alpha)\bar{\boldsymbol{x}}_t, \tag{2.14}$$

where $\boldsymbol{z}_t^b$ are shrinked versions of the predictive particles. We can interpret the shrinkage as follows: For the two dimensional case, if $Z = \alpha X_1 + (1-\alpha)X_2$; $0 \leq \alpha \leq 1$, then Z is a point between X_1 and X_2 on the straight line which connects X_1 and X_2 (see Fig. 2.1). As a result (with a predicted sample point of view), the new predicted sample $\boldsymbol{z}_t^b = \alpha \boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big) + (1-\alpha)\bar{\boldsymbol{x}}_t$ is a sample on the hyperplane which connects $\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big)$ to $\bar{\boldsymbol{x}}_t$, and the position of this sample is between $\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big)$ and $\bar{\boldsymbol{x}}_t$ on the same hyperplane. The interpretation of Eq. 2.14 then is that the predicted particles are shrinked towards the mean. Using $\alpha = 0$, the shrinking is large, and the result is the EnKF. Using $\alpha = 1$, there is no shrinkage, and the result is the GMF. Again, we do not use any clustering (Frei and Künsch, 2013). We have one mixture component for every sample.

Figure 2.2 illustrates the predictive densities of the GMF, EnKF, and RGMF for a particular $0 < \alpha < 1$.

Associated with the shrinkage, we compute a predictive covariance matrix for every mixture component. We construct the predictive covariance matrix such that the total covariance matches that of the Gaussian mixture distribution in Eq. 2.7, just like what is done for the EnKF. That is, the predictive distribution in the RGMF has components $\tilde{\pi}(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1}, b) = N\big(\boldsymbol{x}_t; \tilde{\boldsymbol{z}}_t^b, \tilde{\boldsymbol{P}}_t\big)$, for $b = 1, .., B$, where $\tilde{\boldsymbol{P}}_t$ defines the second-order properties of the predictive distribution. Note that for the shrinked variables we have $\boldsymbol{z}_t^b - \bar{\boldsymbol{z}}_t = \alpha\big(\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big) - \bar{\boldsymbol{x}}_t\big)$. Then, by using the formula for double covariance, conditioning on the mixture components, we can ensure that the following holds:

$$\begin{aligned}\tilde{\boldsymbol{P}}_t \quad &+ \quad \alpha^2 \frac{1}{B}\sum_{b=1}^{B}\big(\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big) - \bar{\boldsymbol{x}}_t\big)\big(\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big) - \bar{\boldsymbol{x}}_t\big)' \\ &= \quad \boldsymbol{P} + \frac{1}{B}\sum_{b=1}^{B}\big(\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big) - \bar{\boldsymbol{x}}_t\big)\big(\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big) - \bar{\boldsymbol{x}}_t\big)'. \qquad (2.15)\end{aligned}$$

For this to be valid, the component-wise covariance matrix for the

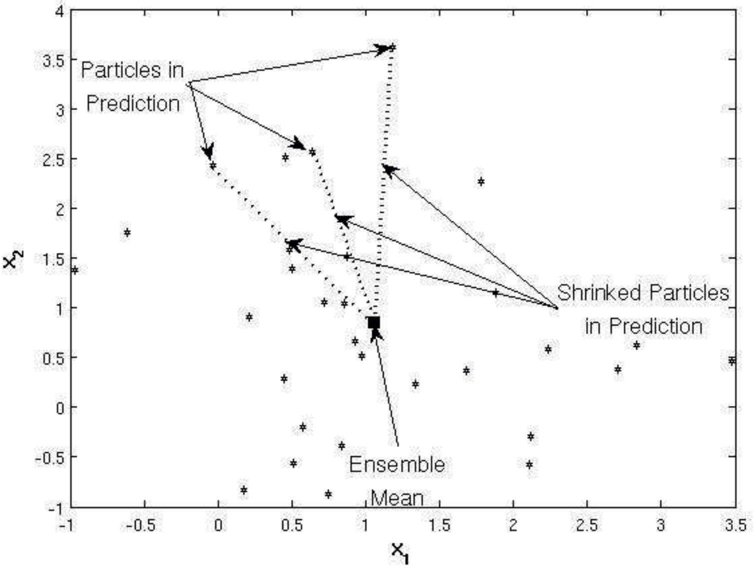

Figure 2.1: A graphicaldescription of shrinkage $\boldsymbol{z}_t^b = \alpha \boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big) + (1-\alpha)\bar{\boldsymbol{x}}_t$; the shrunk samples move on the line (*dotted line*) which connect the ensemble mean (*square*) to the ensemble members (*dotted points*)

shrinked mixture becomes

$$\tilde{\boldsymbol{P}}_t = \boldsymbol{P} + (1-\alpha^2)\frac{1}{B}\sum_{b=1}^{B}\big(\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big) - \bar{\boldsymbol{x}}_t\big)\big(\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big) - \bar{\boldsymbol{x}}_t\big)'. \tag{2.16}$$

The predictive distribution for the robustified version is a mixture over these mean and variances, such that

$$\tilde{\pi}(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1}) = \frac{1}{B}\sum_{b=1}^{B} N\big(\boldsymbol{x}_t; \boldsymbol{z}_t^b, \tilde{\boldsymbol{P}}_t\big). \tag{2.17}$$

The updating proceeds as for the GMF, with

$$\begin{aligned}\tilde{\pi}(\boldsymbol{x}_t|\boldsymbol{Y}_t) &\propto N(\boldsymbol{y}_t; \boldsymbol{H}_t\boldsymbol{x}_t, \boldsymbol{R})\tilde{\pi}(\boldsymbol{x}_t|\boldsymbol{Y}_{t-1}),\\ \tilde{\pi}(\boldsymbol{x}_t|\boldsymbol{Y}_t) &= \sum_{b=1}^{B}\tilde{w}_b N\big(\boldsymbol{x}_t; \tilde{\boldsymbol{x}}_t^b, \tilde{\boldsymbol{S}}_t\big),\end{aligned} \tag{2.18}$$

where $\tilde{\boldsymbol{x}}_t^b$ and $\tilde{S}_t$ are the updated mean and variance, given particle b, i.e.,

$$\begin{aligned}\tilde{\boldsymbol{x}}_t^b &= \boldsymbol{z}_t^b + \tilde{\boldsymbol{P}}_t \boldsymbol{H}_t' \tilde{\boldsymbol{Q}}_t^{-1} \big(\boldsymbol{y}_t - \boldsymbol{H}_t \boldsymbol{z}_t^b\big) \\ \tilde{\boldsymbol{S}}_t &= \tilde{\boldsymbol{P}}_t - \tilde{\boldsymbol{P}}_t \boldsymbol{H}_t' \tilde{\boldsymbol{Q}}_t^{-1} \boldsymbol{H}_t \tilde{\boldsymbol{P}}_t, \ \tilde{\boldsymbol{Q}}_t = \boldsymbol{H}_t \tilde{\boldsymbol{P}}_t \boldsymbol{H}_t' + \boldsymbol{R}. \end{aligned} \tag{2.19}$$

Naturally, all matrices in this expression depend on the shrinkage parameter α. The weights are now given by the following:

$$\tilde{w}_b = \frac{N\big(\boldsymbol{y}_t; \boldsymbol{H}_t \boldsymbol{z}_t^b, \tilde{\boldsymbol{Q}}_t\big)}{\sum_{c=1}^{B} N\big(\boldsymbol{y}_t; \boldsymbol{H}_t \boldsymbol{z}_t^c, \tilde{\boldsymbol{Q}}_t\big)}. \tag{2.20}$$

The GMF gives a very wiggly predictive density plot, while the EnKF is a Gaussian density. Now, if the data match one of the spikes, the particle associated with this spike would get a very large weight w_b in the GMF. This could cause degeneracy. The RGMF is smoother and closer to the Gaussian curve representing the EnKF. If the data match one of the spikes in the GMF representation, the associated increase in the weight for the RGMF, denoted by $\tilde{w}_b$, would not get that much larger than the remaining weights.

The robustification occurs because of the shrinkage. First, the $\boldsymbol{z}_t^b$s are more similar than the $\boldsymbol{g}_t\big(\boldsymbol{x}_{t-1}^b\big)$. Second, the variance $\tilde{\boldsymbol{Q}}_t > \boldsymbol{Q}_t$ and thus the weights become more uniform. The parameter α gives us some flexibility: small α entails a solution close to the Gaussian predictive density (like EnKF), while a large α is close to the GMF solution. One option is to tune α at every time t, i.e., $\alpha \to \alpha_t$. The tuning can be done using the weights $\tilde{w}_b = \tilde{w}_b(\alpha)$. The variance of the weights can be used as an indicator of sample degeneracy. When most samples are more or less equally weighted, we are far away from sample degeneracy, and the variance of the weights is low. If one or a few samples are much more likely to get sampled, some weights are large and the rest, very small. Consequently, the variance of the weights is high. We use the effective sample size (ESS) as an indicator for sample degeneracy. If the ESS is below a specified level, too many weights are negligible. Both the ESS and the variance of the weights capture the sample degeneracy, and they are related. Here, the ESS is defined by the following:

$$\mathrm{ESS}(\tilde{w}) = \sum_{b=1}^{B} \left[B\delta\left(\tilde{w}_b \geq \frac{1}{B}\right) + B\tilde{w}_b\delta\left(\tilde{w}_b \leq \frac{1}{B}\right)\right]. \tag{2.21}$$

The Dirac function is defined as $\delta(\boldsymbol{x}) = 1$ if $\boldsymbol{x} =$ true, and $\delta(\boldsymbol{x}) = 0$ otherwise.

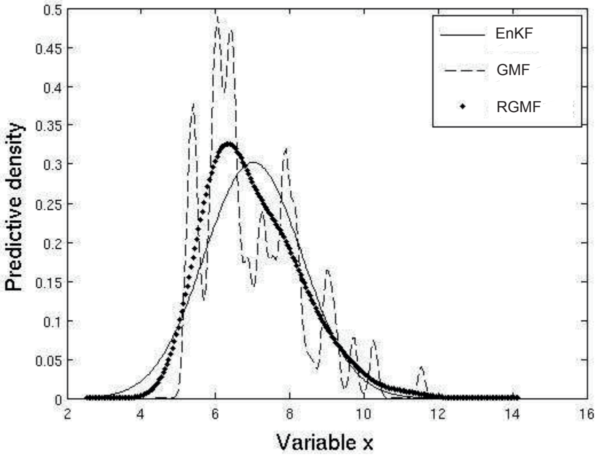

Figure 2.2: The predictive distribution from EnKF (*solid line*), GMF (*dashed line*), and RGMF (*dotted line*)

We can use the ESS to tune the level of shrinkage. We start with small α and increase the α parameter until the ESS is above a threshold: for instance, a fixed fraction, say $t_1 = \frac{B}{5}$.

Pseudo algorithm 1:

```
– Set $\alpha = \epsilon$.
– Tol = 0.
– Repeat until Tol = 1
    1. Compute $\tilde{w} = (\tilde{w}_1, ..., \tilde{w}_B)$.
    2. Compute effective sample size ESS($\tilde{w}$).
    3. If ESS($\tilde{w}$) $< t_1$ set Tol=1 and return $\alpha = \alpha - \epsilon$. Otherwise $\alpha = \alpha + \epsilon$.
```

The tuning parameter ϵ could be 0.1 or similar. Alternatively, one could start with the GMF and decrease α until the ESS goes above the threshold. Trade-off between the two exists too, such as guiding the algorithm by the α from the previous time step.

One could certainly use alternative criteria for tuning the shrinkage parameter α. Since the properties of the PF are well understood, we would

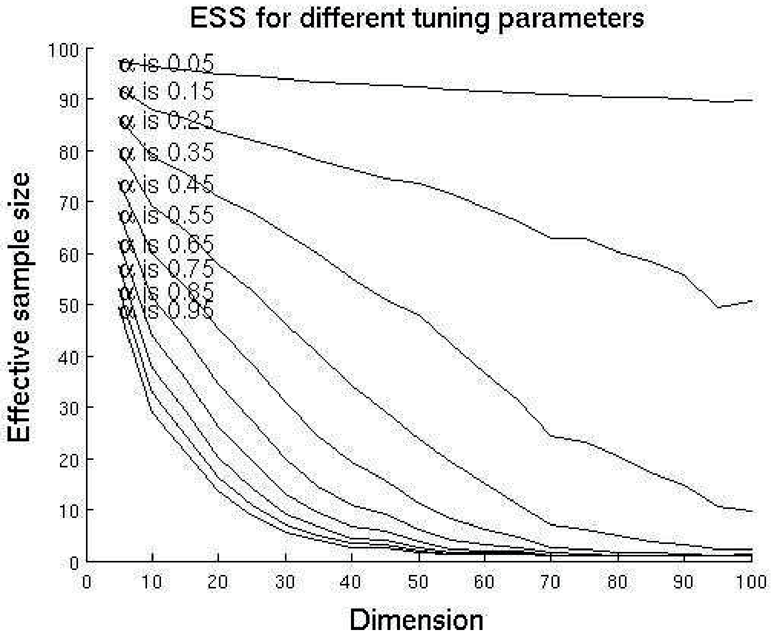

Figure 2.3: Effective sample size (ESS) for different shrinkage values α versus system dimension. For constant α, the ESS decreases as system dimension increases; for constant system dimension, the ESS decreases with increasing α

like to ensure a similar asymptotic behavior for the RGMF. As $B \to \infty$, one expects that the tuned α goes towards 1, and then the asymptotic properties are valid because we are in the PF domain. It is, thus, interesting to study the variability in the weights $\tilde{w}_b$ in Eq. 2.20, as a function of number of ensembles B and tuning parameter α. The weights naturally depend on the model $\boldsymbol{g}_t(\cdot)$, the process noise covariance matrix $\boldsymbol{P}$ and the observation noise covariance matrix $\boldsymbol{R}$, and on the tuning parameter α. Implicitly, the state dimension n and the observation dimension m become very important, like discussed above. Since the different particles have different mean values $\boldsymbol{g}_t\big(x_{t-1}^b\big)$, it is very hard to derive theoretical properties for the weights. Moreover, the particles $\boldsymbol{z}_t^b$ are dependent because they are all shrinked towards the common mean.

Instead, we use simulation to study the variance of the weights. We use a Gaussian model of dimension $n = m$ for both process and data variables. We use one time step $t = 1$ only, generating $\boldsymbol{g}_t(\boldsymbol{x}^b) \sim N\big(\boldsymbol{x}; \underline{0}, k^2\boldsymbol{I}_n\big), b = 1, ..., B,$. We set $k = 1$ and $\boldsymbol{x}_1 \sim \frac{1}{B}\sum_{b=1}^{B} N\big(\boldsymbol{x}_1; \boldsymbol{g}_t(\boldsymbol{x}^b), \boldsymbol{I}_n\big)$. The likelihood is $\boldsymbol{y}_1 \sim N(\boldsymbol{y}_1; \boldsymbol{x}_1, \boldsymbol{I}_n)$. We study the effect of dimension n varying from 5 to 100. The number of particles is , and we repeat the process 25 times to compute an average ESS under each B=100 parameter setting (see Fig. 2.3). For large α, the effective sample size decreases quickly, indicating a degeneracy in the weights. When α is close to 0, the updating is more like an EnKF, and the ESS remains larger; but in high dimensions, an α slightly larger than 0 gives different filtering distributions than the EnKF.

2.4 Simulation

In this part, we compare the proposed algorithm with EnKF and GMF. We consider four different cases: the first two are single-sensor single-target and multi-sensor multi-target tracking examples, where the posterior distributions are multi-modal. The third one is the famous Lorenz 40 model, and the last one is a synthetic reservoir simulator example. For each case, we vary the level of shrinkage and ensemble size to check the performance of the algorithms.

In the examples, we compare the different algorithms in terms of mean square error (MSE),continuously ranked probability score (CRPS), and probability distribution coverage. Here, at any time t, we have $\text{MSE}(t) = \sum_j \big(\hat{x}_{j,t} - x_{j,t}^{\text{true}}\big)^2$, where $\hat{x}_{j,t}$ is the estimated mean of the filtering distribution, and the sum is over all n state dimensions. An integrated MSE is achieved by summing out t. Further, the CRPS is defined by $\text{CRPS}(t) = \sum_j \big(\hat{F}(y_{j,t}) - I\big(y_{j,t} < y_{j,t}^{\text{obs}}\big)\big)^2$ (Gneiting et al., 2005). Here, $\hat{F}(.)$ is the empirical cumulative predictive distribution for data at time t, given all former data $\boldsymbol{Y}_{t-1}$. Smaller values of CRPS means better predictive power. It shows that we often match the observed value, and that we have a narrow prediction band. The sum is over all m observation dimension, and an integrated CRPS is obtained by summing over all times t. The probability distribution coverage is defined as the proportion of times the fitted filtering distribution percentiles cover the true state, over a number of replicates.

2.4.1 Tracking targets with bimodal distributions

These examples describe the position and velocity of planes or ships moving in two dimensions. If we imagine a monitoring system for planes or ships, their positions are measured by radar /sonar. The targets move in a dependent pattern, i.e., if one turns, others likely turn as well. In this simulation, we consider two cases: one-sensor one-target and ten-sensor ten-target. We let $\boldsymbol{x}_t = [x_t \; \dot{x}_t \; y_t \; \dot{y}_t]'$ be the state vector of one target. For one target, $(x_t \; y_t)$ is the north and east positions, and similarly, $(\dot{x}_t \; \dot{y}_t)$ is the north and east velocity. The absolute velocity is $v_t = \sqrt{\dot{x}_t^2 + \dot{y}_t^2}$, while the target is moving at a bearing $\eta_t = \arctan\left(\frac{y_t}{x_t}\right)$.

With constant velocities, a target moves in a straight line, and the dynamical model is linear. We consider a situation where a target maneuvers (30 degrees) to the west whenever the velocity v_t becomes smaller than a threshold c. This model is non-linear, and the dynamics can be phrased by $\pi(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}) \sim N(\boldsymbol{x}_t; \boldsymbol{g}_t(\boldsymbol{x}_{t-1}), \boldsymbol{P})$. Using a time step d$T$, the one-target dynamics for large velocity is

$$\boldsymbol{g}_t(\boldsymbol{x}_{t-1}) = \begin{bmatrix} 1 & \mathrm{d}T & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & \mathrm{d}T \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_{t-1} \\ \dot{x}_{t-1} \\ y_{t-1} \\ \dot{y}_{t-1} \end{bmatrix} \tag{2.22}$$

while for small velocity, it is

$$\boldsymbol{g}_t(\boldsymbol{x}_{t-1}) = \begin{bmatrix} x_{t-1} + \mathrm{d}T\cos(\eta_t)v_{t-1} \\ \cos(\eta_t)v_{t-1} \\ y_{t-1} + \mathrm{d}T\sin(\eta_t)v_{t-1} \\ \\ \sin(\eta_t)v_{t-1} \end{bmatrix} \tag{2.23}$$

$$\eta_t = \frac{\pi}{6} + \eta_{t-1}, \qquad \text{if} \quad v_{t-1} < c. \tag{2.24}$$

Thus, bearing η_t of one target at time t changes westward when the absolute velocity is small. This has an effect on the north and east velocity, whereas the absolute velocity $v_t = v_{t-1}$ remains the same, as expected. As a consequence, the predictions of the north and east positions will tend to be skewed or multi-modal when the distribution for velocity is near the critical velocity c.

The process noise covariance matrix is $\boldsymbol{P} = \mathrm{diag}([0.5^2,\ 2^2,\ 0.5^2,\ 2^2])$, and initial conditions are drawn from $N(\boldsymbol{x}_0; \boldsymbol{\mu}_0, \boldsymbol{P}_0)$, where $\boldsymbol{\mu}_0 = [1000, 75, 1000, 75]'$ and $\boldsymbol{P}_0 = 100\boldsymbol{P}$. We introduce a fixed correlation of 0.9 between all targets, and the joint covariance is a block diagonal in the multi-target situation.

We observe the north and east positions at every time point, with Gaussian additive noise. Thus, the likelihood model for position data is linear and can be phrased by $\pi(\boldsymbol{y}_t|\boldsymbol{x}_t) \sim N(\boldsymbol{y}_t; \boldsymbol{H}_t\boldsymbol{x}_t, \boldsymbol{R})$, where $\boldsymbol{R} = \mathrm{diag}([5^2\ \ 5^2])$ and

$$\boldsymbol{H}_t = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \tag{2.25}$$

Table 2.1: Comparison of different filters for one-sensor one-target case

Values	EnKF	GMF	RGMF
α			0.00001
CRPS(STD)	144(1.5)	140(1.4)	144(1.6)
MSE(STD)	281(26)	87(6.5)	282(26)
Coverage(%)	75	87	75
α			0.2
CRPS(STD)	-	-	143(1.6)
MSE(STD)	-	-	272(26)
Coverage(%)	-	-	74
α			0.5
CRPS(STD)	-	-	140(1.4)
MSE(STD)	-	-	230(14)
Coverage(%)	-	-	76
α			0.8
CRPS(STD)	-	-	131(1.3)
MSE(STD)	-	-	164(8)
Coverage(%)	-	-	85
α			0.99
CRPS(STD)	-	-	140(1.5)
MSE(STD)	-	-	83(7)
Coverage(%)	-	-	89

STD standard deviation

One-sensor one-target case

In this case, we assume that we have one moving target and one radar/sonar sensor measuring the target position in the Cartesian plane.

The main goal of this example is to evaluate the performance of each filter for estimating a general distribution when the system dimension is low (system dimension is $n = 4$ for this case), and there is no sign of sample degeneracy. The number of ensemble members is $B = 500$, and we repeat the simulation for 500 replicates in order to reduce Monte Carlo error. We predict that GMF should be the best, because for low dimensional systems, it can approximate general filtering distributions without sample degeneracy. Moreover, we have a large number of samples here ($B = 500$). We tested with smaller Bs, and sometimes, the GMF diverges, even for such small dimensions.

Table 2.1 and Figs. 2.4 and 2.5 present the simulation results for a selected α. We see that our prediction about GMF is correct, and it is the

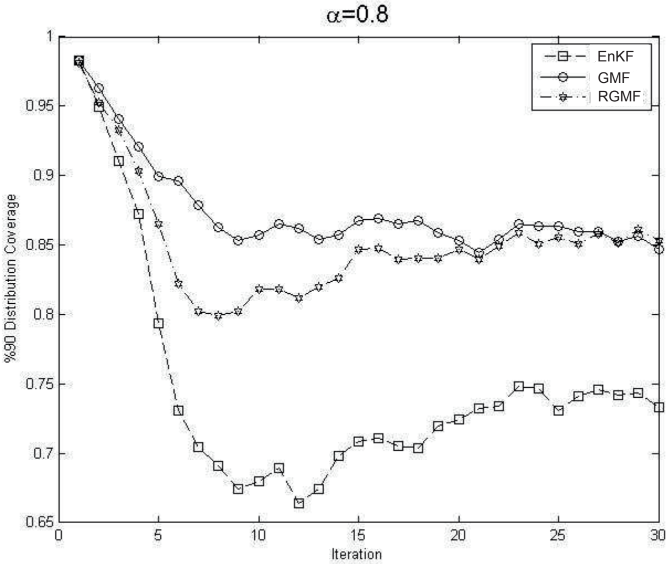

Figure 2.4: One-sensor one-target: distribution coverage for $\alpha = 0.8$, EnKF (*dashed square line*) has the lowest coverage, GMF (*solid circle line*) has the highest coverage, and the RGMF (*dashed dotted star line*) is in between but reaches the coverage of GMF as time goes

best in terms of MSE, CRPS, and probability distribution coverage (nominal level is 90%). When α is close to 1, the RGMFs performance is close to GMF. This result is natural since it is close to the GMF for large values of α. By looking at Table 2.1, we also see that when $\alpha \approx 0$, the RGMFs estimation accuracy is close to EnKF. This result is also predictable because the shrinked samples goes to the predicted ensemble mean.

Figure 2.4 tells us that the percentile coverage distribution starts from a high value (98%) then reduces rapidly. The high initial coverage is caused by the initial state variables. It is noticeable that after some time steps, the percentile coverage for GMF and RGMF reduces with a lower rate than EnKF. This occurs because the posterior distribution is non-Gaussian, and the EnKF does not approximate it very well.

Figure 2.5 shows the MSE and associated confidence intervals, which are calculated using bootstrapping of samples. The GMF has the lowest trend. We can see similar results for CRPS. Simulation results show that

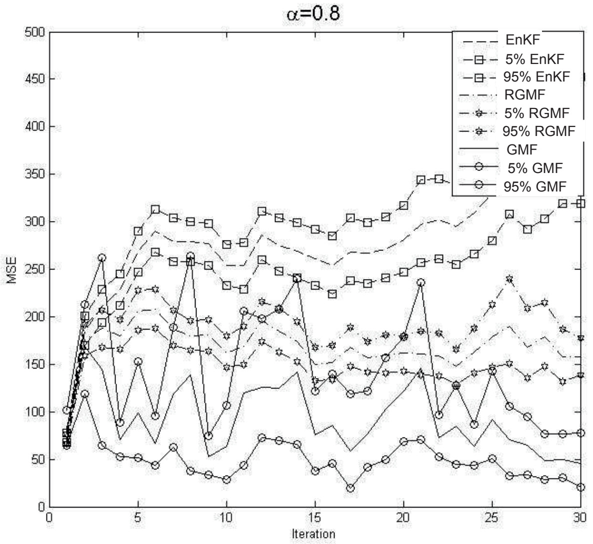

Figure 2.5: One-sensor one-target: MSE for $\alpha = 0.8$, EnKF (*dashed square line*) has the highest MSE with 95% confidence interval, GMF (*solid circle line*) has the lowest MSE, and the RGMF (*dashed dotted star line*) is in between EnKF and GMF

by choosing different values for α, we can change the performance of the filter: being close to GMF when α increases to 1 and close to EnKF for α near 0. It has different performance from EnKF and GMF for $0 < \alpha < 1$.

Ten-sensor ten-target case

In this case, we increase the system dimension by increasing the number of targets and sensors. Now, we are going to check the proposed method for dealing with sample degeneracy in higher dimensional system. We know that the GMF suffers from sample degeneracy when the system dimension increases, and we predict that it should have the worst performance in this example. We also know that EnKF tends to work well in high dimensional systems. In Table 2.2 and Figs. 2.6 and 2.7, we see that the GMF now has the worst performance as we predicted, while EnKF works well. We see that the proposed method works well for this case without sample degeneracy.

Table 2.2: Comparison of different filters for ten-sensor ten-target case

Values	EnKF	GMF	RGMF
α			0
CRPS(STD)	1389(6.9)	2311(12.7)	1389(6.6)
MSE(STD)	257(8.1)	581(28.5)	259(7.3)
Coverage(%)	87	25	87
α			0.1
CRPS(STD)	-	-	1390(7.3)
MSE(STD)	-	-	259(8.7)
Coverage(%)	-	-	86
α			0.4
CRPS(STD)	-	-	1389(7.2)
MSE(STD)	-	-	246(6.7)
Coverage(%)	-	-	87
α			0.7
CRPS(STD)	-	-	1513(8.8)
MSE(STD)	-	-	264(13.3)
Coverage(%)	-	-	83
α			1
CRPS(STD)	-	-	2319(12)
MSE(STD)	-	-	594(27.3)
Coverage(%)	-	-	26

STD standard deviations

Figure 2.6 shows that the percentile coverage (nominal 90%) for the GMF is high at the beginning, because the initial samples are independent, and the distributions are close to the Gaussian distribution. But after some time steps, the coverage drops because of sample collapse (the ESS is low). We have a similar behavior for the MSE (Fig. 2.7).

For small to medium values of α, the performance of the RGMF is as good as the EnKF. According to Table 2.2, for some values (such as $\alpha = 0.4$), the RGMF is the best in the MSE sense. Besides, Table 2.2 shows that the RGMF is close to the EnKF for a wide range of shrinkage levels ($0 < \alpha < 0.8$). For larger αs, the RGMF goes to the GMF. When $\alpha = 0$, the performance of RGMF is a little bit different from EnKF. This is because of Monte Carlo error. We have a similar story for the performance of GMF and RGMF when $\alpha = 1$.

Note that the coverage is better for this ten-sensor ten-target than for single-target situation. We believe this occurs because the ten-target model

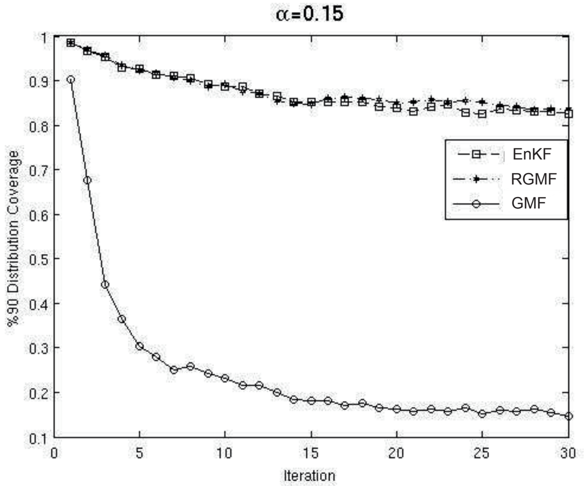

Figure 2.6: Ten-sensor ten-target: distribution coverage for $\alpha = 0.15$, GMF (*solid circle line*) has the lowest coverage, and also the coverage of EnKF (*dash square line*) and RGMF (*dashed dotted star line*) is almost equal

contains more Gaussianity, in the sense that only one or some targets maneuver at a time, and thus the non-linearity is averaged over more dimensions.

2.4.2 Lorenz 40 model

The Lorenz40 model is a highly non-linear model with state dimension $n = 40$ (Sakov and Oke, 2008). It consists of 40 ordinary differential equations with cyclic boundary condition as follows:

$$\begin{aligned} \dot{\lambda}_i &= (\lambda_{i+1} - \lambda_{i-2})\lambda_{i-1} - \lambda_i + 8, \quad i = 1, ..., 40; \\ \lambda_0 &= \lambda_{40}, \quad \lambda_{-1} = \lambda_{39}, \quad \lambda_{41} = \lambda_1. \end{aligned} \tag{2.26}$$

This model is discretized by the standard fourth-order Runge–Kutta scheme, and the system states $\boldsymbol{x}_t$ relate to the solution of the above system as $\boldsymbol{x}_t = \boldsymbol{\lambda}_{0.05t}, t = 1, 2, ...$ where $\boldsymbol{\lambda}_t = [\lambda_1(t), \lambda_2(t), ..., \lambda_{40}(t)]'$. The system dynamics are perturbed with Gaussian noise with mean 0 and standard deviation 0.05. The system state is measured by

$$\boldsymbol{y}_t = \boldsymbol{x}_t + N(0, I_{40}). \tag{2.27}$$

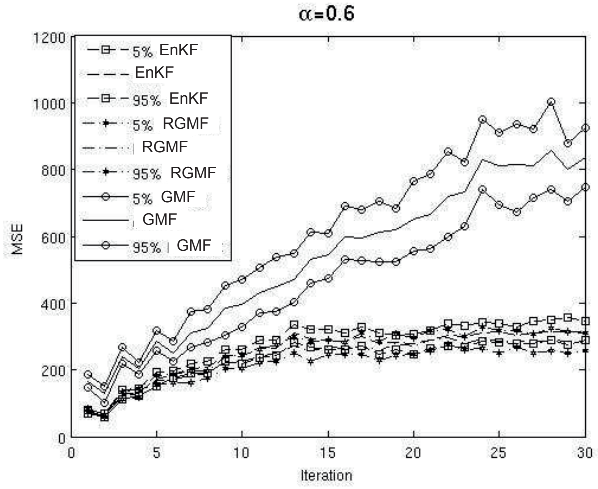

Figure 2.7: Ten-sensor ten-target: MSE for $\alpha = 0.6$, GMF (*solid circle line*) has the highest MSE with 95% credible interval, and also the MSE of EnKF (*dashed square line*) and RGMF (*dashed dotted star line*) is almost equal

We have these measurements at each time step. The initial ensemble members are selected randomly from a set of 10,000 model states obtained from one continuous integration at $t = 1000, 1001, ..., 11000$ (Sakov and Oke, 2008).

In this simulation, we evaluate the proposed method on this famous test in data assimilation. For this example, we choose a constant value for α (instead of adaptive α selection) in order to simultaneously check the sensitivity of the algorithm with respect to the value of α and also the number of ensembles (Figs. 2.8 and 2.9). The comparison of GMF, EnKF, and RGMF is done for different cost functions and evaluation criteria such as the total MSE, percentile posterior distribution coverage, and the variance of the weights.

In Fig. 2.8, we display a three-dimensional plot of total MSE versus ensemble size and the tuning parameter α, averaged for the time interval [900 1000] of the Lorenz model. We note that the MSE tends to decrease with larger ensemble size. According to Fig. 2.8, for small sample size, the smallest MSE occurs for the EnKF ($\alpha = 0$). The GMFs result ($\alpha = 1$) is the worst because the sample size B must be very high for the GMF to obtain the best result. It is not surprising because we know that GMF diverges

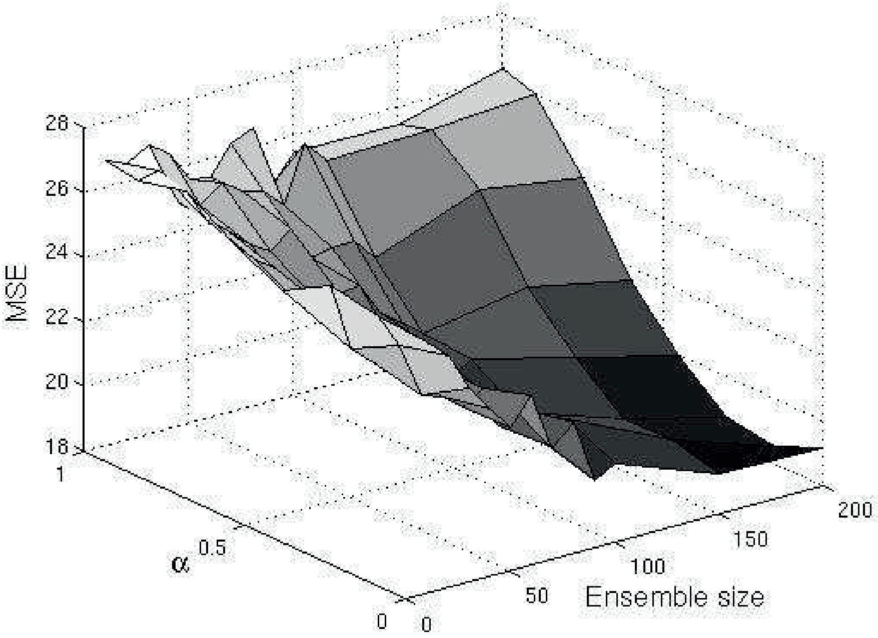

Figure 2.8: Lorenz 40: total MSE versus the ensemble size and tuning parameter α, averaged for the interval [900 1000]

when system dimension is proportionally high (40 for this case), since it suffers from sample degeneracy.

In some cases, when we increase the number of samples, the number of collinear samples increases too, and this results in poorer estimation accuracy. For instance, when $\alpha = 0$ (EnKF), by changing the number of ensembles from 90 to 100, the total MSE increases. Also, for the GMF case (i.e., $\alpha = 1$), the MSE increases as the number of samples goes from 150 to 200. According to Fig. 2.8, the best result in total MSE sense is when $\alpha = 0.15$ and the number of samples is 200. Choosing the optimal value for the tuning parameter α is highly dependent on sample size, and its optimal value depends on both system complexity and number of samples.

Figure 2.9 shows the filter results in the percentile posterior distribution coverage sense. The distribution coverages are too small, but they increase as a function of sample size B. Based on this figure, we have the best coverage for $0.15 \leq \alpha \leq 0.3$. Thus, the RGMF performs the best in this range of shrinkage values. For constant sample size, the coverage distribu-

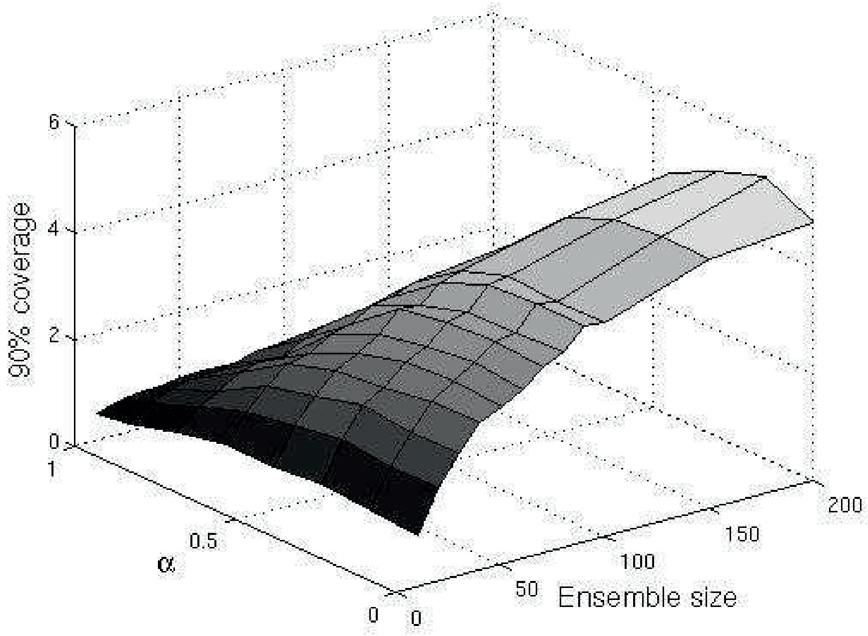

Figure 2.9: Lorenz 40: distribution coverage versus the ensemble size and tuning parameter α, averaged for the interval [900 1000]

tion is like a convex polynomial function of the tuning parameter α (such as second-order polynomial function), where its maximum value is around $\alpha = 0.15$. Also, increasing the number of ensembles does not always lead to increased distribution coverage because of sample collinearity. For instance, when $\alpha = 0.45$, and we increase the number of samples from 70 to 80 or when $\alpha = 0$ (EnKF) and the number of samples is changed from 90 to 100, the posterior distribution coverage decreases.

The coverage is so low because of small process noise covariance. It increases rapidly when we set the process noise to a larger value. For instance, when we increase the process noise standard deviation from 0.05 to 0.1, the percentile coverage distribution increases to about 80%. The parameters used in Sakov and Oke (2008) describe a hard case.

We know that the variance of the weights is 0 for the EnKF, because all weights are equal. For high α, the variance of the weights is much higher. However, when the ensemble size is high, the variance of the weights does not change so much for different αs.

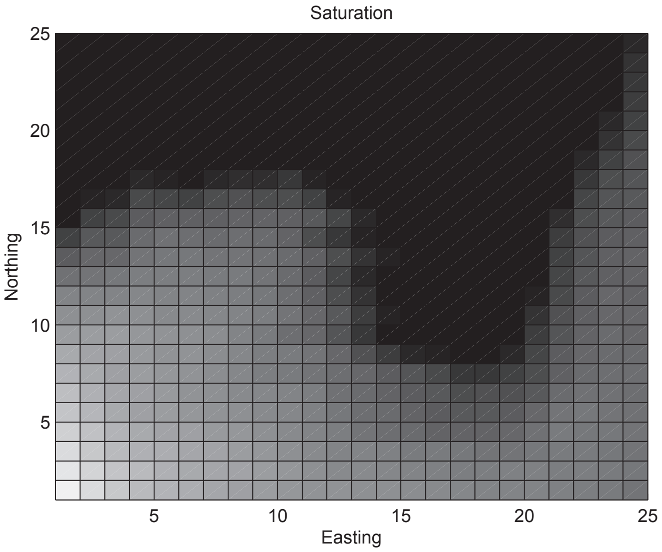

Figure 2.10: Saturation for one run at the third time step (after 300 days)

2.4.3 Saturation estimation based on seismic and electromagnetic data

We next explore the performance of the proposed method on a high dimensional system. Besides, the proposed adaptive method for choosing α is used. The example we consider comes from a reservoir simulation model with a single injection well and a single production well (Lie et al., 2010). The injection well is located at the lower left corner (Fig. 2.10), and water is pumped into this well for replacing and moving oil to the production well, which is located at the upper right corner. Based on fluid dynamics, the flow is faster where the permeability (porosity) is high. There is some heterogeneity in the initial oil saturation in the reservoir. Figure 2.10 shows the saturation profile after 300 days for one realization of the reservoir simulator, while Fig. 2.11 shows the associated oil production over the first 300 days.

The simulator works on fluid dynamics laws, and it contains some parameters such as porosity, permeability, and saturation. The values of these parameters depend on the geophysical characteristics of the reservoir. Assuming the whole reservoir is a combination of cells in a lattice, we are interested in estimating the porosity, permeability, and saturation for each cell. We use a 25×25 lattice in our example. The system state vector

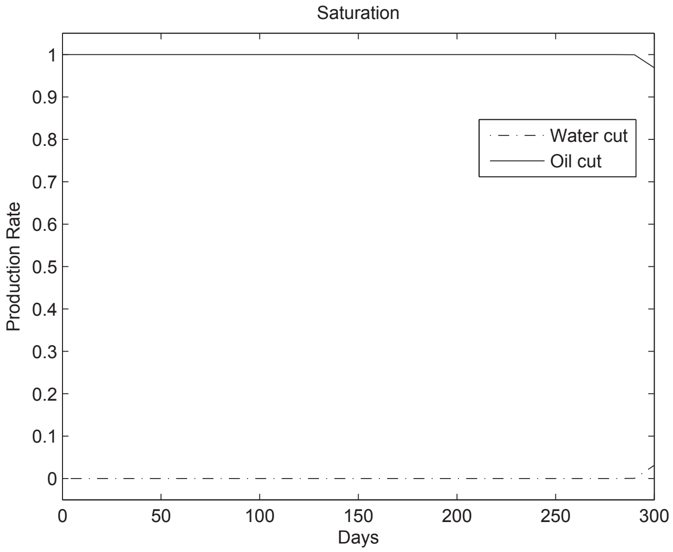

Figure 2.11: Production rate over 300 days at the well in the *upper right* corner of the grid

consists of porosity, permeability, and saturation for each cell. For this example, we assume that the porosity and permeability are known and constant, and we consider the spatial saturation distribution as a random vector developing over time. The noise covariance in the dynamic model is $\boldsymbol{P} = \boldsymbol{P}_t = \boldsymbol{P}(\boldsymbol{x}_t, \boldsymbol{x}_{t-1}) = \text{diag}(|\boldsymbol{x}_t - \boldsymbol{x}_{t-1}|)$, where $\boldsymbol{x}_t$ is the state vector of logistic saturations at all $25 \times 25 = 625$ cells. Logistic saturation transforms the variables on the real line. The initial distribution for the logistic saturation is a Gaussian random field, with a correlation length of about ten cells. At the first time step, this distribution is updated using data from the baseline seismic survey. The dynamic model $\boldsymbol{g}_t$ is the numerical solution of the PDE for fluid flow, and $t = 1, 2, 3$ are the three discretization points, using a 100-day time step, i.e., $t = 1$ means 100 days, etc. (Lie et al., 2010). More details about the simulator are available at http://www.sintef.no/Projectweb/MRST/.

The data consists of seismic and electromagnetic observations (see Fig. 2.12), and these are repeated over time. We assume that initially at time 0, a baseline survey is performed, and a monitoring survey is performed at days 100, 200, and 300 (i.e., at time steps 1, 2, and 3). Two seismic attributes and electromagnetic resistivity are obtained and interpolated along

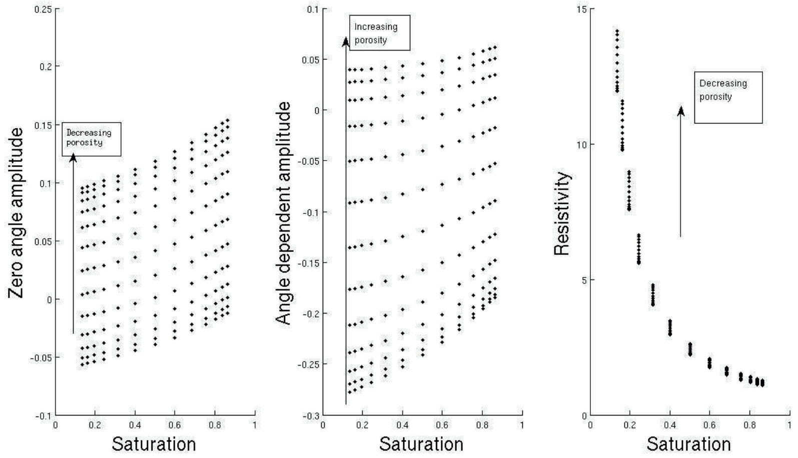

Figure 2.12: Expected response of seismic amplitude data (*two left plots*) and electromagnetic resistivity data (*right plot*) for different values of saturation (*first axis*) and for changing porosity

the top reservoir described by the lattice. The data are informative of the reservoir variables. Seismic attributes are very informative about the porosity and partly saturation, while the electromagnetic data carries coarse scale information about saturation. The likelihood is a function of saturation (Mavko et al., 2003), and the expected response is shown in Fig. 2.12. These expected seismic responses are obtained using a Reuss fluid mixing model and Gassmann's formula for fluid substitution. The electromagnetic resistivity level is calculated by Archie's polynomial law (Eidsvik et al., 2008). The observations are made at all lattice cells, and its dimension is $m = 3 \times 25^2$. The likelihood is slightly non-linear, but we linearize it by using first-order Taylor series expansion. This means that we fit Gaussian approximations at each sample member in all of our algorithms.

In order to check the effect of the number of ensembles on the estimation performance, we consider two different ensemble sizes, namely $B = 100$ and $B = 200$. A comparison between EnKF and RGMF is done over MSE, CRPS, and percentile coverage distribution, averaged over 40 replicates. This entails that we run a simulator 40 times, and for each of these runs, we compute the EnKF and RGMF filtering distributions at all time steps. First, we use only the seismic data. Consequently, the expected response only contains the seismic equations which relate the saturations, as state vector, to the observations of size 2×625. We run the simulator and filters,

Table 2.3: Filter performance for estimating saturation conditioned on seismic amplitude data

Values	EnKF	RGMF
Number of ensembles(B)	100	100
CRPS(STD)	4.1e4(4.2)	4.0e4(3.3)
MSE(STD)	2.6e-2(1.3e-4)	2.9e-2(1.6e-4)
Coverage(%)	76.4	84.0
Number of ensembles(B)	200	200
CRPS(STD)	4.1e4(3.8)	4.0e4(4.3)
MSE(STD)	2.5e-2(2.1e-4)	2.7e-2(1.4e-4)
Coverage(%)	77.0	85.7

We used an adaptive shrinkage parameter α
STD standard deviations

and every 100 days, when the seismic data are available, we condition the saturation variables on this data set.

Table 2.3 shows the MSE, CRPS, and coverage for ensemble sizes $B = 100$ and $B = 200$ for the EnKF and RGMF. The RGMF is based on the proposed adaptive α selection method. The results are averaged over all 625 variables on the lattice and over the three time steps. When we increase B, the MSE and CRPS decreases, while the coverage increases and gets closer to the nominal level of 95%. This holds for both EnKF and RGMF. However, doubling the ensembles size gives twice the computation burden. In comparison, the RGMF is not a large burden, since the reservoir simulator is the most time-consuming part of the algorithms. The CRPS is slightly smaller for the RGMF; this means that the RGMF has a sharper prediction of the true saturation. The total MSE of EnKF is actually smaller than that of RGMF. A more detailed analysis of the MSE trend of both filters over the three time steps shows that the MSE of the RGMF is higher only at time step $t = 1$. Both at times 2 and 3, the EnKF has a higher MSE than RGMF (see Fig. 2.13). It seems the EnKF in this case works as overshrinking in the first step and possibly pays for this in the later stages. When we turn to the distributional point of view by comparing distribution coverage values, the RGMF is clearly better than the EnKF (the distribution coverage of the RGMF is higher than the EnKF by a factor of 8%). The coverage probability is, in some senses, more important than the previous indicators, because it tells us how much we can rely on our filtering distributions. It measures how often the true state is covered by our estimate and uncertainty bounds. Moreover, the coverage calculation for the RGMF is based on using the

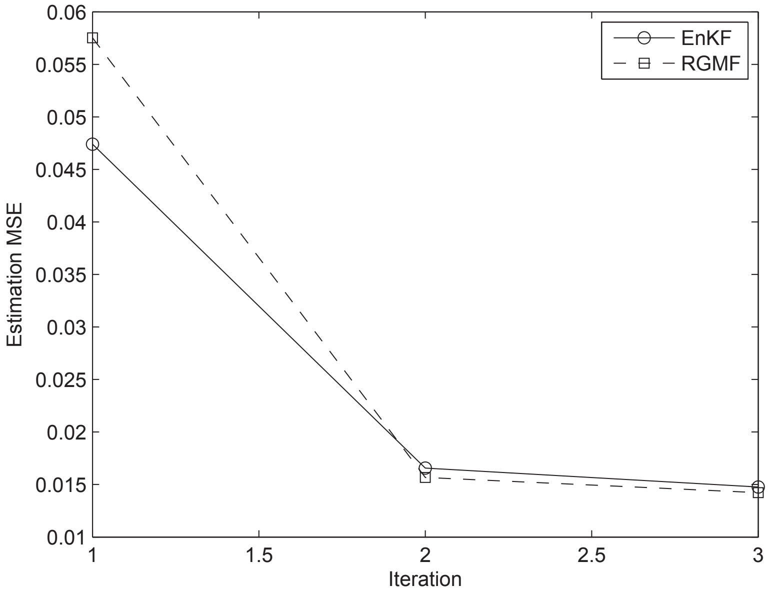

Figure 2.13: Estimation MSE when the ensemble size is $B = 100$ and the estimation is only based on seismic data

mixture of Gaussian, and this would capture skewed posterior distributions.

We use the estimated saturation from each filter to estimate the oil production rate. The results show that the estimated production rate by RGMC has lower MSE and better coverage than EnKF. Inflation would likely improve the average coverage probabilities of the EnKF (Evensen, 2009), but this heuristic approach gives too large variance for some parameters and still too small for others, so the marginal coverage would perhaps not improve much.

Here, the α value is tuned at each step of the RGMF algorithm. Sometimes, the shrinkage parameter goes down to almost 0 (we set a threshold at 0.005), and this corresponds to the EnKF algorithm. At other times, it adapts to values of 0.4–0.1, but levels of about 0.01–0.05% are most common. Just like for the Lorenz 40 example, we see a tendency of higher α values for $B = 200$ than for $B = 100$. In particular, the upper tail of the α distribution over time and replicates is much heavier for $B = 200$, i.e., it is more common to get shrinkage in the range of $\alpha = 0.4 - 0.1$. Nevertheless, in this high dimensional example, B must probably be much higher to get high α values closer to that of the GMF.

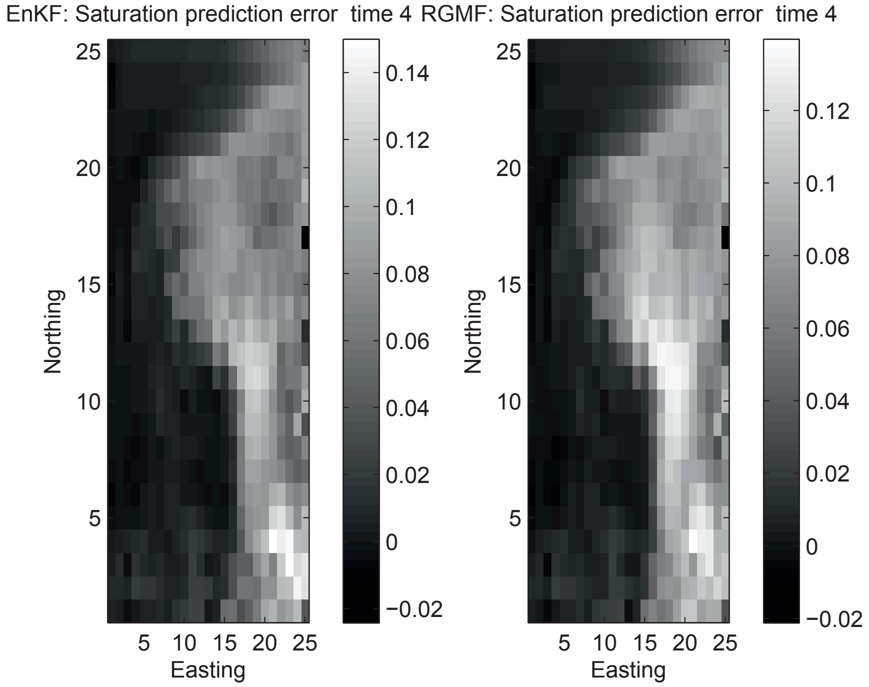

Figure 2.14: Saturation prediction error at the fourth time step (after 400 days) when the ensemble size is $B = 100$ and only seismic observations used

Figure 2.14 shows the saturation prediction error at the fourth step (after 400 days). That is, we propagate sample into the future. Notably, for both methods, the errors are highest near the front of the water flow. The display shows that the saturation prediction error, which is the difference between predicted saturation and the truth, has smaller range for the RGMF ($-0.021 < \mathrm{Err}^{\mathrm{RGMF}}_{\mathrm{prediction}} < 0.139$) than for the EnKF ($-0.024 < \mathrm{Err}^{\mathrm{EnKF}}_{\mathrm{prediction}} < 0.149$). This means that predictive capability of the RGMF is good. There are some small artifacts in the prediction plots, caused by the truncation and noise level in the dynamical model. This is a common challenge when merging numerical solution of PDEs and statistical Monte Carlo sampling.

For the second case, we now use our observations on seismic and electromagnetic data sets. Thus, the response vector consists of both seismic and electromagnetic measurements, at all lattice cells, and the likelihood (Fig. 2.12) relates the saturation (state vector) to the size $m = 3 \times 25^2$ observation vector at each time step. Simulation results show that the MSE and CRPS are clearly reduced for both EnKF and RGMF (the MSE for the EnKF reduces to 1.6e–2 and for the RGMF, to 1.5e–2). This is natural because we have more information. The CRPS value of EnKF is now slightly lower than for RGMF. The coverage goes down from before (76% for the EnKF and 77% for the RGMC). This seems to be a result of high correlation and much collinearity. The adaptive α now typically ends up near $0 - 0.05$, i.e., a slightly lower level than for the case without electromagnetic data.

2.5 Closing remarks

We have studied a filtering method going between the EnKF and the particle filter. Our modeling setup allows us to phrase the sample approximation as a Gaussian mixture. Nevertheless, the Gaussian mixture filter resulting from this approach will degenerate because of the sampling approximation. The robustified procedure we outline in this paper shrinks samples towards the ensemble mean, while maintaining some flexibility in the distribution. The degree of shrinkage is the only tuning parameter in our algorithm. The mixture mean, covariance, and weights follow automatically. Our suggested algorithm is a natural intermediate version of the EnKF and the Gaussian mixture filter in this situation.

We tested the robustified Gaussian mixture filter on a simulation study for target tracking, on the Lorenz 40 model, and on a reservoir simulation example. Results indicate that the robustified filter works better or as well as the Gaussian mixture particle filter and ensemble Kalman filter for systems with different dimensions and complexities.

In the reservoir simulation example, we used the (default) permeability field for the reservoir simulator developed by SINTEF http://www.sintef.no /Projectweb/MRST/ . For a real situation, the permeability field could show more heterogeneity. The differences between the EnKF and RGMF would become larger if we, for instance, have large uncertainty in the saturation variables near channel zones. Intuitively, this causes some realizations to move much faster ahead than others. This induces relevant bimodality or skewness that is not captured by a Gaussian predictive distribution.

In future work, we plan to apply the methods to reservoir models that parametrically characterize features of the reservoir rather than work on a grid. This feature approach gives enormous dimension reduction, but typically induces severe non-Gaussianity. We believe our methods could have the largest potential in such applications.

Acknowledgments

We thank the sponsors of the Uncertainty in Reservoir Evaluation (URE) project at the NTNU.

We further thank Hans Kunsch for helpful discussions. We thank Sintef for the MATLAB reservoir simulation toolbox and NERSC for the Lorenz 40 scripts.

Bibliography

Anderson, B. D. and Moore, J. B. (1979). *Optimal filtering*, volume 11. Prentice-hall Englewood Cliffs, NJ.

Doucet, A., d. Freitas, N., and Gordon, N. (2001). *Sequential Monte-Carlo Methods in Practice*. Springer.

Doucet, A. and Johansen, A. M. (2011). *A tutorial on particle filtering and smoothing: fifteen years later. In: Oxford Handbook of Nonlinear Filtering, 2011.* Oxford University Press.

Dovera, L. and Rossa, E. D. (2011). Multimodal ensemble Kalman filtering using Gaussian mixture models. *Computational Geosciences*, 15(2):307–323.

Eidsvik, J., Bhattacharjya, D., and Mukerji, T. (2008). Value of information of seismic amplitude and CSEM resistivity. *Geophysics*, 73(4):59–69.

Evensen, G. (2003). The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean dynamics*, 53(4):343–367.

Evensen, G. (2009). *Data assimilation, The Ensemble Kalman Filter*. Springer.

Fishman, G. (1996). *Monte Carlo: concepts, algorithms, and applications*. Springer.

Frei, M. and Künsch, H. R. (2013). Mixture ensemble Kalman filters. *Computational Statistics & Data Analysis*, 58(0):127 – 138.

Gneiting, T., Raftery, A. E., Westveld III, A. H., and Goldman, T. (2005). Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review*, 133(5):1098–1118.

Hoteit, I., Luo, X., and Pham, D.-T. (2011). Particle Kalman filtering: A nonlinear Bayesian framework for ensemble Kalman filters. *Monthly Weather Review*.

Hoteit, I., Pham, D., Triantafyllou, G., and Korres, G. (2008). A new approximate solution of the optimal nonlinear filter for data assimilation in meteorology and oceanography. *Monthly Weather Review*, 136(1):317–334.

Ito, K. and Xiong, K. (2000). Gaussian filters for nonlinear filtering problems. *Automatic Control, IEEE Transactions on*, 45(5):910–927.

Jazwinsky, A. (1970). *Stochastic Processes and Filtering Theory*. Academic Press, New York.

Julier, S. J. and Uhlmann, J. K. (1997). A new extension of the Kalman filter to nonlinear systems. *SPIE - Int. Soc. Opt. Eng.*, 3068(1):182–193.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *ASME Journal of Basic Engineering*, 1(1):35–45.

Leeuwen, P. J. V. (2010). Nonlinear data assimilation in geosciences: an extremely efficient particle filter. *Quarterly Journal of the Royal Meteorological Society*, 136(653):1991–1999.

Leeuwen, P. J. V. (2011). Efficient nonlinear data-assimilation in geophysical fluid dynamics. *Computers & Fluids*, 46(1):52 – 58.

Lie, K., Krogstad, S., Ligaarden, I., Natvig, J., Nilsen, H., and Skaflestad, B. (2010). Discretisation on complex grids–open source MATLAB implementation. In *Proceedings of ECMOR XII–12th European Conference on the Mathematics of Oil Recovery, EAGE, Oxford, UK.*

Mavko, G., Mukerji, T., and Dvorkin, J. (2003). *Rock Physics Handbook - Tools for Seismic Analysis in Porous Media.* Cambridge University Press.

Merwe, R. V. D. and Wan, E. (2001). The square-root unscented Kalman filter for state and parameter estimation. In *International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Salt Lake City, UT*, pages 3461–3464. IEEE.

Nørgaard, M., Poulsen, N. K., and Ravn, O. (2000). New developments in state estimation for nonlinear systems. *Automatica*, 36(11):1627 – 1638.

Pitt, M. K. and Shephard, N. (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association*, pages 590–599.

Sætrom, J. and Omre, H. (2011). Ensemble Kalman filtering with shrinkage regression techniques. *Computational Geosciences*, 15(2):271–292.

Sakov, P. and Oke, P. R. (2008). Implications of the form of the ensemble transformation in the ensemble square root filters. *Monthly Weather Review*, 136:1042–1053.

Stordal, A. S., Karlsen, H. A., Nævdal, G., Skaug, H. J., and Valls, B. (2010). Bridging the ensemble Kalman filter and particle filters: the adaptive gaussian mixture filter. *Computational Geosciences*, 15(2):293–305.

Wan, E. A., Merwe, R. V. D., and Nelson, A. T. (2000). Dual estimation and the unscented transformation. *Advances in Neural Information Processing Systems*, 12:666–672.

Paper II

Kalman Filter Variants in the Closed Skew Normal Setting

J. Rezaie and J. Eidsvik.
In revision

Is not included due to copyright

Paper III

Reducing the Dimensionality of Geophysical Data in Conjunction with Seismic History Matching

J. Rezaie, J. Sætrom and E. Smørgrav
Published in the proceedings of the 74^{th} EAGE Conference & Exhibition
incorporating SPE EUROPEC 2012, Copenhagen, Denmark, SPE153924.

Is not included due to copyright

Paper IV

Value of Information Analysis and Bayesian Inversion for Closed Skew-Normal Distributions: Applications to Seismic Amplitude Versus Offset Data

J. Rezaie, J. Eidsvik and T. Mukerji

In revision

Is not included due to copyright