Maxime Conjard

The Selection Kalman Model - Data assimilation for spatio-temporal variables with multimodal spatial histograms

Maxime Conjard

Doctoral thesis

NTNU
Norwegian University of
Science and Technology

NTNU
Norwegian University of
Science and Technology

NTNU

Maxime Conjard

The Selection Kalman Model - Data assimilation for spatio-temporal variables with multimodal spatial histograms

Thesis for the degree of Philosophiae Doctor

Trondheim, June 2020

NTNU

Norwegian University of
Science and Technology

PRINTED IN
NORWAY
NO - 1598

NORDIC SWAN ECOLABEL

Printed matter
2041 0731

The Selection Kalman Model - Data assimilation for spatio-temporal variables with multimodal spatial histograms

Maxime Conjard

May 17, 2021

Preface

This thesis is submitted in partial fulfillment of the requirements for the degree of Philosophiae Doctor (PhD) at the Norwegian University of Science and Technology (NTNU). The research is funded by the Uncertainty in Reservoir Evaluation (URE) consortium, and carried out at the Department of Mathematical Sciences (IMF). Supervisor on the project was Professor Henning Omre, IMF, NTNU and co-supervisor Kjartan Rimstad, DNB, Oslo.

I would like to thank my supervisor Henning Omre, for his time and patience. I would also like to thank my co-authors: Dario Grana, Olaf Cirpka and Emilio Sanchéz-Léon for their contributions to my thesis and to my understanding of data assimilation.

Finally a few words for the people in my life. To my partner, Aiga, putting up with me must not have been an easy task. To the kids, Margaux and Louis, the daily dose of love, hugs and kisses was much appreciated. To my parents, grandparents and brother, for supporting me throughout. And to my old friends, Guillaume, Pierre-Antoine and Quentin - for the good banter.

<div align="right">

Maxime Conjard
Trondheim, May 17, 2021

</div>

Contents

Introduction

Data assimilation is the process of predicting the state of a quantity of interest using observations collected as,

$$\boldsymbol{d} = g(\boldsymbol{r}) + \boldsymbol{\epsilon}, \tag{1}$$

where the observed data vector $\boldsymbol{d}$ is a function g of an unknown quantity of interest $\boldsymbol{r}$ called state vector and of an observation error denoted $\boldsymbol{\epsilon}$. Bayes' theorem,

$$f(\boldsymbol{r}|\boldsymbol{d}) = \frac{f(\boldsymbol{d}|\boldsymbol{r})f(\boldsymbol{r})}{f(\boldsymbol{d})}, \tag{2}$$

represents what the observed data tell us about the quantity of interest, where $f(\boldsymbol{r})$, $f(\boldsymbol{d}|\boldsymbol{r})$, $f(\boldsymbol{d})$ and $f(\boldsymbol{r}|\boldsymbol{d})$ are the prior distribution, the likelihood model, the marginal likelihood and the posterior distribution, respectively. In the general case, the posterior distribution $f(\boldsymbol{r}|\boldsymbol{d})$ is not analytically tractable and must be approximated. The approximation itself can be very computationally demanding.

In this thesis work, we focus on data assimilation for spatio-temporal phenomena. When cast in a Bayesian setting, we specify an initial distribution, a forward model and a likelihood model which define the posterior distribution. The forward model represents the evolution of the spatio-temporal variables at play while the likelihood model describes the data acquisition procedure. When the initial distribution is Gaussian and the forward and likelihood models are linear with additive Gaussian noise (Gauss-linear), the posterior distribution is Gaussian and analytically tractable. It can be assessed using the traditional Kalman model (Kalman, 1960). In this settting, one usually distinguishes between the filtering and smoothing distributions. The filtering distribution predicts the current state of the quantity of interest while the smoothing distribution predicts previous states of the same quantity of interest. They are assessed with the Kalman Filter (Kalman, 1960) and the Kalman smoother (Rauch et al., 1965), respectively.

When the forward and/or likelihood models are non-linear, analytical tractability is lost and the extended Kalman filter (Jazwinski, 1970) and

1

smoother (Yu et al., 2004) can be used to approximate the posterior distribution. They however require repeated evaluations of sensitivity matrices which can be costly for high-dimensional problems. Ensemble methods such as the ensemble Kalman Filter (EnKF) (Evensen, 1994, 2009) and smoother (EnKS) (Evensen and van Leeuwen, 2000) provide an alternative approach that can be used to assimilate spatio-temporal data and evaluate the posterior distribution of the quantity of interest. The EnKF is a Monte-Carlo implementation of the Bayesian update. Ensemble members are simulated from the initial distribution and serve as inputs to the forward model. The likelihood model is then used to condition the ensemble on the collected data. The EnKF has successfully been applied to numerical weather prediction (Sakov et al., 2018), oceanography (Bertino et al., 2003), reservoir simulation (Aanonsen et al., 2009) and groundwater flow (Hendricks Franssen and Kinzelbach, 2008).

However, owing to the linearized conditioning step, the EnKF is not well suited to represent spatio-temporal variables that display multimodal spatial histograms. A regression towards the mean occurs during the conditioning steps and the posterior marginal distributions drift towards Gaussianity. This is a challenging problem in subsurface modeling: subsurface properties, such as porosity and saturation in reservoir simulation and log-conductivity in groundwater flow, often appear as non-Gaussian owing to the underlying geology. Statistical methods, such as Ensemble Randomized Likelihood (EnRML) (Chen and Oliver, 2012), Gaussian anamorphosis (Bertino et al., 2003; Zhou et al., 2012), Gaussian mixture models (Dovera and Della Rossa, 2011) and truncated pluri-Gaussian (Oliver and Chen, 2018), have been developed to address this issue.

An alternative to Gaussian mixture models is the selection-Gaussian distribution (Arellano-Valle and del Pino, 2004; Arellano-Valle et al., 2006; Omre and Rimstad, 2021). The class of selection-Gaussian distributions can represent multimodaliy, skewness and peakedness and is closed under Gauss-linear operations. In addition, the selection-Gaussian distribution is a conjugate prior to Gauss-linear likelihood models. The posterior distribution is therefore selection-Gaussian and analytically tractable when the forward and likelihood models are Gauss-linear. When they are non-linear, it is possible to use the selection-Gaussian distribution as an initial distribution within the ensemble Kalman filter framework and conserve non-Gaussian features. These attributes make the selection-Gaussian distribution well suited for data assimilation when the quantity of interest is suspected to display a multimodal, skewed or heavy-tailed spatial histogram for linear and non-linear models.

We define the hidden Markov model that serves as a basis for the rest

of the thesis in the first section. In the next section, we focus on linear models and present the traditional Kalman model. In the third section, we introduce the major contribution of this thesis, the selection Kalman model, as an extension of the traditional Kalman model that allows for multimodality in linear models. In the fourth section, non-linear models are considered and ensemble methods such as the EnKF are presented. Methods to address the shortcomings of the EnKF are reviewed in the fifth section. Finally, a summary of the papers constituting this PhD thesis is presented.

In this introduction, $f(\boldsymbol{y})$ denotes the probability density function (pdf) of a random variable $\boldsymbol{y}$, $\varphi_n(\boldsymbol{y}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the pdf of the Gaussian n-vector $\boldsymbol{y}$ with expectation n-vector $\boldsymbol{\mu}$ and covariance $(n \times n)$-matrix $\boldsymbol{\Sigma}$. Furthermore, $\Phi_n(A; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the probability of the aforementioned Gaussian n-vector $\boldsymbol{y}$ to be in $A \subset \mathbb{R}^n$. We also use $\boldsymbol{I}_n$ to denote the identity $(n \times n)$-matrix.

Hidden Markov model

Consider the real valued state n-vector $\boldsymbol{r}_t$ for discretized times $t \in \mathcal{T}_r :$ $\{0, 1, \ldots, T, T+1\}$. Let $\boldsymbol{r}$ denote $\{\boldsymbol{r}_0, \boldsymbol{r}_1, \ldots, \boldsymbol{r}_T, \boldsymbol{r}_{T+1}\}$ and $\boldsymbol{r}_{i:j}$ denote $\{\boldsymbol{r}_i, \boldsymbol{r}_{i+1}, \ldots, \boldsymbol{r}_j\}, \forall (i, j) \in \mathcal{T}_r^2, i \leq j$. Assume that the temporal m-vectors of observations $\boldsymbol{d}_t$ for $t \in \mathcal{T}_d : \{0, 1, \ldots, T\}$ are available, and define $\boldsymbol{d} = \{\boldsymbol{d}_0, \boldsymbol{d}_1, \ldots, \boldsymbol{d}_T\}$ and $\boldsymbol{d}_{i:j} = \{\boldsymbol{d}_i, \ldots, \boldsymbol{d}_j\}$ accordingly. The model specified thereafter defines a hidden Markov (HM) model (Cappé et al., 2005), as displayed in Figure 1, with conditional independence and single state dependence.


Figure 1: Graph of the hidden Markov model

Prior model: The prior model on $\boldsymbol{r}$ consists of an initial distribution and a forward model,

$$f(\boldsymbol{r}) = f(\boldsymbol{r}_0)f(\boldsymbol{r}_{1:T+1}|\boldsymbol{r}_0), \tag{3}$$

where $f(\boldsymbol{r}_0)$ is the pdf of the initial state and $f(\boldsymbol{r}_{1:T+1}|\boldsymbol{r}_0)$ defines the forward model.

The forward model given the initial state $[\boldsymbol{r}_{1:T+1}|\boldsymbol{r}_0]$ is defined as

$$f(\boldsymbol{r}_{1:T+1}|\boldsymbol{r}_0) = \prod_{t=0}^{T} f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t), \tag{4}$$

with

$$[\boldsymbol{r}_{t+1}|\boldsymbol{r}_t] = \omega_t(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^r) \sim f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t), \tag{5}$$

where $\omega_t(\cdot, \cdot) \in \mathbb{R}^n$ is the forward function with random n-vector $\boldsymbol{\epsilon}_t^r$. This forward model only involves the variable at the previous time step $\boldsymbol{r}_t$, hence it defines a first-order Markov chain.

Likelihood model: The likelihood model for $[\boldsymbol{d}|\boldsymbol{r}]$ is defined as conditional independent with single-site response,

$$f(\boldsymbol{d}|\boldsymbol{r}) = \prod_{t=0}^{T} f(\boldsymbol{d}_t|\boldsymbol{r}_t), \tag{6}$$

with

$$[\boldsymbol{d}_t|\boldsymbol{r}_t] = \psi_t(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^d) \sim f(\boldsymbol{d}_t|\boldsymbol{r}_t), \tag{7}$$

where $\psi_t(\cdot, \cdot) \in \mathbb{R}^m$ is the likelihood function with random m-vector $\boldsymbol{\epsilon}_t^d$.

Posterior model: The posterior model for the HM model in Figure 1 is also a Markov chain (Cappé et al., 2005; Moja et al., 2018), and is given by

$$\begin{aligned}
[\boldsymbol{r}|\boldsymbol{d}] \sim f(\boldsymbol{r}|\boldsymbol{d}) =& const \times f(\boldsymbol{d}|\boldsymbol{r})f(\boldsymbol{r}) \\
=& const \times f(\boldsymbol{d}_0|\boldsymbol{r}_0)f(\boldsymbol{r}_0) \prod_{t=1}^{T} f(\boldsymbol{d}_t|\boldsymbol{r}_t)f(\boldsymbol{r}_t|\boldsymbol{r}_{t-1})f(\boldsymbol{r}_{T+1}|\boldsymbol{r}_T) \\
=& f(\boldsymbol{r}_0|\boldsymbol{d}) \prod_{t=1}^{T} f(\boldsymbol{r}_t|\boldsymbol{r}_{t-1}, \boldsymbol{d}_{t:T})f(\boldsymbol{r}_{T+1}|\boldsymbol{r}_T).
\end{aligned} \tag{8}$$

Oftentimes in data assimilation, the quantity of interest does not require the evaluation of the full posterior model. It might be sufficient to evaluate the filtering or the smoothing distributions detailed hereafter.

Filtering distribution: The filtering distribution $f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t})$ can be evaluated recursively using:

$$\begin{aligned}
f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t}) =& \int f(\boldsymbol{r}_t, \boldsymbol{r}_{t-1}|\boldsymbol{d}_{0:t})\mathrm{d}\boldsymbol{r}_{t-1} \\
=& const \times \int f(\boldsymbol{d}_t|\boldsymbol{r}_t, \boldsymbol{r}_{t-1}, \boldsymbol{d}_{0:t-1})f(\boldsymbol{r}_t, \boldsymbol{r}_{t-1}|\boldsymbol{d}_{0:t-1})\mathrm{d}\boldsymbol{r}_{t-1} \\
=& const \times \int f(\boldsymbol{d}_t|\boldsymbol{r}_t)f(\boldsymbol{r}_t|\boldsymbol{r}_{t-1})f(\boldsymbol{r}_{t-1}|\boldsymbol{d}_{0:t-1})\mathrm{d}\boldsymbol{r}_{t-1}.
\end{aligned} \tag{9}$$

Smoothing distribution: The smoothing distribution is denoted $f(\boldsymbol{r}_s|\boldsymbol{d}_{0:t})$, $t > s$. Consider,

$$f(\boldsymbol{r}_s|\boldsymbol{d}_{0:t}) = \int f(\boldsymbol{r}_s, \boldsymbol{r}_{s+1}|\boldsymbol{d}_{0:t})\mathrm{d}\boldsymbol{r}_{s+1} \tag{10}$$

$$= const \times \int f(\boldsymbol{r}_s|\boldsymbol{r}_{s+1}, \boldsymbol{d}_{0:t})f(\boldsymbol{r}_{s+1}|\boldsymbol{d}_{0:t})\mathrm{d}\boldsymbol{r}_{s+1}.$$

Since

$$f(\boldsymbol{r}_s|\boldsymbol{r}_{s+1}, \boldsymbol{d}_{0:t}) = \frac{f(\boldsymbol{d}_{s+1:t}|\boldsymbol{r}_s, \boldsymbol{r}_{s+1}, \boldsymbol{d}_{0:s})}{f(\boldsymbol{d}_{s+1:t}|\boldsymbol{r}_{s+1}, \boldsymbol{d}_{0:s})}f(\boldsymbol{r}_s|\boldsymbol{r}_{s+1}, \boldsymbol{d}_{0:s}) \tag{11}$$

$$= f(\boldsymbol{r}_s|\boldsymbol{r}_{s+1}, \boldsymbol{d}_{0:s}),$$

Equation 10 therefore gives,

$$f(\boldsymbol{r}_s|\boldsymbol{d}_{0:t}) = const \times \int f(\boldsymbol{r}_s|\boldsymbol{r}_{s+1}, \boldsymbol{d}_{0:s})f(\boldsymbol{r}_{s+1}|\boldsymbol{d}_{0:t})\mathrm{d}\boldsymbol{r}_{s+1} \tag{12}$$

$$= const \times \int \frac{f(\boldsymbol{r}_{s+1}|\boldsymbol{r}_s)f(\boldsymbol{r}_s|\boldsymbol{d}_{0:s})}{f(\boldsymbol{r}_{s+1}|\boldsymbol{d}_{0:s})}f(\boldsymbol{r}_{s+1}|\boldsymbol{d}_{0:t})\mathrm{d}\boldsymbol{r}_{s+1},$$

the smoothing distribution can thus be recursively evaluated.

The HM model defined in this section gives a very general framework for the evaluation of the posterior distribution of the quantity of interest. Its analytical tractability depends however on the assumptions that are made on the prior and likelihood models. The different models presented hereafter all use the HM model defined in this section.

Traditional Kalman model

The traditional Kalman model (Kalman, 1960) is defined assuming that the initial state is distributed according to a Gaussian distribution,

$$f(\boldsymbol{r}_0) = \varphi_n(\boldsymbol{r}_0; \boldsymbol{\mu}_{0|.}^r, \boldsymbol{\Sigma}_{0|.}^r), \tag{13}$$

with parametrization $\Theta^G = (\boldsymbol{\mu}_{0|.}^r, \boldsymbol{\Sigma}_{0|.}^r)$. It also assumes that the forward and likelihood models are Gauss-linear:

$$[\boldsymbol{r}_{t+1}|\boldsymbol{r}_t] = \omega_t(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^r) = \boldsymbol{A}_t\boldsymbol{r}_t + \boldsymbol{\epsilon}_t^r \rightarrow \varphi_n(\boldsymbol{r}; \boldsymbol{A}_t\boldsymbol{r}_t, \boldsymbol{\Sigma}_t^r) \tag{14}$$

$$[\boldsymbol{d}_t|\boldsymbol{r}_t] = \psi_t(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^d) = \boldsymbol{H}_t\boldsymbol{r}_t + \boldsymbol{\epsilon}_t^d \rightarrow \varphi_m(\boldsymbol{d}; \boldsymbol{H}_t\boldsymbol{r}_t, \boldsymbol{\Sigma}_t^d), \tag{15}$$

where $\boldsymbol{A}_t$ is a $(n \times n)$-matrix and $\boldsymbol{H}_t$ is a $(m \times n)$-matrix, $\boldsymbol{\epsilon}_t^r$ is a centered Gaussian n-vector with covariance $(n \times n)$-matrix $\boldsymbol{\Sigma}_t^r$ and $\boldsymbol{\epsilon}_t^d$ a centered Gaussian m-vector with covariance $(m \times m)$-matrix $\boldsymbol{\Sigma}_t^d$. Both $\boldsymbol{\epsilon}_t^r$ and $\boldsymbol{\epsilon}_t^d$ are assumed to be independent of $\boldsymbol{r}_t$. Under these assumptions, the posterior model in Equation 8 is Gaussian and analytically tractable.

Kalman filter

Since the prior distribution for the initial state is Gaussian, the filtering recursion given in Equation 9 can be evaluated using the Kalman Filter (Kalman, 1960) with the following recursion formulas:

$$f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t-1}) = \varphi_n(\boldsymbol{r}_t; \boldsymbol{\mu}_{t|t-1}^r, \boldsymbol{\Sigma}_{t|t-1}^r). \tag{16}$$

Conditioning:

$$f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t}) = \varphi_n(\boldsymbol{r}_t; \boldsymbol{\mu}_{t|t}^r, \boldsymbol{\Sigma}_{t|t}^r), \tag{17}$$

with,

$$\boldsymbol{\mu}_{t|t}^r = \boldsymbol{\mu}_{t|t-1}^r + \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T (\boldsymbol{H}_t \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T + \boldsymbol{\Sigma}_t^d)^{-1} (\boldsymbol{d}_t - \boldsymbol{H} \boldsymbol{\mu}_{t|t-1}^r) \tag{18}$$
$$\boldsymbol{\Sigma}_{t|t}^r = \boldsymbol{\Sigma}_{t|t-1}^r - \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T (\boldsymbol{H}_t \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T + \boldsymbol{\Sigma}_t^d)^{-1} \boldsymbol{H}_t \boldsymbol{\Sigma}_{t|t-1}^r.$$

Forwarding:
$$f(\boldsymbol{r}_{t+1}|\boldsymbol{d}_{0:t}) = \varphi_n(\boldsymbol{r}_{t+1}; \boldsymbol{\mu}_{t+1|t}^r, \boldsymbol{\Sigma}_{t+1|t}^r), \tag{19}$$
with,

$$\boldsymbol{\mu}_{t+1|t}^r = \boldsymbol{A}_t \boldsymbol{\mu}_{t|t}^r \tag{20}$$
$$\boldsymbol{\Sigma}_{t+1|t}^r = \boldsymbol{A}_t \boldsymbol{\Sigma}_{t|t}^r \boldsymbol{A}_t^T + \boldsymbol{\Sigma}_t^r.$$

The $(n \times m)$-matrix $\boldsymbol{K}_t = \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T (\boldsymbol{H}_t \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T + \boldsymbol{\Sigma}_t^d)^{-1}$ is denoted the Kalman gain. Equations 16–20 can be used recursively to evaluate the parameters of the Gaussian posterior distribution $f(\boldsymbol{r}_T|\boldsymbol{d}_{0:T}) = \varphi_n(\boldsymbol{r}; \boldsymbol{\mu}_{T|T}^r, \boldsymbol{\Sigma}_{T|T}^r)$.

Kalman smoother

The smoothing distribution in Equation 12 can be evaluated with the Kalman smoother (Rauch et al., 1965). Since the prior distribution for

the initial state is Gaussian, the joint distribution of $[\boldsymbol{r}_0, \boldsymbol{d}_0]$ is jointly Gaussian,

$$f(\boldsymbol{r}_0, \boldsymbol{d}_0) = \varphi_{n+m} \left(\begin{bmatrix} \boldsymbol{r}_0 \\ \boldsymbol{d}_0 \end{bmatrix} ; \begin{bmatrix} \boldsymbol{\mu}_{0|.}^r \\ \boldsymbol{H}_0 \boldsymbol{\mu}_{0|.}^r \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{0|.}^r & \boldsymbol{\Sigma}_{0|.}^r \boldsymbol{H}_0^T \\ \boldsymbol{H}_0 \boldsymbol{\Sigma}_{0|.}^r & \boldsymbol{H}_0 \boldsymbol{\Sigma}_{0|.}^r \boldsymbol{H}_0^T + \boldsymbol{\Sigma}_0^d \end{bmatrix} \right) . \tag{21}$$

Assume $[\boldsymbol{r}_{0:t}, \boldsymbol{d}_{0:t}]$ is jointly Gaussian. Consider $[\boldsymbol{r}_{0:t+1}, \boldsymbol{d}_{0:t}]$, then

$$f(\boldsymbol{r}_{0:t+1}, \boldsymbol{d}_{0:t}) = f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t) f(\boldsymbol{r}_{0:t}, \boldsymbol{d}_{0:t}), \tag{22}$$

is the product of a Gauss-linear forward model and a Gaussian distribution which shows that $[\boldsymbol{r}_{0:t+1}, \boldsymbol{d}_{0:t}]$ is jointly Gaussian. Further, consider $[\boldsymbol{r}_{0:t+1}, \boldsymbol{d}_{0:t+1}]$, then

$$f(\boldsymbol{r}_{0:t+1}, \boldsymbol{d}_{0:t+1}) = f(\boldsymbol{d}_{t+1}|\boldsymbol{r}_{t+1}) f(\boldsymbol{r}_{0:t+1}, \boldsymbol{d}_{0:t}), \tag{23}$$

is the product of a Gauss-linear likelihood model and a Gaussian distribution which shows that $[\boldsymbol{r}_{0:t+1}, \boldsymbol{d}_{0:t+1}]$ is jointly Gaussian. By recursion, $[\boldsymbol{r}_{0:t}, \boldsymbol{d}_{0:t}]$ is jointly Gaussian for all t. By assembling sequentially the mean and covariance matrix of $[\boldsymbol{r}_{0:T}, \boldsymbol{d}_{0:T}]$ during the forward run, the smoothing distribution $f(\boldsymbol{r}_s|\boldsymbol{d}_{0:T})$, $s < T$, can be evaluated by marginalizing out $[\boldsymbol{r}_{0:s-1}, \boldsymbol{r}_{s+1:T}]$, and conditioning on $\boldsymbol{d}_{0:T}$. Conceptually, this approach is the simplest way to define Kalman smoothing but it might have a prohibitive storage cost for high-dimensional problems. The Rauch–Tung–Striebel smoother (Rauch et al., 1965) can achieve more efficiently the same outcome with a forward and a backward run using Equation 12.

Kalman filter extensions

The extended Kalman filter (Jazwinski, 1970) is an extension of the Kalman filter that allows for non-linearity in the forward and likelihood models with a first-order Taylor expansion. Assume that the forward model is non linear with additive Gaussian noise,

$$\boldsymbol{r}_{t+1} = \omega_t(\boldsymbol{r}_t) + \boldsymbol{\epsilon}_t^r, \tag{24}$$

the forward model ω_t is then linearized such that:

$$\omega_t(r_t) = \omega_t(\boldsymbol{\mu}_{t|t}^r) + \boldsymbol{\Omega}^T(\boldsymbol{r}_t - \boldsymbol{\mu}_{t|t}^r) + o(\boldsymbol{r}_t - \boldsymbol{\mu}_{t|t}^r), \tag{25}$$

where $\boldsymbol{\Omega} = \nabla \omega_t(\boldsymbol{\mu}_{t|t}^r)$ thereby approximating,

$$\boldsymbol{\mu}_{t+1|t}^r \approx \omega_t(\boldsymbol{\mu}_{t|t}^r) \tag{26}$$
$$\boldsymbol{\Sigma}_{t+1|t}^r \approx \boldsymbol{\Omega} \boldsymbol{\Sigma}_{t|t}^r \boldsymbol{\Omega}^T + \boldsymbol{\Sigma}_t^r.$$

However evaluating $\boldsymbol{\Omega}$ at every time steps can prove costly for high-dimensional problems. In addition, if the forward model is more than weakly non-linear for the considered time steps, the extended Kalman filter is likely to behave poorly.

The unscented Kalman filter (Julier and Uhlmann, 1997) offers more robust performances by considering weighted deterministic particles $\boldsymbol{\chi}_t^i, i = 0, \ldots, 2n$, distributed along a covariance contour. The weights are denoted $\mathrm{w}_t^i, i = 0, \ldots, 2n$ and are assigned to their respective particles such that,

$$\boldsymbol{\chi}_t^0 = \boldsymbol{\mu}_{t|t}^r, \quad \mathrm{w}_t^0 = \kappa/(n+\kappa) \tag{27}$$
$$\boldsymbol{\chi}_t^i = \boldsymbol{\mu}_{t|t}^r + \left((\kappa+n)^{1/2} [\boldsymbol{\Sigma}_{t|t}^r]^{1/2} \right)_i, \quad \mathrm{w}_t^i = 1/2(n+\kappa)$$
$$\boldsymbol{\chi}_t^{i+n} = \boldsymbol{\mu}_{t|t}^r - \left((\kappa+n)^{1/2} [\boldsymbol{\Sigma}_{t|t}^r]^{1/2} \right)_i, \quad \mathrm{w}_t^{i+n} = 1/2(n+\kappa),$$

where $\kappa \in \mathbb{R}$, $i = 1, \ldots, n$ and $\left((\kappa+n)^{1/2} [\boldsymbol{\Sigma}_{t|t}^r]^{1/2} \right)_i$ is the i-th row of the square root matrix of $(\kappa+n)\boldsymbol{\Sigma}_{t|t}^r$

The particles are then propagated forward in time as:

$$\boldsymbol{\chi}_{t+1}^i = \omega_t(\boldsymbol{\chi}_t^i) \tag{28}$$
$$\overline{\boldsymbol{\chi}}_{t+1} = \sum_0^{2n} \mathrm{w}_t^i \boldsymbol{\chi}_{t+1}^i.$$

Then instead of approximating $\boldsymbol{\Omega}$ to estimate $\boldsymbol{\Omega} \boldsymbol{\Sigma}_{t|t}^r \boldsymbol{\Omega}^T$ as in the extended Kalman filter, we consider the $(n \times n)$-matrix $\boldsymbol{\Sigma}_{\chi,t+1}$ which is defined as:

$$\boldsymbol{\Sigma}_{\chi,t+1} = \sum_0^{2n} \mathrm{w}_t^i (\boldsymbol{\chi}_{t+1}^i - \overline{\boldsymbol{\chi}}_{t+1})(\boldsymbol{\chi}_{t+1}^i - \overline{\boldsymbol{\chi}}_{t+1})^T, \tag{29}$$

thereby giving,

$$\boldsymbol{\mu}_{t+1|t}^r \approx \overline{\boldsymbol{\chi}}_{t+1} \tag{30}$$
$$\boldsymbol{\Sigma}_{t+1|t}^r \approx \boldsymbol{\Sigma}_{\chi,t+1} + \boldsymbol{\Sigma}_t^r.$$

The unscented Kalman Filter has a comparable cost to the extended Kalman filter but better estimates the non-linearity in the forward model (Julier and Uhlmann, 1997).

Note that if the likelihood model is non-linear, similar derivations can be presented for both the extended Kalman filter and the unscented Kalman filter. An extended Kalman smoother and an unscented Kalman smoother can also be derived (Yu et al., 2004; Särkkä, 2008).

Selection Kalman model

In the traditional Kalman model, the prior distribution for the initial state is assumed to be Gaussian. Under this assumption, the posterior distribution is Gaussian and therefore cannot represent non-Gaussian features in the quantity of interest. The selection Kalman model defined hereafter provides a solution to the issue. In the selection Kalman model, the distribution for the initial state $f(\boldsymbol{r}_0)$ is assumed to be in the class of selection-Gaussian distributions (Arellano-Valle et al., 2006; Omre and Rimstad, 2021). Consider an auxiliary Gaussian $(n+n)$-vector $[\tilde{\boldsymbol{r}}_0, \boldsymbol{\nu}]$,

$$
\begin{bmatrix} \tilde{\boldsymbol{r}}_0 \\ \boldsymbol{\nu} \end{bmatrix} \sim \varphi_{2n} \left(\begin{bmatrix} \tilde{\boldsymbol{r}}_0 \\ \boldsymbol{\nu} \end{bmatrix}; \begin{bmatrix} \boldsymbol{\mu}_{0|\cdot}^{\tilde{r}} \\ \boldsymbol{\mu}_{\nu} \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}} & \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}} \boldsymbol{\Gamma}_{\nu|\tilde{r}}^T \\ \boldsymbol{\Gamma}_{\nu|\tilde{r}} \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}} & \boldsymbol{\Sigma}_{\nu} \end{bmatrix} \right), \qquad (31)
$$

with expectation n-vectors $\boldsymbol{\mu}_{0|\cdot}^{\tilde{r}}$ and $\boldsymbol{\mu}_{\nu}$, correlation $(n \times n)$-matrix $\boldsymbol{\Gamma}_{\nu|\tilde{r}}$, and where $\boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}}$, $\boldsymbol{\Sigma}_{\nu}$, and $\boldsymbol{\Sigma}_{\nu|\tilde{r}}$ are covariance $(n \times n)$-matrices with $\boldsymbol{\Sigma}_{\nu} = \boldsymbol{\Gamma}_{\nu|\tilde{r}} \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}} \boldsymbol{\Gamma}_{\nu|\tilde{r}}^T + \boldsymbol{\Sigma}_{\nu|\tilde{r}}$. Define a selection set $A \subset \mathbb{R}^n$ of dimension n and let $\boldsymbol{r}_0 = [\tilde{\boldsymbol{r}}_0|\boldsymbol{\nu} \in A]$; then, $\boldsymbol{r}_0$ is in the class of selection-Gaussian distributions and its pdf is,

$$
\begin{aligned}
f(\boldsymbol{r}_0) = &\left[\Phi_n(A; \boldsymbol{\mu}_{\nu}, \boldsymbol{\Sigma}_{\nu}) \right]^{-1} \\
&\times \Phi_n(A; \boldsymbol{\mu}_{\nu} + \boldsymbol{\Gamma}_{\nu|\tilde{r}}(\boldsymbol{r}_0 - \boldsymbol{\mu}_{0|\cdot}^{\tilde{r}}), \boldsymbol{\Sigma}_{\nu|\tilde{r}}) \times \varphi_n(\boldsymbol{r}_0; \boldsymbol{\mu}_{0|\cdot}^{\tilde{r}}, \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}}),
\end{aligned} \qquad (32)
$$

and is parametrized with $\Theta^{SG} = (\boldsymbol{\mu}_{0|\cdot}^{\tilde{r}}, \boldsymbol{\mu}_{\nu}, \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}}, \boldsymbol{\Sigma}_{\nu|\tilde{r}}, \boldsymbol{\Gamma}_{\nu|\tilde{r}}, A)$. Note that the class of Gaussian distributions constitutes a subset of the class of selection-Gaussian distributions with $\boldsymbol{\Gamma}_{\nu|\tilde{r}} = 0 \times \boldsymbol{I}_n$. The dependence in $[\tilde{\boldsymbol{r}}, \boldsymbol{\nu}]$ represented by $\boldsymbol{\Gamma}_{\nu|\tilde{r}}$ and the selection subset A are crucial user-defined parameters. The selection-Gaussian model may represent multi-modal, skewed, and/or peaked marginal distributions (Omre and Rimstad, 2021) as shown in Figure 2. As the traditional Kalman model, the selection Kalman model also assumes the forward and likelihood models to be Gauss-linear, see Equation 14. Under these assumptions, the posterior model in Equation 8 is selection-Gaussian and analytically tractable (Omre and Rimstad, 2021) for two reasons. First, the class of selection-Gaussian distributions is closed under Gauss-linear transformations and second, the selection-Gaussian distribution is a conjugate
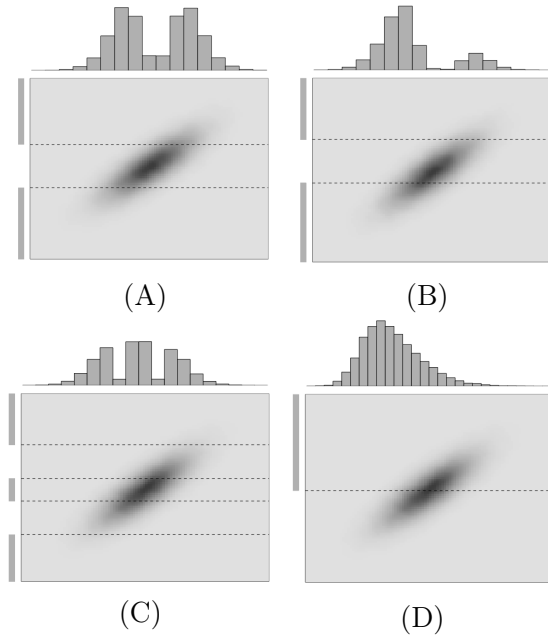
Figure 2: Realizations of 1D selection-Gaussian pdfs (histogram) with varying selection sets $A \subset \mathbb{R}^n$ (solid gray bars) for a bi-Gaussian pdf $[\tilde{r}, \nu]$ (dark gray)

prior to Gauss-linear likelihood models. The selection Kalman model can therefore represent non-Gaussian features in the posterior distribution of the quantity of interest.

Selection Kalman Filter

The selection Kalman filter evaluates the filtering distribution $f(\tilde{\boldsymbol{r}}_T | \boldsymbol{\nu} \in A, \boldsymbol{d}_{0:T})$. The posterior distribution given by the selection Kalman filter is calculated by first defining a Gaussian augmented state $2n$-vector $\boldsymbol{u}_t = [\tilde{\boldsymbol{r}}_t, \boldsymbol{\nu}]$. The forward model in Equation 14 is then adjusted to account for the augmented state vector,

$$[\boldsymbol{u}_{t+1} | \boldsymbol{u}_t] = \begin{bmatrix} \boldsymbol{A}_t & 0 \\ 0 & \boldsymbol{I}_n \end{bmatrix} \boldsymbol{u}_t + \begin{bmatrix} \boldsymbol{\epsilon}_t^r \\ 0 \end{bmatrix}, \tag{33}$$

and so is the likelihood model in Equation 15,

$$[\boldsymbol{d}_t | \boldsymbol{u}_t] = \begin{bmatrix} H & 0 \\ 0 & 0 \end{bmatrix} \boldsymbol{u}_t + \begin{bmatrix} \boldsymbol{\epsilon}_t^d \\ 0 \end{bmatrix}. \tag{34}$$

These forward and likelihood models are Gauss-linear, and the recursion for the augmented state vector is,

$$f(\boldsymbol{u}_{t+1}|\boldsymbol{d}_{0:t}) = \int f(\boldsymbol{u}_{t+1}|\boldsymbol{u}_t)[f(\boldsymbol{d}_t|\boldsymbol{d}_{0:t-1})]^{-1}f(\boldsymbol{d}_t|\boldsymbol{u}_t)f(\boldsymbol{u}_t|\boldsymbol{d}_{0:t-1})\mathrm{d}\boldsymbol{u}_t.$$
(35)

This recursion is identical to Equation 9, and hence it can be evaluated using the Kalman filter recursions in Equations 16–20. The filtering distribution,

$$f(\boldsymbol{u}_T|\boldsymbol{d}_{0:T}) = \varphi_{2n}(\boldsymbol{u}; \boldsymbol{\mu}^u_{T|T}, \boldsymbol{\Sigma}^u_{T|T}),$$
(36)

yielded by the Kalman filter recursions is a Gaussian $2n$-vector with mean $2n$-vector $\boldsymbol{\mu}^u_{T|T}$ and covariance $(2n \times 2n)$-matrix $\boldsymbol{\Sigma}^u_{T|T}$ representing $f(\tilde{\boldsymbol{r}}_T, \boldsymbol{\nu}|\boldsymbol{d}_{0:T})$. The posterior distribution $f(\tilde{\boldsymbol{r}}_T|\boldsymbol{\nu} \in A, \boldsymbol{d}_{0:T})$ is then assessed using Markov chain Monte Carlo (McMC) simulation (Omre and Rimstad, 2021).

Selection Kalman Smoother

The selection Kalman smoother aims to evaluate $[\boldsymbol{r}_s|\boldsymbol{\nu} \in A, \boldsymbol{d}_{0:T}]$, $s < T$. As in the selection Kalman filter, we consider the augmented forward and likelihood models presented in Equations 33 and 34, respectively. Similarly to the Kalman smoother, one can show that $[\boldsymbol{r}_{0:t}, \boldsymbol{\nu}, \boldsymbol{d}_{0:t}]$ is jointly Gaussian for all t. The smoothing distribution $f(\boldsymbol{r}_s|\boldsymbol{\nu} \in A, \boldsymbol{d}_{0:T})$, $s < T$, is obtained by marginalizing out $[\boldsymbol{r}_{0:s-1}, \boldsymbol{r}_{s+1:T}]$, and conditioning on $\boldsymbol{d}_{0:T}$ and then on $\boldsymbol{\nu} \in A$ using McMC simulations as in the selection Kalman filter.

Non-linear models: Ensemble methods

When the forward and/or likelihood models are nonlinear, Equations 5 and 7 cannot be written in closed form. The analytical tractability of the posterior model in Equation 8 is lost and the traditional Kalman model cannot be used. When that is the case, ensemble methods are a popular alternative to the extended Kalman filter in the field of data assimilation. A collection of realizations, called an ensemble, are generated from the initial distribution. The distribution from which the initial ensemble is simulated is not limited to Gaussian distributions, it is therefore possible to represent non-Gaussian features in the prior. Each data assimilation cycle consists of two steps, conditioning and forwarding, represented by

the following recursions,

$$f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t}) = [f(\boldsymbol{d}_t|\boldsymbol{d}_{0:t-1})]^{-1} f(\boldsymbol{d}_t|\boldsymbol{r}_t) f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t-1}) \tag{37}$$

$$f(\boldsymbol{r}_{t+1}|\boldsymbol{d}_{0:t}) = \int f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t) f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t}) \mathrm{d}\boldsymbol{r}_t,$$

that in general are not analytically tractable. The ensemble members are propagated from one time step to the next by the forward model. When conditioning on the data, ensemble methods circumvent the need to evaluate the sensitivity matrix $\boldsymbol{\Omega}$ at each assimilation, contrary to the extended Kalman filter. The added cost of having to run the forward model for each ensemble member is mitigated by the fact that these forward runs can be easily parallelized.

Particle Filter

Particle filters (Gordon et al., 1993; Doucet et al., 2001) are a class of Monte Carlo methods for filtering that represent the filtering distribution as a weighted ensemble. The particle filter starts with an initial ensemble of size n_p, where all ensemble members $\boldsymbol{r}_0^i, i = 1, \ldots, n_p$, are generated from the initial distribution $f(\boldsymbol{r}_0)$ and assigned the same weight $w_0^i = 1/n_p, i = 1, \ldots, n_p$. During the conditioning step, the particle filter updates the weights according to:

$$w_t^i = \frac{w_{t-1}^i f(\boldsymbol{d}_t|\boldsymbol{r}_t^i)}{\sum_{k=1}^{n_p} w_{t-1}^k f(\boldsymbol{d}_t|\boldsymbol{r}_t^k)}, i = 1, \ldots, n_p. \tag{38}$$

During the forwarding step, the particles are propagated in time according to,

$$\boldsymbol{r}_{t+1}^i = \omega_t(\boldsymbol{r}_t^i, \boldsymbol{\epsilon}_t^r), i = 1, \ldots, n_p. \tag{39}$$

Asymptotically, the posterior ensemble $\{\boldsymbol{r}_T^i, w_T^i, i = 1, \ldots, n_p\}$ weakly represents the filtering distribution in the sense that,

$$\mathbb{E}(g(\boldsymbol{r}_T)) = \int g(\boldsymbol{r}_T) f(\boldsymbol{r}_T|\boldsymbol{d}_{0:T}) d\boldsymbol{r}_T \tag{40}$$

$$= \lim_{n_p \to \infty} \sum_{i=1}^{n_p} \int g(\boldsymbol{r}_T) w_t^i \delta(\boldsymbol{r}_T^i - \boldsymbol{r}_T) d\boldsymbol{r}_T$$

$$= \sum_{i=1}^{\infty} w_t^i g(\boldsymbol{r}_T^i),$$

where δ denotes the Dirac pdf. The use of particle filter is limited by one major issue : as data is assimilated, the weight is distributed among fewer

and fewer particles. This eventually leads to almost all the weight being concentrated on only one particle. This collapse is even more acute in high-dimensional problems. It possible to somewhat mitigate this effect by reducing the variance of the weights using resampling (Doucet et al., 2001; van Leeuwen, 2009).

Ensemble Kalman Filter

The ensemble Kalman filter (EnKF) (Evensen, 1994) is another Monte-Carlo method that bears some similarity with the particle filter. The EnKF also generates an ensemble from the initial distribution $f(\boldsymbol{r}_0)$ denoted $\boldsymbol{e}_0 = \{\boldsymbol{r}_0^{u(i)}, i = 1, \ldots, n_e\}$, where n_e is the ensemble size. Note that no weights are assigned to the ensemble members. At time t, the ensemble $\boldsymbol{e}_t^r = \{\boldsymbol{r}_t^{u(i)}, i = 1, \ldots, n_e\}$ is assumed to represent $f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t-1})$. The conditioning step is initiated by generating pseudo observations using the likelihood model $\boldsymbol{d}_t^{(i)} = \psi_t(\boldsymbol{r}_t^{u(i)}, \boldsymbol{\epsilon}_t^d), i = 1, \ldots, n_e$, and defining $\boldsymbol{e}_t = \{(\boldsymbol{r}_t^{u(i)}, \boldsymbol{d}_t^{(i)}), i = 1, \ldots, n_e\}$. The conditioning step is given by,

$$\text{Assess } \boldsymbol{\Sigma}_{rd} \text{ from } \boldsymbol{e}_t \rightarrow \hat{\boldsymbol{\Sigma}}_{rd} \rightarrow \hat{\boldsymbol{K}}_t = \hat{\boldsymbol{\Gamma}}_{rd}[\hat{\boldsymbol{\Sigma}}_d]^{-1} \tag{41}$$

$$\boldsymbol{r}_t^{c(i)} = \boldsymbol{r}_t^{u(i)} + \hat{\boldsymbol{K}}_t(\boldsymbol{d}_t - \boldsymbol{d}_t^{(i)}), i = 1, \ldots, n_e.$$

The ensemble $\{\boldsymbol{r}_t^{c(i)}, i = 1, \ldots, n_e\}$ is assumed to represent $f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t})$. The Kalman gain $\boldsymbol{K}_t$ is inferred by calculating the sample covariance matrix of the ensemble $\boldsymbol{e}_t$. The forwarding step is given by,

$$\boldsymbol{r}_{t+1}^{u(i)} = \omega_t(\boldsymbol{r}_t^{c(i)}, \boldsymbol{\epsilon}_t^r), i = 1, \ldots, n_e, \tag{42}$$

and the ensemble $\boldsymbol{e}_{t+1}^r = \{\boldsymbol{r}_{t+1}^{u(i)}, i = 1, \ldots, n_e\}$ represents $f(\boldsymbol{r}_{t+1}|\boldsymbol{d}_{0:t})$. Note that if the initial distribution is Gaussian and the forward and likelihood models are Gauss-linear, the posterior distribution given by the EnKF is asymptotically correct as $n_e \rightarrow \infty$. The EnKF is also subject to ensemble collapse (Sætrom and Omre, 2013) in part because of the coupling between ensemble members introduced by the conditioning in which the same estimated Kalman gain is used to update all the ensemble members.

Even though the analytical properties of the particle filter may appear superior to that of the EnKF, the latter is preferable to the former in high dimensional models (Li et al., 2016; Katzfuss et al., 2020).

Ensemble Kalman smoother

The ensemble Kalman smoother (EnKS) (Evensen and van Leeuwen, 2000) evaluates the smoothing distribution $f(\boldsymbol{r}_s|\boldsymbol{d}_{0:T})$, $s < T$. Consider

the interpolation pdf $f(\boldsymbol{r}_{0:T}|\boldsymbol{d}_{0:T})$ with corresponding HM model recursions (Cappé et al., 2005). The recursion is initiated by

$$[\boldsymbol{r}_0|\boldsymbol{d}_0] \sim f(\boldsymbol{r}_0|\boldsymbol{d}_0) = [f(\boldsymbol{d}_0)]^{-1}f(\boldsymbol{d}_0|\boldsymbol{r}_0)f(\boldsymbol{r}_0). \tag{43}$$

Then, for $t = 1, \ldots, T$,

$$\begin{aligned}[\boldsymbol{r}_{0:t}|\boldsymbol{d}_{0:t}] \sim & f(\boldsymbol{r}_{0:t}|\boldsymbol{d}_{0:t}) \\ = & [f(\boldsymbol{d}_t|\boldsymbol{d}_{0:t-1})]^{-1}f(\boldsymbol{d}_t|\boldsymbol{r}_t)f(\boldsymbol{r}_t|\boldsymbol{r}_{t-1})f(\boldsymbol{r}_{0:t-1}|\boldsymbol{d}_{0:t-1}).\end{aligned} \tag{44}$$

The EnKS is implemented in practice in a similar manner to the EnKF, with a conditioning and a forwarding step. The main difference is that at a given time t, the state vector includes the ensemble representation of all the states up to time t. At time t, the ensemble $\boldsymbol{e}_t^r = \{\boldsymbol{r}_{0:t}^{u(i)} = (\boldsymbol{r}_0^{u(i)}, \ldots, \boldsymbol{r}_t^{u(i)}), i = 1, \ldots, n_e\}$ represents $f(\boldsymbol{r}_{0:t}|\boldsymbol{d}_{0:t-1})$. As in the EnKF, the conditioning step is initiated by generating pseudo observations using the likelihood model $\boldsymbol{d}_t^{(i)} = \psi_t(\boldsymbol{r}_t^{u(i)}, \boldsymbol{\epsilon}_t^d), i = 1, \ldots, n_e$, and defining $\boldsymbol{e}_t = \{(\boldsymbol{r}_{0:t}^{u(i)}, \boldsymbol{d}_t^{(i)}), i = 1, \ldots, n_e\}$. The conditioning step is then given by

$$\boldsymbol{r}_{0:t}^{c(i)} = \boldsymbol{r}_{0:t}^{u(i)} + \hat{\boldsymbol{\Gamma}}_{r_{0:t}d_t}\hat{\boldsymbol{\Sigma}}_d^{-1}(\boldsymbol{d}_t - \boldsymbol{d}_t^i), i = 1, \ldots, n_e, \tag{45}$$

where then ensemble $\{\boldsymbol{r}_{0:t}^{c(i)}, i = 1, \ldots, n_e\}$ represents $f(\boldsymbol{r}_{0:t}|\boldsymbol{d}_{0:t})$ and where $\hat{\boldsymbol{\Gamma}}_{r_{0:t}d_t}$ and $\hat{\boldsymbol{\Sigma}}_d$ are inferred by calculating the sample covariance matrix of the ensemble $\boldsymbol{e}_t$. The forward step is given by,

$$\boldsymbol{r}_{0:t+1}^{u(i)} = [\boldsymbol{r}_{0:t}^{c(i)}, \omega_t(\boldsymbol{r}_t^{c(i)}, \boldsymbol{\epsilon}_t^r)], i = 1, \ldots, n_e, \tag{46}$$

where then ensemble $\boldsymbol{e}_{t+1}^r = \{\boldsymbol{r}_{0:t+1}^{u(i)}, i = 1, \ldots, n_e\}$ represents $f(\boldsymbol{r}_{0:t+1}|\boldsymbol{d}_{0:t})$. Note that the size of the state-space vector increases at every time step. The target smoothing distribution $f(\boldsymbol{r}_s|\boldsymbol{d}_{0:T})$, $s < T$, can be assessed by extracting the relevant ensemble from $\{\boldsymbol{r}_{0:T}^{c(i)}, i = 1, \ldots, n_e\}$.

Improving the EnKF

In this section we propose a solution to four of the major shortcomings of the EnKF:

1. *Localization:* Against estimation uncertainty in the covariance matrix estimate

2. *Inflation:* Against systematic underestimation of the uncertainty in the posterior ensemble

3. *Iterative methods:* Improve data match when the non-linearity is too dominant

4. *Selection Ensemble Kalman Filter:* Mitigates drift towards Gaussianity

Localization

The conditioning step in the EnKF is based on the estimation of $\boldsymbol{\Sigma}_{rd}$ from which we calculate the estimated Kalman gain $\hat{\boldsymbol{K}}_t = \hat{\boldsymbol{\Gamma}}_{rd}[\hat{\boldsymbol{\Sigma}}_d]^{-1}$. The ensemble size is commonly of the order of 10^2, while the dimension of the state vector is in many applications of the order of $10^3 - 10^5$. Estimating $\boldsymbol{K}_t$ therefore leads to estimation uncertainty, whereby causing unwarranted reduction in variance away from where the data is collected. When the dimension of the data is small, for instance oil production rate measured at a few wells, the estimation uncertainty will mostly affect the estimate of $\boldsymbol{\Gamma}_{rd}$. However, when the dimension of the data is large, for instance when assimilating seismic amplitudes, the estimation uncertainty will affect the estimates of both $\boldsymbol{\Gamma}_{rd}$ and $\boldsymbol{\Sigma}_d$. Either way, there are not enough ensemble members to accurately estimate $\boldsymbol{K}_t$. To remedy the issue, covariance localization is used to ensure that physically distant points are uncorrelated. To that end, we define a damping $(n + m) \times (n + m)$-matrix $\boldsymbol{\Theta}$ that we multiply element-wise with the estimate of $\boldsymbol{\Sigma}_{rd}$ from the ensemble $\hat{\boldsymbol{\Sigma}}_{rd}$,

$$\tilde{\boldsymbol{\Sigma}}_{rd} = \boldsymbol{\Theta} \circ \hat{\boldsymbol{\Sigma}}_{rd}, \tag{47}$$

where $\circ$ is the element-wise matrix product. The piecewise rational function presented in Gaspari and Cohn (1999) is often used to define $\boldsymbol{\Theta}$. Note that because the size of the ensemble is usually smaller than the dimension of the state vector, the estimate of $\boldsymbol{\Sigma}_{rd}$ is rank deficient. Using localization will helped improve the rank of $\hat{\boldsymbol{\Sigma}}_{rd}$.

Example: The example depicted in Figure 3 is adapted from Paper III. The ensemble (size 100) representing the porosity field (size 4096) is conditioned on the bottom hole pressure and the oil production rate collected at the data collection points. The figure shows that localization helps conserve variability away from where the data was collected when the ensemble size is much smaller than the dimension of the grid size.

Inflation

In the conditioning step of the EnKF, see Equation 41, the Kalman gain $\boldsymbol{K}_t$ is estimated using the ensemble which introduces coupling between

15

Figure 3: Marginal Variance: prior ensemble (left), the black dots denote the data collection points, marginal variance of the posterior ensemble with localization (center), posterior ensemble without localization (right)

ensemble members (Sætrom and Omre, 2013). This in turn leads to an underestimation of the posterior marginal variances. Ensemble inflation (Anderson and Anderson, 1999) is a pragmatic approach that addresses the consequence of the problem (underestimated variance) rather than its cause (coupled ensemble members) by increasing the spread of the ensemble prior to conditioning without altering the correlation between ensemble members. The inflated ensemble is given by,

$$\boldsymbol{r}_I^{u(i)} = \alpha \times (\boldsymbol{r}^{u(i)} - \bar{\boldsymbol{r}}^u) + \bar{\boldsymbol{r}}^u, i = 1, \ldots, n_e, \tag{48}$$

where $\bar{\boldsymbol{r}}^u$ is the ensemble mean and $\alpha \geq 1$ is the inflation factor.

Example: The example depicted in Figure 4 is taken from the Report. The ensemble representing the log-conductivity field is conditioned on drawdown measurements from real pumping test data. The updated ensemble members are used to simulate drawdown observations that are then compared to the actual drawdown measurements. Without inflation, the predicted drawdowns fail to cover the measured data, while with inflation, the predicted drawdowns cover the measured data more convincingly. However, one must be careful not to read too much into the uncertainty quantification in the posterior ensemble as it has been arbitrarily altered.

In our experience, results obtained with inflation have often been underwhelming and hard to interpret, while localization appeared to be more robust and often made resorting to inflation unnecessary.

Iterative methods

When the non-linearity in the forward and/or likelihood models becomes too dominant for the EnKF to accommodate, matching the observed

Figure 4: Data match using the EnKF (left) and the EnKF with inflation (right). The red markers represent the measured data, the black line the mean prediction, the dark gray area represents the 70% confidence interval, the light gray area represents the 90% confidence interval

data becomes difficult. Indeed, the conditioning step of the EnKF is an approximation whose validity is challenged when the forward and likelihood models deviate too far from Gauss-linearity.

Iterative EnKF Because the conditioning step can be seen as the first step in a Gauss-Newton algorithm, it is intuitive to replace the conditioning step by an iterative method such as Gauss-Newton or Levenberg-Marquardt. Assume the likelihood model is defined with additive Gaussian error term, $[\boldsymbol{d}_t|\boldsymbol{r}_t] = \psi_t(\boldsymbol{r}_t) + \boldsymbol{\epsilon}_t^d$, with m-vector $\boldsymbol{\epsilon}_t^d$ being centred Gaussian with covariance matrix $\boldsymbol{\Sigma}_t^d$, the conditioning step of the iterative EnKF (IEnKF) (Gu and Oliver, 2007; Sakov et al., 2012) is then given by,

$$\boldsymbol{r}_t^{c(i)} = \mathrm{argmax}_{\boldsymbol{r}}\{(\boldsymbol{r}_t^{u(i)} - \boldsymbol{r})^T\left[\hat{\boldsymbol{\Sigma}}_r\right]^{-1}(\boldsymbol{r}_t^{u(i)} - \boldsymbol{r}) + \ldots \qquad (49)$$

$$(\psi_t(\boldsymbol{r}) + \boldsymbol{\epsilon}_t^d - \boldsymbol{d}_t)^T\left[\boldsymbol{\Sigma}_t^d\right]^{-1}(\psi_t(\boldsymbol{r}) + \boldsymbol{\epsilon}_t^d - \boldsymbol{d}_t)\}, i = 1, \ldots, n_e,$$

where $\hat{\boldsymbol{\Sigma}}_r$ is inferred by calculating the sample covariance matrix of the ensemble $\boldsymbol{e}_t^r = \{\boldsymbol{r}_t^{u(i)}, i = 1, \ldots, n_e\}$ and where the gradients used for the maximization are either calculated by solving the adjoint problem or estimated using the ensemble.

Multiple data assimilation   The idea behind multiple data assimilation (MDA) (Emerick and Reynolds, 2012, 2013) is to split the conditioning step in several smaller substeps to mitigate the non-linearity and retrieve a better data match. Assume the likelihood model is defined with additive Gaussian error term, $[\boldsymbol{d}_t|\boldsymbol{r}_t] = \psi_t(\boldsymbol{r}_t) + \boldsymbol{\epsilon}_t^d$, with m-vector $\boldsymbol{\epsilon}_t^d$ being centred Gaussian with covariance matrix $\boldsymbol{\Sigma}_t^d$. This likelihood model can therefore be decomposed as,

$$f(\boldsymbol{d}_t|\boldsymbol{r}_t)) = \prod_{i=1}^{L} [(f(\boldsymbol{d}_t|\boldsymbol{r}_t))]^{w_i} \tag{50}$$

$$\propto \prod_{i=1}^{L} \exp\left(-\frac{1}{2}(\boldsymbol{d}_t - \psi_t(\boldsymbol{r}_t))^T \left[\frac{1}{w_i}\boldsymbol{\Sigma}_t^d\right]^{-1}(\boldsymbol{d}_t - \psi_t(\boldsymbol{r}_t^u))\right),$$

with positive decomposition factors $\boldsymbol{w} : \{w_i, i = 1, \ldots, L\}$ such that $\sum_{i=1}^{L} w_i = 1$. Equation 50 gives the blueprint for the MDA EnKF algorithm - the conditioning step is repeated L times using the following likelihood model:

$$[\boldsymbol{d}_t|\boldsymbol{r}_t] = \psi_t(\boldsymbol{r}_t) + \frac{1}{\sqrt{w_i}}\boldsymbol{\epsilon}_t^d, i = 1, \ldots, L. \tag{51}$$

Because the measurement error of each sub-likelihood model is larger than in the actual likelihood model, each sub-steps results in a smaller update.

Example: The example depicted in Figure 5 is taken from the Report and shows the effect of iterative methods on data match. The ensemble representing the log-conductivity field is conditioned on drawdown measurements from real pumping test data. The updated ensemble members are used to simulate drawdown observations that are then compared to the actual drawdown measurements. Data match for the EnKF is poor while it is satisfactory for the IEnKF and the MDA EnKF. The spread obtained using the MDA EnKF is larger than the IEnKF because the MDA EnKF increases the measurement error variance to account for the repeated conditioning steps within each data assimilation cycle, while the IEnKF does not.

Selection EnKF

The selection EnKF (SEnKF) extends the selection Kalman Filter to non-linear forward and likelihood models. The goal of the SEnKF is to allow for multimodality in the posterior ensemble, which is difficult to
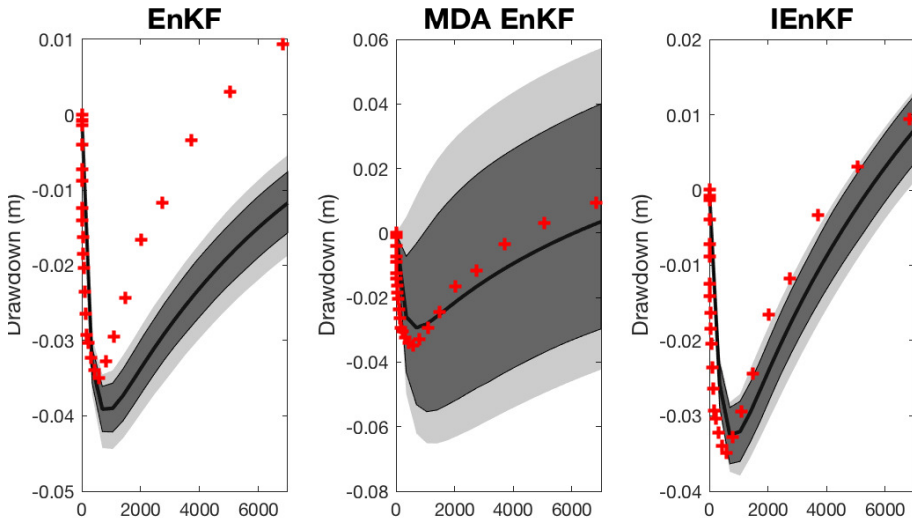
Figure 5: Data match using the EnKF (left), MDA EnKF (center) and IEnKF (right). The red dots represent the measured data, the black line the mean prediction, the dark gray area represents the 70% confidence interval, the light gray area represents the 90% confidence interval

achieve in the EnKF: when specifying a non-Gaussian prior distribution, the successive conditioning steps render the posterior ensemble Gaussian-like.

Let the initial distribution $f(\boldsymbol{r}_0)$ be a selection-Gaussian. We take advantage of the structure of the selection-Gaussian distribution and consider the augmented state vector $[\tilde{\boldsymbol{r}}_0, \boldsymbol{\nu}]$ which is jointly Gaussian. Conditioning on $\boldsymbol{\nu} \in A$ is then done after conditioning the augmented state vector on the data, therefore allowing for multimodality in the posterior ensemble. We further define the appropriate forward and likelihood models for the augmented state vector. The forward model is given by

$$\begin{bmatrix} \tilde{\boldsymbol{r}}_{t+1} \\ \boldsymbol{\nu}_{t+1} \end{bmatrix} \begin{vmatrix} \tilde{\boldsymbol{r}}_t \\ \boldsymbol{\nu}_t \end{vmatrix} = \begin{bmatrix} \omega_t(\tilde{\boldsymbol{r}}_t, \boldsymbol{\epsilon}_t^r) \\ \boldsymbol{\nu}_t \end{bmatrix}, \tag{52}$$

while the likelihood model is given by

$$\begin{bmatrix} \boldsymbol{d}_t \end{vmatrix} \begin{matrix} \tilde{\boldsymbol{r}}_t \\ \boldsymbol{\nu}_t \end{matrix} \end{bmatrix} = \psi_t(\tilde{\boldsymbol{r}}_t, \boldsymbol{\epsilon}_t^d). \tag{53}$$

The SEnKF is a two step algorithm. The first step is to run the EnKF: the initial ensemble $\boldsymbol{e}_0$ of contains realizations from the aug-

19

mented state vector $[\tilde{\boldsymbol{r}}_0, \boldsymbol{\nu}]$ that is jointly Gaussian, the forward and likelihood model are defined by Equations 52 and 53, the posterior ensemble $\boldsymbol{e}_T = \{(\tilde{\boldsymbol{r}}_T^{u(i)}, \boldsymbol{\nu}_T^{u(i)}), i = 1, \ldots, n_e\}$ therefore represents $f(\tilde{\boldsymbol{r}}_T, \boldsymbol{\nu}_T | \boldsymbol{d}_{0:T})$. The second step of the SEnKF is to use MCMC to assess $f(\boldsymbol{r}_T | \boldsymbol{d}_{0:T}) = f(\tilde{\boldsymbol{r}}_T | \boldsymbol{\nu} \in A, \boldsymbol{d}_{0:T})$: the posterior ensemble is assumed to be jointly Gaussian $\varphi_{2n}((\tilde{\boldsymbol{r}}, \boldsymbol{\nu}); \hat{\boldsymbol{\mu}}_{\tilde{r}\nu}, \hat{\boldsymbol{\Sigma}}_{\tilde{r}\nu})$ with expectation vector $\boldsymbol{\mu}_{\tilde{r}\nu}$ and covariance matrix $\boldsymbol{\Sigma}_{\tilde{r}\nu}$ estimated from $\boldsymbol{e}_T$, $f(\boldsymbol{r}_T | \boldsymbol{d}_{0:T}) = f(\tilde{\boldsymbol{r}}_T | \boldsymbol{\nu} \in A, \boldsymbol{d}_{0:T})$ is therefore selection-Gaussian and estimated using the Metropolis-Hastings algorithm detailed in Omre and Rimstad (2021). A selection ensemble Kalman smoother (SEnKS) and selection ensemble smoother (S-ES) can be defined similarly. The SEnKS is a straightforward extension of the EnKS using the same augmented state vector, forward and likelihood models as the SEnKF.
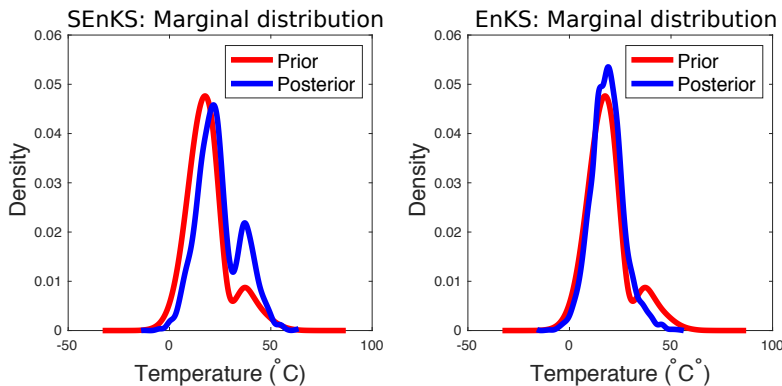

Figure 6: Example of a marginal distribution of the SEnKS (left) and EnKS (right)

Example: The example depicted in Figure 6 is taken from Paper II. It show the posterior marginal distribution of the initial temperature field at a given location overlain with the prior marginal distribution at the same location. It illustrates the SEnKS's ability to conserve bimodality in the posterior ensemble.

Summary of the papers

The objective of this thesis work is to develop a new methodology for sequential data assimilation that can represent multimodality in the posterior distribution. In Paper I, the methodology is defined for linear problems and Paper II extends it non-linear problems. Case studies based on the advection-diffusion equation and the diffusion equation, respectively, illustrate the approach. The methodology is then tested on subsurface modeling problems where the quantities of interest often display multimodal spatial histograms. In Paper III, we apply the methodology to reservoir modeling and the assimilation of seismic and production data. The goal is to estimate porosity and water saturation. In the Report, we consider groundwater flow and the assimilation of real pumping test data with the purpose of estimating hydraulic conductivity. A large variety of EnKF algorithms are presented, evaluated and compared.

Paper 1: Spatio-temporal Inversion using the Selection Kalman Model

Maxime Conjard and Henning Omre

Abstract: Data assimilation in models representing spatio-temporal phenomena poses a challenge, particularly if the spatial histogram of the variable appears with multiple modes. The traditional Kalman model is based on a Gaussian initial distribution and Gauss-linear forward and observation models. This model is contained in the class of Gaussian distribution and is therefore analytically tractable. It is however unsuitable for representing multimodality. We define the selection Kalman model that is based on a selection-Gaussian initial distribution and Gauss-linear forward and observation models. The selection-Gaussian distribution can be seen as a generalization of the Gaussian distribution and may represent multimodality, skewness and peakedness. This selection Kalman model is contained in the class of selection-Gaussian distributions and therefore it is analytically tractable. An efficient recursive algorithm for assessing the selection Kalman model is specified. The synthetic case study of spatio-temporal inversion of an initial state, inspired by pollution monitoring, suggests that the use of the selection Kalman model offers significant improvements compared to the traditional Kalman model when reconstructing discontinuous initial states.
Major contribution: In this paper, we define the selection Kalman model for sequential data assimilation of linear problems. The selection Kalman model is a generalization of the traditional Kalman model. In the tradi-

tional Kalman model, the posterior distribution is Gaussian while in the selection Kalman model, the posterior distribution is selection-Gaussian thereby allowing for multimodality to be represented in the posterior distribution. The results from the synthetic case study based on the advection-diffusion equation illustrate the benefits of using the selection Kalman model when reconstructing an initial temperature field that displays a bimodal spatial histogram.

Paper 2: Data Assimilation in Spatio-Temporal Models with Non-Gaussian Initial States: The Selection Ensemble Kalman Model

Maxime Conjard and Henning Omre

Abstract: Assimilation of spatio-temporal data poses a challenge when allowing non-Gaussian features in the prior distribution. It becomes even more complex with nonlinear forward and likelihood models. The ensemble Kalman model and its many variants have proven resilient when handling nonlinearity. However, owing to the linearized updates, conserving the non-Gaussian features in the posterior distribution remains an issue. When the prior model is chosen in the class of selection-Gaussian distributions, the selection Ensemble Kalman model provides an approach that conserves non-Gaussianity in the posterior distribution. The synthetic case study features the prediction of a parameter field and the inversion of an initial state for the diffusion equation. By using the selection Kalman model, it is possible to represent multimodality in the posterior model while offering a 20 to 30% reduction in root mean square error relative to the traditional ensemble Kalman model.

Major contribution: In this paper, we extend the selection Kalman model to sequential data assimilation of non-linear problems by defining the selection ensemble Kalman filter and the selection ensemble Kalman smoother. In the traditional ensemble Kalman filter/smoother, the ensemble drifts towards Gaussianity as data are assimilated. In the selection ensemble Kalman filter/smoother, the posterior distribution is assumed to be selection-Gaussian therefore allowing for multimodality to be represented in the posterior distribution. The case studies based on the diffusion equation, similar to the one in Paper I, show that the selection ensemble Kalman filter/smoother can improve the estimation of quantities of interest that display bimodal spatial histograms.

Paper 3: Ensemble-based seismic and production data assimilation using the selection Kalman model

Maxime Conjard and Dario Grana

Abstract: Data assimilation in reservoir modeling often involves model variables that are multimodal, such as porosity and permeability. Well established data assimilation methods such as ensemble Kalman filter and ensemble smoother approaches, are based on Gaussian assumptions that are not applicable to multimodal random variables. The selection ensemble smoother is introduced as an alternative to traditional ensemble methods. In the proposed method, the prior distribution of the model variables, for example the porosity field, is a selection-Gaussian distribution, which allows to model the multimodal behavior of the posterior ensemble. The proposed approach is applied on a two-dimensional synthetic channelized reservoir for validation. In the application, the unknown reservoir model of porosity and permeability is estimated from the measured data. Seismic and production data are assumed to be repeatedly measured in time and the reservoir model is updated every time new data are assimilated. The example shows that the selection ensemble Kalman model improves the characterisation of the bimodality of the model parameters compared to the results of the ensemble smoother.

Major contribution: In this paper, we define the selection ensemble smoother based on the selection Kalman model. The selection ensemble smoother is used to assimilate seismic and production data from a synthetic reservoir model. The data is assimilated in one step contrary to the selection ensemble Kalman smoother in which data is assimilated sequentially. The results from this realistic case study shows that the selection ensemble smoother provides a credible alternative to the ensemble smoother when the quantities of interest display bimodal spatial histograms.

Report: Ensemble methods applied to groundwater flow

Maxime Conjard, Emilio Sanchéz-Léon, Olaf Cirpka and Henning Omre

Abstract: Assimilation of spatio-temporal data is challenging especially when parameters are suspected to display a non-Gaussian spatial histogram. When that is the case, the selection ensemble Kalman filter has been shown to produce encouraging results on synthetic tests cases, the goal is therefore to situate its performance on a real data application when compared to established methods. To that end, we first present a review of existing ensemble Kalman filtering methods. We then present a

synthetic and real data study where hydraulic conductivity is estimated using pumping test data. The synthetic study confirms the suitability of the selection ensemble Kalman filter when the conductivity field displays a bimodal spatial histogram. The real data study shows that the selection ensemble Kalman filter provides robust results, but its relevance is questioned when considering the added cost and the comparively good results provided by the ensemble Kalman filter with multiple data assimilation. *Major contribution:* This report determines the applicability of the selection Kalman model to real data experiments in which the quantity of interest displays a bimodal spatial histogram. A review of various versions of the EnKF algorithm is presented and their respective performances are discussed on a real data case study. The latter shows that the SEnKF provides robust results when compared to other established methods such as the IEnKF and the EnKF MDA.

Synthesis

The objective of this thesis work is to develop a methodology for spatio-temporal data assimilation when the variables of interest display multi-modal spatial histograms. Data assimilation is cast in a Bayesian setting, and we specify a hidden Markov model. We define the selection Kalman model in which the initial distribution is selection-Gaussian. The selection-Gaussian distribution can represent multimodality, skewness and heavytailedness.

For Gauss-linear models, the selection Kalman model is a generalization of the traditional Kalman model (Paper I). The posterior distribution is selection-Gaussian and analytically tractable.

For non-linear models, analytical tractability is lost. We extend the use of the selection Kalman model to ensemble methods by defining the selection ensemble Kalman filter (Paper II, Report), the selection ensemble Kalman smoother (Paper II) and the selection smoother (Paper III). These three algorithms mitigate the regression towards the mean that occurs in traditional ensemble methods, thereby allowing for multimodality to be represented in the posterior distribution.

Synthetic (Paper I, II, III, Report) and real data (Report) case studies show that the selection Kalman model provides robust results that can improve the estimation of quantities of interest that display multimodal spatial histograms for Gauss-linear and non-linear models.

Further research should consider high-dimensional problems and focus on improving the McMC sampling to accomodate larger grid size. A comparative study comparing the merits of the selection Kalman model to that of Gaussian mixture models would also be of interest.

Bibliography

Aanonsen, S., Nævdal, G., Oliver, D., Reynolds, A., and Vallès, B. (2009). The ensemble Kalman filter in reservoir engineering–A review. *SPE Journal*, 14:393–412.

Anderson, J. and Anderson, S. L. (1999). A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review*, 127:2741–2758.

Arellano-Valle, R. B., Branco, M. D., and Genton, M. G. (2006). A unified view on skewed distributions arising from selections. *Canadian Journal of Statistics*, 34(4):581–601.

Arellano-Valle, R. B. and del Pino, G. (2004). *From symmetric to asymmetric distributions: A unified approach*, pages 113–133. Chapman & Hall / CRC, New York.

Bertino, L., Evensen, G., and Wackernagel, H. (2003). Sequential data assimilation techniques in oceanography. *International Statistical Review*, 71(2):223–241.

Cappé, O., Moulines, E., and Ryden, T. (2005). *Inference in Hidden Markov Models (Springer Series in Statistics)*. Springer-Verlag, Berlin, Heidelberg.

Chen, Y. and Oliver, D. (2012). Ensemble randomized maximum likelihood method as an iterative ensemble smoother. *Mathematical Geosciences*, 44(1):1–26.

Doucet, A., de Freitas, N., and Gordon, N. (2001). *An Introduction to Sequential Monte Carlo Methods*, pages 3–14. Springer New York, New York.

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

Emerick, A. and Reynolds, A. (2012). History matching time-lapse seismic data using the ensemble Kalman filter with multiple data assimilations. *Computational Geosciences*, 16:639–659.

Emerick, A. and Reynolds, A. (2013). Ensemble smoother with multiple data assimilation. *Computers & Geosciences*, 55:3–15.

Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99(C5):101–143.

Evensen, G. (2009). *Data assimilation: the ensemble Kalman filter*. Springer Science & Business Media, Berlin Heidelberg.

Evensen, G. and van Leeuwen, P. J. (2000). An Ensemble Kalman Smoother for Nonlinear Dynamics. *Monthly Weather Review*, 128(6):1852–1867.

Gaspari, G. and Cohn, S. (1999). Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, 125:723–757.

Gordon, N. J., Salmond, D. J., and Smith, A. F. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)*, 140(2):107–113.

Gu, Y. and Oliver, D. (2007). An iterative ensemble Kalman filter for multiphase fluid flow data assimilation. *SPE Journal*, 12:438–446.

Hendricks Franssen, H.-J. and Kinzelbach, W. (2008). Real-time groundwater flow modeling with the ensemble Kalman filter: Joint estimation of states and parameters and the filter inbreeding problem. *Water Resources Research*, 44(9).

Jazwinski, A. H. (1970). *Stochastic processes and filtering theory*. Acad. Press, New York.

Julier, S. J. and Uhlmann, J. K. (1997). New extension of the Kalman filter to nonlinear systems. In Kadar, I., editor, *Signal Processing, Sensor Fusion, and Target Recognition VI*, volume 3068, pages 182 – 193. SPIE.

Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1):35–45.

Katzfuss, M., Stroud, J. R., and Wikle, C. K. (2020). Ensemble Kalman methods for high-dimensional hierarchical dynamic space-time models. *Journal of the American Statistical Association*, 115(530):866–885.

Li, R., Prasad, V., and Huang, B. (2016). Gaussian mixture model-based ensemble Kalman filtering for state and parameter estimation for a PMMA process. *Processes*, 4(2):9.

Moja, S., Asfaw, Z., and Omre, H. (2018). Bayesian inversion in hidden Markov models with varying marginal proportions. *Mathematical Geosciences*, 51:463–484.

Oliver, D. and Chen, Y. (2018). Data assimilation in truncated plurigaussian models: Impact of the truncation map. *Mathematical Geosciences*, 50:867–893.

Omre, H. and Rimstad, K. (2021). Bayesian spatial inversion and conjugate selection Gaussian prior models. *Journal of Uncertainty Quantification*, To appear.

Rauch, H. E., Tung, F., and Striebel, C. T. (1965). Maximum likelihood estimates of linear dynamic systems. *AIAA Journal*, 3(8):1445–1450.

Sætrom, J. and Omre, H. (2013). Uncertainty quantification in the ensemble Kalman filter. *Scandinavian Journal of Statistics*, 40(4):868–885.

Sakov, P., Haussaire, J.-M., and Bocquet, M. (2018). An iterative ensemble Kalman filter in the presence of additive model error. *Quarterly Journal of the Royal Meteorological Society*, 144(713):1297–1309.

Sakov, P., Oliver, D., and Bertino, L. (2012). An iterative EnKF for strongly nonlinear systems. *Monthly Weather Review*, 140(6):1988–2004.

Särkkä, S. (2008). Unscented Rauch–Tung–Striebel smoother. *IEEE Transactions on Automatic Control*, 53(3):845–849.

van Leeuwen, P. J. (2009). Particle filtering in geophysical systems. *Monthly Weather Review*, 408:4089–4114.

Yu, B. M., Shenoy, K. V., and Sahani, M. (2004). Derivation of extended Kalman filtering and smoothing equations. *Technical Report, Stanford*.

Zhou, H., Li, L., Franssen, H.-J., and Gomez-Hernandez, J. (2012). Pattern recognition in a bimodal aquifer using the normal-score ensemble Kalman filter. *Mathematical Geosciences*, 44:1–17.

Spatio-temporal Inversion using the Selection Kalman Model

Maxime Conjard and Henning Omre

Spatio-Temporal Inversion Using the Selection Kalman Model

Maxime Conjard *[†] **and Henning Omre**[†]

Department of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway

Data assimilation in models representing spatio-temporal phenomena poses a challenge, particularly if the spatial histogram of the variable appears with multiple modes. The traditional Kalman model is based on a Gaussian initial distribution and Gauss-linear forward and observation models. This model is contained in the class of Gaussian distribution and is therefore analytically tractable. It is however unsuitable for representing multimodality. We define the selection Kalman model that is based on a selection-Gaussian initial distribution and Gauss-linear forward and observation models. The selection-Gaussian distribution can be seen as a generalization of the Gaussian distribution and may represent multimodality, skewness and peakedness. This selection Kalman model is contained in the class of selection-Gaussian distributions and therefore it is analytically tractable. An efficient recursive algorithm for assessing the selection Kalman model is specified. The synthetic case study of spatio-temporal inversion of an initial state, inspired by pollution monitoring, suggests that the use of the selection Kalman model offers significant improvements compared to the traditional Kalman model when reconstructing discontinuous initial states.

Keywords: inverse problem, spatio-temporal variables, Kalman model, multimodality, data assimilation

1 INTRODUCTION

Data assimilation in models representing spatio-temporal phenomena is challenging. Most statistical spatio-temporal models are based on assumptions of temporal stationarity, possibly with a parametric, seasonal trend model [1]. We consider spatio-temporal phenomena where the dynamic spatial variables evolve according to a set of differential equations. Such phenomena will, in statistics, normally be modeled as hidden Markov models [2]. The celebrated Kalman model [3] is one of the most frequently used hidden Markov models.

In studies of hidden Markov models, it is natural to distinguish between filtering and smoothing [2]. Filtering entails predicting the spatial variable at a given time with observations up to that point in time. Smoothing entails predicting the spatial variable given observations both at previous and later times. Filtering is naturally based on recursive temporal updating while smoothing appears as more complicated since updating must also be made backwards in time. We focus on a particular smoothing challenge, namely to assess the initial state given observations at later times and we denote the task spatio-temporal inversion.

Spatio-temporal inversion is of interest in many applications. In petroleum engineering, initial water saturation is often unknown. Ensemble smoothing techniques [4–6] are commonly used to evaluate this parameter and improve fluid flow prediction. In air pollution monitoring [7], inverse trajectory methods are used to identify potential source contribution. Source mapping of wildfire

origin from airborne smoke observations is a spectacular example [8]. Evaluation of groundwater pollution mostly focuses on the future pollution of the pollutant, but as emphasized in [9], the identification of the heterogeneous source may be complicated.

We study a continuous spatial variable, a random field (RF), with temporal behavior governed by a set of differential equations. The spatio-temporal variable is discretized in space and time, and the hidden Markov model is cast in a Bayesian framework. The prior model consists of an initial spatial model and a forward spatio-temporal model, representing the evolution of the spatio-temporal phenomenon. The likelihood model represents the observation acquisition procedure. The corresponding posterior model, fully defined by the prior and likelihood models, represents the spatio-temporal phenomenon given the available observations. The traditional Kalman model constitutes a very particular hidden Markov model [3] with a Gaussian initial model and a linear forward function with Gaussian error term (Gauss-linear) forward model, and a Gauss-linear likelihood model. Since the class of Gaussian models is closed under linear operations, the posterior distribution is also Gaussian in the Kalman model, and the posterior model parameters are analytically tractable. Based on this posterior Gaussian model, both filtering and smoothing can easily be performed. In particular, the spatio-temporal inversion can be obtained by integrating out the spatial variables at all time points except the initial one, which is a simple task in Gaussian models. Most spatio-temporal models used in statistical studies are defined in the traditional Kalman model framework [10, 11]. Moreover, most of these models are based on spatial stationarity and consider filtering. Their focus is primarily on computational efficiency, not on model flexibility.

The fundamental Gauss-linear assumptions of the traditional Kalman model are often not suitable in real studies. The initial spatial variable may appear as non-Gaussian and/or the forward and/or the likelihood functions are non-linear. In the control theory community, linearizations such as the extended Kalman filter [12] or quantile-based representation such as the unscented Kalman filter [13] are recommended in these cases. These approaches are suitable for models with mild deviations from Gauss-linearity. Statisticians will more naturally use various Monte-Carlo based approaches such as the particle filter [2] or the ensemble Kalman Filter (EnKF) [14]. The particle filter is a sequential Monte Carlo algorithm with data assimilation made by reweighting the particles. To avoid singular solutions, resampling is usually required during the assimilation. The need for resampling makes the definition of an efficient corresponding particle smoother difficult. The EnKF is also a sequential Monte Carlo algorithm with data assimilation based on linear updates of each ensemble member. The sequential linear updates cause the ensemble distribution to drift toward Gaussianity. A corresponding ensemble smoother [15] is available but the ensemble drift toward Gaussianity makes it difficult to preserve non-Gaussianity in the posterior distribution. The discrete representation of the spatial variable makes the spatio-temporal model high-dimensional, and according to [16, 17], the EnKF is preferable to the particle filter in high dimensional models. Lastly, brute force Markov chain Monte Carlo

(McMC) [18] algorithms may be used for spatio-temporal inversion, but the increasing coupling of the temporal observations makes these algorithms inefficient. Focus in our study is on the spatial initial state of the spatio-temporal phenomenon, and we aim at reproducing clearly non-Gaussian marginal features, such as multi-modality, skewness and peakedness. Several models with such features are presented in the literature.

In [19], a hidden Markov model with a skew-Gaussian initial model is defined, and for Gauss-linear forward and likelihood models, it is demonstrated that the filtering is analytically tractable. The skew-Gaussian model is based on a selection concept, and the current spatio-temporal model will later be defined along these lines.

In [16, 20–24], the initial model in the hidden Markov model is defined to be a Gaussian mixture model representing multimodality. These studies all consider filtering problems and the filter algorithms are based on a combination of clustering/particle filter and Kalman filter/EnKF. The Gaussian mixture model contains a latent categorical mode indicator, which in a spatial setting must have spatial coupling, for example in the form of a Markov RF [25]. Data assimilation in such categorical Markov RFs, either by particle filter or EnKF, appears as very complicated [23, 24], particularly in a smoothing setting.

The spatio-temporal case we consider in the current study has a non-Gaussian spatial initial model, while both the forward and likelihood models are Gauss-linear. We study this special case since it has a particularly elegant analytical solution. Further, we study spatio-temporal inversion, which entails smoothing to assess the initial spatial variable given observations up to current time, as it constitutes a particularly challenging problem. To our knowledge, no reliable methodology exists for solving such a spatio-temporal inverse problem.

In the hidden Markov model considered in this study, the initial spatial model is assigned a selection-Gaussian RF [26], which may capture multi-modality, skewness and/or peakedness in the spatial histogram of the initial spatial variable. Recall that the forward and likelihood models are assumed to be Gauss-linear. Since the class of selection-Gaussian models is closed under linear operations [26] the posterior model will also be selection-Gaussian. The posterior model parameters are then analytically tractable. A general algorithm for assessing this posterior selection-Gaussian model is defined. Based on this posterior model, both smoothing and filtering can be performed. We denote this special hidden Markov model the selection Kalman model. The class of Gaussian models is a central member in the class of selection-Gaussian models [26], hence one may consider the selection Kalman model to be a generalization of the Kalman model. We develop the results presented above and demonstrate the use of the selection Kalman model on a synthetic case study of spatio-temporal inversion. This entails assessing the initial spatial variable in a dynamic model based on a limited set of observations.

The characteristics of the class of selection-Gaussian models are central in the development of the selection Kalman model properties. These characteristics are thoroughly discussed in [26],
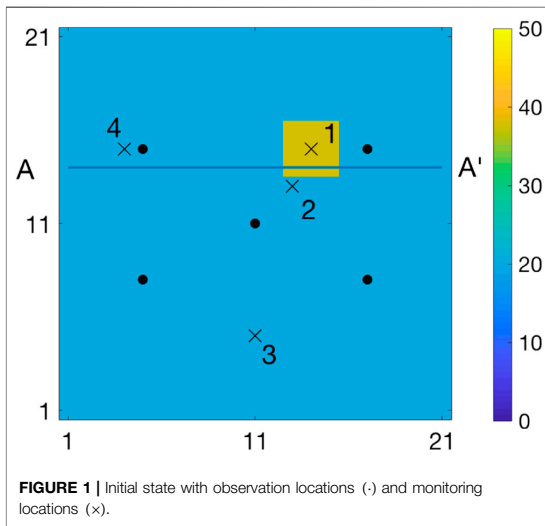
FIGURE 1 | Initial state with observation locations (·) and monitoring locations (×).

which is inspired by the results presented in [27, 28]. In these papers, the general concept of the selection-Gaussian pdf is defined in a probabilistic setting. In [19], the one-sided selection concept is used to define a skew Kalman filter in a non-spatial setting.

In this paper $y \sim f(y)$ denotes a random variable y distributed according to the probability density function (pdf) $f(y)$, or alternatively according to the corresponding cumulative distribution function (cdf) $F(y)$. Moreover, $\varphi_n(y; \mu, \Sigma)$ denotes the pdf of the Gaussian n-vector y with expectation n-vector μ and covariance $(n \times n)$-matrix Σ. Further $\Phi_n(A; \mu, \Sigma)$ denotes the probability of the aforementioned Gaussian n-vector y to be in $A \subset \mathbb{R}^n$. We also use i_n to denote the all-ones n-vector and I_n to denote the identity $(n \times n)$-matrix.

In **Section 2**, the problem is set. In **Section 3**, the traditional Kalman model is cast in a Bayesian hidden Markov model framework. The generalization to the selection Kalman model is then defined, and the analytical tractability is investigated. Further a general recursive algorithm for assessing the posterior distribution is specified. In **Section 3**, a synthetic case study of the advection-diffusion equation is discussed to showcase the ability of the selection Kalman model to solve the spatio-temporal inversion problem. The goal is to reconstruct the initial state. Results from the selection Kalman model and the traditional Kalman model are compared. In **section 4**, conclusions are presented.

2 PROBLEM SETTING

The case is defined in a spatio-temporal setting. Consider the variable $\{r_t(x); x \in \mathcal{L}_r, t \in \mathcal{T}\}; r.(\cdot) \in \mathbb{R}$, with $\mathcal{L}_r$ a grid of size n over a two-dimensional spatial area of interest while $\mathcal{T}$: $\{0, 1, \ldots, T\}$ is a regular discretization in time. Let $t = T$ represent current time while $t = 0$ represents the initial time. The spatial variable $\{r_0(x); x \in \mathcal{L}_r\}$ is a discretized representation of the initial state which later will be assumed to be unknown. **Figure 1**

displays the initial state that we evaluate in the case study. It is a spatial field with two areas: the blue area is at low value and the yellow area is at much higher value. One may consider the yellow area as the release of a pollutant at time $t = 0$.

The spatio-temporal variable evolves in time, $\{r_{t+1}(x); x \in \mathcal{L}_r\} = \omega_t[\{r_t(x); x \in \mathcal{L}_r\}]$ where $\omega_t(\cdot)$ is a dynamic function usually represented by a set of discretized differential equations. **Figure 2** shows the temporal evolution of the field presented in **Figure 1** according to a set of differential equations. The spatio-temporal variable is not fully observable, it can only be measured at a number of observation sites. The observations at the observation sites are collected with some measurement error, they appear as time series denoted $\{d_t = (d_t^1, \ldots, d_t^m), t \in \mathcal{T}\}$ where m is the number of observation sites. The five observation locations in the case study are represented by dots in **Figure 1**. **Figure 3** displays the observations collected at these observation locations. The typical challenge is to infer the spatio-temporal variable $\{r_t(x); x \in \mathcal{L}_r, t \in \mathcal{T}\}$ based on the observed time series $\{d_t; t \in \mathcal{T}\}$. It constitutes a complex spatio-temporal inverse problem. In the current study we focus on assessing the initial spatial variable $\{r_0(x); x \in \mathcal{L}_r\}$ from the observed time series $\{d_t; t \in \mathcal{T}\}$.

3 MODEL DEFINITION

Consider the unknown temporal n-vector r_t, representing the discretized spatial variable $\{r_t(x); x \in \mathcal{L}_r\}$, for $t \in \mathcal{T}_r$: $\{0, 1, \ldots, T, T+1\}$. Define the variable $r = \{r_0, r_1, \ldots, r_T, r_{T+1}\}$ and let $r_{i:j}$ denote $\{r_i, r_{i+1}, \ldots, r_j\}, \forall (i, j) \in \mathcal{T}_r^2, i \le j$. Moreover assume that the temporal m-vectors of observations d_t for $t \in \mathcal{T}_d$: $\{0, 1, \ldots, T\}$ are available, and define $d = \{d_0, d_1, \ldots, d_T\}$ and $d_{i:j} = \{d_i, \ldots, d_j\}$ accordingly. The objective of this study is to assess $[r_0|d]$. To that end, we define a Kalman type model, represented as a hidden Markov model in a Bayesian inversion framework.

3.1 Bayesian Inversion

The Kalman type model, phrased as Bayesian inversion, requires the specification of a prior model for r and a likelihood model for $[d|r]$. The model specified below defines a hidden Markov model as displayed in **Figure 4**.

3.1.1 Prior Model

The prior model on r synthesizes the knowledge and experience with the spatial variable of interest, and it consists of an initial distribution and a forward model:

Initial Distribution

The prior distribution for the initial state r_0 is denoted $f(r_0)$.

Forward Model

The forward model given the initial state $[r_{1:T+1}|r_0]$ is defined as,

$$f(r_{1:T+1}|r_0) = \prod_{t=0}^{T} f(r_{t+1}|r_t) \tag{1}$$

with,

$$[r_{t+1}|r_t] = \omega_t(r_t, \varepsilon_t^r) \sim f(r_{t+1}|r_t) \tag{2}$$

FIGURE 2 | Spatio-temporal diffusion.

where $\omega_t(\cdot, \cdot) \in \mathbb{R}^n$ is the function that propagates r_t forward in time, with ε_t^r a random component. Since $\omega_t(\cdot, \cdot)$ only involves the variable at the previous time step, r_t, the forward model is a Markov chain.

3.1.2 Likelihood Model

The likelihood model on $[d|r]$ provides a link between the variable of interest r and the observations d and is defined as,

$$f(d|r) = \prod_{t=0}^{T} f(d_t|r_t) \tag{3}$$

with,

$$[d_t|r_t] = \psi_t\left(r_t, \varepsilon_t^d\right) \sim f(d_t|r_t) \tag{4}$$

where $\psi_t(\cdot, \cdot) \in \mathbb{R}^m$ is the likelihood function with ε_t^d a random component. The likelihood model is defined assuming conditional independence and single state response and is thus in factored form.

3.1.3 Posterior Model

Bayesian inversion endeavors to assess the posterior distribution of $[r|d]$,

$$
\begin{aligned}
f(r|d) &= \left[\int f(d|r) f(r) dr \right]^{-1} \times f(d|r) f(r) \\
&= const \times f(d_0|r_0) f(r_0) \\
&\quad \times \prod_{t=1}^{T} f(d_t|r_t) f(r_t|r_{t-1}) f(r_{T+1}|r_T) \\
&= f(r_0|d) \prod_{t=1}^{T} f(r_t|r_{t-1}, d_{t:T}) f(r_{T+1}|r_T)
\end{aligned} \tag{5}
$$

which is a non-stationary Markov chain for the hidden Markov model with a likelihood model in factored form as defined above [29]. Assessing such a posterior distribution is usually difficult as both the normalizing constant and the conditional transition matrices are challenging to calculate.

3.2 Kalman Type Models

The current study is limited to Kalman type models. They comprise an initial and a process part.

Initial Distribution

   The initial distribution is identical to the initial distribution of the prior model $f(r_0)$, and as such captures the characteristics of the initial state of the process. Two model classes are later discussed: the Gaussian and the selection-Gaussian classes.

Process Model

   The process model includes the forward model and likelihood models defined in **Section 3.1**. It thus characterizes the process dynamics and the observation acquisition procedure. The forward model is defined by,

$$
\begin{aligned}
[r_{t+1}|r_t] &= A_t r_t + \varepsilon_t^r \\
f(r_{t+1}|r_t) &= \varphi_n\left(r_{t+1}; A_t r_t, \Sigma_t^{r|r}\right)
\end{aligned} \tag{6}
$$

with forward $(n \times n)$-matrix A_t and n-vector error term ε_t^r defined as centered Gaussian with covariance $(n \times n)$-matrix $\Sigma_t^{r|r}$. The forward model is therefore Gauss-linear. The likelihood component is defined by,
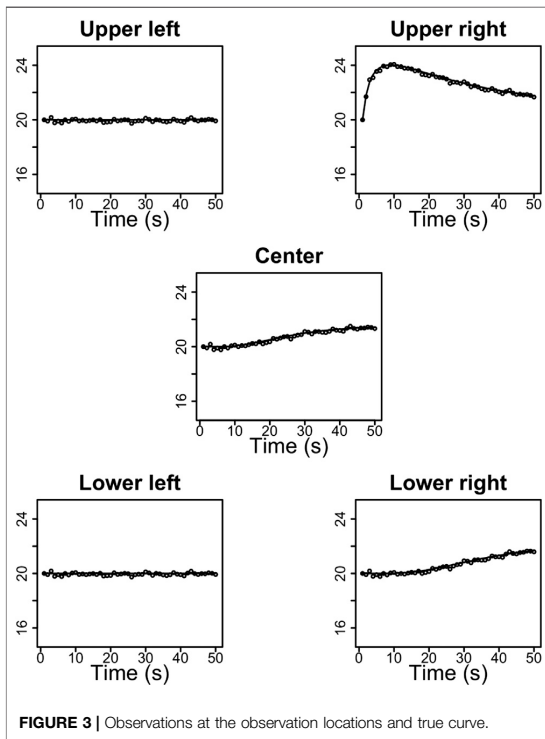
FIGURE 3 | Observations at the observation locations and true curve.

$$[d_t|r_t] = Hr_t + \varepsilon_t^d$$
$$f(d_t|r_t) = \varphi_p\left(d_t; Hr_t, \Sigma_t^{d|r}\right) \tag{7}$$

with the observation $(m \times n)$-matrix H and the m-vector error term ε_t^d defined as centered Gaussian with covariance $(m \times m)$-matrix $\Sigma_t^{d|r}$. The likelihood model is also Gauss-linear. This process model coincides with the frequently used traditional Kalman model [3].

3.3 Traditional Kalman Model

The traditional Kalman model is defined by letting the initial distribution be in the class of Gaussian pdfs,

$$r_0 \sim f(r_0) = \varphi_n\left(r_0; \mu_0^r, \Sigma_0^r\right) \tag{8}$$

with initial expectation n-vector μ_0^r and positive definite covariance $(n \times n)$-matrix Σ_0^r. The Gaussian initial distribution is parametrized by $\Theta^G = (\mu_0^r, \Sigma_0^r)$. In our spatial study, this initial distribution will be a discretized stationary Gaussian RF. The process model is Gauss-linear and identical to the traditional Kalman type.

This traditional Kalman model is analytically tractable. The posterior distribution $f(r|d)$ is Gaussian and the posterior distribution parameters can be calculated by algebraic operations on the parameters of the initial distribution, process model and the observed data. Therefore the assessment of the posterior distribution does not require computationally demanding integrals. The analytical tractability follows from the recursive reproduction of Gaussian pdfs:

- The initial model $f(r_0)$ is Gaussian and the likelihood model $f(d_0|r_0)$ is Gauss-linear, hence the joint model $f(r_0, d_0)$ is Gaussian. Consequently, the conditional model $f(r_0|d_0)$ is Gaussian.
- The conditional model $f(r_0|d_0)$ is Gaussian and the dynamic model $f(r_1|r_0)$ is Gauss-linear, hence the joint conditional model $f(r_1, r_0|d_0)$ is Gaussian.

By recursion, we obtain that $f(r|d) = f(r_0, \ldots, r_{T+1}|d_0, \ldots, d_T)$ is Gaussian. Note in particular that since $f(r|d)$ is Gaussian, so is $f(r_0|d)$. This pdf is obtained by marginalization of $f(r|d)$ which, for the Gaussian case, amounts to removing rows from the expectation vector and rows and columns from the covariance matrix. Additionally, the joint pdf $f(r, d)$ can be assessed using a simple recursive algorithm, see **Supplementary Appendix Algorithm A1** in **Supplementary Appendix A**.

From the joint Gaussian pdf $f(r, d)$, the posterior distribution $f(r|d)$ can be analytically assessed. In spatial models, the grid dimension n may be large while the number of data collection sites m usually is small. **Supplementary Appendix Algorithm A1** requires storing the covariance $[n(T + 2) + m(T + 1)] \times [n(T + 2) + m(T + 1)]$-matrix of the Gaussian vector $[r, d]$ which is hardly ever necessary in practice where the target distribution is clearly identified. Only the spatial variables of interest need to be stored, which entails that only the covariance $[n + m(T + 1)] \times [n + m(T + 1)]$-matrix of $[r_0, d]$ need to be stored in our spatio-temporal inversion study.


FIGURE 4 | Graph of the hidden Markov model.

3.4 Selection Kalman Model

The selection Kalman model is defined by letting the initial distribution be in the class of selection-Gaussian pdfs [27, 28]. This class is defined by considering an auxilary n-vector $\tilde{r}$ with pdf from the Gaussian class,

$$f(\tilde{r}) = \varphi_n(\tilde{r}; \mu_{\tilde{r}}, \Sigma_{\tilde{r}}) \tag{9}$$

with expectation n-vector $\mu_{\tilde{r}}$ and covariance $(n \times n)$-matrix $\Sigma_{\tilde{r}}$. In our spatial study this pdf will represent a discretized stationary Gaussian RF. Define further an auxiliary q-vector ν by a Gauss-linear extension,

$$[\nu|\tilde{r}] = \mu_\nu + \Gamma_{\nu|\tilde{r}}(\tilde{r} - \mu_{\tilde{r}}) + \varepsilon_{\nu|\tilde{r}} \tag{10}$$

with the expectation q-vector μ_ν, and the regression $(q \times n)$-matrix $\Gamma_{\nu|\tilde{r}}$ and the centered Gaussian q-vector $\varepsilon_{\nu|\tilde{r}}$, independent of $\tilde{r}$, with covariance $(q \times q)$-matrix $\Sigma_{\nu|\tilde{r}}$. In the current spatial study the dimension of $\tilde{r}$ and ν will be identical. Generally, we have,

$$f(\nu|\tilde{r}) = \varphi_q(\nu; \mu_{\nu|\tilde{r}}, \Sigma_{\nu|\tilde{r}}) \tag{11}$$

with $\mu_{\nu|\tilde{r}} = \mu_\nu + \Gamma_{\nu|\tilde{r}}(\tilde{r} - \mu_{\tilde{r}})$. As a consequence, $[\tilde{r}, \nu]$ is jointly Gaussian,

$$\begin{bmatrix} \tilde{r} \\ \nu \end{bmatrix} \sim f(\tilde{r}, \nu) = \varphi_{n+q}\left(\begin{bmatrix} \tilde{r} \\ \nu \end{bmatrix}; \begin{bmatrix} \mu_{\tilde{r}} \\ \mu_\nu \end{bmatrix}, \begin{bmatrix} \Sigma_{\tilde{r}} & \Sigma_{\tilde{r}}\Gamma_{\nu|\tilde{r}}^T \\ \Gamma_{\nu|\tilde{r}}\Sigma_{\tilde{r}} & \Sigma_\nu \end{bmatrix} \right) \tag{12}$$

with the covariance $(q \times q)$-matrix $\Sigma_\nu = \Gamma_{\nu|\tilde{r}}\Sigma_{\tilde{r}}\Gamma_{\nu|\tilde{r}}^T + \Sigma_{\nu|\tilde{r}}$. Define a selection subset $A \subset \mathbb{R}^q$, and define the class of selection-Gaussian pdfs by $r_A = [\tilde{r}|\nu \in A]$. In the current spatial study the set A will be separable in $\mathbb{R}^q$. Generally, it follows that,

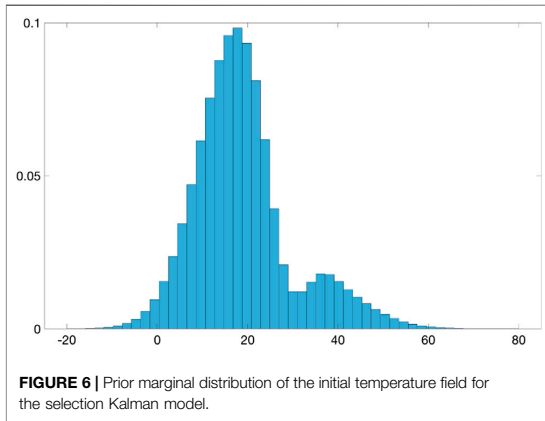$$\begin{aligned} f(r_A) &= f(\tilde{r}|\nu \in A) \\ &= [\Phi_q(A; \mu_\nu, \Sigma_\nu)]^{-1} \\ &\quad \times \Phi_q(A; \mu_{\nu|\tilde{r}}, \Sigma_{\nu|\tilde{r}}) \times \varphi_n(\tilde{r}; \mu_{\tilde{r}}, \Sigma_{\tilde{r}}) \end{aligned} \tag{13}$$

This class of pdfs is parametrized by $\Theta^{SG} = (\mu_{\tilde{r}}, \Sigma_{\tilde{r}}, \mu_\nu, \Gamma_{\nu|\tilde{r}}, \Sigma_{\nu|\tilde{r}}, A)$ for all valid parameter sets. The class of selection-Gaussian pdfs is very flexible and may represent multi-modality, skewness and peakedness [26].

Four one-dimensional selection-Gaussian pdfs are displayed in **Figure 5** in order to demonstrate the influence of the selection set $A \subset \mathbb{R}$. The bivariate variable $[\tilde{r}, \nu]$ is bi-Gaussian and identical in all displays, while the selection sets are marked as solid gray bars along the vertical ν-axis. **Figure 5A** contains a selection set comprised of two segments symmetric about the expectation of ν, making the selection-Gaussian pdf along the horizontal axis bimodal and symmetric. **Figure 5B** contains a selection set of two asymmetric segments, making the selection-Gaussian pdf bimodal and asymmetric. **Figure 5C** contains a selection set of three segments symmetric about the expectation of ν, making the selection-Gaussian pdf trimodal and symmetric.
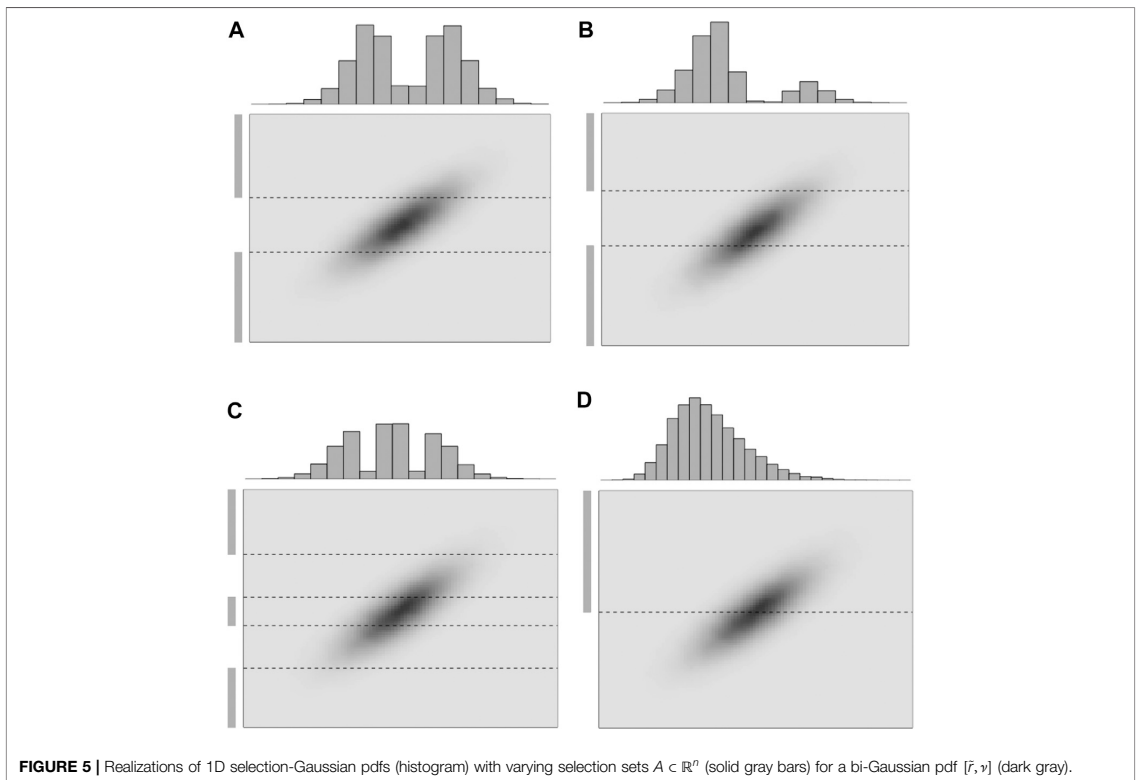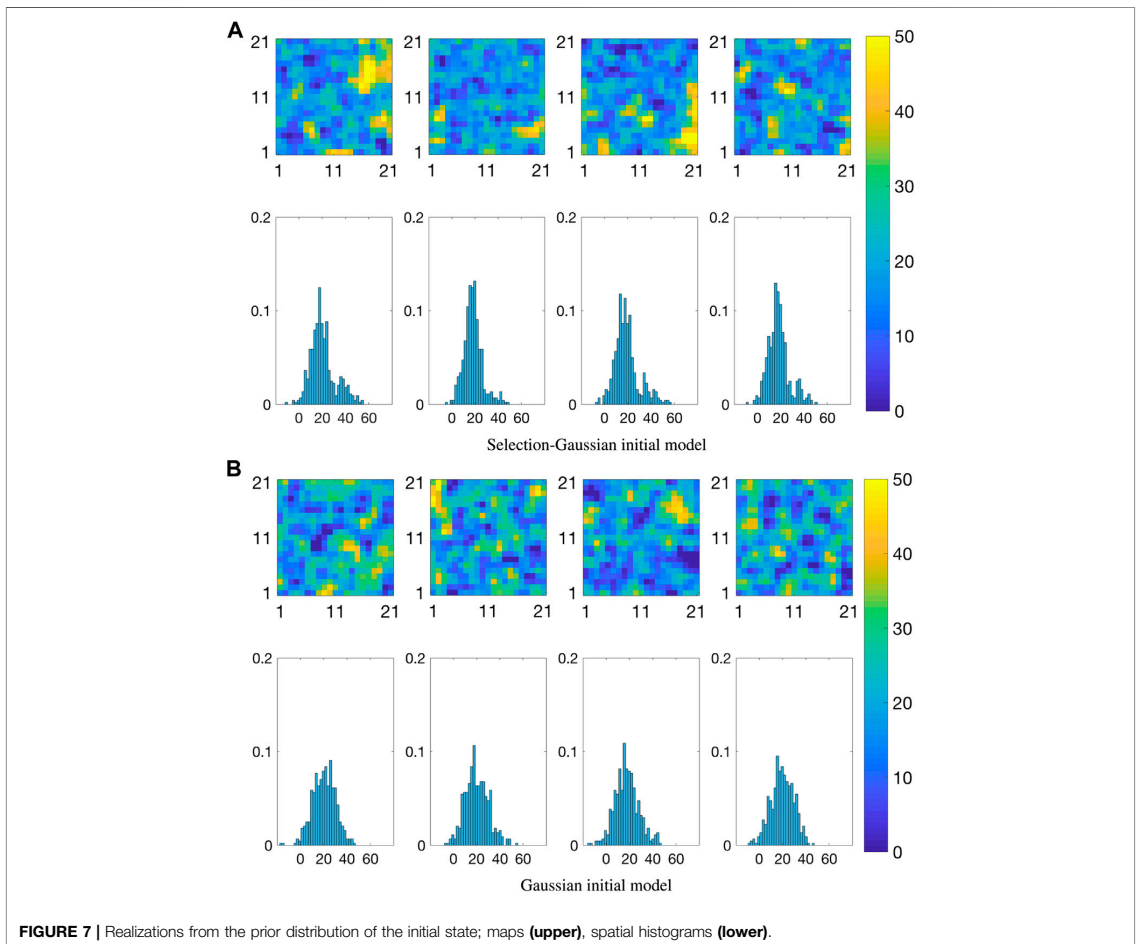

FIGURE 5 | Realizations of 1D selection-Gaussian pdfs (histogram) with varying selection sets $A \subset \mathbb{R}^n$ (solid gray bars) for a bi-Gaussian pdf $[\tilde{r}, \nu]$ (dark gray).

FIGURE 6 | Prior marginal distribution of the initial temperature field for the selection Kalman model.

Lastly, **Figure 5D** contains a selection set comprised of only one segment, making the selection-Gaussian pdf skewed. This selection concept can be extended to higher dimensions and even to discretized spatial models.

Note that assigning a null-matrix to $\Gamma_{\nu|\tilde{r}}$ entails that $f(\tilde{r}, \nu) = f(\tilde{r})f(\nu)$ and selection on ν does not influence $\tilde{r}$. It follows that $f(r_A) = f(\tilde{r})$ is Gaussian. The selection-Gaussian model can therefore be seen as a generalization of the Gaussian one. Assume that the conditional independence, $f(d_t, \nu|\tilde{r}_t) = f(d_t|\tilde{r}_t)f(\nu|\tilde{r}_t)$, holds for all t, it can then be demonstrated [26] that the following recursive reproduction of selection-Gaussian pdfs holds:

- The initial model $f(r_{A,0})$ is selection-Gaussian and the likelihood model $f(d_0|r_{A,0})$ is Gauss-linear, hence the joint model $f(r_{A,0}, d_0)$ is selection-Gaussian. Moreover, the conditional model $f(r_{A,0}|d_0)$ is selection-Gaussian.

FIGURE 7 | Realizations from the prior distribution of the initial state; maps **(upper)**, spatial histograms **(lower)**.

- The conditional model $f(r_{A,0}|d_0)$ is selection-Gaussian and the dynamic model $f(r_{A,1}|r_{A,0})$ is Gauss-linear, the joint conditional model $f(r_{A,1}, r_{0,A}|d_0)$ is therefore selection-Gaussian.

By recursion, we obtain that $f(r_A|d) = f(r_{A,0}, \ldots, r_{A,T+1}|d_0, \ldots, d_T)$ is selection-Gaussian. Recall that these characteristics are similar to those of the class of Gaussian pdfs that make the traditional Kalman model analytically tractable. The selection Kalman model is defined with an initial distribution from the class of selection-Gaussian pdfs and a process model which is Gauss-linear and identical to the traditional Kalman type. From the characteristics of the class of selection-Gaussian distributions, it

follows that the posterior distribution $f(r_A|d)$ is in the class of selection-Gaussian distributions and so is $f(r_{A,0}|d)$.

Consider the augmented $(n + q)$-vector $[\tilde{r}_0, \nu]$ which together with the selection set $A \in \mathbb{R}^q$ defines the initial state $r_{A,0} = [\tilde{r}_0||\nu \in A]$. The recursive algorithm, see **Algorithm 1**, is initiated with this augmented vector which is Gaussian. The conditional independence $f(d_t, \nu|\tilde{r}_t) = f(d_t|\tilde{r}_t)f(\nu|\tilde{r}_t)$ entails that $f(d_t|r_{A,t}) = f(d_t|r_t)$, which is Gauss-linear for all t.

Algorithm 1 provides the Gaussian pdf of the $[n(T+2) + q + m(T+1)]$-vector $[\tilde{r}, \nu, d]$. From the joint Gaussian pdf $f(\tilde{r}, \nu, d)$, the pdf $f(r_{A,0}|d) = f(\tilde{r}_0|\nu \in A, d)$ can be assessed by first marginalizing $\tilde{r}$ and thereafter sequentially

ALGORITHM 1 | *Joint Selection Kalman Model*

- *Define*

$$\mu_t^{\tilde{r}} = \mathbb{E}[\tilde{r}_t]$$
$$\mu_0^{\nu} = \mathbb{E}[\nu]$$
$$\mu_t^{d} = \mathbb{E}[d_t]$$
$$\Sigma_{ts}^{\tilde{r}\tilde{r}} = \text{Cov}(\tilde{r}_t, \tilde{r}_s) = \Sigma_{st}^{\tilde{r}\tilde{r}T}$$
$$\Sigma_{ts}^{d\tilde{d}} = \text{Cov}(d_t, d_s) = \Sigma_{st}^{ddT}$$
$$\Sigma_{00}^{\nu\nu} = \text{Cov}(\nu, \nu)$$
$$\Gamma_{ts}^{\tilde{r}d} = \text{Cov}(\tilde{r}_t, d_s) = \Gamma_{st}^{d\tilde{r}T}$$
$$\Gamma_{t0}^{\tilde{r}\nu} = \text{Cov}(\tilde{r}_t, \nu) = \Gamma_{0t}^{\nu\tilde{r}T}$$
$$\Gamma_{t0}^{d\nu} = \text{Cov}(d_t, \nu) = \Gamma_{0t}^{\nu dT}$$

- *Initiate*

$$\begin{bmatrix} \tilde{r}_0 \\ \nu \end{bmatrix} \sim \varphi_{n+q}\left(\begin{bmatrix} \tilde{r} \\ \nu \end{bmatrix}; \begin{bmatrix} \mu_{\tilde{r}} \\ \mu_{\nu} \end{bmatrix}, \begin{bmatrix} \Sigma_{\tilde{r}} & \Sigma_{\tilde{r}}\Gamma_{\nu|\tilde{r}}^T \\ \Gamma_{\nu|\tilde{r}}\Sigma_{\tilde{r}} & \Sigma_{\nu} \end{bmatrix} \right)$$
$$\mu_0^{\tilde{r}} = \mu_{\tilde{r}}$$
$$\mu_0^{\nu} = \mu_{\nu}$$
$$\Sigma_{00}^{\tilde{r}\tilde{r}} = \Sigma_{\tilde{r}}$$
$$\Sigma_{00}^{\tilde{r}\nu} = \Sigma_{\tilde{r}}\Gamma_{\nu|\tilde{r}}^T$$
$$\Sigma_{00}^{\nu\nu} = \Sigma_{\nu}$$

- *Iterate* $t = 0, \ldots, T$

  Likelihood model:
  $$\mu_t^d = H\mu_t^{\tilde{r}}$$
  $$\Sigma_{tt}^{dd} = H\Sigma_{tt}^{\tilde{r}\tilde{r}}H^T + \Sigma_t^{d|r}$$
  $$\Gamma_{t0}^{d\nu} = H\Gamma_{t0}^{\tilde{r}\nu}$$
  Iterate $s = 0, \ldots, t$
  $$\Gamma_{ts}^{\tilde{r}d} = \Sigma_{ts}^{\tilde{r}\tilde{r}}H^T$$
  End iterate s

  If $t > 0$: Iterate $s = 0, \ldots, t-1$
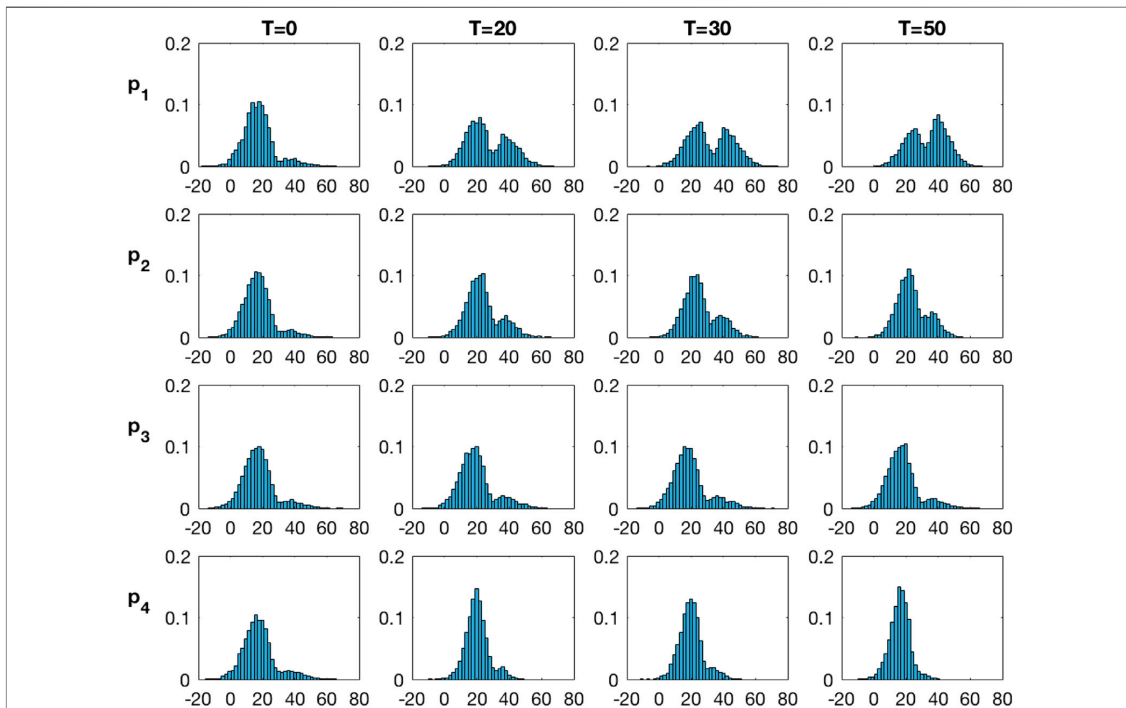 $$\Sigma_{t,s}^{dd} = H\Gamma_{ts}^{\tilde{r}d}$$
  End iterate s

  Forwarding model:
  $$\mu_{t+1}^{\tilde{r}} = A_t\mu_t^{\tilde{r}}$$
  $$\Sigma_{(t+1)(t+1)}^{\tilde{r}\tilde{r}} = A_t\Sigma_{tt}^{\tilde{r}\tilde{r}}A_t^T + \Sigma_t^{r|r}$$
  $$\Gamma_{t+1,0}^{\tilde{r}\nu} = A_t\Gamma_{t0}^{\tilde{r}\nu}$$
  Iterate $s = 0, \ldots, t$
  $$\Sigma_{(t+1)s}^{\tilde{r}\tilde{r}} = A_t\Sigma_{ts}^{\tilde{r}\tilde{r}}$$
  $$\Gamma_{(t+1)s}^{\tilde{r}d} = A_t\Gamma_{ts}^{\tilde{r}d}$$
  End iterate s

- *End iterate t :*
  The Gaussian $(n(T+2) + q + m(T+1))$-vector $[\tilde{r}, \nu, d]$ is then fully assessed by the algorithm.

FIGURE 8 | Marginal pdfs at monitoring locations for increasing current time T from the inversion with the selection Kalman model.

conditioning on d and then on v. The final step, conditioning on $v \in A$, is computer demanding even though v has only dimension q. Rejection sampling is only possible for very low values of q. We therefore use the Metropolis-Hastings algorithm, a McMC method, detailed in [26] and extended from [30]. **Algorithm 1** requires storing the covariance $[n(T+2) + q + m(T+1)] \times [n(T+2) + q + m(T+1)]$ -matrix of the augmented Gaussian vector $[\tilde{r}, v, d]$ which can usually be avoided in practice. Only the spatial variables of interest need to be stored, which entails that only the covariance $[n + q + m(T+1)] \times [n + q + m(T+1)]$-matrix of $[\tilde{r}_0, v, d]$ need to be stored in our spatio-temporal inversion study.

4 SYNTHETIC STUDY

The synthetic study is introduced in **Section 2**, and we discuss it in larger detail in this section.

4.1 Model
Consider a discretized spatio-temporal continuous RF representing the evolution of a temperature field $\{r_t(x), x \in \mathcal{L}_r\}$, $t \in \mathcal{T}_r : \{0, 1, \ldots, T, T+1\}$; $r_t(x) \in \mathbb{R}$, as defined in **Section 2**. The number of spatial grid nodes is $n = 21 \times 21$, while temporal reference T is the current time up to $T = 50$. The discretized spatial field at time t is represented by the n-vector r_t.

The initial temperature field r_0, given in **Figure 1**, is assumed to be unknown. It is divided into two distinct areas: the blue area where the temperature is set at 20°C and the yellow area where the temperature is set at 45°C. Assume that, given the initial temperature field, the field evolves according to the advection-diffusion equation, a linear partial differential equation,

$$\frac{\partial r_t(x)}{\partial t} - \lambda \nabla^2 r_t(x) + c \cdot \nabla r_t(x) = 0 \qquad (14)$$

$$\nabla r_t(x) \cdot n = 0 \qquad (15)$$

with $\lambda \in \mathbb{R}_+$ the known diffusivity coefficient, n the outer normal to the domain and $c = [c_1, c_2]$ the known velocity field. The forward model is defined as,

$$[r_{t+1}|r_t] = A r_t + \varepsilon_t^r \qquad (16)$$

$$f(r_{t+1}|r_t) = \varphi_n(r_{t+1}; A r_t, \Sigma_t^{r|r}) \qquad (17)$$

where the $(n \times n)$-matrix A is obtained by discretizing the advection-diffusion equation using finite differences, see **Supplementary Appendix B** for finite differences scheme and parameter values, while the centered Gaussian n-vector ε_t^r, with covariance $(n \times n)$-matrix $\Sigma_t^{r|r} = 0 \times I_n$ represents model error. Under these assumptions, the forward model is exact which constitutes a limiting case to Gauss-linear models. The evolution of the temperature field is described in **Figure 2**.
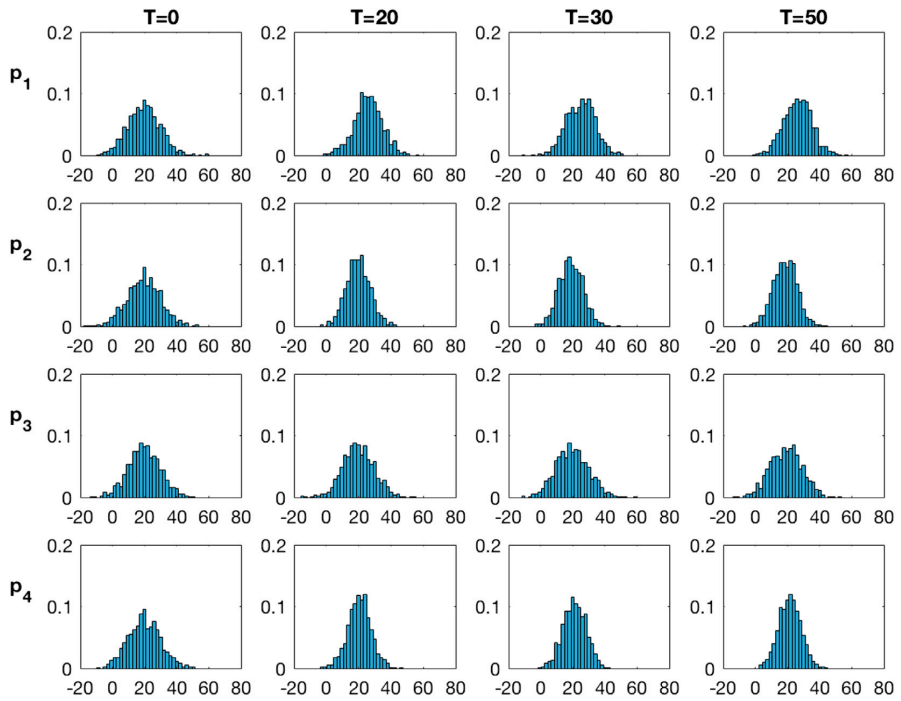
FIGURE 9 | Marginal pdfs at monitoring locations for increasing current time T from the inversion with the traditional Kalman model.
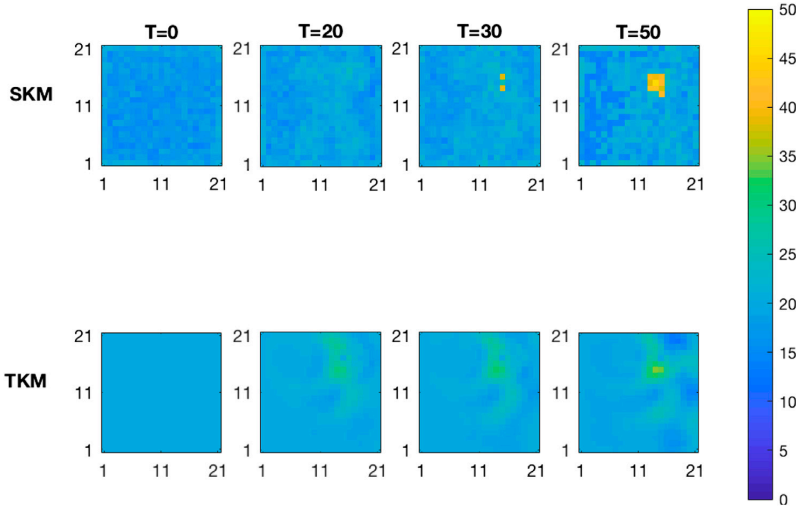
FIGURE 10 | MMAP predictions of the initial state for increasing current time T from the inversion with the selection Kalman model (SKM-upper) and with the traditional Kalman model (TKM-lower).
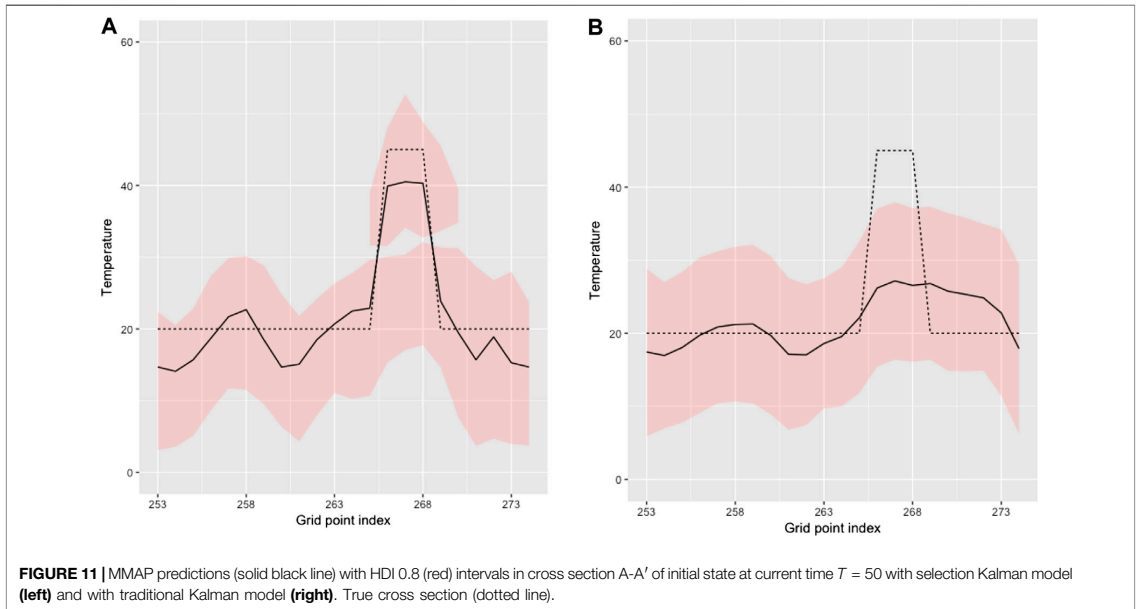
FIGURE 11 | MMAP predictions (solid black line) with HDI 0.8 (red) intervals in cross section A-A′ of initial state at current time $T = 50$ with selection Kalman model **(left)** and with traditional Kalman model **(right)**. True cross section (dotted line).

The observations are acquired at $m = 5$ different locations on the spatial grid $\mathcal{L}_r$ at each temporal node in $\mathcal{T}_d$ providing the set of m-vectors $\{d_t, t \in \mathcal{T}_d\}$. The observation locations are represented with dots in **Figure 1**. The corresponding likelihood model is defined as,

$$[d_t|r_t] = Hr_t + \varepsilon_t^d \tag{18}$$

$$f(d_t|r_t) = \varphi_m(d_t; Hr_t, \Sigma_t^{d|r}) \tag{19}$$

where the observation $(m \times n)$-matrix H is a binary selection matrix, see **Supplementary Appendix B**, while the centered Gaussian m-vector ε_t^d with covariance $(m \times m)$-matrix $\Sigma_t^{d|r} = \sigma_{d|r}^2 \times I_m$ with $\sigma_{d|r} = 0.1$, represents independent observation errors. Under these assumptions, the likelihood model is Gauss-linear. The observations are displayed as time series in **Figure 3**. Note that $\Sigma_t^{r|r}$ and $\Sigma_t^{d|r}$ are in this example constant through time.

The forward and likelihood models are Gauss-linear. In order to fully defined the selection Kalman model and traditional Kalman model, we must specify the prior distribution for the initial temperature field for both approaches.

We assume we know the initial temperature field has large areas with low temperatures in the range $[5°C, 25°C]$ and possibly, smaller areas with high temperatures in the range $[40°C, 55°C]$. The exact location, extent and temperature of these smaller areas are unknown. The prior is therefore spatially stationary in both models.

The prior distribution is set to be selection-Gaussian for the selection Kalman model. Such a prior model can represent multimodality. The model is constructed according to [26] and is defined considering an auxiliary discretized stationary Gaussian RF,

$$f\left(\tilde{r}_0\right) = \varphi_n\left(\tilde{r}_0; \mu_{\tilde{r}} i_n, \sigma_{\tilde{r}}^2 \Sigma_{\tilde{r}}^\rho\right) \tag{20}$$

with expectation and variance levels, $\mu_{\tilde{r}}$ and $\sigma_{\tilde{r}}^2$ respectively. The spatial correlation $(n \times n)$-matrix $\Sigma_{\tilde{r}}^\rho$ is defined by an isotropic second order exponential spatial correlation function $\rho_{\tilde{r}}(\tau) = \exp(-\tau^2/\delta^2); \tau \in \mathbb{R}_+$. Define the auxiliary n-vector ν given $\tilde{r}_0$,

$$[\nu|\tilde{r}_0] = \gamma\left(\tilde{r}_0 - \mu_{\tilde{r}} i_n\right) + \varepsilon^\nu \tag{21}$$

$$f(\nu|\tilde{r}_0) = \varphi_n\left[\nu; \gamma\left(\tilde{r}_0 - \mu_{\tilde{r}} i_n\right), (1 - \gamma^2)I_n\right] \tag{22}$$

$$= \prod_{i=1}^n \varphi_1\left[\nu_i; \gamma\left(\tilde{r}_{0,i} - \mu_{\tilde{r}}\right), (1 - \gamma^2)\right] \tag{23}$$

with coupling parameter $\gamma \in \mathbb{R}_{[-1,1]}$ and centered Gaussian independent n-vector ε^ν with variance $(1 - \gamma^2)$. Note that this pdf is in factored form. Consequently the joint pdf of $[\tilde{r}_0, \nu]$ is,

$$\begin{bmatrix} \tilde{r}_0 \\ \nu \end{bmatrix} \sim \varphi_{2n}\left(\begin{bmatrix} \tilde{r}_0 \\ \nu \end{bmatrix}; \begin{bmatrix} \mu_{\tilde{r}} i_n \\ 0 i_n \end{bmatrix}, \begin{bmatrix} \sigma_{\tilde{r}}^2 \Sigma_{\tilde{r}}^\rho & \sigma_{\tilde{r}}^2 \gamma \Sigma_{\tilde{r}}^\rho \\ \sigma_{\tilde{r}}^2 \gamma \Sigma_{\tilde{r}}^\rho & \sigma_{\tilde{r}}^2 \gamma^2 \Sigma_{\tilde{r}}^\rho + (1 - \gamma^2)I_n \end{bmatrix}\right) \tag{24}$$

Define a separable selection set $A \subset \mathbb{R}^n$ such that $A = B^n, B \subset \mathbb{R}$. Therefore, the prior distribution is represented by the discretized selection-Gaussian RF $r_{A,0}$ defined as,

$$r_{A,0} = [\tilde{r}_0|\nu \in A] \tag{25}$$

$$f(r_{A,0}) = \left[\Phi_n\left(A; 0i_n, \sigma_{\tilde{r}}^2 \gamma^2 \Sigma_{\tilde{r}}^\rho + (1 - \gamma^2)I_n\right)\right]^{-1}$$

$$\times \prod_{i=1}^n \Phi_1\left[A_i; \gamma\left(\tilde{r}_i - \mu_{\tilde{r}}\right), (1 - \gamma^2)\right] \tag{26}$$

$$\times \varphi_n\left(r_{A,0}; \mu_r i_n, \sigma_{\tilde{r}}^2 \Sigma_{\tilde{r}}^\rho\right)$$

Note that after selection on the auxiliary variable ν is made, the expectation and variance of the resulting $r_{A,0}$ will no longer be $\mu_{\tilde{r}} i_n$ and $\sigma_{\tilde{r}}^2 \Sigma_{\tilde{r}}^\rho$.
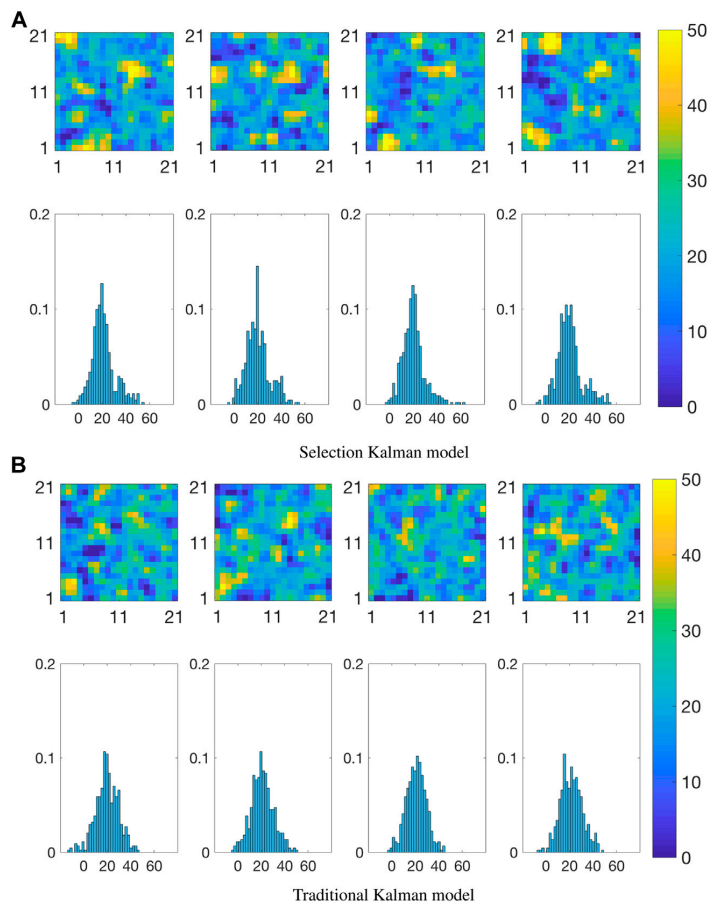
FIGURE 12 | Realizations from the posterior distribution of the initial state at current time $T = 50$.

The parameters values are listed in **Table 1** and they are chosen to reflect what is assumed about the initial temperature field. The prior marginal distribution is bimodal, the dominant mode is centered about 18°C while the smaller mode is centered about 40°C as shown in **Figure 6**. The spread of the dominant mode covers the assumed temperature range for the low temperature areas while the spread of the smaller mode covers the assumed range for the high temperature areas. Realizations from the prior distribution and associated spatial histograms are shown in **Figure 7A**. They exhibit large areas at low temperatures and smaller areas at higher temperatures. Similarly to the marginal distribution, the spatial histograms cover the assumed range for high and low temperature areas.

The prior distribution for the traditional Kalman model is Gaussian and is defined as,

$$f(r_0) = \varphi_n(r_0; \mu_r i_n, \sigma_r^2 \Sigma_r^\rho) \qquad (27)$$

with expectation and variance levels, μ_r and σ_r^2, respectively and spatial correlation $(n \times n)$-matrix Σ_r^ρ defined by a second order spatial correlation function $\rho_r(\tau) = \exp(-\tau^2/\delta^2); \tau \in \mathbb{R}_+$. The parameter values are listed in **Table 1**.

Figure 7B displays four realizations with associated spatial histograms from the prior distribution for the traditional Kalman model. The mean and variance levels are chosen so that the prior covers the assumed range for the high and low temperature areas as can be seen in the spatial histograms.

TABLE 1 | Parameter values for the selection Gaussian initial model (SKM) and the Gaussian initial model (TKM).

SKM	$\mu_{\tilde{r}}$	$\sigma_{\tilde{r}}$	δ	Γ	A
	28.75	10	0.15	0.95	$[[-\infty, -0.2] \cup [0.5, +\infty)]^n$
TKM	μ_r	σ_r	δ		
	20	10	0.15		

Figure 7B can be compared to **Figure 7A**, and one observes that only the selection-Gaussian distribution can capture bi-modality in the spatial histograms. In studies with real data, the prior model specification must of course be based on experience with the phenomenon under study.

In the next section, we demonstrate the effect of specifying different initial models in the spatio-temporal inversion model.

4.2 Results

The challenge is to restore r_0 based on the observations $d = \{d_0, \ldots, d_T\}$ by evaluating the posterior distribution in the selection Kalman model $f(r_{A,0}|d_0, \ldots, d_T)$ and in the traditional Kalman filter $f(r_0|d_0, \ldots, d_T)$. We compare the results from these two models that have been specified in the previous section. The posterior distributions are analytically tractable for both the selection Kalman model and the traditional Kalman model. They are calculated using **Algorithm 1** and **Supplementary Appendix Algorithm A1** respectively. In order to evaluate the results, we present various characteristics of the posterior distributions for increasing values of current time T:

1. Marginal pdfs at four monitoring locations represented by crosses and numbered $1, 2, 3, 4$ in **Figure 1**,

$$f(r_{A,0,i}|d_{0:T}) = \int f(r_{A,0}|d_{0:T})dr_{A,0,-i} \quad i = 1, \ldots, 4 \qquad (28)$$

and similarly for $f(r_{0,i}|d_{0:T})$ based on $f(r_0|d_{0:T})$. The index $-i$ stands for all the indices in $1, \ldots, n$ but the ith index.

2. Spatial prediction based on a marginal maximum a posteriori (MMAP) criterion,

$$\hat{r}_{A,0} = MMAP\{r_{A,0}|d_{0:T}\} = \{MAP\{r_{A,0,j}|d_{0:T}\}; j = 1, 2, \ldots, n\}$$
$$= \{\arg\max\{f(r_{A,0,j}|d_{0:T})\}, j = 1, 2, \ldots, n\} \qquad (29)$$

and similarly for $\hat{r}_0$ based on $f(r_0|d_{0:T})$. This MMAP criterion is used as the marginal posterior model may be multi-modal. For uni-modal symmetric posterior distributions such as the Gaussian one, the MMAP predictor coincides with the expectation predictor.

3. The MMAP prediction and the associated 0.80 prediction interval along a horizontal profile A-A', see **Figure 1**. The prediction interval is computed as the highest density interval (HDI) [31], which entails that the prediction intervals may consist of several intervals for multimodal posterior pdfs.

4. Realizations From the Posterior Pdfs $f(r_{A,0}|d_{0:T})$ and $f(r_0|d_{0:T})$.

Figure 8 displays the marginal posterior pdfs based on the selection Kalman model at the four monitoring locations,
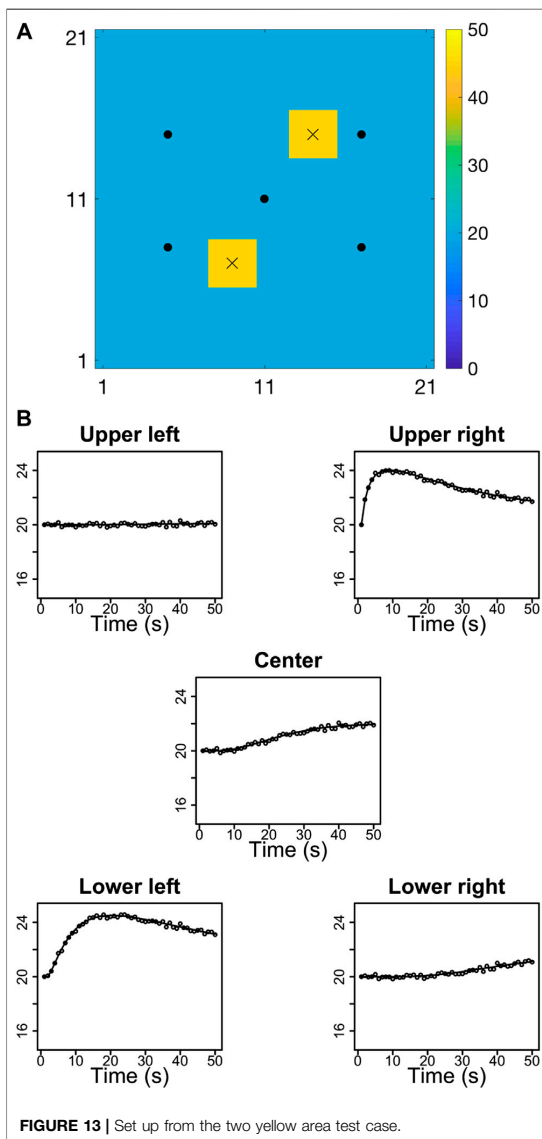

FIGURE 13 | Set up from the two yellow area test case.

vertically, for increasing current time T, horizontally. At current time $T = 0$, all pdfs are virtually identical to the marginal pdf of the stationary initial model. As current time T increases, and the observations are assimilated, one observes substantial differences between the marginal pdfs at the monitoring locations. The height of the high-value mode increases depending on the proximity of monitoring location to the yellow area, as expected. The posterior marginal pdf at observation location one clearly indicates that it lies in the yellow area already at current time $T = 20$ as the high-value mode is increasing. At location two the high-value mode also

FIGURE 14 | Results from the two-yellow-area test case.
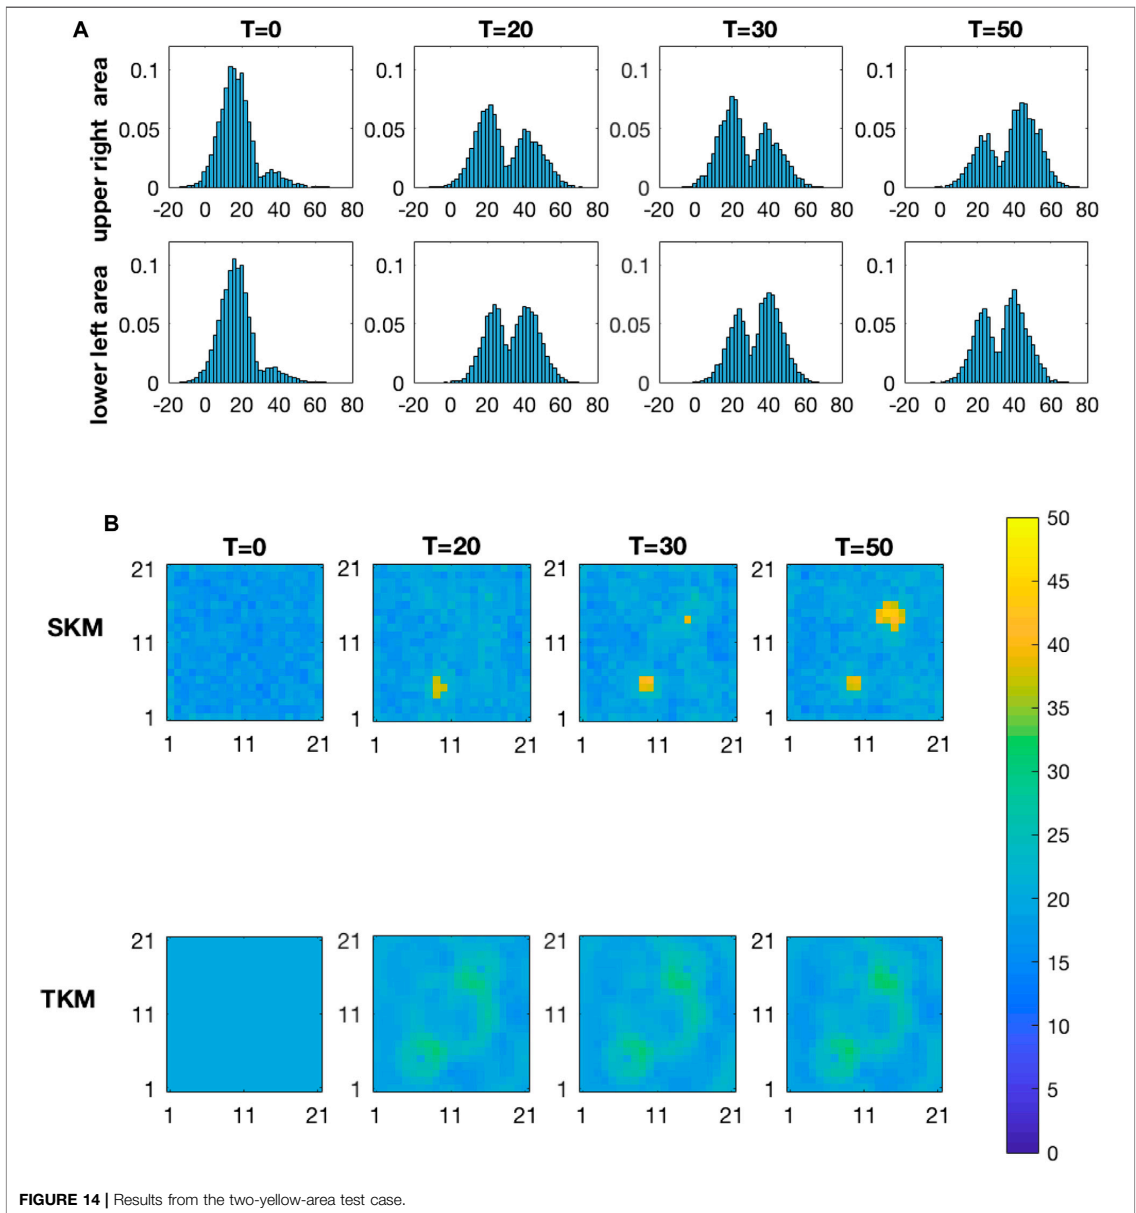
increases somewhat at $T = 20$, but does not increase more thereafter. This monitoring location is outside the yellow area, although fairly close to it. Location three is far from both the yellow area and observation locations and the posterior marginal pdf remains almost identical to the prior model. Lastly, location four is far from the yellow area but close to an observation location at which the observations remain

stationary, hence the low-value mode grows to be completely dominant.

Figure 9 displays the marginal pdfs from the traditional Kalman model. These marginal posterior pdfs are also virtually identical at current time $T = 0$. As current time T increases the marginal pdfs at the monitoring locations are indeed different as they are shifting. However, this shift is

difficult to observe. By using the selection Kalman model, the indications of a yellow area at the correct location can be observed from current time $T = 20$, while one can hardly observe any indications of it if the traditional Kalman model is used.

The upper panels of **Figure 10** display the MMAP spatial prediction based on the selection Kalman model for increasing current time T. At current time $T = 0$, the predictions are virtually constant bar some boundary effect as the initial prior model is stationary. As current time T increases, indications of the yellow appear at $T = 30$, it is however at $T = 50$ that correct location and spatial extent are identified. The prediction value of the yellow area is very close to the correct value of 45. The blue area value is predicted with some variability around the expected 20. The lower panels of **Figure 10** present the corresponding spatial predictions based for the traditional Kalman model. As current time T increases, indications of something occurring in the yellow area appears, but the location is uncertain and the spatial extent only vaguely outlined. Moreover the predicted value in the yellow area is much lower than the correct value 45. The background value is however fairly precisely predicted around the expected 20. The circular features centered about the observation locations that appear on the predictions based on the selection Kalman model in **Figure 10** are not artifacts. These features are also present on the predictions based on the traditional Kalman model, although less prominent.

We evaluate the root mean square error (RMSE) values of the two models at time $T = 50$. The RMSE criterion is used to quantify the difference between the MMAP predictions in **Figure 10** and the truth in **Figure 1**. The RMSE for the selection Kalman model is 2.76 while the RMSE of the traditional Kalman model is 3.33 The selection Kalman model therefore offers a 18% reduction in RMSE compared to the RMSE of the traditional Kalman model.

Figure 11 displays the MMAP predictions with associated 0.80 prediction intervals along the horizontal profile A-A'. The prediction from the selection Kalman model captures the yellow area while the prediction from the traditional Kalman model barely indicates the yellow area. The prediction intervals follow the same pattern. Note, however, that the prediction intervals of the selection Kalman model may appear as two intervals close to the yellow area since the marginal posterior models are bimodal. By using the selection Kalman model, the location, spatial extent and value of the yellow area is very precisely predicted at current time $T = 50$. Predictions based on the traditional Kalman model are less precise and rather blurred.

Figure 12 displays realizations from the posterior pdfs at $T = 50$. For the selection Kalman model, see **Figure 12A**, the yellow area is precisely reproduced in the majority of realizations while for traditional Kalman model, see **Figure 12B**, the yellow area is only vaguely indicated. Note however that the realizations from the selection Kalman model reflect the bimodality of the prior model outside the central area where the five spot observation design provides the most information. These observations are consistent with the results observed in **Figures 8**, **9**.

Conditioning on the observed data takes the same time for both methods but the selection Kalman model requires sampling from a high dimensional truncated Gaussian pdf in order to evaluate the posterior distribution which means that the computational demand for the selection Kalman model is higher than that of the traditional Kalman model. For $n = 441$, as in our study, it takes an average of $7.4s$ to generate 100 realizations from a selection-Gaussian on our laptop computer, so a little over 12 min to generate the 10,000 realizations used to estimate the MMAP for the selection Kalman model. Note that the sampling becomes increasingly more resource consuming as the grid dimension increases and the computational time can be reduced by introducing parallelization in the algorithm.

To demonstrate the generality of the selection Kalman model, we define an alternative true initial state with two yellow areas, see **Figure 13A**. We used the same model parameters as in the primary case. The prior distribution for both the selection Kalman model and the traditional Kalman model are identical to the first case. Note in particular that the number of yellow areas is not specified. The observed time series will of course be different, see **Figure 13B**. These time series have many similarities with the ones from the primary case. We inspect the marginal pdfs at two monitoring locations, one inside each yellow area, as they evolve with current time T, see **Figure 14A**. Both marginal pdfs are identical at current time T = 0, and as current time T increases the height of the high-value mode increases, indicating that both monitoring locations are within the yellow areas. In **Figure 14B** the corresponding MMAP predictions are displayed for increasing current time T. We observe that location, areal extent and value of both yellow areas are well reproduced, but not as well as for the first case since identifying two sources is more complicated. The identification challenge is of course increasing with an increasing number of yellow areas. **Figure 14B** also displays the MMAP prediction for the traditional Kalman model for the two-yellow-area case, where location, areal extent and value are hard to evaluate, similarly to the first test case.

5 CONCLUSION

We define a selection Kalman model based on a selection-Gaussian initial distribution and Gauss-linear dynamic and observation models. This model may represent spatial phenomena with initial states with spatial histograms that are multimodal, skewed and/or peaked. The selection Kalman model is demonstrated to be contained in the class of selection-Gaussian distributions and hence analytically tractable. The analytical tractability makes the assessment of the selection-Gaussian posterior model fast and reliable. Moreover, an efficient recursive algorithm for assessing the selection Kalman model is specified. Note that the traditional Kalman model is a special case of the selection Kalman model, hence the latter can be seen as a generalization of the former.

A synthetic spatio-temporal inversion case study with Gauss-linear forward and observation models is used to demonstrate the characteristics of the methodology. We specify both a selection Kalman model and a traditional Kalman model and evaluate their ability to restore the initial state based on the observed time series. The time series are noisy observations of the variable of interest collected at

a set of sites. The selection Kalman model outperforms the traditional Kalman model. The former model identifies location, areal extent and value of the yellow area very reliably. The traditional Kalman model only provides blurry indications with severe under-prediction of the yellow area. We conclude that for spatio-temporal inversion where the initial spatial state has bimodal or multimodal spatial histograms, the selection Kalman model is far more suitable than the traditional Kalman model.

The selection Kalman model has potential applications far beyond the simple case evaluated in this case study. For all spatio-temporal problems where multimodal spatial histograms appear, the selection Kalman model should be considered. The model can easily be extended to a selection extended Kalman model, along the lines of the extended Kalman model. A more challenging and interesting extension is the definition of a selection ensemble Kalman model including non-linear dynamic and observation models. Research along these lines is currently taking place.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/**Supplementary Material**, further inquiries can be directed to the corresponding author.

REFERENCES

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fams.2021.636524/full#supplementary-material.

1. Handcock MS, Wallis JR. An approach to statistical spatial-temporal modeling of meteorological fields. *J Am Stat Assoc* (1994). 89:368–78. doi:10.1080/01621459.1994.10476754

2. Cappé O, Moulines E, Ryden T. *Inference in hidden Markov models (Springer series in statistics)*. Berlin, Heidelberg: Springer-Verlag (2005).

3. Kalman RE. A new approach to linear filtering and prediction problems. *J Basic Eng* (1960). 82:35–45. doi:10.1115/1.3662552

4. Emerick AA, Reynolds AC. Ensemble smoother with multiple data assimilation. *Comput Geosci* (2013). 55:3–15. doi:10.1016/j.cageo.2012.03.011

5. Evensen G, Raanes PN, Stordal AS, Hove J. Efficient implementation of an iterative ensemble smoother for data assimilation and reservoir history matching. *Front Appl Math Stat* (2019). 5:47. doi:10.3389/fams.2019.00047

6. Liu M, Grana D. Time-lapse seismic history matching with iterative ensemble smoother and deep convolutional autoencoder. *Geophysics* (2019). 85:1–63. doi:10.1190/geo2019-0019.1

7. Hopke PK. Review of receptor modeling methods for source apportionment. *J Air Waste Manag Assoc* (2016). 66:237–59. doi:10.1080/10962247.2016.1140693

8. Cheng MD, Lin CJ. Receptor modeling for smoke of 1998 biomass burning in Central America. *J Geophys Res Atmos* (2001). 106:22871–86. doi:10.1029/2001JD900024

9. Kjeldsen P, Bjerg PL, Rügge K, Christensen TH, Pedersen JK. Characterization of an old municipal landfill (Grindsted, Denmark) as a groundwater pollution source: landfill hydrology and leachate migration. *Waste Manage Res* (1998). 16:14–22. doi:10.1177/0734242X9801600103

10. Wikle CK, Cressie N. A dimension-reduced approach to space-time Kalman filtering. *Biometrika* (1999). 86:815–29. doi:10.1093/biomet/86.4.815

11. Todescato M, Carron A, Carli R, Pillonetto G, Schenato L. Efficient spatio-temporal Gaussian regression via Kalman filtering. *Automatica* (2020). 118:109032. doi:10.1016/j.automatica.2020.109032

12. Jazwinski AH. *Stochastic processes and filtering theory*. New York, NY: Academic Press (1970).

13. Julier SJ, Uhlmann JK. New extension of the Kalman filter to nonlinear systems. In: I Kadar, editor. *Signal processing, sensor fusion, and target recognition VI (SPIE)*, 3068 (1997). p. 182–93. doi:10.1117/12.280797

14. Evensen G. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J Geophys Res* (1994). 99:101–43. doi:10.1029/94JC00572

15. Evensen G, van Leeuwen PJ. An ensemble Kalman smoother for nonlinear dynamics. *Monthly Weather Rev* (2000). 128:1852–67. doi:10.1175/1520-0493(2000)128<1852:AEKSFN>2.0.CO;2

16. Li R, Prasad V, Huang B. Gaussian mixture model-based ensemble Kalman filtering for state and parameter estimation for a PMMA process. *Processes* (2016). 4:9. doi:10.3390/pr4020009

17. Katzfuss M, Stroud JR, Wikle CK. Ensemble Kalman methods for high-dimensional hierarchical dynamic space-time models. *J Am Stat Assoc* (2020). 115:866–85. doi:10.1080/01621459.2019.1592753

18. Robert C, Casella G. *Monte Carlo statistical methods. Springer Texts in Statistics*. New York, NY: Springer (2005).

19. Naveau P, Genton MG, Shen X. A skewed Kalman filter. *J Multivariate Anal* (2005). 94:382–400. doi:10.1016/j.jmva.2004.06.002

20. Ackerson G, Fu K. On state estimation in switching environments. *IEEE Trans Automatic Control* (1970). 15:10–7. doi:10.1109/TAC.1970.1099359

21. Chen R, Liu JS. Mixture Kalman filters. *J R Stat Soc Ser B (Statistical Methodology)* (2000). 62:493–508. doi:10.1111/1467-9868.00246

22. Smith KW. Cluster ensemble kalman filter. *Tellus A* (2007). 59:749–57. doi:10.1111/j.1600-0870.2007.00246.x

23. Dovera L, Della Rossa E. Multimodal ensemble Kalman filtering using Gaussian mixture models. *Comput Geosciences* (2011). 15:307–23. doi:10.1007/s10596-010-9205-3

24. Bengtsson T, Snyder C, Nychka D. Toward a nonlinear ensemble filter for high-dimensional systems. *J Geophys Res Atmos* (2003). 108. doi:10.1029/2002JD002900

25. Ulvmoen M, Omre H. Improved resolution in Bayesian lithology/fluid inversion from prestack seismic data and well observations: Part 1—methodology. *Geophysics* (2010). 75(2):R21–R35. doi:10.1190/1.3294570

26. Omre H, Rimstad K. Bayesian spatial inversion and conjugate selection Gaussian prior models. *Methodology*. Preprint: arXiv:1812.01882 (2018).

27. Arellano-Valle RB, Branco MD, Genton MG. A unified view on skewed distributions arising from selections. *Can J Stat* (2006). 34:581–601. doi:10.1002/cjs.5550340403

28. Arellano-Valle R, del Pino G. From symmetric to asymmetric distributions: a unified approach. In: M Genton, editor. *Skew-elliptical distributions and their applications: a journey beyond normality.* New York, NY: Chapman and Hall/CRC) (2004). p. 113–33.

29. Moja S, Asfaw Z, Omre H. Bayesian inversion in hidden Markov models with varying marginal proportions. *Math Geosci* (2018). 51:463–84. doi:10.1007/s11004-018-9752-z

30. Robert CP. Simulation of truncated normal variables. *Stat Comput* (1995). 121–5. doi:10.1007/BF00143942

31. Hyndman RJ. Computing and graphing highest density regions. *The Am Statistician* (1996). 50:120–6. doi:10.2307/2684423

Supplementary Material

1   RECURSIVE ALGORITHM FOR TRADITIONAL KALMAN MODEL

ALGORITHM 1.1. *Joint Traditional Kalman Model*

- *Define*

$$\boldsymbol{\mu}_t^r = \mathbb{E}[\boldsymbol{r}_t]$$

$$\boldsymbol{\mu}_t^d = \mathbb{E}[\boldsymbol{d}_t]$$

$$\boldsymbol{\Sigma}_{ts}^{rr} = \mathbf{Cov}(\boldsymbol{r}_t, \boldsymbol{r}_s) = \boldsymbol{\Sigma}_{st}^{rr\,T}$$

$$\boldsymbol{\Sigma}_{ts}^{dd} = \mathbf{Cov}(\boldsymbol{d}_t, \boldsymbol{d}_s) = \boldsymbol{\Sigma}_{st}^{dd\,T}$$

$$\boldsymbol{\Gamma}_{ts}^{rd} = \mathbf{Cov}(\boldsymbol{r}_t, \boldsymbol{d}_s) = \boldsymbol{\Gamma}_{st}^{dr\,T}$$

- *Initiate*

$$\boldsymbol{r}_0 \sim \varphi_n(\boldsymbol{r}_0; \boldsymbol{\mu}_0^r, \boldsymbol{\Sigma}_0^r)$$

$$\boldsymbol{\mu}_0^r = \boldsymbol{\mu}_0^r$$

$$\boldsymbol{\Sigma}_{00}^{rr} = \boldsymbol{\Sigma}_0^r$$

- *Iterate* $t = 0, ..., T$

  Likelihood model:

  $$\boldsymbol{\mu}_t^d = \boldsymbol{H}\boldsymbol{\mu}_t^r$$

  $$\boldsymbol{\Sigma}_{tt}^{dd} = \boldsymbol{H}\boldsymbol{\Sigma}_{tt}^{rr}\boldsymbol{H}^T + \boldsymbol{\Sigma}_t^{d|r}$$

  Iterate $s = 0, ..., t$

  $$\boldsymbol{\Gamma}_{ts}^{rd} = \boldsymbol{\Sigma}_{ts}^{rr}\boldsymbol{H}^T$$

  End iterate s

  If $t > 0$: *Iterate* $s = 0, ..., t-1$

  $$\boldsymbol{\Sigma}_{ts}^{dd} = \boldsymbol{H}\boldsymbol{\Gamma}_{ts}^{rd}$$

  End iterate s

  Forwarding model:

  $$\boldsymbol{\mu}_{t+1}^r = \boldsymbol{A}_t\boldsymbol{\mu}_t^r$$

  $$\boldsymbol{\Sigma}_{(t+1)(t+1)}^{rr} = \boldsymbol{A}_t\boldsymbol{\Sigma}_{tt}^{rr}\boldsymbol{A}_t^T + \boldsymbol{\Sigma}_t^{r|r}$$

  Iterate $s = 0, ..., t$

  $$\boldsymbol{\Sigma}_{(t+1)s}^{rr} = \boldsymbol{A}_t\boldsymbol{\Sigma}_{ts}^{rr}$$

  $$\boldsymbol{\Gamma}_{(t+1)s}^{rd} = \boldsymbol{A}_t\boldsymbol{\Gamma}_{ts}^{rd}$$

  End iterate s

- *End iterate t*

$$f\left(\begin{bmatrix}\boldsymbol{r}\\\boldsymbol{d}\end{bmatrix}\right) = \varphi_{n(T+2)+m(T+1)}\left(\begin{bmatrix}\boldsymbol{r}\\\boldsymbol{d}\end{bmatrix}; \begin{bmatrix}\boldsymbol{\mu}_r\\\boldsymbol{\mu}_d\end{bmatrix}, \begin{bmatrix}\boldsymbol{\Sigma}^{rr} & \boldsymbol{\Gamma}^{rd}\\\boldsymbol{\Gamma}^{dr} & \boldsymbol{\Sigma}^{dd}\end{bmatrix}\right),$$

is then fully defined by the algorithm.

2 PARAMETERS IN THE DYNAMIC AND LIKELIHOOD MODELS

Forward $(n \times n)$-matrix A is derived from the following finite difference scheme:

$$
\begin{aligned}
r_{i,j}^{t+1} =\, & r_{i,j}^t + \Delta t \Big(-c_2 \frac{r_{i,j+1}^{t+1} - r_{i,j}^{t+1}}{\Delta x} \\
& + \lambda \frac{r_{i+1,j}^{t+1} + r_{i-1,j}^{t+1} + r_{i,j+1}^{t+1} + r_{i,j-1}^{t+1} + r_{i,j}^{t+1} - 4r_{i,j}^{t+1}}{\Delta x^2} \Big),
\end{aligned}
$$

where $\Delta x = 0.1m$, $\Delta t = 0.5s$, $\lambda = 1.43 \times 10^{-2} m^2.s^{-1}$ and $[c_1, c_2] = [0, -0.1]m.s^{-1}$. When (i,j) is on the no-flow boundary, we use

$$
\frac{r_{i,j+1}^{t+1} - r_{i,j-1}^{t+1}}{2\Delta x} = 0,
$$

$$
\frac{r_{i+1,j}^{t+1} - r_{i-1,j}^{t+1}}{2\Delta x} = 0.
$$

Both equations are necessary at the corners, only one on the rest of the boundary. Observation $(m \times n)$-matrix $\boldsymbol{H}$ is a binary selection matrix defined as,

$$
\boldsymbol{H} = \begin{bmatrix}
0 & \dots & 0 & 1 & 0 & \dots & & & & & & \dots & 0 \\
0 & \dots & & \dots & 0 & 1 & 0 & \dots & & & & \dots & 0 \\
0 & \dots & & & & \dots & 0 & 1 & 0 & \dots & & \dots & 0 \\
0 & \dots & & & & & \dots & 0 & 1 & 0 & \dots & \dots & 0 \\
0 & \dots & & & & & & \dots & 0 & 1 & 0 & \dots & 0
\end{bmatrix}.
$$

Data Assimilation in Spatio-Temporal Models with Non-Gaussian Initial States: The Selection Ensemble Kalman Model

Maxime Conjard and Henning Omre

Data Assimilation in Spatio-Temporal Models with Non-Gaussian Initial States—The Selection Ensemble Kalman Model

Maxime Conjard *,† [ID] **and Henning Omre** †

Department of Mathematical Sciences, Norwegian University of Science and Technology,
7491 Trondheim, Norway; henning.omre@ntnu.no
* Correspondence: maxime.conjard@ntnu.no
† These authors contributed equally to this work.

check for
updates

Abstract: Assimilation of spatio-temporal data poses a challenge when allowing non-Gaussian features in the prior distribution. It becomes even more complex with nonlinear forward and likelihood models. The ensemble Kalman model and its many variants have proven resilient when handling nonlinearity. However, owing to the linearized updates, conserving the non-Gaussian features in the posterior distribution remains an issue. When the prior model is chosen in the class of selection-Gaussian distributions, the selection Ensemble Kalman model provides an approach that conserves non-Gaussianity in the posterior distribution. The synthetic case study features the prediction of a parameter field and the inversion of an initial state for the diffusion equation. By using the selection Kalman model, it is possible to represent multimodality in the posterior model while offering a 20 to 30% reduction in root mean square error relative to the traditional ensemble Kalman model.

1. Introduction

Data assimilation of spatio-temporal models is a challenge in many fields of study, including, but not limited to, air pollution mapping, weather forecast, petroleum engineering, and ground water flow assessment. Over the years, methods have been developed to handle increasingly complex problems. It started with the Kalman filter as presented in the seminal publication [1]. The Kalman filter is based on a Gaussian initial model and Gauss-linear forward and observation models. It defined the foundation for data assimilation and is still used in many assimilation studies. The extended Kalman filter (EKF) [2] appeared as a natural methodological extension that allowed for nonlinearity in the Kalman filter framework by linearization. The ensemble Kalman filter (EnKF) [3,4] defined a Monte Carlo approach to the filter and it became popular as it allowed for nonlinearity in the forward and observation models without having to evaluate analytical gradients. The EnKF and its variants have proven to be efficient in solving high-dimensional and nonlinear problems, see [5,6]. In the EnKF, the initial ensemble members represent the initial state which may not have an analytical expression. The forward model then propagates the ensemble members forward in time. Pseudo observations are generated using the observation model. The conditioning of each ensemble member is made with the Kalman weights estimated from the ensemble to give the best linear update. In cases where the initial

model is non-Gaussian, the distribution of the variable of interest conditioned on the data will tend toward Gaussianity as observations are assimilated due to the linear assimilation rule.

Non-Gaussian initial distributions may be conserved by using a univariate transform into Gaussian marginals while assuming multi-Gaussianity in the transformed space. A univariate back transform is then used to return to the original space. This approach has a long history in traditional statistics, geostatistics, and more recently in ensemble methods for data assimilation, which is referred to as copulas [7], normal score transform [8], and Gaussian anamorphosis [9], respectively. The latter has shown to improve the performance of the EnKF in many applications [10,11]. There are however some unresolved issues since Gaussian anamorphosis transforms the marginal distributions rather than the full distribution, and the effect on the resulting variables interdependence is uncertain.

The Ensemble Randomized Maximum Likelihood Filter (EnRML) [12] and its close relative the Iterative EnKF (IEnKF) [6] are primarily used to handle nonlinearities in the forward and observation models, but they will also retain certain non-Gaussian features in the filtering distribution. These filters require gradient evaluations to execute the update which can be complicated even if the adjoint state method is used. One alternative is to evaluate the gradient using the ensemble itself [13], but this approach introduces an approximation with unclear consequences, particularly in models with multimodal marginals.

Multimodality in the prior model can be represented using categorical auxiliary variables to construct Gaussian mixture prior models [14–16]. In a spatial setting, these models appear as a combination of Gaussian random fields whose parameters depend on the value taken by the categorical variable, but in order to retain spatial dependence, the categorical variable must also have a spatial dependence. This indicator spatial variable can be modeled as a Markov [17] or truncated pluri-Gaussian [18] random field. For both of these models, there are challenges related to temporal data assimilation, although some encouraging examples have been developed [19].

We define and study an alternative prior model, the selection-Gaussian random field [20,21], which may represent multimodality, skewness, and peakedness. This random field model is conjugate with respect to Gauss-linear forward and observation models, similarly to the Gaussian random field model. The posterior distribution is therefore analytically tractable under these assumptions [22]. For general forward and observation models, ensemble based algorithms along the lines of the EnKF can be designed. Such selection ensemble Kalman algorithms are the focus of this study, and they are evaluated on a couple of examples.

In Section 2, we introduce the selection ensemble Kalman model. It provides a framework for the use of the selection-Gaussian distribution as a prior in data assimilation. This framework is then used for ensemble filtering and smoothing through the selection EnKF (SEnKF) and the selection EnKS (SEnKS) algorithms. In Section 3, a synthetic case study of the diffusion equation, with two distinct test cases, showcases the ability of the proposed approaches to assess a parameter field and the initial state of a dynamic field. Results from the SEnKF and the SEnKS are compared to that of the traditional EnKF and the EnKS, respectively. In Section 4, potential shortcomings are discussed and the results are put into perspective with respect to applicability in more realistic applications. In Section 5, conclusions are presented.

In this paper, $f(y)$ denotes the probability density function (pdf) of a random variable y, $\varphi_n(y; \mu, \Sigma)$ denotes the pdf of the Gaussian n-vector y with expectation n-vector μ and covariance $(n \times n)$-matrix Σ. Furthermore, $\Phi_n(A; \mu, \Sigma)$ denotes the probability of the aforementioned Gaussian n-vector y to be in $A \subset \mathbb{R}^n$. We also use i_n to denote the all-ones n-vector, I_n to denote the identity $(n \times n)$-matrix and $\mathbf{1}(S)$ to denote the indicator function that equals 1 when S is true and 0 otherwise. We consider log-diffusivity to be an adimensional quantity and it will therefore not be given a unit.

2. Materials and Methods

Consider the unknown temporal n-vector r_t for $t \in \mathcal{T}_r : \{0, 1, \ldots, T, T+1\}$. Let $r = \{r_0, r_1, \ldots, r_T, r_{T+1}\}$ denote the variable of interest and let $r_{i:j}$ denote $\{r_i, r_{i+1}, \ldots, r_j\}, \forall (i,j) \in \mathcal{T}_r^2, i \leq j$.

Assume that the temporal *m*-vectors of observations d_t for $t \in \mathcal{T}_d : \{0, 1, \ldots, T\}$ are available, and define $d = \{d_0, d_1, \ldots, d_T\}$ and $d_{i:j} = (d_i, \ldots, d_j)$ accordingly. The model specified hereafter defines a hidden Markov (HM) model [23] as displayed in Figure 1.

Prior model: The prior model on *r* consists of an initial and a forward model,

$$f(\boldsymbol{r}) = f(\boldsymbol{r}_0)f(\boldsymbol{r}_{1:T+1}|\boldsymbol{r}_0), \tag{1}$$

where $f(\boldsymbol{r}_0)$ is the pdf of the initial state and $f(\boldsymbol{r}_{1:T+1}|\boldsymbol{r}_0)$ defines the forward model.

(a) *Initial distribution:* The distribution for the initial state $f(\boldsymbol{r}_0)$ is assumed to be in the class of selection-Gaussian distributions [20,21]. Consider a Gaussian $(n + n)$-vector $[\tilde{\boldsymbol{r}}, \boldsymbol{v}]$,

$$\begin{bmatrix} \tilde{\boldsymbol{r}} \\ \boldsymbol{v} \end{bmatrix} \sim \varphi_{2n} \left(\begin{bmatrix} \tilde{\boldsymbol{r}} \\ \boldsymbol{v} \end{bmatrix} ; \begin{bmatrix} \boldsymbol{\mu}_{\tilde{\boldsymbol{r}}} \\ \boldsymbol{\mu}_{v} \end{bmatrix} , \begin{bmatrix} \boldsymbol{\Sigma}_{\tilde{\boldsymbol{r}}} & \boldsymbol{\Sigma}_{\tilde{\boldsymbol{r}}} \boldsymbol{\Gamma}_{v|\tilde{\boldsymbol{r}}}^{T} \\ \boldsymbol{\Gamma}_{v|\tilde{\boldsymbol{r}}} \boldsymbol{\Sigma}_{\tilde{\boldsymbol{r}}} & \boldsymbol{\Sigma}_{v} \end{bmatrix} \right), \tag{2}$$

with *n*-vectors $\boldsymbol{\mu}_{\tilde{\boldsymbol{r}}}$ and $\boldsymbol{\mu}_{v}$, $(n \times n)$-matrix $\boldsymbol{\Gamma}_{v|\tilde{\boldsymbol{r}}}$, and where $\boldsymbol{\Sigma}_{\tilde{\boldsymbol{r}}}$, $\boldsymbol{\Sigma}_{v}$, and $\boldsymbol{\Sigma}_{v|\tilde{\boldsymbol{r}}}$ are all three covariance $(n \times n)$-matrices with $\boldsymbol{\Sigma}_{v} = \boldsymbol{\Gamma}_{v|\tilde{\boldsymbol{r}}} \boldsymbol{\Sigma}_{\tilde{\boldsymbol{r}}} \boldsymbol{\Gamma}_{v|\tilde{\boldsymbol{r}}}^{T} + \boldsymbol{\Sigma}_{v|\tilde{\boldsymbol{r}}}$. Define a selection set $A \subset \mathbb{R}^n$ of dimension *n* and let $\boldsymbol{r}_0 = [\tilde{\boldsymbol{r}}|\boldsymbol{v} \in A]$; then, $\boldsymbol{r}_0$ is in the class of selection-Gaussian distribution and its pdf is,

$$f(\boldsymbol{r}_0) = [\Phi_n(A; \boldsymbol{\mu}_{v}, \boldsymbol{\Sigma}_{v})]^{-1} \tag{3}$$
$$\times \Phi_n(A; \boldsymbol{\mu}_{v} + \boldsymbol{\Gamma}_{v|\tilde{\boldsymbol{r}}}(\boldsymbol{r}_0 - \boldsymbol{\mu}_{\tilde{\boldsymbol{r}}}), \boldsymbol{\Sigma}_{v|\tilde{\boldsymbol{r}}}) \times \varphi_n(\boldsymbol{r}_0; \boldsymbol{\mu}_{\tilde{\boldsymbol{r}}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{r}}}).$$

Note that the class of Gaussian distributions constitutes a subset of the class of selection-Gaussian distributions with $\boldsymbol{\Gamma}_{v|\tilde{\boldsymbol{r}}} = 0 \times \boldsymbol{I}_n$. The dependence in $[\tilde{\boldsymbol{r}}, \boldsymbol{v}]$ represented by $\boldsymbol{\Gamma}_{v|\tilde{\boldsymbol{r}}}$ and the selection subset A are crucial user-defined parameters with the latter being temporally constant. The selection-Gaussian model may represent multimodal, skewed, and/or peaked marginal distributions, see [21]. In this study, the initial distribution is defined to be a discretized stationary selection-Gaussian random field with parametrization,

$$\begin{aligned} \boldsymbol{\mu}_{\tilde{\boldsymbol{r}}} &= \mu_{\tilde{r}} \boldsymbol{i}_n \\ \boldsymbol{\mu}_{v} &= \mu_{v} \boldsymbol{i}_n \\ \boldsymbol{\Sigma}_{\tilde{\boldsymbol{r}}} &= \sigma_{\tilde{r}}^2 \boldsymbol{\Sigma}_{\tilde{\boldsymbol{r}}}^{\rho} \\ \boldsymbol{\Sigma}_{v} &= \gamma^2 \boldsymbol{\Sigma}_{\tilde{\boldsymbol{r}}}^{\rho} + (1 - \gamma^2) \boldsymbol{I}_n \\ \boldsymbol{\Gamma}_{v|\tilde{\boldsymbol{r}}} &= \gamma \sigma_{\tilde{r}}^{-1} \boldsymbol{I}_n. \end{aligned} \tag{4}$$

For a given spatial correlation $(n \times n)$-matrix $\boldsymbol{\Sigma}_{\tilde{\boldsymbol{r}}}^{\rho}$, a stationary selection-Gaussian random field is fully parametrized by $\boldsymbol{\Theta}^{SG} = (\mu_{\tilde{r}}, \mu_{v}, \sigma_{\tilde{r}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{r}}}^{\rho}, \gamma, A)$. Similarly, a stationary Gaussian random field is parametrized by $\boldsymbol{\Theta}^{G} = (\mu_r, \sigma_r, \boldsymbol{\Sigma}_{r}^{\rho})$.

(b) *Forward model:* The forward model given the initial state $[\boldsymbol{r}_{1:T+1}|\boldsymbol{r}_0]$ is defined as

$$f(\boldsymbol{r}_{1:T+1}|\boldsymbol{r}_0) = \prod_{t=0}^{T} f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t), \tag{5}$$

with

$$[\boldsymbol{r}_{t+1}|\boldsymbol{r}_t] = \omega_t(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^r) \sim f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t), \tag{6}$$

where $\omega_t(\cdot, \cdot) \in \mathbb{R}^n$ is the forward model with random *n*-vector $\boldsymbol{\epsilon}_t^r$, independent and identically distributed (iid) for each *t*. This forward model may be nonlinear, but, since it only involves the variable at the previous time step $\boldsymbol{r}_t$, it defines a first-order Markov chain. Note that $f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t)$ cannot generally be written in closed form.

Figure 1. Graph of the hidden Markov model.

Likelihood model: The likelihood model for $[\boldsymbol{d}|\boldsymbol{r}]$ is defined as conditional independent with single-site response,

$$f(\boldsymbol{d}|\boldsymbol{r}) = \prod_{t=0}^{T} f(\boldsymbol{d}_t|\boldsymbol{r}_t), \tag{7}$$

with

$$[\boldsymbol{d}_t|\boldsymbol{r}_t] = \psi_t(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^d) \sim f(\boldsymbol{d}_t|\boldsymbol{r}_t), \tag{8}$$

where $\psi_t(\cdot, \cdot) \in \mathbb{R}^m$ is the likelihood function with random m-vector $\boldsymbol{\epsilon}_t^d$, iid for each t. Note that $f(\boldsymbol{d}_t|\boldsymbol{r}_t)$ cannot generally be written in closed form.

Posterior model: The posterior model for the HM model in Figure 1 is given by

$$
\begin{aligned}
[\boldsymbol{r}|\boldsymbol{d}] \sim f(\boldsymbol{r}|\boldsymbol{d}) &= \ const \times f(\boldsymbol{d}|\boldsymbol{r})f(\boldsymbol{r}) \\
&= \ const \times f(\boldsymbol{d}_0|\boldsymbol{r}_0)f(\boldsymbol{r}_0)\prod_{t=1}^{T} f(\boldsymbol{d}_t|\boldsymbol{r}_t)f(\boldsymbol{r}_t|\boldsymbol{r}_{t-1})f(\boldsymbol{r}_{T+1}|\boldsymbol{r}_T) \\
&= \ f(\boldsymbol{r}_0|\boldsymbol{d})\prod_{t=1}^{T} f(\boldsymbol{r}_t|\boldsymbol{r}_{t-1}, \boldsymbol{d}_{t:T})f(\boldsymbol{r}_{T+1}|\boldsymbol{r}_T),
\end{aligned} \tag{9}
$$

and is also a Markov chain, see [23,24]. This model is denoted the selection ensemble Kalman model. If the forward and likelihood models are Gauss-linear, the posterior model is also selection-Gaussian and analytically tractable, see [22]. When the forward and/or likelihood models are nonlinear, however, approximate or sampling based assessment of the posterior model must be made. For this purpose, we introduce the selection ensemble Kalman filter (SEnKF) and smoother (SEnKS) in the spirit of the traditional ensemble Kalman model [3].

The traditional EnKF algorithm aims at assessing the forecast pdf $f(\boldsymbol{r}_{T+1}|\boldsymbol{d}_{0:T})$, and it is justified by general HM model recursions, see [23]. The algorithm is initiated by

$$[\boldsymbol{r}_1|\boldsymbol{d}_0] \sim f(\boldsymbol{r}_1|\boldsymbol{d}_0) = \int f(\boldsymbol{r}_1|\boldsymbol{r}_0)[f(\boldsymbol{d}_0)]^{-1}f(\boldsymbol{d}_0|\boldsymbol{r}_0)f(\boldsymbol{r}_0)\mathrm{d}\boldsymbol{r}_0, \tag{10}$$

and utilizes the recursion for $t = 1, \ldots, T$,

$$[\boldsymbol{r}_{t+1}|\boldsymbol{d}_{0:t}] \sim f(\boldsymbol{r}_{t+1}|\boldsymbol{d}_{0:t}) = \int f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t)[f(\boldsymbol{d}_t|\boldsymbol{d}_{0:t-1})]^{-1}f(\boldsymbol{d}_t|\boldsymbol{r}_t)f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t-1})\mathrm{d}\boldsymbol{r}_t. \tag{11}$$

The expressions are represented by an ensemble of realizations, which in each recursion is conditioned using a linearized approximation with Kalman weights estimated from the ensemble. Thereafter, the ensemble is forwarded to the next time step. The SEnKF introduced in this study relies on the same relations as above, but it operates on the augmented $(n + n)$-vector $[\tilde{\boldsymbol{r}}., \boldsymbol{\nu}]$, see Equation (2). Hence, the forward model is defined as

$$\begin{bmatrix} \tilde{\boldsymbol{r}}_{t+1} \\ \boldsymbol{\nu}_{t+1} \end{bmatrix} \bigg| \begin{bmatrix} \tilde{\boldsymbol{r}}_t \\ \boldsymbol{\nu}_t \end{bmatrix} = \begin{bmatrix} \omega_t(\tilde{\boldsymbol{r}}_t, \boldsymbol{\epsilon}_t^r) \\ \boldsymbol{\nu}_t \end{bmatrix}. \tag{12}$$

where the auxiliary n-vector $\boldsymbol{\nu}_t$ is temporally constant.

The likelihood model is defined as

$$\left[\boldsymbol{d}_t \middle| \begin{matrix} \tilde{\boldsymbol{r}}_t \\ \boldsymbol{\nu}_t \end{matrix} \right] = \psi_t(\tilde{\boldsymbol{r}}_t, \epsilon_t^d). \tag{13}$$

The SEnKF algorithm provides an ensemble representation of

$$\left[\begin{matrix} \tilde{\boldsymbol{r}}_{T+1} \\ \boldsymbol{\nu} \end{matrix} \middle| \boldsymbol{d}_{0:T} \right] \sim f(\tilde{\boldsymbol{r}}_{T+1}, \boldsymbol{\nu} | \boldsymbol{d}_{0:T}), \tag{14}$$

and, based on this ensemble, empirical sampling based inference, see [21], is used to obtain the forecast of interest:

$$[\boldsymbol{r}_{T+1} | \boldsymbol{d}_{0:T}] \sim f(\boldsymbol{r}_{T+1} | \boldsymbol{d}_{0:T}) = f(\tilde{\boldsymbol{r}}_{T+1} | \boldsymbol{d}_{0:T}, \boldsymbol{\nu} \in A). \tag{15}$$

The SEnKF algorithm is specified in Algorithm A1 in Appendix A.

The traditional EnKS algorithm aims at evaluating the interpolation pdf $f(\boldsymbol{r}_{0:T} | \boldsymbol{d}_{0:T})$ with corresponding HM model recursions, see [23]. The algorithm is initiated by

$$[\boldsymbol{r}_0 | \boldsymbol{d}_0] \sim f(\boldsymbol{r}_0 | \boldsymbol{d}_0) = [f(\boldsymbol{d}_0)]^{-1} f(\boldsymbol{d}_0 | \boldsymbol{r}_0) f(\boldsymbol{r}_0), \tag{16}$$

and the recursions for $t = 1, \ldots, T$,

$$\begin{aligned} [\boldsymbol{r}_{0:t} | \boldsymbol{d}_{0:t}] \sim \; & f(\boldsymbol{r}_{0:t} | \boldsymbol{d}_{0:t}) \\ = \; & [f(\boldsymbol{d}_t | \boldsymbol{d}_{0:t-1})]^{-1} f(\boldsymbol{d}_t | \boldsymbol{r}_{0:t}, \boldsymbol{d}_{0:t-1}) f(\boldsymbol{r}_t | \boldsymbol{r}_{0:t-1}, \boldsymbol{d}_{0:t-1}) f(\boldsymbol{r}_{0:t-1} | \boldsymbol{d}_{0:t-1}) \\ = \; & [f(\boldsymbol{d}_t | \boldsymbol{d}_{0:t-1})]^{-1} f(\boldsymbol{d}_t | \boldsymbol{r}_t) f(\boldsymbol{r}_t | \boldsymbol{r}_{t-1}) f(\boldsymbol{r}_{0:t-1} | \boldsymbol{d}_{0:t-1}). \end{aligned} \tag{17}$$

The expressions are represented by an ensemble of realizations. Forwarding is made on the ensemble and the conditioning is empirically linearized. Note that the dimension of the model increases very fast, one may therefore only store the interpolation pdf $f(\boldsymbol{r}_s | \boldsymbol{d}_{0:T})$ at the time point s of interest. The SEnKS introduced in this study relies on the relations defined above and uses an extended $(n+n)$-vector $[\tilde{\boldsymbol{r}}, \boldsymbol{\nu}]$ as defined in Equation (2). The forward and likelihood models are identical to those defined for the filter. The SEnKS algorithm provides an ensemble representation of

$$\left[\begin{matrix} \tilde{\boldsymbol{r}}_{0:T} \\ \boldsymbol{\nu} \end{matrix} \middle| \boldsymbol{d}_{0:T} \right] \sim f(\tilde{\boldsymbol{r}}_{0:T}, \boldsymbol{\nu} | \boldsymbol{d}_{0:T}), \tag{18}$$

and by using empirical sampling based inference, see [21], the interpolation of interest is assessed,

$$[\boldsymbol{r}_{0:T} | \boldsymbol{d}_{0:T}] \sim f(\boldsymbol{r}_{0:T} | \boldsymbol{d}_{0:T}) = f(\tilde{\boldsymbol{r}}_{0:T} | \boldsymbol{d}_{0:T}, \boldsymbol{\nu} \in A). \tag{19}$$

The SEnKS algorithm is specified in Algorithm A2 in the Appendix A. Both algorithms, SEnKF and SEnKS, contain empirically linearized conditioning and asymptotic results, when the ensemble size goes to infinity, and are consistent only for Gauss-linear forward and likelihood models. Under these assumptions, the model is analytically tractable; however, see [22]. In spite of this lack of asymptotic consistency for general HM models, the ensemble Kalman scheme has proven surprisingly reliable for high-dimensional, weakly nonlinear models even with very modest ensemble sizes [25].

3. Results

We consider two test cases to illustrate the relevance of the selection ensemble Kalman algorithms presented in Section 2. The model, common to both test cases, is based on the diffusion equation.

The test cases are designed such that it will be opportune to consider bi-modal initial distributions. In the first test case, we compare the SEnKF to the traditional EnKF with a focus on predicting the diffusivity field that contains a high diffusivity channel. In the second test case, we compare the SEnKS to the traditional EnKS with a focus on evaluating the initial temperature field that is divided into two distinct areas where the initial temperature is substantially higher in one than in the other.

3.1. Model

Consider a discretized spatio-temporal random field, $\{r_t(x), x \in \mathcal{L}_r \subset \mathbb{R}^2\}$ where $t \in \mathcal{L}_t$: $\{0, 1, \ldots, T\}$ and $r_t(\cdot) \in \mathbb{R}$ that represents temperature (°C). Let a discretized spatial random field, $\{\lambda(x), x \in \mathcal{L}_r \subset \mathbb{R}^2\}$; with $\lambda(\cdot) \in \mathbb{R}_\oplus$ representing diffusivity ($\mathrm{m}^2\,\mathrm{s}^{-1}$). Let x be the spatial reference on the regular spatial grid $\mathcal{L}_r$ on the domain D, while t is the temporal reference on the regular temporal grid $\mathcal{L}_t$. The number of spatial grid nodes is $n = 21 \times 21$, and they are placed every 10 cm vertically and horizontally. The discretized temperature field at time t may be represented by the n-vector r_t and the diffusivity field by the n-vector λ. Both are assumed to be unknown. Note that the Kalman models are defined on the joint variable $[r_t, \lambda]$.

Assume that, given the initial temperature field, the field evolves according to the diffusion equation:

$$\frac{\partial r_t(x)}{\partial t} - \nabla \cdot (\lambda(x)\nabla r_t(x)) = q \tag{20}$$
$$\nabla r_t(x) \cdot n = 0,$$

with n the outer normal to the domain and q a source term. The expression in Equation (20) is discretized using finite differences and the forward model is defined as

$$[r_{t+1}|r_t, \lambda] = \omega^*(r_t, \lambda), \tag{21}$$

with $\omega^*(\cdot, \cdot) \in \mathbb{R}^n$. Convergence and stability of the numerical method are easily ensured for the finite difference scheme that is used. The initial temperature field r_0 is considered unknown in the test cases.

The forward model is assumed to be perfect in the sense that there is no model error. The forward model in Equation (6) then takes the form,

$$\omega([r_t, \lambda], 0i_n) = \begin{bmatrix} \omega^*(r_t, \lambda) \\ \lambda \end{bmatrix}. \tag{22}$$

This forward model is nonlinear due to the product of r_t and λ in Equation (20). Consequently, the assumption of Gauss-linearity required for both the traditional Kalman model [1] and the selection Kalman model [22] is violated and necessitates ensemble based algorithms.

The observations are acquired in a $m = 5$ location pattern on the spatial grid $\mathcal{L}_r$ at each temporal node in $\mathcal{L}_t$, providing the set of observations m-vectors d_t, $t \in \mathcal{L}_t$. The corresponding likelihood model is defined as

$$[d_t|r_t] = \psi_t(r_t, \epsilon_t^d) = Hr_t + \epsilon_t^d \tag{23}$$
$$f(d_t|r_t) = \varphi_m(d_t; Hr_t, \Sigma_{d|r}),$$

where the observation $(m \times n)$-matrix H is a binary selection matrix, while the centered Gaussian m-vector ϵ_t^d with the covariance $(m \times m)$-matrix $\Sigma_{d|r} = \sigma_{d|r}^2 I_m$, and $\sigma_{d|r}^2 = 0.1$ represents independent observation errors. This likelihood model is in Gauss-linear form.

3.2. Test Case 1: Predicting the Parameter Field

The focus of this test case is to predict the unknown diffusivity field λ based on the observations d. Because diffusivity is constant in time, smoothing and filtering give an identical prediction of the field. However, filtering is preferred because it does not require updating the ensemble at all future times in addition to the previous one, see [26]. The posterior model is evaluated using the SEnKF, see Appendix A and the results are compared to those from the traditional EnKF algorithm.

The true diffusivity n-vector λ is displayed in Figure 2. The diffusivity λ is always positive. To ensure that ensemble updates do not lead to negative diffusivity values, we work on $\log(\lambda)$. The figure shows a channel in which the diffusivity is higher than in the rest of the field. The diffusivity field is formally defined as

$$\lambda(x) = \lambda_1 \mathbf{1}(x \in D_1) + \lambda_2 \mathbf{1}(x \in D_2), \tag{24}$$

where $D_1 \subset D$ is the low diffusivity area and $D_2 \subset D$ is the high diffusivity channel. The parameter values are $\lambda_1 = e^{-12} \text{ m}^2 \text{ s}^{-1}$ and $\lambda_2 = e^{-5} \text{ m}^2 \text{ s}^{-1}$. The true temperature field is initially at $20\,^\circ\text{C}$ and the heat source on the lower border of the high diffusivity channel starts pumping in heat at $T = 0$ at a constant volumetric rate $q = 15 \text{ W m}^{-3}$, see Figure 2. The temporal evoluation of the temperature field, shown in Figure 3, is obtained by solving the diffusion equation in Equation (20) for the log-diffusivity field in Figure 2 and the initial temperature field defined above. The temperature observations d, see Figure 4, are then collected from the five locations shown in Figure 2 using the likelihood model defined in Equation (23). The measurements are taken every second from $T = 0$ to $T = 100$. As the heat from the source diffuses mostly along the high diffusivity channel, the observed temperature increases substantially at observation locations within the channel.

The unknown initial field for log-diffusivity $\log(\lambda)$ is assigned a stationary selection-Gaussian random field prior model with parameters $\Theta_\lambda^{SG} = (\mu_{\tilde{r}}^\lambda, \mu_v^\lambda, \sigma_{\tilde{r}}^\lambda, \Sigma_{\tilde{r}}^\rho, \gamma^\lambda, A)$, see [21] and Equation (2). The parameter values for the prior model are listed in Table 1.
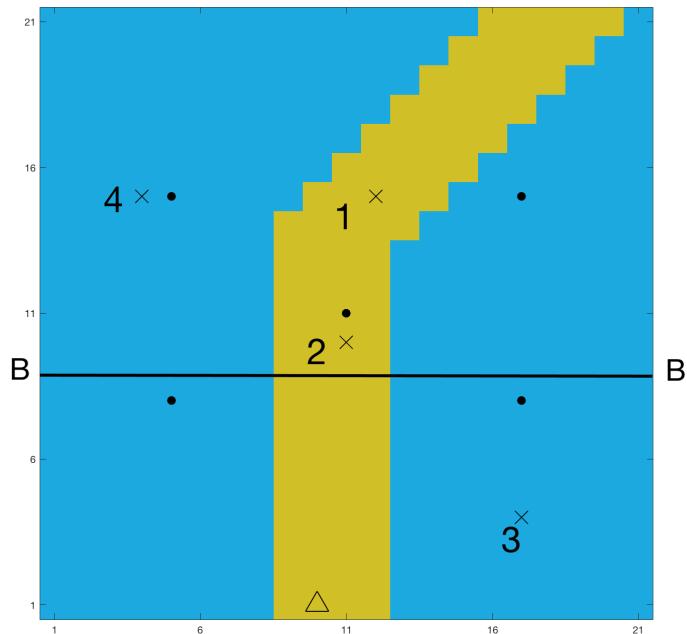

Figure 2. Initial log-diffusivity field with observation locations ·, monitoring locations ×, and heat source △.

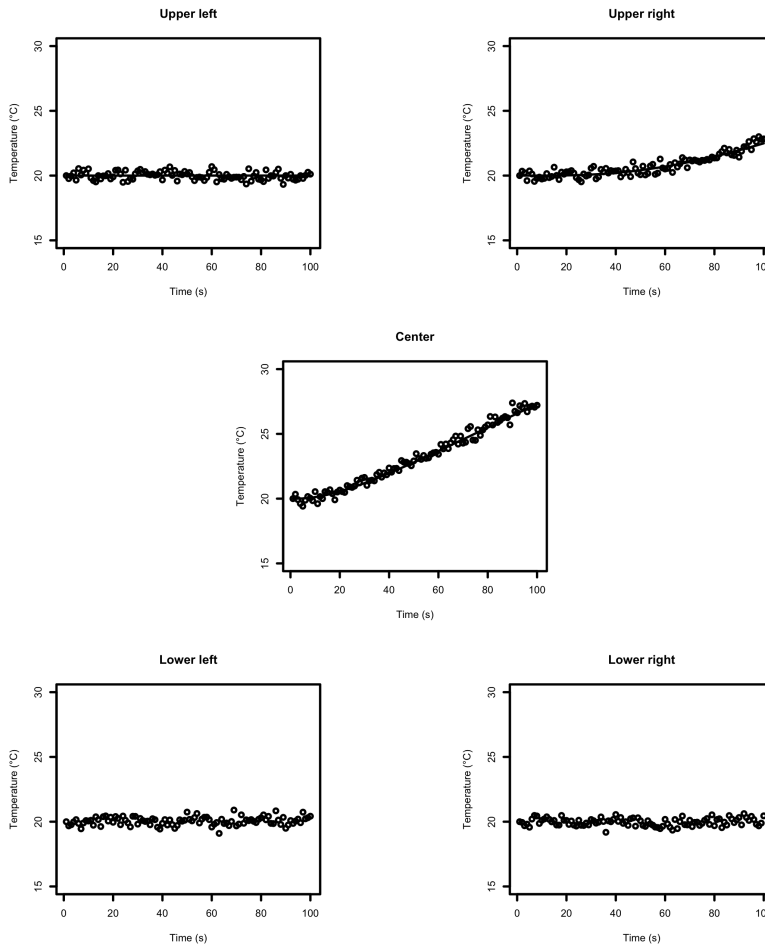**Figure 3.** True temperature (°C) field evolution over time.

Figure 4. Data collected over time (·) and true temperature evolution (line) at the data collection points.

Table 1. Parameter values for the selection-Gaussian initial distribution for the initial log-diffusivity field.

Parameters	Values
$\mu_{\tilde{r}}^{\lambda}$	-8.5
μ_{ν}^{λ}	0
$\sigma_{\tilde{r}}^{\lambda}$	$\sqrt{1.6}$
γ^{λ}	0.9
A	$]-\infty, -0.3] \cup [0.5, +\infty[^{n}$

The unknown initial temperature field r_0 is assigned a stationary Gaussian random field prior model with parameters $\Theta_r^G = (\mu_r, \sigma_r, \Sigma_r^\rho)$ with expectation and variance levels $\mu_r = 20$ and $\sigma_r^2 = 2$, respectively. The variance level is relatively large as we assume little prior knowledge of the initial temperature field. For both prior models, the spatial correlation $(n \times n)$-matrix $\Sigma_.^\rho$ is defined by the second order exponential spatial correlation function $\rho_.(\tau) = \exp(-\tau^2/\delta^2); \delta = 0.15$, with interdistance τ.

Figure 5 contains realizations from the prior model of the log-diffusivity field and their associated spatial histograms. The prior model is specified to be spatially stationary except for boundary effects with bi-modal spatial histograms. The selection set $A \subset \mathbb{R}^n$ for the prior model is chosen to obtain bi-modal marginal distributions with a very dominant mode centered slightly above the value for λ_1 and a very small mode centered slightly below the value for λ_2. The prior is therefore not centered at the true values. Note that the joint random field $[\log(\lambda), r_0]$ will appear as a bi-variate selection-Gaussian random field, see [21].
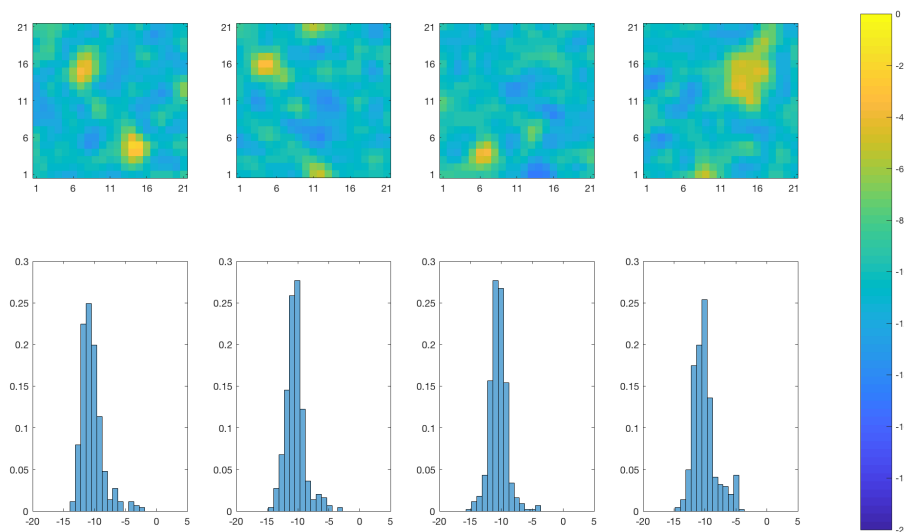
Figure 5. Realizations from the initial selection-Gaussian distribution of the log diffusivity $f(\log(\lambda))$ at time $T = 0$ (upper panels) and associated spatial histogram (lower panels). Lower panels: the horizontal axes represent the log-diffusivity, the vertical axes represent the relative prevalence of each log-diffusivity value for the realization in the panel right above.

The SEnKF operates on the $3n$-vector $[\log(\tilde{\lambda}), \nu, r_0]$, and therefore we generate an initial ensemble with $n_e = 10{,}000$ ensemble members that are sampled from the Gaussian $3n$-vector $[\log(\tilde{\lambda}), \nu, r_0]$ with pdf,

$$\begin{bmatrix} \log(\tilde{\lambda}) \\ \nu \\ r_0 \end{bmatrix} \sim \varphi_{3n} \left(\begin{bmatrix} \log(\tilde{\lambda}) \\ \nu \\ r_0 \end{bmatrix}; \begin{bmatrix} \mu_{\tilde{r}}^\lambda \\ \mu_\nu^\lambda \\ \mu_r \end{bmatrix}, \begin{bmatrix} \sigma_{\tilde{r}}^{\lambda\,2}\Sigma_{\tilde{r}}^\rho & \gamma^\lambda \sigma_{\tilde{r}}^\lambda \Sigma_{\tilde{r}}^{\rho\,T} & 0 \\ \gamma^\lambda \sigma_{\tilde{r}}^\lambda \Sigma_{\tilde{r}}^\rho & \gamma^{\lambda 2}\Sigma_{\tilde{r}}^\rho + (1-\gamma^{\lambda 2})I_n & 0 \\ 0 & 0 & \sigma_r^2 \Sigma_r^\rho \end{bmatrix} \right). \quad (25)$$

The EnKF operates on the $2n$-vector $[\log(\lambda), r_0]$, and therefore we generate an initial ensemble with $n_e = 10{,}000$ ensemble members that are sampled from the selection-Gaussian distribution $f(\log(\lambda), r_0)$. The variables $\log(\lambda)$ and r_0 are independent, so we generate them independently: 10,000 samples from the selection-Gaussian n-vector $\log(\lambda)$ with parameters Θ_λ^{SG} and 10,000 samples from the Gaussian n-vector r_0 with parameters Θ_r^G. It is important to understand that both ensemble algorithms are initiated with an ensemble from an identical selection-Gaussian random field prior

model for $[\log(\boldsymbol{\lambda}), \boldsymbol{r}_0]$ at $T = 0$, which reflects the bi-modality of the prior model. Due to the size of the ensemble relative to the dimension of the problem, we are using neither localization nor inflation in the algorithms.

To illustrate the differences between the SEnKF and the EnKF, we present the following results for both algorithms:

1. The marginal posterior distributions $f(\log(\lambda_i)|\boldsymbol{d}_{0:T})$ of the log-diffusivity field at four monitoring locations denoted $1, 2, 3, 4$ on Figure 2, at time $T = 0, 50, 80, 100$.
2. The marginal maximum a posteriori (MMAP) prediction of the log-diffusivity field at time at time $T = 0, 50, 80, 100$.
3. Realizations from the posterior distribution $f(\log(\boldsymbol{\lambda})|\boldsymbol{d}_{0:T})$ at time $T = 100$.
4. The root mean square errors (RMSE) of the MMAP prediction of the log-diffusivity field relative to the true log-diffusivity field at time $T = 100$.

Figures 6 and 7 show the marginal posterior pdfs $f(\log(\lambda_i)|\boldsymbol{d}_{0:T})$ at the four monitoring locations at time $T = 0, 50, 80, 100$ for the SEnKF and EnKF algorithms, respectively. Monitoring locations 1 and 2 are placed within the high diffusivity area while the two other locations are placed far into the low diffusivity area. At $T = 0$, all pdfs are identical, in all locations due to the stationary prior model and for both algorithms due to identical prior models. The SEnKF results appear to preserve bi-modality as observations are assimilated. As more data are made available, the high value mode increases at monitoring locations 1 and 2 that are inside the high diffusivity area. The low value mode remains dominant at monitoring locations within the low diffusivity areas. These results reflect expected behaviors. The traditional EnKF results are significantly different since the bi-modality of the marginal pdfs disappears already at $T = 50$. The marginal pdfs are Gaussian-like and are gently moved toward high and low values depending on which diffusivity area the monitoring locations are in. This regression toward the mean effect of the EnKF is generally recognized as it gives the best prediction in the squared error sense [27].
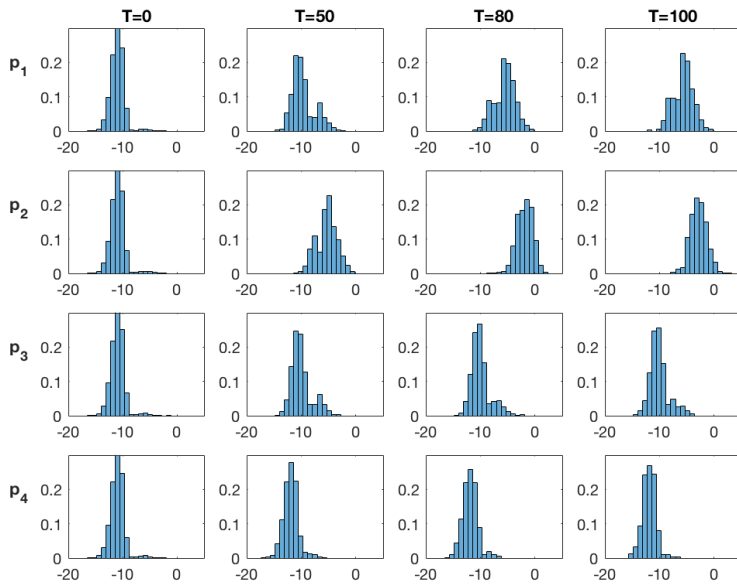
Figure 6. SEnKF approach: Marginal posterior distribution of the log diffusivity $f(\log(\lambda_i)|\boldsymbol{d}_{0:T})$ at time $T = 0, 50, 80, 100$ at the monitoring locations $(1, 2, 3, 4)$ denoted (p_1, p_2, p_3, p_4).

Figure 7. EnKF approach: Marginal posterior distribution of the log diffusivity $f(\log(\lambda_i)|\boldsymbol{d}_{0:T})$ at time $T = 0, 50, 80, 100$ at the monitoring locations $(1, 2, 3, 4)$ denoted (p_1, p_2, p_3, p_4).

Figure 8 displays the MMAP predictions based on the SEnKF and the traditional EnKF at time $T = 0, 50, 80, 100$. At $T = 0$, the predictions from the two algorithms are identical since they use identical prior models. As observations are assimilated, the SEnKF predictions reproduce the high diffusivity area relatively well, with clear contrast. The traditional EnKF predictions also indicate the diffusivity areas, but with less contrast. Figure 9 displays the MMAP prediction at $T = 100$ along the section B-B' shown in Figure 2. The high contrast reliable reconstruction of the high diffusivity channel by the SEnKF algorithm is confirmed. The traditional EnKF predictions appear less reliable. The 80% highest density interval (HDI) [28] covers the true diffusivity values for the SEnKF while these values are far outside the interval for the traditional EnKF results. The results are consistent with the observations made regarding the marginal posterior pdfs in Figures 6 and 7.

Figures 10 and 11 show realizations and spatial histograms from the posterior distribution of the log-diffusivity at time $T = 100$ for the SEnKF and traditional EnKF algorithms, respectively. The realizations from the SEnKF largely reproduce the channel with clear contrast while the realizations from the EnKF also reproduce the channel, but with much less contrast. The spatial histograms also underline the difference in contrast in that they are clearly bi-modal for the SEnKS and much more Gaussian-like for the EnKS.

Table 2 shows that the RMSE of the MMAP prediction relative to the true diffusivity field for the SEnKF is approximately 30% lower than for the EnKF .

Table 2. RMSE comparing the MMAP prediction and the true log diffusivity field at time $T = 100$.

	SEnKF	ENKF
$RMSE_{T=100}$	2.72	3.76

This test case clearly illustrates the SEnKF's ability to conserve multimodality in the posterior distribution and it leads to predictions with better constrast and accuracy. We conclude that the reconstruction of the true diffusivity field is done more reliably by the SEnKF algorithm than by the EnKF algorithm.
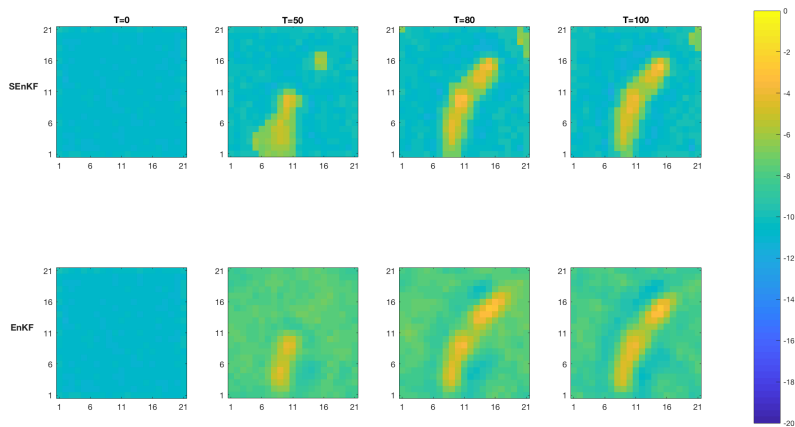
**Figure 8.** MMAP predictions of the log diffusivity field $(\log(\boldsymbol{\lambda})|\boldsymbol{d}_{0:T})$ at time $T = 0, 50, 80, 100$ (upper panels—SEnKF approach, lower panels—EnKF approach).

Figure 9. MMAP predictions of the log diffusivity field with 80% HDI in cross section B-B' at time $T = 100$ with SEnKF (**left**) and with EnKF (**right**).

Figure 10. SEnKF approach: Realizations of the posterior distribution of the log diffusivity $f(\log(\boldsymbol{\lambda})|\boldsymbol{d}_{0:T})$ at time $T = 100$ (**upper** panels) and associated spatial histogram (**lower** panels). Lower panels: the horizontal axes represent the log-diffusivity, the vertical axes represent the relative prevalence of each log-diffusivity value for the realization in the panel right above.

Figure 11. EnKF approach: Realizations of the posterior distribution of the log diffusivity $f(\log(\lambda)|d_{0:T})$ at time $T = 100$ (**upper** panels) and associated spatial histogram (**lower** panels). Lower panels: the horizontal axes represent the log-diffusivity, the vertical axes represent the relative prevalence of each log-diffusivity value for the realization in the panel right above.

3.3. Test Case 2: Reconstructing the Initial Field

The focus of the study is to evaluate the unknown initial state of the temperature field r_0 based on the observations d. The posterior model $f(r_0|d)$ is assessed using the SEnKS, see Appendix A, and the results are compared to those from the traditional EnKS.

The true initial temperature field r_0 is set at 20 °C except for a square shaped region with temperature set at 45 °C, see Figure 12. The temperature field is formally defined as

$$r_0(x) = \tau_1 \mathbf{1}(x \in D_1) + \tau_2 \mathbf{1}(x \in D_2), \tag{26}$$

where $D_1 \subset D$ is the low temperature area and $D_2 \subset D$ is the high temperature area, and $\tau_1 = 20$ °C and $\tau_2 = 45$ °C. Figure 12 shows the true log-diffusivity n-vector $\log(\lambda)$. The diffusivity λ is always positive. To ensure that ensemble updates do not lead to negative diffusivity values, we work on $\log(\lambda)$. The heat contained in the high temperature area will diffuse towards the rest of the field according to the diffusion equation in Equation (20), see Figure 13. The temporal observations are collected at five different observation locations according to the likelihood model in Equation (23), see Figure 12. Figure 14 displays the observations d where it is clear that the observed temperature increases substantially only at the observation locations close to the high temperature area. The measurements are taken every second from T = 0 to T = 50.

The unknown initial temperature field r_0 is assigned a stationary selection-Gaussian random field prior model with parameters $\Theta_r^{SG} = (\mu_{\tilde{r}}, \mu_{\nu}, \sigma_{\tilde{r}}, \Sigma_{\tilde{r}}^{\rho}, \gamma, A)$. The parameter values are listed in Table 3. The unknown log-diffusivity field $\log(\lambda)$ is assigned a stationary Gaussian random field prior model with parameters $\Theta_\lambda^G = (\mu_\lambda, \sigma_\lambda, \Sigma_\lambda^{\rho})$ with expectation and variance levels $\mu_\lambda = -8.5$ and $\sigma_\lambda^2 = 2$, respectively. For both prior models, the spatial correlation $(n \times n)$-matrix Σ^{ρ} is defined by the second order exponential spatial correlation function $\rho.(\tau) = \exp(-\tau^2/\delta^2); \delta = 0.15$, with interdistance τ.

Figure 15 contains four realizations from the prior model of the temperature field and their spatial histograms. The marginal initial distributions of the realizations are bi-modal and spatially stationary except for boundary effects. The selection set $A \subset \mathbb{R}^n$ in the prior model is chosen to obtain a bi-modal marginal distribution with one large mode approximately centered about 20 °C and a smaller mode centered close to 45 °C.

Table 3. Parameter values for the selection-Gaussian initial distribution for the initial temperature prior model.

Parameters	Values
$\mu_{\tilde{r}}$	28.75
μ_{ν}	0
$\sigma_{\tilde{r}}$	$\sqrt{10}$
γ	0.8
A	$]-\infty, -0.2] \cup [0.5, +\infty[^n$

Figure 12. Initial temperature (°C) field (**left**) with data collection points · and monitoring locations × and reference log-diffusivity field (**right**).

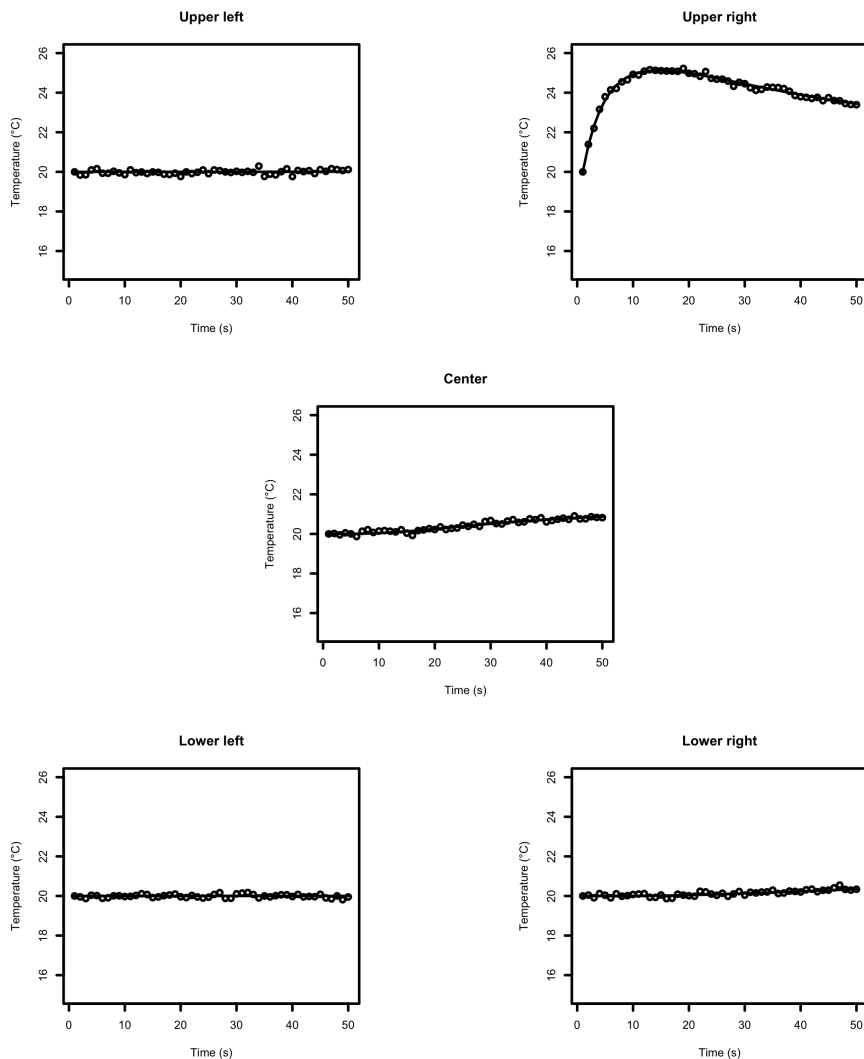**Figure 13.** True temperature (°C) field evolution over time.

The SEnKS operates on the $3n$-vector $[\tilde{r}_0, \nu, \log(\lambda)]$, and therefore we generate an initial ensemble with $n_e = 10{,}000$ ensemble members that are sampled from the Gaussian $3n$-vector $[\tilde{r}_0, \nu, \log(\lambda)]$ with pdf,

$$
\begin{bmatrix} \tilde{\boldsymbol{r}}_0 \\ \boldsymbol{\nu} \\ \log(\boldsymbol{\lambda}) \end{bmatrix} \sim \varphi_{3n} \left(\begin{bmatrix} \tilde{\boldsymbol{r}}_0 \\ \boldsymbol{\nu} \\ \log(\boldsymbol{\lambda}) \end{bmatrix} ; \begin{bmatrix} \boldsymbol{\mu}_{\tilde{r}} \\ \boldsymbol{\mu}_{\nu} \\ \boldsymbol{\mu}_{\lambda} \end{bmatrix} , \begin{bmatrix} \sigma_{\tilde{r}}^2 \boldsymbol{\Sigma}_{\tilde{r}}^{\rho} & \gamma \sigma_{\tilde{r}} \boldsymbol{\Sigma}_{\tilde{r}}^{\rho T} & 0 \\ \gamma \sigma_{\tilde{r}} \boldsymbol{\Sigma}_{\tilde{r}}^{\rho} & \gamma^2 \boldsymbol{\Sigma}_{\tilde{r}}^{\rho} + (1-\gamma^2) \boldsymbol{I}_n & 0 \\ 0 & 0 & \sigma_{\lambda}^2 \boldsymbol{\Sigma}_{\lambda}^{\rho} \end{bmatrix} \right). \tag{27}
$$

The EnKS operates on the $2n$-vector $[\boldsymbol{r}_0, \log(\boldsymbol{\lambda})]$, and therefore we generate an initial ensemble with $n_e = 10{,}000$ ensemble members that are sampled from selection-Gaussian distribution $f(\boldsymbol{r}_0, \log(\boldsymbol{\lambda}))$. The variables $\boldsymbol{r}_0$ and $\log(\boldsymbol{\lambda})$ are independent, so we generate them independently: 10,000 samples from the selection-Gaussian n-vector $\boldsymbol{r}_0$ with parameters $\boldsymbol{\Theta}_r^{SG}$ and 10,000 samples from the Gaussian n-vector $\log(\boldsymbol{\lambda})$ with parameters $\boldsymbol{\Theta}_{\lambda}^{G}$. Due to the size of the ensemble relative to the dimension of the problem, we used neither localization nor inflation in the algorithms.

To illustrate the differences between the SEnKS and the EnKS we present the following results for both algorithms:

1. The marginal posterior distributions $f(r_{0,i}|\boldsymbol{d}_{0:T})$ of the initial temperature field at four monitoring locations denoted $1, 2, 3, 4$ on Figure 12, at time $T = 0, 20, 30, 50$.
2. The marginal maximum a posteriori (MMAP) prediction of the initial temperature field at time at time $T = 0, 20, 30, 50$.
3. Realizations from the posterior distribution $f(\boldsymbol{r}_0|\boldsymbol{d}_{0:T})$ of the initial temperature field at time $T = 50$.
4. The root mean square errors (RMSE) of the MMAP prediction of the initial temperature field relative to the true initial temperature field at time $T = 50$.

Figure 14. Data collected over time (points) and true temperature (°C) evolution at the data collection points (line).

The marginal posterior pdfs $f(r_{0,i}|\boldsymbol{d}_{0:T})$ at the four monitoring locations at time $T = 0,20,30,50$ are displayed in Figures 16 and 17 for the SEnKS and EnKS algorithms, respectively. At $T = 0$, the prior models for both algorithms are identical and so are the marginal pdfs. Monitoring location 1 is placed

inside the high temperature area. As observations are assimilated, the marginal pdf from the SEnKS remain bi-modal, but the high value mode increases steadily. For the other monitoring locations, all placed outside the high temperature area, the bi-modality is reproduced but with a dominant low value mode. The relative size of the modes reflects the distance to the high temperature area and the observation locations. The marginal pdfs from the EnKS lose their bi-modality after a few assimilation steps and from then on the Gaussian-like marginal pdfs are only slightly shifted by the assimilation of observations.

Figure 15. Realizations from the selection-Gaussian initial distribution of the initial temperature field $f(\mathbf{r}_0)$ at time $T = 0$ (upper panels) and associated spatial histogram (lower panels). Upper panels: the colorbar gives the temperature in °C. Lower panels: the horizontal axes represent the temperature (°C), the vertical axes represent the relative prevalence of each temperature value for the realization right above.
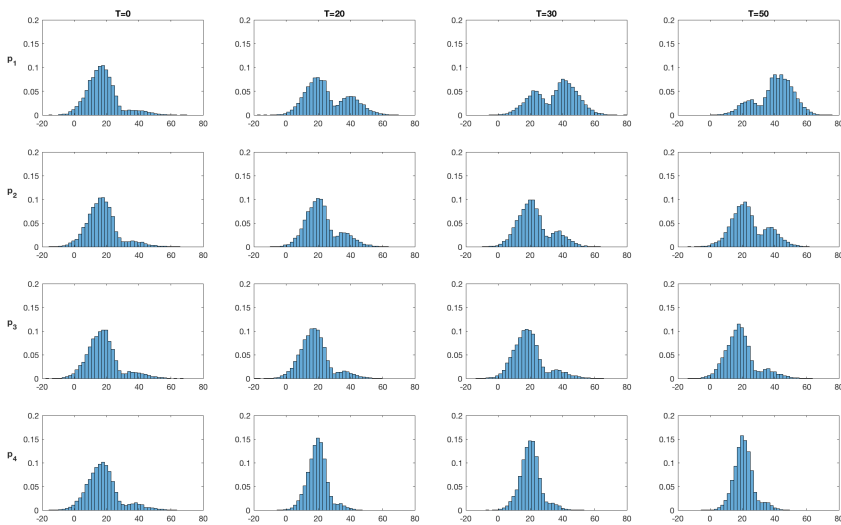
Figure 16. SEnKS approach: Marginal posterior distributions of the initial temperature $f(\mathbf{r}_{0,i}|\mathbf{d}_{0:T})$ at time $T = 0, 20, 30, 50$ at monitoring locations ($i = 1, 2, 3, 4$) denoted (p_1, p_2, p_3, p_4). The horizontal axes representing temperature are expressed in °C.
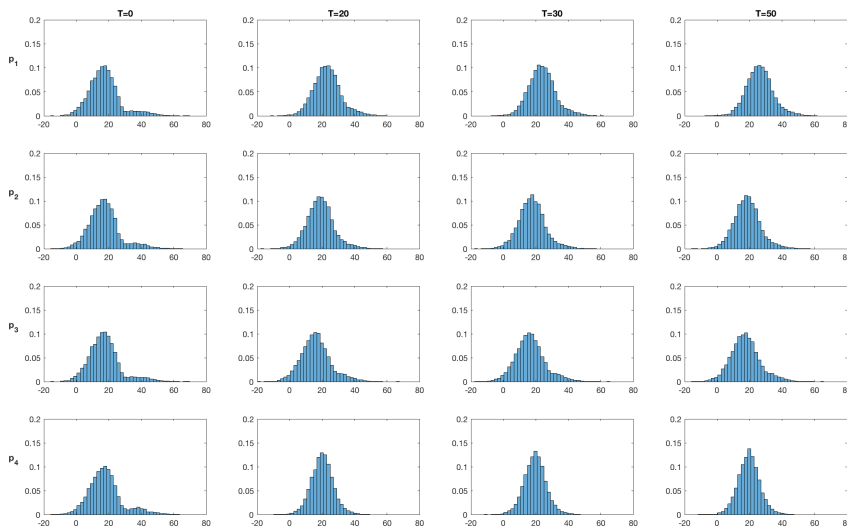
Figure 17. EnKS approach: Marginal posterior distributions of the initial temperature $f(r_{0,i}|d_{0:T})$ at time $T = 0, 20, 30, 50$ at monitoring locations ($i = 1, 2, 3, 4$) denoted (p_1, p_2, p_3, p_4). The horizontal axes representing temperature are expressed in °C.

Figure 18 displays the MMAP predictions of the initial temperature field based on the SEnKS and the traditional EnKS at time $T = 0, 20, 30, 50$. For the SEnKS, the high temperature area is clearly identifiable with clear contrast from time $T = 30$ while for the EnKS the high temperature area is hardly ever identifiable on the MMAP predictions that show little contrast. Figure 19 displays the MMAP prediction of the initial temperature field at $T = 50$ along the section A-A', see Figure 12, for the SEnKS and the traditional EnKS. The SEnKS clearly identifies the high temperature area and the 80% HDI covers the truth, while the EnKS clearly fails to identify the high temperature area and the 80% HDI does not even cover it.
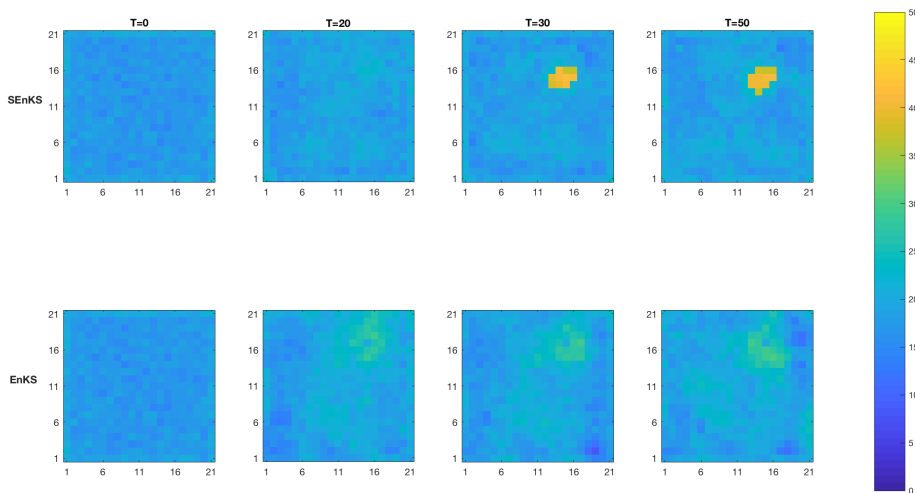
Figure 18. MMAP predictions of the initial temperature (°C) field at time $T = 0, 20, 30, 50$ for the SEnKS approach (**upper**) and the EnKS approach (**lower**).
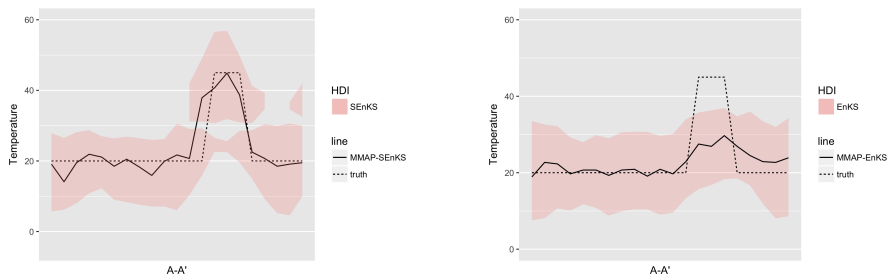
Figure 19. MMAP predictions of the initial temperature (°C) field $(r_0|d_{0:T})$ with 80% HDI in cross section A-A' at time $T = 50$ with SEnKS (**left**) and with EnKS (**right**).

Realizations of the posterior model at $T = 50$ based on the SEnKS and the traditional EnKS algorithms are displayed in Figures 20 and 21, respectively. The SEnKS produces realizations that appear bi-modal while the ensemble members from the EnKS display more symmetric spatial histograms. Even though the differences between the realizations are quite subtle, they are consistent with previous results.

Table 4 shows that RMSE of the MMAP prediction relative to the true initial temperature field for the SEnKS is approximately 20% lower than for the EnKS.

This test case clearly illustrates the ability of the SEnKS to conserve multimodality in the posterior distribution, and it leads to predictions with better contrast and accuracy. We conclude that the SEnKS algorithm provides a more reliable reconstruction of the initial state of the temperature field than the traditional EnKS algorithm. Note that the posterior model for the unknown diffusivity field $f(\log(\lambda)|d)$ can also be assessed with the two algorithms. When comparing the MMAP predictions relative to the true diffusivity field, see Figure 12, we observe that none of the algorithms provide reliable predictions. We conclude that the small scale variations in the field are not sufficiently distinct to be identified.

Figure 20. SEnKS approach: Realizations from the posterior distribution $f(r_0|d_{0:T})$ of the initial temperature field at time $T = 50$. Upper panels: the colorbar gives the temperature in °C. Lower panels: the horizontal axes represent the temperature (°C), the vertical axes represent the relative prevalence of each temperature value for the realization right above.
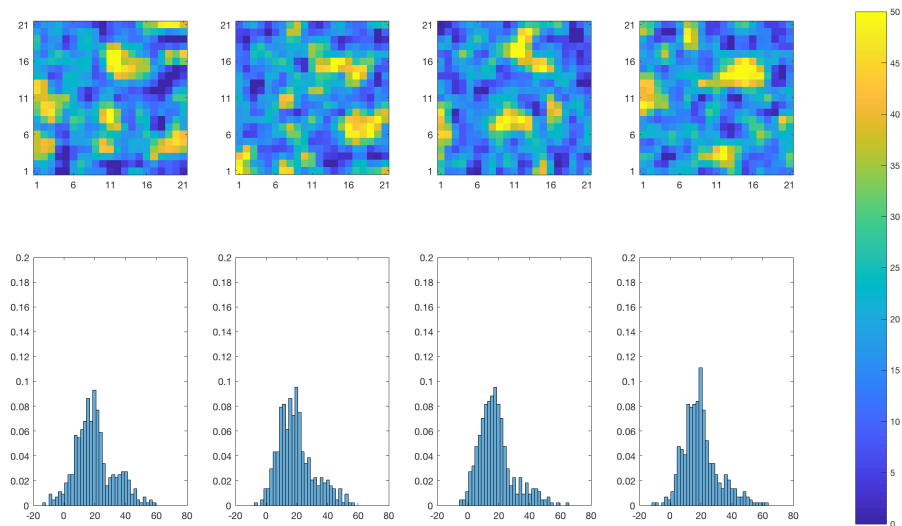
Figure 21. EnKS approach: Realizations from the posterior distribution of the initial temperature field $f(r_0|d_{0:T})$ at time $T = 50$. Upper panels: the colorbar gives the temperature in °C. Lower panels: the horizontal axes represent the temperature (°C), the vertical axes represent the relative prevalence of each temperature value for the realization right above.

Table 4. RMSE comparing the MMAP prediction of the initial temperature field and the initial temperature field at time $T = 50$.

	SEnKS	**ENKS**
$RMSE_{T=50}$	2.92	3.72

4. Discussion

The traditional EnKF and EnKS algorithms provide an ensemble that directly represents the posterior models $f(r_{T+1}|d_{0:T})$ and $f(r_{0:T}|d_{0:T})$, respectively. Hence, the posterior models can be assessed by displaying statistics based on these ensembles. Reliable assessment of the posterior model in the two test cases can be obtained with approximately 1000 ensemble members. The SEnKF and SEnKS algorithms under study provide an ensemble of the augmented posterior models, $f(\tilde{r}_{T+1}, \nu|d_{0:T})$ and $f(\tilde{r}_{0:T}, \nu_{0:T}|d_{0:T})$, respectively. In order to obtain the posterior models of interest, $f(r_{T+1}|d_{0:T})$ and $f(r_{0:T}|d_{0:T})$, the conditioning on $\nu \in A$ must be made by empirical sampling based inference, see [21,22]. This inference requires the estimation of the expectation vector $\mu_{\tilde{r}\nu}$ and covariance matrix $\Sigma_{\tilde{r}\nu}$. The two test cases are defined on a (21×21)-grid for both $\tilde{r}_t$ and ν—hence in dimension 882. The expectation and covariance will have 882 and 389,403 unique entries, respectively. Our experience from this study is that approximately 10,000 ensemble members are required to obtain reliable assessment of the posterior models of interest. To reduce the ensemble size, we have tested various localization approaches [29], without notable success, and leave the subject for further research.

5. Conclusions

Data assimilation of spatio-temporal variables with multimodal spatial histograms is challenging. Traditional ensemble Kalman algorithms enforce a regression towards the mean due to the linearized conditioning on observations, hence the multimodality is averaged out. We introduce the selection ensemble Kalman algorithms, termed SEnKF and SEnKS. These algorithms are based on recursive expressions similar to the ones justifying the traditional ensemble Kalman algorithms, but they are defined in an augmented space including the selection variable. From the two case studies, we conclude

that multimodality is much better represented by the selection ensemble Kalman algorithms than by the traditional ones. We obtain RMSE reductions in the range of 20 to 30%.

The traditional ensemble Kalman algorithms provide an ensemble representation of the posterior model of interest hence making assessment of the posterior pdf simple. The selection ensemble Kalman algorithms are defined in an augmented space and conditioning on the selection variable must be made a posteriori. For this conditioning to be reliable, the ensemble size needs to be much larger than for the traditional algorithms. Hence, there is a trade-off between improved reproduction of multimodal characteristics of the phenomenon under study and the computational demands. In our case study, the ensemble size needed to be increased by approximately a factor of ten.

We have not fully explored the possibilities of robust estimation of model parameters in the conditioning of the selection variable. This robustification may reduce the ensemble size requirements. Note that parallelization in forwarding of the ensemble is possible and it will reduce the computer demands.

Author Contributions: Conceptualization, M.C. and H.O.; Formal analysis, M.C. and H.O.; Funding acquisition, H.O.; Investigation, M.C. and H.O.; Methodology, M.C. and H.O.; Project administration, H.O.; Resources, H.O.; Software, M.C.; Supervision, H.O.; Validation, M.C. and H.O.; Visualization, H.O.; Writing—original draft, M.C. and H.O. All authors have read and agreed to the published version of the manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

Glossary

$\boldsymbol{r}_t \in \mathbb{R}^n$	discretized spatial variable at time t.
$[\tilde{\boldsymbol{r}}_t, \boldsymbol{\nu}] \in \mathbb{R}^{2n}$	Gaussian variables; basis and auxiliary variables, at time t.
$A \subset \mathbb{R}^n$	selection set.
$\boldsymbol{r}_t = [\tilde{\boldsymbol{r}}_t \vert \boldsymbol{\nu} \in A]$	selection Gaussian variable at time t.
$\boldsymbol{\mu}_. \in \mathbb{R}^n$	expectation vector.
$\boldsymbol{\Sigma}_. \in \mathbb{R}^n \times \mathbb{R}^n$	covariance matrix.
$\boldsymbol{\Sigma}_.^\rho \in \mathbb{R}^n \times \mathbb{R}^n$	correlation matrix.
$\boldsymbol{\Gamma}_{\cdot\vert\cdot} \in \mathbb{R}^n \times \mathbb{R}^n$	matrix cross-correlation
$\boldsymbol{d}_t \in \mathbb{R}^m$	observation variable at time t.
$\omega(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^r)$	forward function at time t.
$\psi(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^d)$	observation function at time t.
$\rho_.(\tau) \in \mathbb{R}_{[-1,1]}$	spatial correlation function.

Appendix A

The algorithms detailed in Algorithms A1 and A2 follow the formalism in [4].

Algorithm A1 description: The SEnKF is a two-step algorithm. The first step is a traditional EnKF that evaluates $[\tilde{\boldsymbol{r}}_{T+1}, \boldsymbol{\nu} \vert \boldsymbol{d}_{0:T}]$. The second step consists of a sampling step where the target quantity $[\boldsymbol{r}_{T+1} \vert \boldsymbol{d}_0, ..., \boldsymbol{d}_T]$ is evaluated using $[\tilde{\boldsymbol{r}}_{T+1}, \boldsymbol{\nu} \vert \boldsymbol{d}_{0:T}]$ from the first step.

Algorithm A1 Selection Ensemble Kalman Filter (SEnKF)

A time series of ensembles is defined as $e_t = \{(\tilde{r}_t^{u(i)}, v_t^{u(i)}, d_t^i), i = 1, ..., n_e\}, \forall t = 0, \cdots, T$ and the $(2n + m)$-vector $[\tilde{r}_t, v_t, d_t]$ has the following covariance matrix:

$$\Sigma_{\tilde{r}vd} = \begin{bmatrix} \Sigma_{\tilde{r}v} & \Gamma_{\tilde{r}v,d} \\ \Gamma_{d,\tilde{r}v} & \Sigma_d \end{bmatrix}$$

1. **Initiate:**
2. $n_e =$ No. of ensemble members

3. Generate $\begin{bmatrix} \tilde{r}_0^{u(i)} \\ v_0^{u(i)} \end{bmatrix} \sim f(\tilde{r}_0, v), i = 1, ..., n_e$

4. Generate $\epsilon_0^{d(i)} \sim \mathcal{U}_m[0, 1], i = 1, ..., n_e$
5. $d_0^i = \psi_0(\tilde{r}_0^{u(i)}, \epsilon_0^{d(i)}), i = 1, ..., n_e$
6. $e_0 = \{(\tilde{r}_0^{u(i)}, v_0^{u(i)}, d_0^i), i = 1, ..., n_e\}$
7. **Iterate** $t = 0, ..., T$:
8. Conditioning:
9. Estimate $\Sigma_{\tilde{r}vd}$ from $e_t \longrightarrow \hat{\Sigma}_{\tilde{r}vd}$

10. $\begin{bmatrix} \tilde{r}_t^{c(i)} \\ v_t^{c(i)} \end{bmatrix} = \begin{bmatrix} \tilde{r}_t^{u(i)} \\ v_t^{u(i)} \end{bmatrix} + \hat{\Gamma}_{\tilde{r}v,d} \hat{\Sigma}_d^{-1} (d_t - d_t^i), i = 1, ..., n_e$

11. Forwarding:
12. Generate $\epsilon_t^{\tilde{r}(i)} \sim \mathcal{U}_n[0, 1], i = 1, ..., n_e$

13. $\begin{bmatrix} \tilde{r}_{t+1}^{u(i)} \\ v_{t+1}^{u(i)} \end{bmatrix} = \begin{bmatrix} \omega_t(\tilde{r}_t^{c(i)}, \epsilon_t^{\tilde{r}(i)}) \\ v_t^{c(i)} \end{bmatrix}, i = 1, ..., n_e$

14. If $t < T$
15. Generate $\epsilon_{t+1}^{d(i)} \sim \mathcal{U}_n[0, 1], i = 1, ..., n_e$
16. $d_{t+1}^{(i)} = \psi_{t+1}(\tilde{r}_{t+1}^{u(i)}, \epsilon_{t+1}^{d(i)}), i = 1, ..., n_e$
17. $e_{t+1} = \{(\tilde{r}_{t+1}^{u(i)}, v_{t+1}^{u(i)}, d_{t+1}^i), i = 1, ..., n_e\}$
18. Else
19. $e_{t+1} = \{(\tilde{r}_{t+1}^{u(i)}, v_{t+1}^{u(i)}), i = 1, ..., n_e\}$
20. **End iterate**
21. Estimate $\mu_{\tilde{r}v}, \Sigma_{\tilde{r}v}$ from $e_{T+1} \rightarrow \hat{\mu}_{\tilde{r}v}, \hat{\Sigma}_{\tilde{r}v}$
22. Assess
23. $\hat{f}(r_{T+1}|d_0, ..., d_T) = [\Phi_n(A; \hat{\mu}_v, \hat{\Sigma}_v)]^{-1} \times \Phi_n(A; \hat{\mu}_v + \hat{\Gamma}_{v|\tilde{r}}(r - \hat{\mu}_{\tilde{r}}), \hat{\Sigma}_{v|\tilde{r}}) \times \varphi_n(r; \hat{\mu}_{\tilde{r}}, \hat{\Sigma}_{\tilde{r}})$
24. End Algorithm

The ensemble e_{T+1} represents $[\tilde{r}_{T+1}, v|d_{0:T}]$. To assess $f(r_{T+1}|d_{0:T}) = f(\tilde{r}_{T+1}|d_{0:T}, v \in A)$, the sampling algorithm specified in [21] requires $\mathbf{E}[\tilde{r}_{T+1}, v|d_{0:T}] = \mu_{\tilde{r}v}$ and $\mathbf{Cov}[\tilde{r}_{T+1}, v|d_{0:T}] = \Sigma_{\tilde{r}v}$ which are estimated using the ensemble e_{T+1}.

Algorithm A2 description: The SEnKS is a two-step algorithm. The first step is a traditional EnKS that evaluates $[\tilde{r}_t, v|d_{0:T}]$. The second step consists of a sampling step where the target quantity $[r_t|d_0, ..., d_T]$ is evaluated using $[\tilde{r}_t, v|d_{0:T}]$ from the first step.

Algorithm A2 Selection Ensemble Kalman Smoother (SEnKS)

Two time series of ensemble sets are defined as

$$\left. \begin{aligned} e_t^{\tilde{r}v} &= \{(\tilde{\boldsymbol{r}}_t^i, \boldsymbol{v}_t^i), i = 1, \dots, n_e\} \\ e_t^d &= \{\boldsymbol{d}_t^i, i = 1, \dots, n_e\} \end{aligned} \right\} \text{ for } t = 0, \dots, T$$

and the accumulated ensemble set defined as

$$E_{0:t}^{\tilde{r}v} = \{(\tilde{\boldsymbol{r}}_{0:t}^i, \boldsymbol{v}_{0:t}^i), i = 1, \dots, n_e\} \quad \text{for } t = 0, \dots, T$$

The $(2n + m)$-vector $[\tilde{\boldsymbol{r}}_t, \boldsymbol{v}_t, \boldsymbol{d}_t]$ has covariance matrix $\boldsymbol{\Sigma}_{\tilde{r}vd}^t = \begin{bmatrix} \boldsymbol{\Sigma}_{\tilde{r}v}^t & \boldsymbol{\Gamma}_{\tilde{r}v,d}^t \\ \boldsymbol{\Gamma}_{d,\tilde{r}v}^t & \boldsymbol{\Sigma}_d^t \end{bmatrix}$

The $(2n(t+1) + m)$-vector $[\tilde{\boldsymbol{r}}_{0:t}, \boldsymbol{v}_{0:t}, \boldsymbol{d}_t]$ has covariance matrix $\boldsymbol{\Sigma}_{\tilde{r}vd}^{0:t} = \begin{bmatrix} \boldsymbol{\Sigma}_{\tilde{r}v}^{0:t} & \boldsymbol{\Gamma}_{\tilde{r}v,d}^{0:t} \\ \boldsymbol{\Gamma}_{d,\tilde{r}v}^{0:t} & \boldsymbol{\Sigma}_d^t \end{bmatrix}$

1.   **Initiate**
2. $n_e = $ No. of ensemble members
3.   Generate $\begin{bmatrix} \tilde{\boldsymbol{r}}_0^{u(i)} \\ \boldsymbol{v}_0^{u(i)} \end{bmatrix} \sim f(\tilde{\boldsymbol{r}}_0, \boldsymbol{v}), i = 1, \dots, n_e$
4. $E_{0:0}^{\tilde{r}v} = \{(\tilde{\boldsymbol{r}}_0^{u(i)}, \boldsymbol{v}_0^{u(i)}), i = 1, \dots, n_e\}$
5.   Generate $\boldsymbol{\epsilon}_0^{d(i)} \sim \mathcal{U}_m[0,1], i = 1, \dots, n_e$ iid
6. $\boldsymbol{d}_0^i = \psi_0(\tilde{\boldsymbol{r}}_0^{u(i)}, \boldsymbol{\epsilon}_0^{d(i)}), i = 1, \dots, n_e$
7. $e_0^d = \{\boldsymbol{d}_0^i, i = 1, \dots, n_e\}$
8.   Estimate $\boldsymbol{\Sigma}_{\tilde{r}vd}^0$ from $E_{0:0}^{\tilde{r}v}, e_0^d \longrightarrow \hat{\boldsymbol{\Sigma}}_{\tilde{r}vd}^0$
9. $\begin{bmatrix} \tilde{\boldsymbol{r}}_0^{c(i)} \\ \boldsymbol{v}_0^{c(i)} \end{bmatrix} = \begin{bmatrix} \tilde{\boldsymbol{r}}_0^{u(i)} \\ \boldsymbol{v}_0^{u(i)} \end{bmatrix} + \hat{\boldsymbol{\Gamma}}_{\tilde{r}v,d}^0 \hat{\boldsymbol{\Sigma}}_d^{0 \; -1}(\boldsymbol{d}_0 - \boldsymbol{d}_0^i), \quad i = 1, \dots, n_e$
10.   **Iterate** $t = 1, \dots, T$:
11.   Fowarding
12.   Generate $\boldsymbol{\epsilon}_t^{\tilde{r}(i)} \sim \mathcal{U}_n[0,1], i = 1, \dots, n_e$
13. $\begin{bmatrix} \tilde{\boldsymbol{r}}_t^{u(i)} \\ \boldsymbol{v}_t^{u(i)} \end{bmatrix} = \begin{bmatrix} \omega_t(\tilde{\boldsymbol{r}}_{t-1}^{c(i)}, \boldsymbol{\epsilon}_t^{\tilde{r}(i)}) \\ \boldsymbol{v}_{t-1}^{c(i)} \end{bmatrix}, i = 1, \dots, n_e$
14. $E_{0:t}^{\tilde{r}v} = \{E_{0:t-1}^{\tilde{r}v}, (\tilde{\boldsymbol{r}}_t^{u(i)}, \boldsymbol{v}_t^{u(i)}), i = 1, \dots, n_e\}$
15.   Generate $\boldsymbol{\epsilon}_t^{d(i)} \sim \mathcal{U}_m[0,1], i = 1, \dots, n_e$ iid
16. $\boldsymbol{d}_t^i = \psi_0(\tilde{\boldsymbol{r}}_t^{u(i)}, \boldsymbol{\epsilon}_t^{d(i)}), i = 1, \dots, n_e$
17. $e_t^d = \{\boldsymbol{d}_t^i, i = 1, \dots, n_e\}$
18.   Estimate $\boldsymbol{\Sigma}_{\tilde{r}vd}^{0:t}$ from $E_{0:t}^{\tilde{r}v}, e_t^d \longrightarrow \hat{\boldsymbol{\Sigma}}_{\tilde{r}vd}^{0:t}$
19. $\begin{bmatrix} \tilde{\boldsymbol{r}}_{0:t}^{c(i)} \\ \boldsymbol{v}_{0:t}^{c(i)} \end{bmatrix} = \begin{bmatrix} \tilde{\boldsymbol{r}}_{0:t}^{u(i)} \\ \boldsymbol{v}_{0:t}^{u(i)} \end{bmatrix} + \hat{\boldsymbol{\Gamma}}_{\tilde{r}v,d}^{0:t} \hat{\boldsymbol{\Sigma}}_d^{t \; -1}(\boldsymbol{d}_t - \boldsymbol{d}_t^i), \quad i = 1, \dots, n_e$
20.   **End iterate**
21. $E_{0:t}^S = \{(\tilde{\boldsymbol{r}}_{0:T}^{c(i)}, \boldsymbol{v}_{0:T}^{c(i)}), i = 1, \dots, n_e\}$
22.   Select
23.   For arbitrary $t \in [0, T]$, select corresponding ensemble e_t^S from $E_{0:T}^S$
24.   Estimate $\boldsymbol{\mu}_{\tilde{r}v}^t, \boldsymbol{\Sigma}_{\tilde{r}v}^t$ from $e_t^S \longrightarrow \hat{\boldsymbol{\mu}}_{\tilde{r}v}^t, \hat{\boldsymbol{\Sigma}}_{\tilde{r}v}^t$
25.   Assess
26. $\hat{f}(\boldsymbol{r}_t | \boldsymbol{d}_{0:T}) = \left[\Phi_n(A; \hat{\boldsymbol{\mu}}_v^t, \hat{\boldsymbol{\Sigma}}_v^t) \right]^{-1} \times \Phi_n(A; \hat{\boldsymbol{\mu}}_v^t + \hat{\boldsymbol{\Gamma}}_{v|\tilde{r}}^t(\boldsymbol{r} - \hat{\boldsymbol{\mu}}_{\tilde{r}}^t), \hat{\boldsymbol{\Sigma}}_{v|\tilde{r}}^t) \times \varphi_n(\boldsymbol{r}; \hat{\boldsymbol{\mu}}_{\tilde{r}}^t, \hat{\boldsymbol{\Sigma}}_{\tilde{r}}^t)$
27.   End Algorithm

The ensemble $E_{0:T}^S$ represents $[\tilde{\boldsymbol{r}}_{0:T}, \boldsymbol{v}_{0:T} | \boldsymbol{d}_{0:T}]$. To assess $f(\boldsymbol{r}_t | \boldsymbol{d}_{0:T}) = f(\tilde{\boldsymbol{r}}_t | \boldsymbol{d}_{0:T}, \boldsymbol{v} \in A)$, the sampling algorithm in [21] requires $E(\tilde{\boldsymbol{r}}_t, \boldsymbol{v} | \boldsymbol{d}_{0:T}) = \boldsymbol{\mu}_{\tilde{r}v}^t$ and $\textbf{Cov}(\tilde{\boldsymbol{r}}_t, \boldsymbol{v} | \boldsymbol{d}_{0:T}) = \boldsymbol{\Sigma}_{\tilde{r}v}^t$, which are estimated using the sub-ensemble e_t^S of $E_{0:t}^S$.

References

1. Kalman, R.E. A new approach to linear filtering and prediction problems. *Trans. ASME-J. Basic Eng.* **1960**, *82*, 35–45. [CrossRef]

2. McElhoe, B.A. An Assessment of the Navigation and Course Corrections for a Manned Flyby of Mars or Venus. *IEEE Trans. Aerosp. Electron. Syst.* **1966**, *AES-2*, 613–623. [CrossRef]

3. Evensen, G. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.* **1994**, *99*, 10143. [CrossRef]

4. Myrseth, I.; Omre, H. The Ensemble Kalman Filter and Related Filters. In *Large Scale Inverse Problems and Quantification of Uncertainty*; John Wiley & Sons, Ltd.: London, UK, 2010; chapter 11, pp. 217–246.

5. Houtekamer, P.L.; Mitchell, H.L.; Pellerin, G.; Buehner, M.; Charron, M.; Spacek, L.; Hansen, B. Atmospheric Data Assimilation with an Ensemble Kalman Filter: Results with Real Observations. *Mon. Weather. Rev.* **2005**, *133*, 604–620. [CrossRef]

6. Sakov, P.; Oliver, D.; Bertino, L. An Iterative EnKF for Strongly Nonlinear Systems. *Mon. Weather. Rev.* **2012**, *140*, 1988–2004. [CrossRef]

7. Sklar, A. Random variables, joint distribution functions, and copulas. *Kybernetika* **1973**, *9*, 449–460.

8. Isaaks, E.H.; Srivastava, R.M. *Applied Geostatistics*; Oxford University Press: New York, NY, USA, 1989.

9. Bertino, L.; Evensen, G.; Wackernagel, H. Sequential Data Assimilation Techniques in Oceanography. *Int. Stat. Rev.* **2003**, *71*, 223–241. [CrossRef]

10. Simon, E.; Bertino, L. Application of the Gaussian anamorphosis to assimilation in a 3D coupled physical-ecosystem model of the North Atlantic with the EnKF: a twin experiment. *Ocean. Sci.* **2009**, *5*, 495–510. [CrossRef]

11. Xu, T.; Gomez-Hernandez, J. Characterization of non-Gaussian conductivities and porosities with hydraulic heads, solute concentrations, and water temperatures. *Water Resour. Res.* **2016**, *52*, 6111–6136. [CrossRef]

12. Gu, Y.; Oliver, D. An Iterative Ensemble Kalman Filter for Multiphase Fluid Flow Data Assimilation. *SPE J.* **2007**, *12*, 438–446. [CrossRef]

13. Evensen, G. Analysis of iterative ensemble smoothers for solving inverse problems. *Comput. Geosci.* **2018**, *22*, 885–908. [CrossRef]

14. Dovera, L.; Della Rossa, E. Multimodal ensemble Kalman filtering using Gaussian mixture models. *Comput. Geosci.* **2010**, *15*, 307–323. [CrossRef]

15. Rimstad, K.; Omre, H. Approximate posterior distributions for convolutional two-level hidden Markov models. *Comput. Stat. Data Anal.* **2013**, *58*, 187–200. [CrossRef]

16. Grana, D.; Fjeldstad, T.; Omre, H. Bayesian Gaussian Mixture Linear Inversion for Geophysical Inverse Problems. *Math. Geosci.* **2017**, *49*, 493–515. [CrossRef]

17. Besag, J. Spatial interaction and the statistical analysis of lattice systems. *J. R. Stat. Soc. Ser.* **1974**, *36*, 192–236. [CrossRef]

18. Le Loc'h, G.; Beucher, H.; Galli, A.; Doligez, B. Improvement In The Truncated Gaussian Method: Combining Several Gaussian Functions. In *Proceedings of the ECMOR IV—4th European Conference on the Mathematics of Oil Recovery*; European Association of Geoscientists & Engineers: Røros, Norway, 1994. [CrossRef]

19. Oliver, D.; Chen, Y. Data Assimilation in Truncated Plurigaussian Models: Impact of the Truncation Map. *Math. Geosci.* **2018**, *50*, 867–893. [CrossRef]

20. Arellano-Valle, R.B.; Branco, M.D.; Genton, M.G. A unified view on skewed distributions arising from selections. *Can. J. Stat.* **2006**, *34*, 581–601. [CrossRef]

21. Omre, H.; Rimstad, K. Bayesian Spatial Inversion and Conjugate Selection Gaussian Prior Models. *arXiv* **2018**, arXiv:1812.01882.

22. Conjard, M.; Omre, H. Spatio-temporal Inversion using the Selection Kalman Model. *arXiv* **2020**, arXiv:stat.ME/2006.14343

23. Cappé, O.; Moulines, E.; Ryden, T. *Inference in Hidden Markov Models (Springer Series in Statistics)*; Springer-Verlag: Berlin/Heidelberg, Germany, 2005.

24. Moja, S.; Asfaw, Z.; Omre, H. Bayesian Inversion in Hidden Markov Models with Varying Marginal Proportions. *Math. Geosci.* **2018**, *51*, 463–484. [CrossRef]

25. Evensen, G. *Data Assimilation*; Springer-Verlag: Berlin Heidelberg, Germany, 2006; Volume 307. [CrossRef]

26. Evensen, G. The Ensemble Kalman filter: Theoretical Formulation and Practical Implementation. *Ocean. Dyn.* **2003**, *53*, 343–367. [CrossRef]

27. Burgers, G.; Van Leeuwen, P.J.; Evensen, G. On the Analysis Scheme in the Ensemble Kalman Filter. *Mon. Weather. Rev.* **1998**, *126*, 1719–1724. [CrossRef]

28. Hyndman, R. Computing and Graphing Highest Density Regions. *Am. Stat.* **1996**, *50*, 120–126.

29. Gaspari, G.; Cohn, S. Quarterly Journal of the Royal Meteorological Society. *J. Comput. Graph. Stat.* **1999**, *125*, 723–757.

Ensemble-based seismic and production data assimilation using the selection Kalman model

Maxime Conjard and Dario Grana

Ensemble-Based Seismic and Production Data Assimilation Using Selection Kalman Model

Maxime Conjard[1] · **Dario Grana**[2]

Abstract Data assimilation in reservoir modeling often involves model variables that are multimodal, such as porosity and permeability. Well established data assimilation methods such as ensemble Kalman filter and ensemble smoother approaches, are based on Gaussian assumptions that are not applicable to multimodal random variables. The selection ensemble smoother is introduced as an alternative to traditional ensemble methods. In the proposed method, the prior distribution of the model variables, for example the porosity field, is a selection-Gaussian distribution, which allows modeling of the multimodal behavior of the posterior ensemble. The proposed approach is applied for validation on a two-dimensional synthetic channelized reservoir. In the application, an unknown reservoir model of porosity and permeability is estimated from the measured data. Seismic and production data are assumed to be repeatedly measured in time and the reservoir model is updated every time new data are assimilated. The example shows that the selection ensemble Kalman model improves the characterisation of the bimodality of the model parameters compared to the results of the ensemble smoother.

Keywords Selection Kalman model · Seismic inversion · History matching · Ensemble smoother · Channelized reservoir · Multimodality

Maxime Conjard
maxime.conjard@ntnu.no

[1]   Department of Mathematical Sciences, NTNU, Trondheim, Norway

[2]   Department of Geology and Geophysics, School of Energy Resources, University of Wyoming, Laramie, WY, USA

🖄 Springer

1 Introduction

The predictions of fluid flow reservoir models are generally affected by the uncertainty in the initial reservoir model of petrophysical properties, such as porosity and saturation. This model is usually built by integrating different types of data (core measurements, well logs, and surface seismic measurements) and by using geophysical data and statistical algorithms. The so-obtained static reservoir model is then used as an input for fluid flow simulators to predict the dynamic changes in fluid saturation during injection and production. However, there are several sources of uncertainties that might impact the accuracy of the static reservoir model, and as a consequence they affect the accuracy of the fluid flow predictions. Probability random fields can be used to represent the model uncertainty. When additional data become available, including production data or repeated seismic surveys, the static model is updated to reduce the initial uncertainty and improve the reservoir description (Oliver et al. 2008; Evensen 2009).

During injection and production, direct observations of the fluid saturation can be acquired at the well locations. However, this information is available only at sparse locations and cannot fully explain the local variations far away from the wells. Seismic data, repeatedly acquired at multiple times, are used to monitor the changes in seismic response during injection and production (Landrø 2001; Landrø et al. 2003; Dadashpour et al. 2010; Trani et al. 2012). The changes in fluid saturation affect the elastic response of the reservoir. Therefore, by measuring seismic and well data at different times, rock and fluid properties can be predicted by solving inverse problems. In reservoir modeling, this process is called history matching (Williams et al. 1998), which is a data assimilation problem where a set of unknown model variables are updated every time new data are available (Evensen 2009). Time-lapse seismic data have been previously used in history matching in reservoir applications (Huang et al. 1997; Landa et al. 1997; Aanonsen et al. 2003; Kretz et al. 2002; Dong et al. 2005). The integration of seismic data leads to several challenges associated with the resolution of the data and the dimensionality of the problem. In this work, the measured data include oil production and bottom hole pressure at the wells and surface seismic data, and the unknown model variables are the porosity and saturation fields, but the process can be extended to other rock and fluid properties.

Several methods have been presented to solve data assimilation problems, such as gradient-based techniques (e.g. Gauss–Newton), derivative-free optimization methods (e.g. Hooke–Jeeves direct search, genetic algorithm, and particle swarm optimization) and Monte Carlo methods (e.g. ensemble Kalman filter and particle filtering) as described in Oliver and Chen (2011). However, one of the main challenges in seismic and production data assimilation is the representation of the uncertainty. The Ensemble Kalman filter (EnKF) is a Bayesian updating method that relies on the Kalman filter predictor-corrector structure (Evensen 2009). This method has been successfully applied to several geoscience problems including subsurface monitoring (Nævdal et al. 2003; Oliver et al. 2008; Evensen 2009; Oliver and Chen 2011). Skjervheim et al. (2005) applied the EnKF approach to time-lapse seismic and production data in a three-dimensional field case.

Several variations of the ensemble approach have been presented to improve the accuracy of the predictions, including the ensemble smoother and the ensemble smoother with multiple data assimilation (Emerick and Reynolds 2013). Numerous attempts focusing on reproducing the uncertainty in the Kalman gain estimate have also been presented (Myrseth et al. 2010; Sætrom and Omre 2011). Mathematical formulations including localization, covariance inflation, and dimensionality reduction have also been tested for preventing overfitting, ensemble collapsing and spurious correlations and to make the approach applicable to practical applications, including non-linear models in large grids. However, ensemble based methods are generally based on Gaussian assumptions for the conditional distribution of the model, and therefore it is not theoretically applicable to models with multimodal random variables. A regression towards the mean occurs during the conditioning steps, thereby rendering posterior marginal distributions Gaussian. This is a challenging problem in reservoir simulation because subsurface properties, such as porosity and saturation, often show non-Gaussian distributions owing to the underlying geology. The focus on this study is on representing non-Gaussian features in the posterior ensemble of the variables of interest.

Statistical methods, such as Ensemble Randomized Likelihood (EnRML) (Chen and Oliver 2012), Gaussian anamorphosis (Bertino et al. 2003; Zhou et al. 2012), Gaussian mixture models (Dovera and Della Rossa 2011) and truncated pluri-Gaussian (Oliver and Chen 2018), have been developed to address this issue. EnRML and Gaussian anamorphosis attempt to modify the conditioning step to preserve the desirable features in the posterior ensemble, and they are compatible with any prior ensemble. Gausssian mixture and truncated pluri-Gaussian models assume specific prior distribution, and the ensemble algorithm is adapted to the prior model. Recent developments also include the indicator based data assimilation proposed in Kumar and Srinivasan (2018) and the hierarchical truncated pluri-Gaussian simulation (Silva and Deutsch 2019).

An alternative to Gaussian mixture models is the selection-Gaussian models (Arellano-Valle et al. 2006; Arellano-Valle and del Pino 2004; Omre and Rimstad 2018). These models have been applied to spatio-temporal inversion in Conjard and Omre (2020) and validated using synthetic examples. In this manuscript, the compatibility of the selection Kalman model framework detailed in Conjard and Omre (2020) is extended by integrating the selection Kalman model with the ensemble smoother. The resulting approach is named 'selection ensemble smoother'. The prior for the initial state vector is a selection-Gaussian distribution, but instead of sampling the initial ensemble directly from the selection-Gaussian distribution, the proposed approach takes advantage of the structure of the selection-Gaussian to sample from a Gaussian augmented state vector, thereby preventing unwanted regression towards the mean while keeping the computational cost to a minimum. A similar approach has been proposed for asymmetric priors using the closed skew Gaussian distribution (Naveau et al. 2005; Rezaie and Eidsvik 2014).

The proposed methodology is then tested on a two-dimensional channelized reservoir, with the goal of estimating the porosity and saturation fields based on seismic and production data. In Sect. 2, the problem is defined in a Bayesian formulation, the novel method is presented and the forward operator for the reservoir model is defined.

In Sect. 3, a synthetic two-dimensional case study featuring a channelized reservoir with two facies is presented. In Sect. 4, conclusions are presented.

2 Methodology

The Methodology section describes the proposed framework for data assimilation in reservoir modeling. The Bayesian formulation of the selection ensemble smoother is first presented. Then, the forward operators for the geophysics and fluid flow modeling are summarized.

In this paper, $f(\boldsymbol{y})$ denotes the probability density function (pdf) of a random variable $\boldsymbol{y}$, $\varphi_n(\boldsymbol{y}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the pdf of the Gaussian n-vector $\boldsymbol{y}$ with expectation n-vector $\boldsymbol{\mu}$ and covariance $(n \times n)$-matrix $\boldsymbol{\Sigma}$. Furthermore, $\Phi_n(A; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the probability of the aforementioned Gaussian n-vector $\boldsymbol{y}$ to be in $A \subset \mathbb{R}^n$. Further, $\boldsymbol{i}_n$ is used to denote the all-ones n-vector.

2.1 Hidden Markov Model

From a Bayesian perspective, the data assimilation can be assimilated in a hidden Markov model formulation, where at a given time the current state is updated based on the previous state and the current measurements. The state of the chain represents the values of the variables of interest. If a variable of interest is $\boldsymbol{r}$ and the observable measurement is $\boldsymbol{d}$, the goal of data assimilation is to predict the initial value of the variable of interest, i.e. the state of the chain $\boldsymbol{r}_0 = \boldsymbol{r}_{t=0}$ based on the data $\boldsymbol{d}_t$ available at different times $t = 0, \ldots, T$, assuming that the forward model and the likelihood function of the measurements are known.

Given the prior model of the initial state vector $\boldsymbol{r}_0 \in \mathbb{R}^n$, the distribution of the vector of the following states $[\boldsymbol{r}_{1:T+1}|\boldsymbol{r}_0] = [\boldsymbol{r}_1, \ldots, \boldsymbol{r}_{T+1}|\boldsymbol{r}_0]$ is defined as

$$f(\boldsymbol{r}_{1:T+1}|\boldsymbol{r}_0) = \prod_{t=0}^{T} f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t), \tag{1}$$

with

$$[\boldsymbol{r}_{t+1}|\boldsymbol{r}_t] = \omega_t(\boldsymbol{r}_t) + \boldsymbol{\epsilon}_t^r \sim f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t), \tag{2}$$

where $\omega_t(\cdot) \in \mathbb{R}^n$ is the forward operator with random Gaussian n-vector $\boldsymbol{\epsilon}_t^r$, independent and identically distributed (iid) for each time t. The forward model defines a first-order Markov chain.

Consider, then, a set of measurements $\boldsymbol{d} = \{\boldsymbol{d}_1, \ldots, \boldsymbol{d}_T\}$ where $\boldsymbol{d}_t \in \mathbb{R}^m, \forall t = 1, \ldots, T$. It is assumed that the measured data depend on the variables of interest according to a likelihood model. The likelihood model for $[\boldsymbol{d}|\boldsymbol{r}]$ is conditionally independent and defined as

$$f(\boldsymbol{d}|\boldsymbol{r}) = \prod_{t=0}^{T} f(\boldsymbol{d}_t|\boldsymbol{r}_t), \tag{3}$$

with

$$[\boldsymbol{d}_t|\boldsymbol{r}_t] = \psi_t(\boldsymbol{r}_t) + \boldsymbol{\epsilon}_t^d \sim f(\boldsymbol{d}_t|\boldsymbol{r}_t), \tag{4}$$

where $\psi_t(\cdot, \cdot) \in \mathbb{R}^m$ is the likelihood function with random Gaussian m-vector $\boldsymbol{\epsilon}_t^d$, iid for each t.

Let $\Phi_t = \omega_t \circ \omega_{t-1} \circ \ldots \omega_1$ such that

$$[\boldsymbol{r}_{t+1}|\boldsymbol{r}_0] = \Phi_t(\boldsymbol{r}_0) + \boldsymbol{\epsilon}_t^r, \tag{5}$$

then, it is possible to write $[\boldsymbol{d}_t|\boldsymbol{r}_0]$ as

$$[\boldsymbol{d}_t|\boldsymbol{r}_0] = \psi_t \circ \Phi_{t-1}(\boldsymbol{r}_0) + \boldsymbol{\epsilon}_t^d, \tag{6}$$

and consequently $[\boldsymbol{d}|\boldsymbol{r}_0]$ as

$$[\boldsymbol{d}|\boldsymbol{r}_0] = \begin{bmatrix} \psi_1 \circ \Phi_0(\boldsymbol{r}_0) \\ \ldots \\ \psi_T \circ \Phi_{T-1}(\boldsymbol{r}_0) \end{bmatrix} + \begin{bmatrix} \boldsymbol{\epsilon}_1^d \\ \ldots \\ \boldsymbol{\epsilon}_T^d \end{bmatrix}. \tag{7}$$

Equations 1 and 4 provide the framework for the EnKF, whereas Eq. 7 gives the framework for the ensemble smoother using the same hidden Markov model formulation.

2.2 Ensemble Smoother

The ensemble smoother is an inverse method that assimilates the measured data $\boldsymbol{d}$ in a unique step by evaluating the likelihood $f(\boldsymbol{r}_0|\boldsymbol{d})$. The likelihood model described in Eq. 7 is represented as

$$\boldsymbol{d} = g(\boldsymbol{r}_0) + \boldsymbol{\delta}, \tag{8}$$

where $g(\boldsymbol{r}_0) = \begin{bmatrix} \psi_1 \circ \Phi_0(\boldsymbol{r}_0) \\ \ldots \\ \psi_T \circ \Phi_{T-1}(\boldsymbol{r}_0) \end{bmatrix}$ and $\boldsymbol{\delta} = \begin{bmatrix} \boldsymbol{\epsilon}_1^d \\ \ldots \\ \boldsymbol{\epsilon}_T^d \end{bmatrix}$.

In this data assimilation approach, an ensemble of n_e members $\boldsymbol{r}^{p(i)}$, $i = 1, \ldots, n_e$ is generated from the prior distribution of the initial state $f(\boldsymbol{r}_0)$ and the ensemble members are updated as

$$\boldsymbol{r}^{u(i)} = \boldsymbol{r}^{p(i)} + \hat{\boldsymbol{\Gamma}}_{r,d} \hat{\boldsymbol{\Sigma}}_d^{-1} (\boldsymbol{d} - \boldsymbol{d}^i), i = 1, \ldots, n_e, \tag{9}$$

where $\boldsymbol{d}^i$ represents the pseudo-data (i.e. the data predicted from the model according to the forward operator) generated for the ensemble member i as $\boldsymbol{d}^{(i)} = g(\boldsymbol{r}^{p(i)}) + \boldsymbol{\delta}$, and where the crosscovariance and covariance matrices, $\hat{\boldsymbol{\Gamma}}_{r,d}$ and $\hat{\boldsymbol{\Sigma}}_d$, are computed as

$$\hat{\boldsymbol{\Gamma}}_{r,d} = \frac{1}{n_e} \sum_{i=1}^{n_e} (\boldsymbol{r}^{p(i)} - \overline{\boldsymbol{r}}^p)(\boldsymbol{d}^i - \overline{\boldsymbol{d}})^T, \tag{10}$$

$$\hat{\boldsymbol{\Sigma}}_d = \frac{1}{n_e} \sum_{i=1}^{n_e} (\boldsymbol{d}^i - \overline{\boldsymbol{d}})(\boldsymbol{d}^i - \overline{\boldsymbol{d}})^T, \tag{11}$$

with $\overline{\boldsymbol{r}}^p$ being the empirical mean from the initial ensemble and $\overline{\boldsymbol{d}}$ the empirical mean of the pseudo-data. Localization has been proposed in data assimilation problems to avoid spurious correlations and ensemble collapse (Houtekamer and Mitchell 1998; Chen and Oliver 2011). The method is based on distance-dependent functions to constrain the updates of the model variables and their variability to a specific area based on the observed data in that region. Localization should be applied on a scale larger than the correlation length of the spatial variograms of the model variables to avoid introducing biases on the spatial correlation and resolution of the model. In the case study, localization is used on both $\hat{\boldsymbol{\Gamma}}_{r,d}$ and $\hat{\boldsymbol{\Sigma}}_d$ because, as is often the case, the size of the ensemble is small compared to the dimension of the state vector.

One of the most popular inverse methods in reservoir modeling simulation is the ensemble smoother with multiple data assimilation (ES-MDA) (Emerick and Reynolds 2013). In the MDA version of the ensemble smoother, the data are assimilated multiple times to improve convergence. This approach requires a limited number of runs of the forward operator and handles the non-linearity in the forward and likelihood models. The mathematical formulation of the MDA approach is detailed in Emerick and Reynolds (2013).

In the ensemble smoother, the prior distribution for the initial state can be chosen freely, however it is difficult to preserve non-Gaussian features in the posterior distribution owing to the linearized update (Eq. 9). The selection ensemble smoother attempts to represent non-Gaussian features in the posterior distribution.

2.3 Selection Ensemble Smoother

The selection-Gaussian pdf can represent multimodality, skewness, and heavy-tailed distributions; it is an extension of the skew-Gaussian distribution (Omre and Rimstad 2018). Consider a Gaussian $(n + n)$-vector $[\tilde{\boldsymbol{r}}, \boldsymbol{\kappa}]$

$$\begin{bmatrix} \tilde{\boldsymbol{r}} \\ \boldsymbol{\kappa} \end{bmatrix} \sim \varphi_{2n} \left(\begin{bmatrix} \tilde{\boldsymbol{r}} \\ \boldsymbol{\kappa} \end{bmatrix}; \begin{bmatrix} \boldsymbol{\mu}_{\tilde{r}} \\ \boldsymbol{\mu}_{\kappa} \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{\tilde{r}} & \boldsymbol{\Sigma}_{\tilde{r}} \boldsymbol{\Gamma}_{\kappa|\tilde{r}}^T \\ \boldsymbol{\Gamma}_{\kappa|\tilde{r}} \boldsymbol{\Sigma}_{\tilde{r}} & \boldsymbol{\Sigma}_{\kappa} \end{bmatrix} \right), \tag{12}$$

where $\boldsymbol{\mu}_{\tilde{r}}$ and $\boldsymbol{\mu}_{\kappa}$ are n-vectors, $\boldsymbol{\Gamma}_{\kappa|\tilde{r}}$ is a $(n \times n)$-matrix, and $\boldsymbol{\Sigma}_{\tilde{r}}$, $\boldsymbol{\Sigma}_{\kappa}$, and $\boldsymbol{\Sigma}_{\kappa|\tilde{r}}$ are three covariance $(n \times n)$-matrices with $\boldsymbol{\Sigma}_{\kappa} = \boldsymbol{\Gamma}_{\kappa|\tilde{r}} \boldsymbol{\Sigma}_{\tilde{r}} \boldsymbol{\Gamma}_{\kappa|\tilde{r}}^T + \boldsymbol{\Sigma}_{\kappa|\tilde{r}}$.

If $\boldsymbol{r}_0$ is defined as $\boldsymbol{r}_0 = [\tilde{\boldsymbol{r}}|\boldsymbol{\kappa} \in A]$, where $A \subset \mathbb{R}^n$ is a selection set of dimension n, then $\boldsymbol{r}_0$ is distributed according to a selection-Gaussian pdf given by

$$f(\boldsymbol{r}_0) = \left[\Phi_n(A; \boldsymbol{\mu}_\kappa, \boldsymbol{\Sigma}_\kappa)\right]^{-1}$$
$$\Phi_n(A; \boldsymbol{\mu}_\kappa + \boldsymbol{\Gamma}_{\kappa|\tilde{r}}(\boldsymbol{r}_0 - \boldsymbol{\mu}_{\tilde{r}}), \boldsymbol{\Sigma}_{\kappa|\tilde{r}}) \times \varphi_n(\boldsymbol{r}_0; \boldsymbol{\mu}_{\tilde{r}}, \boldsymbol{\Sigma}_{\tilde{r}}). \qquad (13)$$

The structure of the selection-Gaussian model provides a suitable model for non-Gaussian features of the posterior distribution of the inverse problem.

A selection-Gaussian random variable $\boldsymbol{r}_A \in \mathbb{R}^n$ can be written as $[\boldsymbol{r}|\boldsymbol{\kappa} \in A]$ with $\boldsymbol{r} \in \mathbb{R}^n$, $\boldsymbol{\kappa} \in \mathbb{R}^n$, $A \subset \mathbb{R}^n$, and where $[\boldsymbol{r}, \boldsymbol{\kappa}]$ is jointly Gaussian. When conditioning on data, it is always possible to consider $[\boldsymbol{r}, \boldsymbol{\kappa}]$, condition first on the data $\boldsymbol{d}$ and then on $\boldsymbol{\kappa} \in A$. Similarly when transforming $\boldsymbol{r}_A :\mapsto g(\boldsymbol{r}_A)$, one can show that the distribution of $g(\boldsymbol{r}_A))$ is equal to that of $[g(\boldsymbol{r})|\boldsymbol{\kappa} \in A]$. In any data assimilation process where a selection-Gaussian distribution is chosen as a prior, it is therefore possible to work on the augmented Gaussian vector $[\boldsymbol{r}, \boldsymbol{\kappa}]$ throughout the assimilation step(s), and condition on the latent variable $\boldsymbol{\kappa}$ being in A afterwards. The selection Ensemble Kalman filter detailed in (Conjard and Omre 2020) adopts this approach.

The proposed selection ensemble smoother for data assimilation problems is defined by a prior distribution $f(\boldsymbol{r}_0)$ that is selection-Gaussian such that $\boldsymbol{r}_0 = [\tilde{\boldsymbol{r}}_0|\boldsymbol{\kappa} \in A]$, $A \subset \mathbb{R}^n$, and where the prior state vector is represented by the augmented vector $[\tilde{\boldsymbol{r}}_0, \boldsymbol{\kappa}]$. When assessing $[\boldsymbol{r}_0|\boldsymbol{d}_{1:T}]$, the selection ensemble smoother first evaluates $[\tilde{\boldsymbol{r}}_0, \boldsymbol{\kappa}|\boldsymbol{d}_{1:T}]$ and then adopts a Markov chain Monte-Carlo (MCMC) algorithm (Omre and Rimstad 2018) to calculate $[\boldsymbol{r}_0|\boldsymbol{d}_{1:T}] = [\tilde{\boldsymbol{r}}_0|\boldsymbol{\kappa} \in A, \boldsymbol{d}_{1:T}]$. There are two advantages in using the augmented state vector. First, in ensemble methods, the update step is optimal if the prior is Gaussian and the prior and the data are jointly Gaussian. Therefore, by enforcing a Gaussian initial state vector, the optimal update conditions are ensured. Then, since the update is optimal only if the prior is Gaussian and the forward and likelihood models linear and Gaussian, the traditional update tends to make the posterior ensemble more and more Gaussian; however, conditioning on the latent variable after the assimilation is performed allows representing non-Gaussian features in the posterior distribution.

The selection ensemble smoother proceeds as follows. First, n_e ensemble members are generated

$$[\tilde{\boldsymbol{r}}^{p(i)}, \boldsymbol{\kappa}^{p(i)}], i = 1, \ldots, n_e, \qquad (14)$$

from the augmented Gaussian prior distribution of the initial state $f(\boldsymbol{r}_0, \boldsymbol{\kappa})$, and they are updated as

$$\begin{bmatrix} \tilde{\boldsymbol{r}}^{u(i)} \\ \boldsymbol{\kappa}^{u(i)} \end{bmatrix} = \begin{bmatrix} \tilde{\boldsymbol{r}}^{p(i)} \\ \boldsymbol{\kappa}^{p(i)} \end{bmatrix} + \hat{\boldsymbol{\Gamma}}_{\tilde{r}\kappa,d}\hat{\boldsymbol{\Sigma}}_d^{-1}(\boldsymbol{d} - \boldsymbol{d}^i), i = 1, \ldots, n_e, \qquad (15)$$

where $\boldsymbol{d}^i$ represents the pseudo-data generated for the ensemble member i as $\boldsymbol{d}^{(i)} = g(\tilde{\boldsymbol{r}}^{p(i)}) + \boldsymbol{\delta}$, and where $\hat{\boldsymbol{\Gamma}}_{\tilde{r}\kappa,d}$ and $\hat{\boldsymbol{\Sigma}}_d$ are calculated as follows

$$\hat{\boldsymbol{\Gamma}}_{\tilde{r}\kappa,d} = \frac{1}{n_e} \sum_{i=1}^{n_e} \left(\begin{bmatrix} \tilde{\boldsymbol{r}}^{p(i)} \\ \kappa^{p(i)} \end{bmatrix} - \begin{bmatrix} \overline{\tilde{\boldsymbol{r}}}^p \\ \overline{\kappa}^p \end{bmatrix} \right) (\boldsymbol{d}^i - \overline{\boldsymbol{d}})^T, \tag{16}$$

$$\hat{\boldsymbol{\Sigma}}_d = \frac{1}{n_e} \sum_{i=1}^{n_e} (\boldsymbol{d}^i - \overline{\boldsymbol{d}})(\boldsymbol{d}^i - \overline{\boldsymbol{d}})^T, \tag{17}$$

with $\overline{\tilde{\boldsymbol{r}}}^p$ and $\overline{\kappa}^p$ being the empirical means from each component of the initial ensemble and $\overline{\boldsymbol{d}}$ being the empirical mean of the pseudo-data. In the case study, localization is used on both $\hat{\boldsymbol{\Gamma}}_{\tilde{r}\kappa,d}$ and $\hat{\boldsymbol{\Sigma}}_d$ because the size of the ensemble is small compared to the dimension of the state vector.

After conditioning on the data, the ensemble represents $[\boldsymbol{r}_0, \kappa | \boldsymbol{d}]$ which is assumed to be Gaussian. The target distribution $f(\boldsymbol{r}_0 | \kappa \in A, \boldsymbol{d})$ can be written as

$$f(\boldsymbol{r}_0 | \kappa \in A, \boldsymbol{d}) = \frac{1}{f(\kappa \in A | \boldsymbol{d})} \int_A f(\boldsymbol{r}_0 | \kappa, \boldsymbol{d}) f(\kappa | \boldsymbol{d}) d\kappa, \tag{18}$$

$$= \int_{\mathbb{R}^n} f(\boldsymbol{r}_0 | \kappa, \boldsymbol{d}) \frac{f(\kappa, \boldsymbol{d}) \boldsymbol{I}_A(\kappa)}{f(\kappa \in A, \boldsymbol{d})} d\kappa, \tag{19}$$

$$= \int_{\mathbb{R}^n} f(\boldsymbol{r}_0 | \kappa, \boldsymbol{d}) f(\kappa | \kappa \in A, \boldsymbol{d}) d\kappa. \tag{20}$$

Equation 20 shows how sampling from the posterior distribution is performed: first from the truncated Gaussian $\kappa^* \sim [\kappa | \kappa \in A, \boldsymbol{d}]$ and then from the Gaussian vector $\boldsymbol{r}^* \sim [\boldsymbol{r}_0 | \kappa^*, \boldsymbol{d}]$. Hence $\boldsymbol{r}^*$ are distributed according to the target distribution $f(\boldsymbol{r}_0 | \kappa \in A, \boldsymbol{d})$.

The challenging step is the sampling from the truncated Gaussian, which is performed using an MCMC sampling algorithm (Omre and Rimstad 2018). The MCMC algorithm, denoted by S, generates κ^*

$$\kappa^* = S(\hat{\boldsymbol{\mu}}_{\kappa | \boldsymbol{d}}, \hat{\boldsymbol{\Sigma}}_{\kappa | \boldsymbol{d}}, A), \tag{21}$$

where $\hat{\boldsymbol{\mu}}_{\kappa | \boldsymbol{d}}$ and $\hat{\boldsymbol{\Sigma}}_{\kappa | \boldsymbol{d}}$ are the empirical mean and covariance matrix of the posterior ensemble for κ. Note that the selection set A is considered constant throughout the algorithm. The empirical covariance matrix estimate $\hat{\boldsymbol{\Sigma}}_{\kappa | \boldsymbol{d}}$ is known to be subject to sampling errors, which are addressed by using localization on $\hat{\boldsymbol{\Sigma}}_{\kappa | \boldsymbol{d}}$. The algorithm itself has to be tailored to the specific application as rank deficiency is to be expected in high-dimensional inverse problems. The original algorithm expected a positive definite covariance matrix, and even though localization does increase the rank of the localised matrix, it is unlikely that it becomes full rank after the transformation. Therefore, singular value decomposition is used to evaluate the pseudo inverse instead of outright inverting the matrix in the algorithm. The singular value decomposition of

$\hat{\Sigma}_{\kappa|d}$ gives,

$$\hat{\Sigma}_{\kappa|d} = U \Sigma V^T, \tag{22}$$

where U and V are $(n \times n)$ orthogonal matrices, and Σ is an $(n \times n)$ diagonal matrix whose diagonal entries are non-negative. The pseudo inverse $\hat{\Sigma}_{\kappa|d}^+$ of $\hat{\Sigma}_{\kappa|d}$ is then given by

$$\hat{\Sigma}_{\kappa|d}^+ = V \Sigma^+ U^T, \tag{23}$$

where Σ^+ is the pseudoinverse of Σ, which is formed by replacing every non-zero diagonal entry by its inverse and transposing the resulting matrix.

2.4 Forward Operator

The data assimilation problem discussed in this work requires two operators: the fluid flow simulation operator used as a forward model and the geophysical operator used for the likelihood model of the measured data.

The forecast of the dynamic behavior of fluid displacement in the subsurface requires a mathematical-physical model generally called fluid flow simulation. For hydrocarbon-water two-phase flow in subsurface reservoirs, the model is based on Darcy's equation and mass balance law and includes two partial differential equations (PDE) in two unknowns, namely water saturation $s_w(t)$ and water pressure $p_w(t)$. This model assumes that there are only two immiscible fluids, the flow is isothermal, and the capillary pressure is a function of saturation. For a system of oil and water, the PDE system can be written as follows

$$\nabla \cdot \left[\frac{k k_w}{B_w \mu_w} (\nabla p_w(t) - \gamma_w \nabla d) \right] + q_w = \frac{\partial}{\partial t} \left(\frac{\phi s_w(t)}{B_t} \right), \tag{24}$$

$$\nabla \cdot \left[\frac{k k_o}{B_o \mu_o} (\nabla (p_w(t) - p_c(t)) - \gamma_o \nabla d) \right] + q_o = \frac{\partial}{\partial t} \left(\frac{\phi (1 - s_w(t))}{B_o} \right), \tag{25}$$

where k is the absolute permeability, ϕ is the porosity, d is the depth, p_c is the capillary pressure, q_w is the water production rate, q_o is the oil production rate and the constants k_w, B_w, μ_w, γ_w, γ_w, k_o, B_o, μ_o, and γ_o are parameters associated with the fluid properties of oil and water. If porosity and permeability are known, then the system can be solved using finite difference methods to compute the water and saturation and pressure, $s_w(t)$ and $p_w(t)$, at any time t. The oil saturation and pressure can then be computed from water saturation and water pressure.

The permeability field is assumed to be a point-wise function of the porosity field given by Kozeny–Carman's equation

$$k = a \times \frac{\phi^3}{(1 - \phi)^2}, \tag{26}$$

where a is a geometrical factor estimated by fitting the true permeability field. This approximation allows removing $\boldsymbol{k}$ from the state vector because it is then a known function of the porosity $\boldsymbol{\phi}$.

Similarly, the prediction of the geophysical response of subsurface models requires a mathematical-physical model that includes rock physics relations and seismic wave propagation models. Rock physics models allow computing P-wave and S-wave velocities and densities given the petrophysical properties, porosity, and fluid saturation. The seismic wave propagation model allows calculating the seismic response, in terms of amplitude and travel time, based on the velocities and densities. For a partially-saturated porous rock of porosity ϕ and water saturation s_w, the density $\rho(\phi, s_w)$ can be computed as

$$\rho(\phi, s_w) = (1 - \phi)\rho_m + \phi\left(\rho_w s_w + \rho_o(1 - s_w)\right), \tag{27}$$

and the P-wave and S-wave velocities, $v_p(\phi, s_w)$ and $v_s(\phi, s_w)$, are given by Raymer–Dvorkin's relations (Dvorkin et al. 2014)

$$v_p(\phi, s_w) = (1 - \phi)^2 \sqrt{\frac{K_m + \frac{4}{3}G_m}{\rho_m}} + \phi \sqrt{\frac{\left(\frac{s_w}{K_w} + \frac{1-s_w}{K_o}\right)^{-1}}{\rho_w s_w + \rho_o(1 - s_w)}}, \tag{28}$$

$$v_s(\phi, s_w) = (1 - \phi)^2 \sqrt{\frac{G_m}{\rho_m}} \sqrt{\frac{(1 - \phi)\rho_m}{\rho(\phi, s_w)}}, \tag{29}$$

where ρ_m is the density of the solid phase, ρ_w is the density of water, ρ_o is the density of oil, K_m is the bulk modulus of the solid phase, G_m is the shear modulus of the solid phase, K_w is the bulk modulus of water, and K_o is the bulk modulus of oil. All these parameters are assumed to be constant and known. The reflection coefficients $r_{pp}(t, \theta)$ associated with seismic wave propagation are a function of the seismic travel time and incident angle θ and depend on the relative change in P-wave and S-wave velocity and density. For small incident angles, the reflection coefficients can be approximated using Aki–Richards equations (Aki and Richards 1980) as

$$r_{pp}(t, \theta) = r_p(\theta)\frac{\partial}{\partial t}v_p(t) + r_s(\theta)\frac{\partial}{\partial t}v_s(t) + r_\rho(\theta)\frac{\partial}{\partial t}\rho(t), \tag{30}$$

where

$$r_p(\theta) = \frac{1}{2}(1 + \tan^2(\theta)), \tag{31}$$
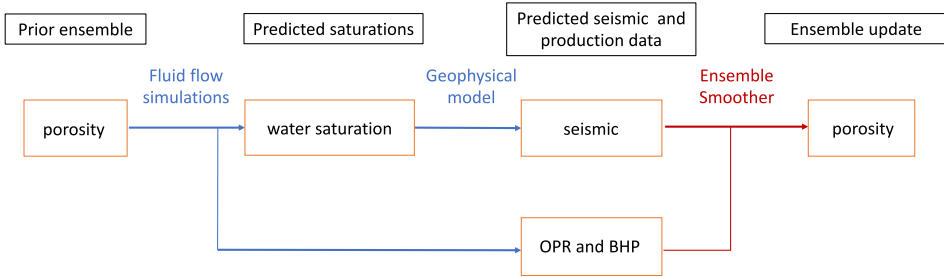
$$r_s(\theta) = -4c\sin^2(\theta), \tag{32}$$

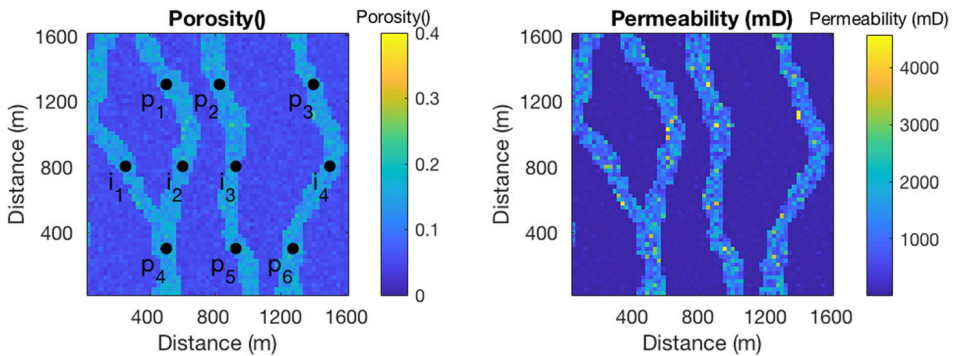$$r_\rho = \frac{1}{2}(1 - 4c\sin^2(\theta)), \tag{33}$$

with c being a constant value equal to the square of the average v_s/v_p ratio.

Fig. 1 Graph of the hidden Markov model

**Fig. 2** Data assimilation flowchart

Fig. 3 Reference porosity and permeability fields, with the location of the injection and producer wells

The amplitudes $s(t, \theta)$ of the seismogram are then approximated as a convolution of a wavelet $w(t, \theta)$ and the reflection coefficients $r_{pp}(t, \theta)$ as

$$s(t, \theta) = \int w(u, \theta) r_{pp}(t - u, \theta) du, \tag{34}$$

where the wavelet $w(t, \theta)$ is generally assumed to be known. As shown in Buland and Omre (2003), the convolution can be discretized by decomposing the matrix associated with the forward operator as $G = WAD$, where W is the wavelet matrix, A is the reflection coefficient matrix, and D is the first order differential matrix.
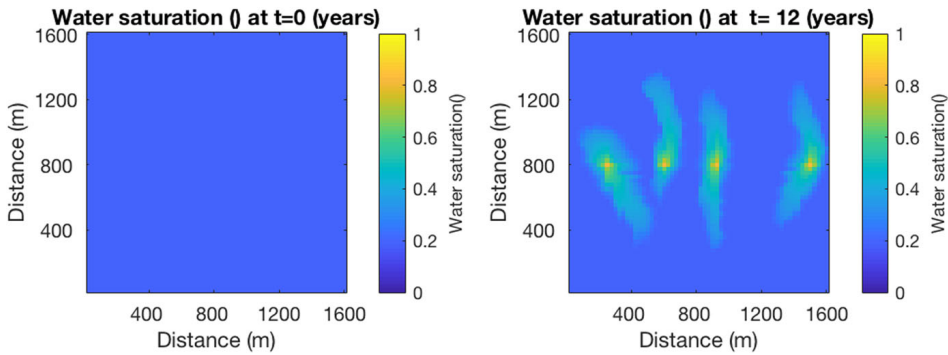
Fig. 4 Water saturation at initial and final times
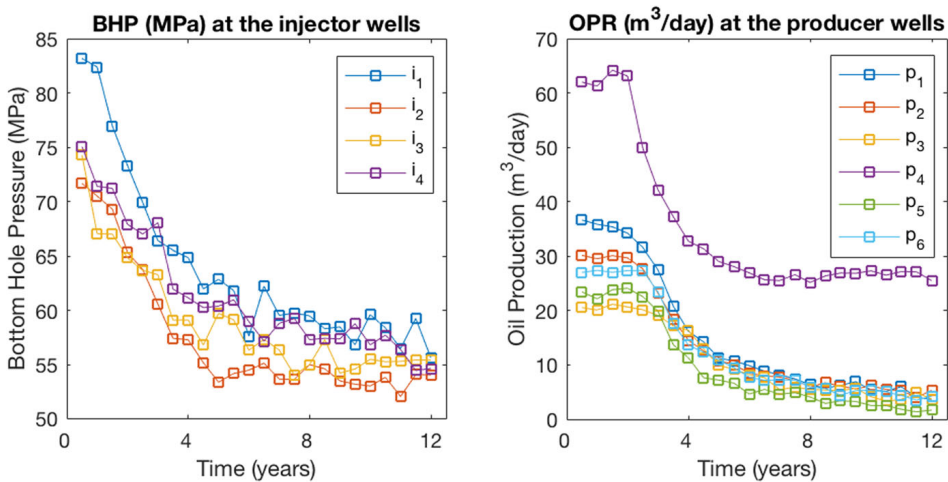
Fig. 5 Well production data: bottom hole pressure and oil production rate

In the current implementation, we adopted a forward model based on the convolution of the wavelet with the linearized approximation of Zoeppritz equations (Shuey 1985; Aki and Richards 1980) because it is mathematically tractable, and the computational cost is limited. The proposed approach can be extended to other forward operations including convolutional models based on full Zoeppritz equations, Kennett's invariant imbedding method and the Born weak scattering approximation (Zoeppritz 1919; Kennett 1984; Russell 1988; Sheriff and Geldart 1995; Yilmaz 2001; Aki and Richards 1980). In theory, the same formulation can also be applied to the full waveform model for wave propagation, however, the computational cost of the approach is still prohibitive for field-scale applications. Applications of ensemble-based methods to full waveform inversion problems have been proposed Thurin et al. (2019) and Gineste et al. (2020), according to Gaussian assumptions of the error and model variables.
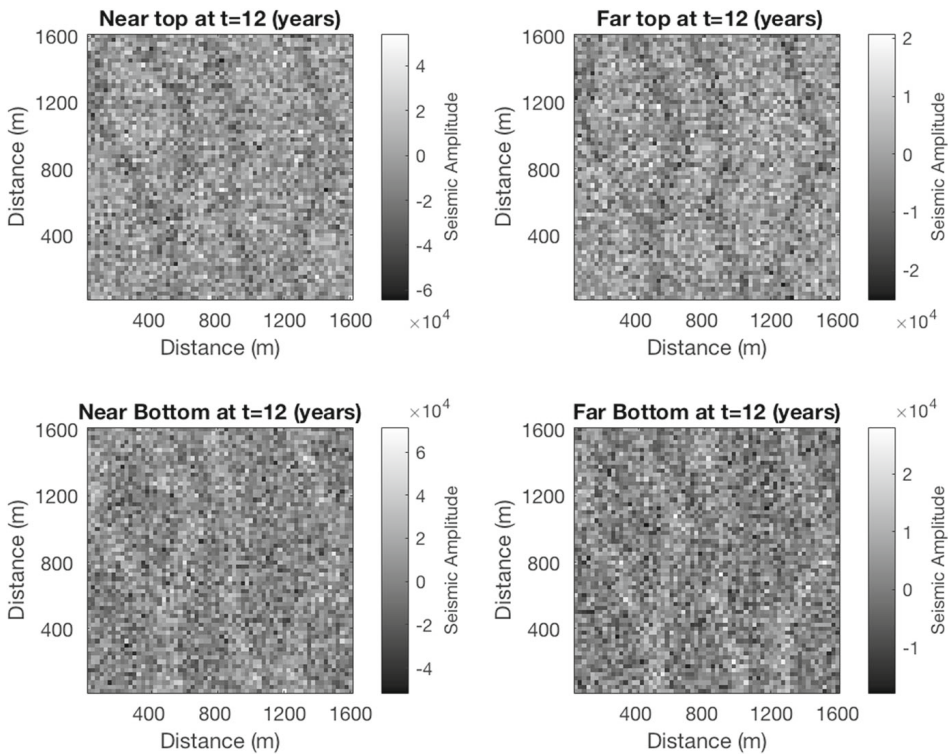
Fig. 6 Seismic data (near and far angle) at the top and bottom of the reservoir after 12 years of injection and production

Table 1 Parameter values

| $\mu_{\bar{r}}^{\phi}()$	$\mu_{\kappa}^{\phi}()$	$\sigma_{\bar{r}}^{\phi}()$	$\gamma^{\phi}()$	A()	$d_x(m)$	$d_y(m)$
-2.5	0	0.2	0.99	$(]-\infty,-0.2] \cap [0.3, 2[)^n$	100	200

3 Application

The proposed method is applied to a two-dimensional synthetic case to demonstrate the validity of the selection ensemble smoother (S-ES) and the advantages compared to the traditional ensemble smoother (ES).

The synthetic example mimics the production of hydrocarbon in a reservoir by water injection as shown in Liu and Grana (2018, 2020). The geology of the reservoir includes four channels in the north-south direction with an average porosity of 0.2. The channels are surrounded by shale with effective porosity close to 0. It is assumed that the entire reservoir is filled by oil with an irreducible water saturation of 0.2. Water is injected in four wells and oil is produced from six wells. All the wells are located within the channels. Fluid flow occurs predominantly within the channels due to the low porosity and permeability of the surrounding shale, and the injected water displaces the oil in place towards the producing wells.
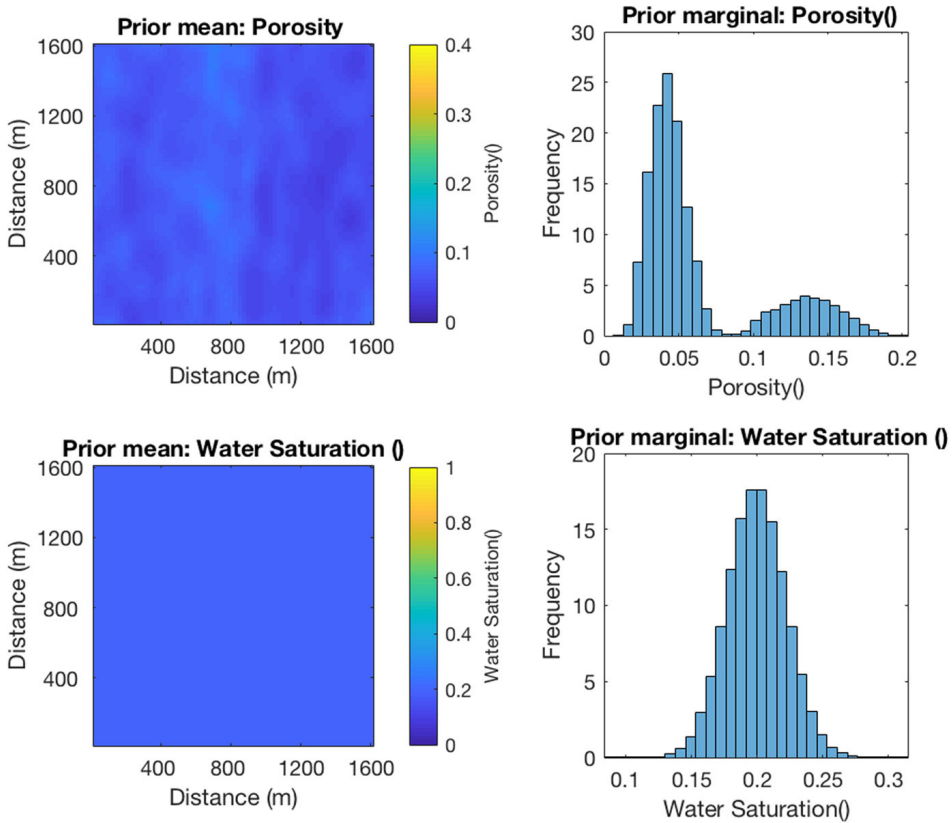
Fig. 7 Prior ensemble of models: mean maps and marginal histograms
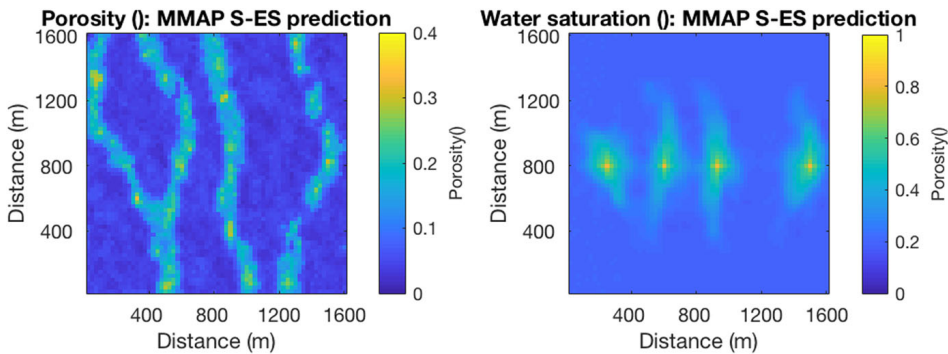
Fig. 8 Posterior ensemble prediction from the selection ensemble smoother: marginal maximum a posteriori (MMAP) predictions

The reservoir model mimics the two-phase (oil and water) flow due to water injection and oil production in a geological system represented in a two-dimensional uniform grid of $n = 64 \times 64$ cells of dimension $40 \times 40 \times 25$ m. Figure 3 shows the true porosity $\boldsymbol{\phi}_r$ and permeability $\boldsymbol{k}_r$ fields. The reference porosity model is generated using the following geostatistical modeling approach: First, a facies model for a channelized
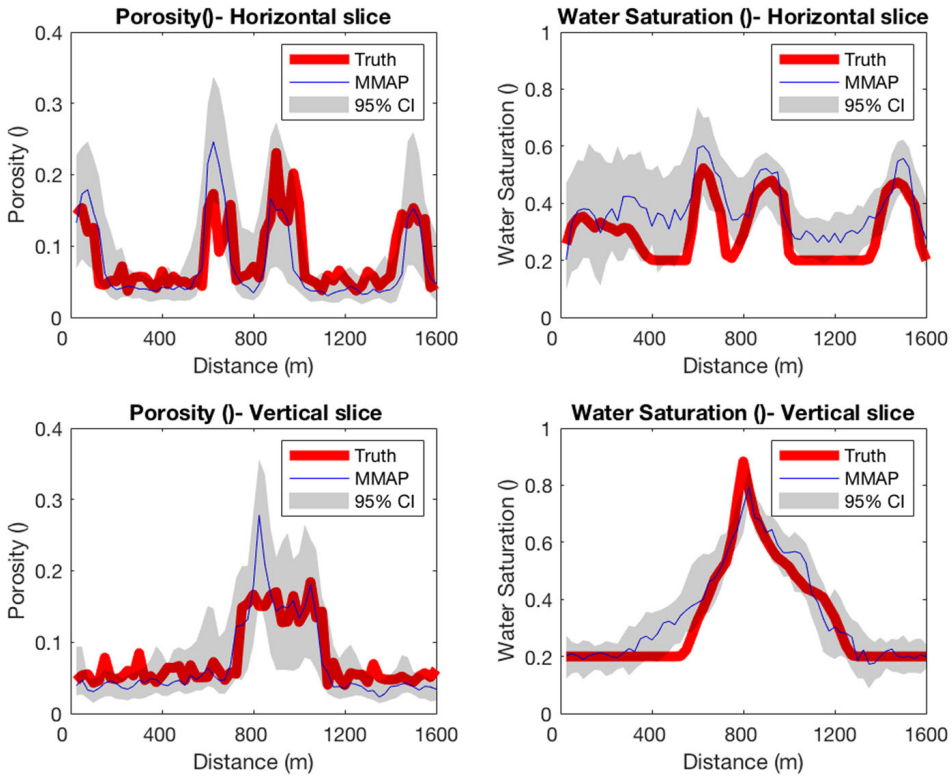
Fig. 9 Posterior ensemble prediction from the selection ensemble smoother: predictions along horizontal and vertical lines

system is generated using object modeling with randomly sampled parameters of the channel geometry (length, width, and tortuosity). Then a porosity model is generated using sequential Gaussian simulation within each facies using facies-dependent prior distributions to mimic medium-high porosity within the channels and low porosity in the background shale. Finally, permeability is simulated in the logarithmic domain, conditioned on the porosity field using sequential Gaussian co-simulation. Figure 3 shows the geostatistical realization used as the reference (true) model for the application. Figure 3 also shows the locations of the four water injection wells (i_1, i_2, i_3, i_4) and the six producer wells ($p_1, \ldots, p_6$).

Fluid flow is then simulated to create the reference (true) saturation maps for a time period of 12 years using the Matlab Reservoir Simulation Toolbox (MRST) (Lie 2019). Figure 4 shows the evolution of the water saturation in the field for the reference porosity and permeability field specified in Fig. 3. The initial water saturation s_{w0} is set constant to 0.2, then water propagates predominantly within the channels until it reaches, after 12 years, one of the producers. Figure 5 shows the bottom hole pressure (BHP) and the oil production rate (OPR) every six months at the injector and producer wells, respectively.

In addition to the saturation and production data, a synthetic geophysical dataset is generated for the data assimilation problem. To generate seismic data, two additional
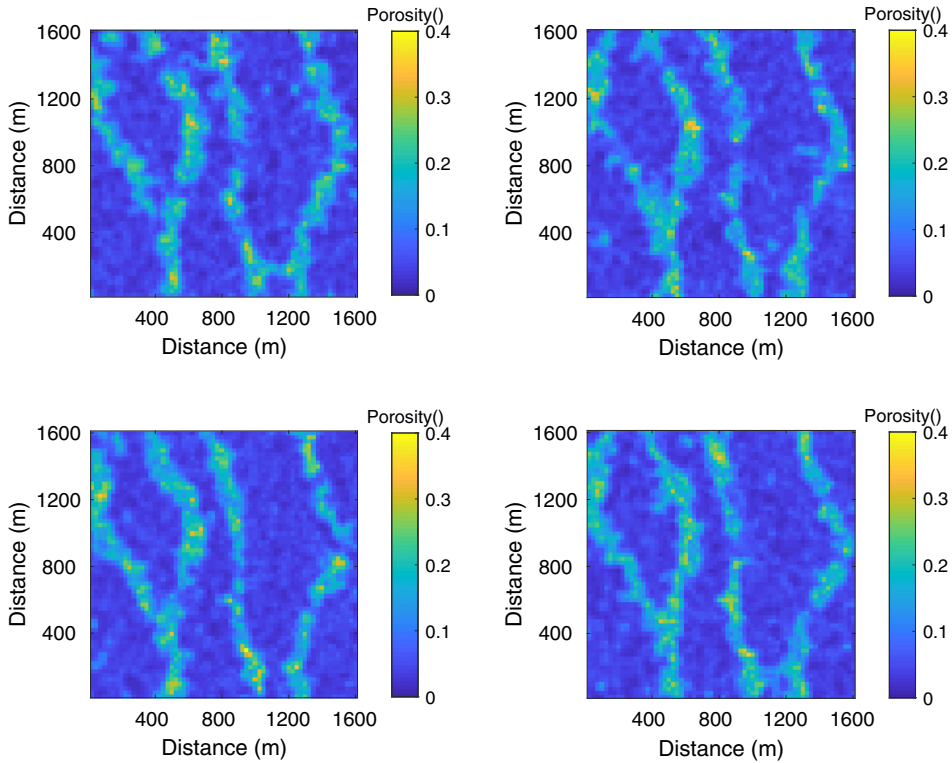
Fig. 10 Posterior realizations from the posterior distribution of the porosity field obtained from the selection ensemble smoother
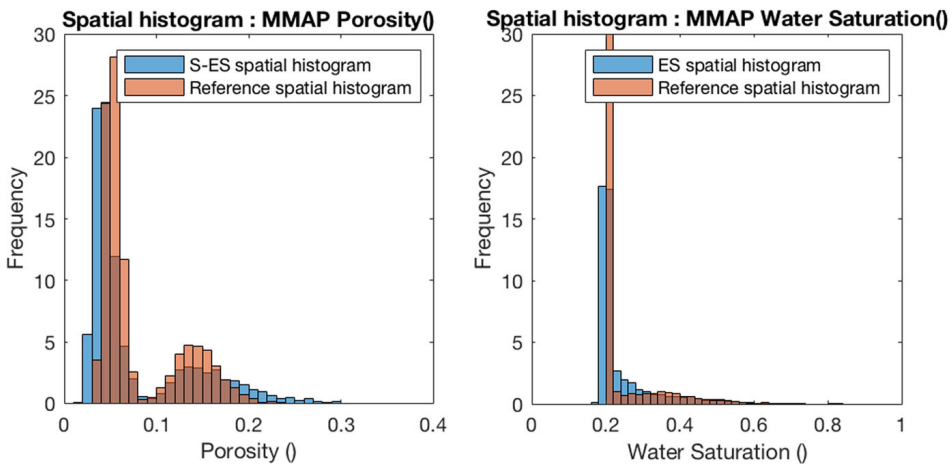
Fig. 11 Posterior statistics of the selection ensemble smoother results compared to the true model

geological layers representing non-permeable shale formations are added above and below the reservoir model. The model then includes three layers (top shale, reservoir, and bottom shale) and two interfaces (top reservoir and bottom reservoir). Seismic data are then generated using the rock physics model with constant mineral properties
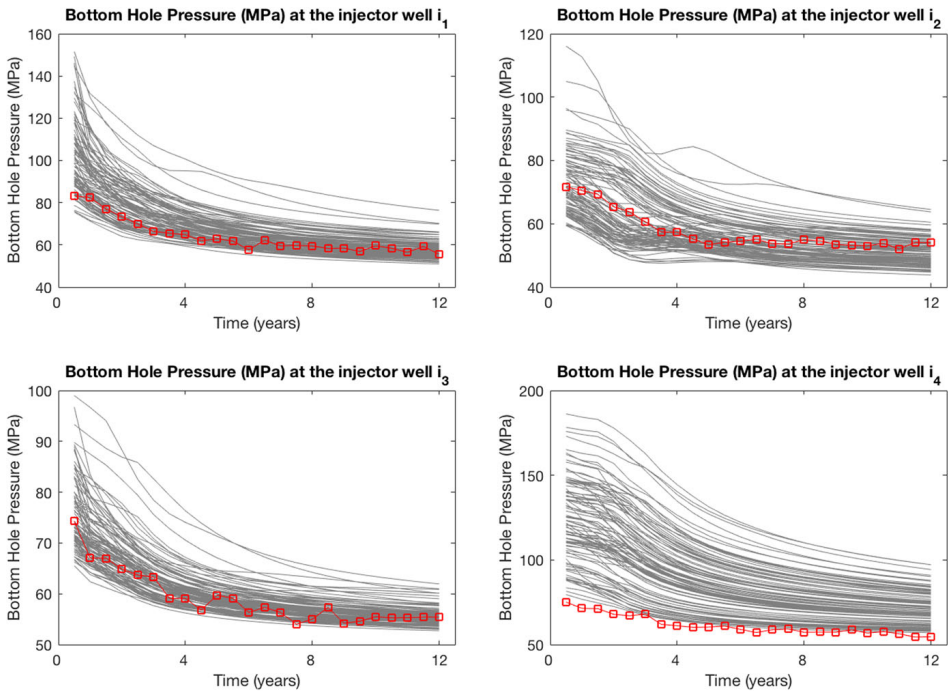
Fig. 12 Predicted bottom hole pressure (BHP) at the 4 injection wells obtained from the selection ensemble smoother; the ensemble members are in black, the measured BHP in red
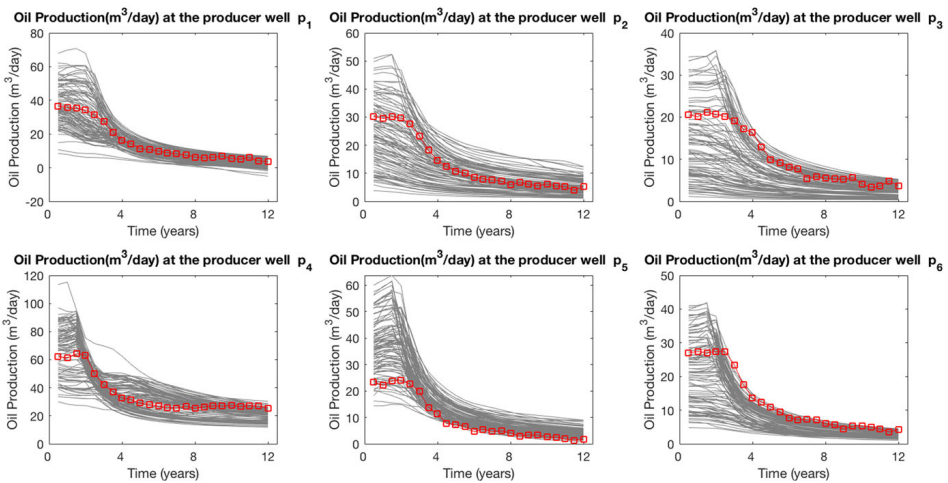
Fig. 13 Predicted oil production rate (OPR) at the 6 producer wells obtained from the selection ensemble smoother; the ensemble members are in grey, the measured OPR in red

in sand and shale and the seismic convolutional operator. The seismic data are calculated for two incident angles of 8° (near) and 24° (far). Figure 6 shows the seismic data computed 12 years after injection and production started. Because the reservoir properties are assumed to be vertically homogeneous within the model, two main
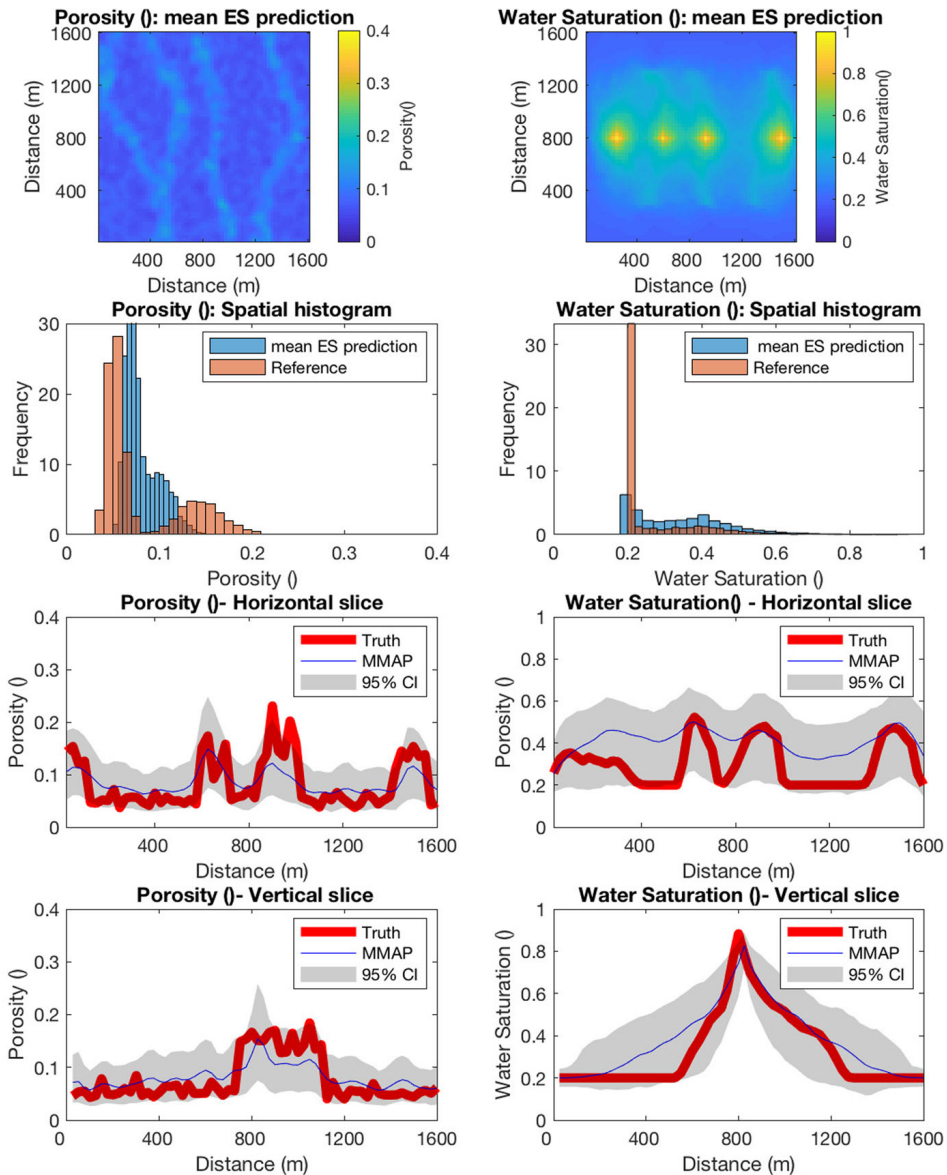
Fig. 14 Ensemble smoother predictions: mean maps, marginal histograms, and predictions along horizontal and vertical lines

reflections are generated in the seismic response, at the top and at the bottom of the reservoir. For this reason, for each trace, the dimension of the data is 4×1 (two angles and two measurements per angle), the dimension of the vector of elastic variables is 9×1 (three properties and three locations per property), and the dimension of the vector of porosity and initial saturation is 6×1 (two properties and three locations per property). The seismic operator is represented by a matrix $\boldsymbol{G} = \boldsymbol{WAD}$ of dimensions

Fig. 15 Predicted bottom hole pressure (BHP) at the 4 injection wells obtained from the ensemble smoother; the ensemble members are in black, the measured BHP in red
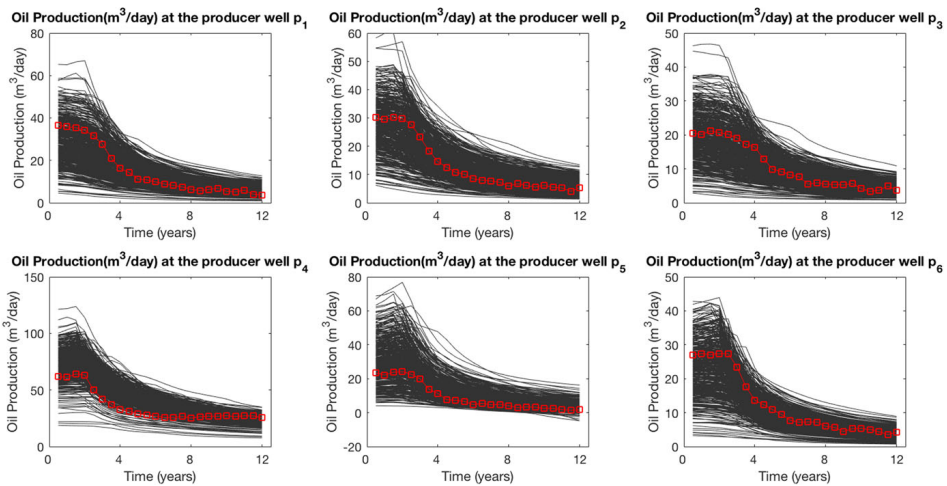
Fig. 16 Predicted oil production rate (OPR) at the 6 producer wells obtained from the ensemble smoother; the ensemble members are in grey, the measured OPR in red
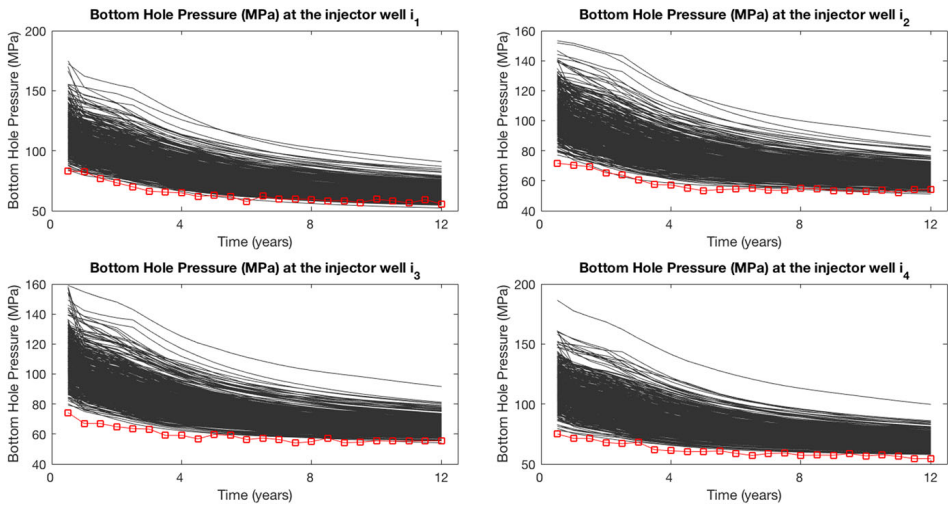
4×9 ($\boldsymbol{W}$ is 4×4, $\boldsymbol{A}$ is 4×6, $\boldsymbol{D}$ is 6×9) and the rock physics model is a system of 3 equations (one for each elastic property) with two unknowns (porosity and saturation).

The focus is on the prediction of the porosity field assumed to be spatially variable but constant in time, and the water saturation field before and after the experiment. Following the formalism in Eq. 8, all the measured data, production and seismic are saved in a vector $\boldsymbol{d}$ that contains OPR and BHP measurements every six months for twelve years and the seismic amplitudes measured after twelve years. Porosity and initial water saturation are stored in $\boldsymbol{r}_0$. The forward model f generates the model predictions according to the discretization of Eqs. 24, 25, 30 and 34 . The measured data are perturbed by adding white noise represented by the Gaussian error term $\boldsymbol{\delta}$ with

standard deviation 1, $\sqrt{2}/2$ for BHP and OPR respectively. The seismic amplitudes have a signal to noise ratio of approximately 2.5.

In this case study, data assimilation is performed using two methods: S-ES and ES. Figure 2 displays a flowchart summarizing the data assimilation procedure. For the S-ES, the prior distribution of the porosity field $\boldsymbol{\phi}$ is defined on log-scale as $f(\log(\boldsymbol{\phi}))$ and is selection-Gaussian with parameters

$$\boldsymbol{\Theta}_{\phi}^{SG} = (\mu_{\tilde{r}}^{\phi}, \mu_{\kappa}^{\phi}, \sigma_{\tilde{r}}^{\phi}, \boldsymbol{\Sigma}_{\tilde{r}}^{\rho}, \gamma^{\phi}, A). \tag{35}$$

The parameters (γ^{ϕ}, A) are chosen so that the prior marginal is bimodal, with modes approximately matching the average porosity values of the two facies in the reference field. The spatial correlation $(n \times n)$-matrix $\boldsymbol{\Sigma}^{\rho}$ is defined by the second order exponential spatial correlation function $\rho(\tau; d_x, d_y) = \exp\left[-\left(\frac{\tau_x^2}{d_x^2} + \frac{\tau_y^2}{d_y^2}\right)\right]$. The parameters values are listed in Table 1. The chosen values for d_x and d_y assume longer correlation along the y-axis than along the x-axis. The prior distribution for the initial water saturation field is Gaussian centered at $\boldsymbol{\mu}_{sw} = 0.2 \times \boldsymbol{i}_n$ with covariance matrix $\boldsymbol{\Sigma}_{sw} = 0.0005 * \boldsymbol{\Sigma}_{\tilde{r}}^{\rho}$, where $\boldsymbol{\Sigma}_{\tilde{r}}^{\rho}$ is the covariance matrix of the porosity field.

An initial ensemble $\boldsymbol{e}_0 = \{\boldsymbol{r}_0^{p(i)}, i = 1, \ldots, n_e\}$ of size $n_e = 500$ is generated. Covariance localization (Gaspari and Cohn 1999) is used for the conditioning. The localization was chosen restrictive enough to preserve an acceptable rank in the estimated covariance matrix, thereby preventing ensemble collapse. The mean of the realizations and the marginal distributions are shown in Fig. 7. The spatial stationarity is shown in the maps, whereas the marginal distribution is displayed by the histograms, which show the bimodality of the selection-Gaussian prior distribution for porosity.

The posterior distributions of porosity and water saturation are then computed using the S-ES method. Figure 8 shows the posterior mean for the porosity field and the water saturation field after 12 years obtained by the S-ES. Comparing the posterior mean of the porosity field to the reference porosity field, it can be observed that channels are correctly identified with good accuracy. The predicted water saturation field is smoother than the reference water saturation, but the overall pattern is aptly reproduced. Figure 9 shows the prediction and 95% confidence interval for porosity and water saturation along a vertical and horizontal line of the two-dimensional field. Even though the uncertainty is slightly underestimated, the S-ES captures the characteristics of the variations of the porosity and water saturation fields. Inflation was not adopted to show the uncertainty quantification of the methods, but the use of inflation factors might mitigate the underestimation in the uncertainty quantification. Figure 10 shows realizations from the posterior distribution of the porosity field obtained with the S-ES. The channels are correctly predicted with some uncertainties on the boundaries. Figure 11 shows the spatial histograms of the predicted porosity field and water saturation field obtained with the S-ES and compares them the reference values. The marginal maximum a posteriori (MMAP) is used as a predictor due to the marginal bimodality of the posterior ensemble. The spatial histogram of the mean porosity field shows an accurate match with the reference porosity field, whereas the water saturation spatial histogram is slightly smoother than the references one. Both findings are consistent

Table 2 Statistics comparing the S-ES prediction to the ES prediction

	S-ES	ES
MAE_ϕ	1.96×10^{-2}	2.6×10^{-2}
MAE_{sw}	6.6×10^{-2}	8.5×10^{-2}
corr_ϕ	0.85	0.82
corr_{sw}	0.84	0.75
80% coverage	0.85	0.71

with the analysis of Fig. 8. Figures 12 and 13 show the predicted OPR and BHP for the S-ES respectively, showing a good agreement with the true data.

For the traditional ES approach, the prior ensemble is generated assuming that the prior distributions for the log-porosity and initial water saturation fields are Gaussian. The prior distribution for the porosity field is defined as $f(\log(\boldsymbol{\phi})) \sim \varphi_n(\log(\boldsymbol{\phi}), \boldsymbol{\mu}_r, \sigma_r^2 \times \boldsymbol{\Sigma}_r^\rho)$, where $\boldsymbol{\mu}_r = -2.5\boldsymbol{i}_n$ and $\sigma_r^2 = 0.2$ are chosen such that the marginal prior distribution approximates the spatial histogram of the reference porosity field. The covariance matrix $\boldsymbol{\Sigma}_r^\rho$ and the prior distribution for the initial water saturation field are the same as in the S-ES case. Figure 14 shows the results for the posterior ensemble for the porosity field and water saturation field obtained using the ES. The posterior mean of the porosity field seems to identify the channels but the porosity values are underestimated within the channels. The predicted water saturation is much smoother than the reference water saturation field and the prediction from the S-ES. The spatial histogram of the mean porosity field shows that the ES fails to capture the bimodal nature of the porosity field, while the spatial histogram of the water saturation field severely underestimates the number of locations where saturation remains equal to the initial value of 0.2 and overpredicts saturation in several locations. By analyzing the predictions of porosity and water saturation along the horizontal and vertical lines, the ES prediction fails to capture the sharp changes, even though the coverage of the confidence intervals is satisfactory. Figures 15 and 16 display the predicted OPR and BHP for the ES, showing a better match for OPR and a worse match for BHP compared to the S-ES results.

The two methods are compared using different statistics: the mean absolute error (MAE), the correlation, and the 80% coverage probability are considered. Let MAE_ϕ and corr_ϕ define, respectively, the MAE and the correlation between the predicted porosity field and the reference porosity field. Further, let MAE_{sw} and corr_{sw} define, respectively, the MAE and the correlation between the predicted water saturation field and the reference water saturation at the final time step. The 80% coverage probability is calculated for the predicted porosity field. The results are detailed in Table 2. The MAE measures for the ES are about 30% higher than those of the S-ES. The correlations for the porosity are comparable. The correlation for the water saturation for the S-ES is about 10% higher than that of the ES. The 80% coverage probability seems to indicates the ES underestimates the uncertainty more than the S-ES overestimates it. Overall these statistics are favorable to the S-ES. The sensitivity to the data does not seem to cause over-fitting. The use of multiple data assimilation and iterative methods could improve the data match for both the ensemble smoother

and selection ensemble smoother. The S-ES has a higher computational demand than the ES because the posterior distribution needs to be sampled. On a 64×64 grid, it takes a few minutes to generate 200 samples, which is considered acceptable in view of the results. When the dimension of the data is very large, such as for field case reservoir models defined on 3D grids with millions of nodes, the matrix $\hat{\mathbf{\Sigma}}_d$ in Eqs. 9 and 15 is too large to be stored, let alone inverted. In that case, one could propose a parametric approach where the covariance is defined by a few model parameters and maximum likelihood estimation is used to estimate those parameters (Skauvold and Eidsvik 2019). Alternatively, a methodological development along the lines of the parametric Kalman filter could be considered (Pannekoucke et al. 2016).

4 Conclusions

Seismic and production data assimilation is performed to predict the spatial distribution of porosity and water saturation using a novel method named 'selection ensemble smoother', a Monte Carlo ensemble method that extends the ensemble smoother to selection-Gaussian models. By using a selection-Gaussian prior and an augmented initial state vector, the posterior distribution can represent multimodal posterior distributions. The main advantage of the selection ensemble smoother is the use of the augmented state vector which makes possible conditioning on the latent variable after data assimilation. The posterior distribution is then sampled assuming that the model is selection-Gaussian, thereby allowing multimodal features in the posterior distribution. The design of the experiment is optimal for the selection ensemble smoother approach since the channelized system leads to a bimodal porosity distribution. The results from the selection ensemble smoother are promising since the predictions are accurate and the data predictions match the measurements. The limited band-with of the seismic measurements, which limits the resolution of the data, prevents the inversion from clearly identifying the non-stationarity in the porosity field but still allows the selection ensemble smoother method to select the right mode, whereas the ensemble smoother results regress towards the mean.

References

Aanonsen SI, Aavatsmark I, Barkve T, Cominelli A, Gonard R, Gosselin O, Kolasinski M, Reme H et al (2003) Effect of scale dependent data correlations in an integrated history matching loop combining production data and 4D seismic data. In: Society of Petroleum Engineers, SPE reservoir simulation symposium

Aki K, Richards P (1980) Quantitative seismology: theory and methods. W. H. Freeman & Co, San Francisco

Arellano-Valle R, del Pino G (2004) From symmetric to asymmetric distributions: a unified approach. In: Genton M (ed) Skew-elliptical distributions and their applications: a journey beyond normality. Chapman & Hall / CRC, New York, pp 113–133

Arellano-Valle RB, Branco MD, Genton MG (2006) A unified view on skewed distributions arising from selections. Can J Stat 34(4):581–601

Bertino L, Evensen G, Wackernagel H (2003) Sequential data assimilation techniques in oceanography. Int Stat Rev 71(2):223–241

Buland A, Omre H (2003) Bayesian linearized AVO inversion. Geophysics 68(1):185–198

Chen Y, Oliver DS (2011) Localization of ensemble-based control-setting updates for production optimization. SPE J 17(01):122–136

Chen Y, Oliver DS (2012) Ensemble randomized maximum likelihood method as an iterative ensemble smoother. Math Geosci 44(1):1–26

Conjard M, Omre H (2020) Data assimilation in spatio-temporal models with non-Gaussian initial states-the selection ensemble Kalman model. Appl Sci 10(17):5742

Dadashpour M, Echeverria Ciaurri D, Mukerji T, Kleppe J, Landrø M (2010) A derivative-free approach for the estimation of porosity and permeability using time-lapse seismic and production data. J Geophys Eng 7(4):351–368

Dong Y, Oliver DS et al (2005) Quantitative use of 4D seismic data for reservoir description. SPE J 10(01):91–99

Dovera L, Della Rossa E (2011) Multimodal ensemble Kalman filtering using Gaussian mixture models. Comput Geosci 15(2):307–323

Dvorkin J, Gutierrez MA, Grana D (2014) Seismic reflections of rock properties. Cambridge University Press, Cambridge

Emerick AA, Reynolds AC (2013) Ensemble smoother with multiple data assimilation. Comput Geosci 55:3–15

Evensen G (2009) Data assimilation: the ensemble Kalman filter. Springer, Berlin

Gaspari G, Cohn S (1999) Construction of correlation functions in two and three dimensions. Q J R Meteorol Soc 125(554):723–757

Gineste M, Eidsvik J, Zheng Y (2020) Ensemble-based seismic inversion for a stratified medium. Geophysics 85(1):29–39

Houtekamer PL, Mitchell HL (1998) Data assimilation using an ensemble Kalman filter technique. Mon Weather Rev 126(3):796–811

Huang X, Meister L, Workman R et al (1997) Reservoir characterization by integration of time-lapse seismic and production data. In: SPE annual technical conference and exhibition, 38695, Society of Petroleum Engineers

Kennett BLN (1984) Guided wave propagation in laterally varying media, I. Theoretical development. Geophys J Int 79(1):235–255

Kretz V, Ravalec-Dupin L, Roggero F, et al (2002) An integrated reservoir characterization study matching production data and 4D seismic. In: SPE annual technical conference and exhibition, 77516, Society of Petroleum Engineers

Kumar D, Srinivasan S (2018) Ensemble-based assimilation of nonlinearly related dynamic data in reservoir models exhibiting non-Gaussian characteristics. Math Geosci 51:75–107

Landa J L, Horne R N, et al. (1997) A procedure to integrate well test data, reservoir performance history and 4-D seismic information into a reservoir description. In: SPE annual technical conference and exhibition, 38653, Society of Petroleum Engineers

Landrø M (2001) Discrimination between pressure and fluid saturation changes from time-lapse seismic data. Geophysics 66(3):836–844

Landrø M, Veire HH, Duffaut K, Najjar N (2003) Discrimination between pressure and fluid saturation changes from marine multicomponent time-lapse seismic data. Geophysics 68(5):1592–1599

Lie KA (2019) An introduction to reservoir simulation using MATLAB/GNU Octave: User guide for the MATLAB reservoir simulation toolbox (MRST). Cambridge University Press, Cambridge

Liu M, Grana D (2018) Ensemble-based seismic history matching with data reparameterization using convolutional autoencoder. In: SEG technical program expanded abstracts 2018, Society of Exploration Geophysicists, pp 3156–3160

Liu M, Grana D (2020) Time-lapse seismic history matching with an iterative ensemble smoother and deep convolutional autoencoder. Geophysics 85(1):15–31

Myrseth I, Omre H et al (2010) Hierarchical ensemble Kalman filter. SPE J 15(02):569–580

Nævdal G, Johnsen L M, Aanonsen S I, Vefring E H, et al. (2003) Reservoir monitoring and continuous model updating using ensemble Kalman filter. In: SPE annual technical conference and exhibition, 84372, Society of Petroleum Engineers

Naveau P, Genton MG, Shen X (2005) A skewed Kalman filter. J Multivar Anal 94(2):382–400

Oliver D, Chen Y (2018) Data assimilation in truncated Plurigaussian models: impact of the truncation map. Math Geosci 50:867–893

Oliver DS, Chen Y (2011) Recent progress on reservoir history matching: a review. Comput Geosci 15(1):185–221

Oliver DS, Reynolds AC, Liu N (2008) Inverse theory for petroleum reservoir characterization and history matching. Cambridge University Press, Cambridge

Omre H, Rimstad K (2018) Bayesian spatial inversion and conjugate selection Gaussian prior models. arXiv preprint arXiv:1812.01882

Pannekoucke O, Ricci S, Barthelemy S, Ménard R, Thual O, Tellus A (2016) Parametric Kalman filter for chemical transport models. Dyn Meteorol Oceanogr 68(1):31547

Rezaie J, Eidsvik J (2014) Kalman filter variants in the closed skew normal setting. Comput Stat Data Anal 75:1–14

Russell B H (1988) Introduction to seismic inversion methods. Society of Exploration Geophysicists

Sætrom J, Omre H (2011) Ensemble Kalman filtering with shrinkage regression techniques. Comput Geosci 15(2):271–292

Sheriff RE, Geldart LP (1995) Exploration seismology, 2nd edn. Cambridge University Press, Cambridge

Shuey RT (1985) A simplification of the Zoeppritz equations. Geophysics 50(4):609–614

Silva D, Deutsch C (2019) Multivariate categorical modeling with hierarchical truncated pluri-Gaussian simulation. Math Geosci 51:527–552

Skauvold J, Eidsvik J (2019) Parametric spatial covariance models in the ensemble Kalman filter. Spatial Stat 29:226–242

Skjervheim J A, Evensen G, Aanonsen S I, Ruud B O, Johansen T A, et al. (2005) Incorporating 4D seismic data in reservoir simulation models using ensemble Kalman filter. In: SPE annual technical conference and exhibition, 95789, Society of Petroleum Engineers

Thurin J, Brossier R, Métivier L (2019) Ensemble-based uncertainty estimation in full waveform inversion. Geophys J Int 219(3):1613–1635

Trani M, Arts R, Leeuwenburgh O et al (2012) Seismic history matching of fluid fronts using the ensemble Kalman filter. SPE J 18(01):159–171

Williams M, Keating J, Barghouty M et al (1998) The stratigraphic method: a structured approach to history matching complex simulation models. SPE Reserv Eval Eng 1(02):169–176

Yilmaz Ö (2001) Seismic data analysis: processing, inversion, and interpretation of seismic data. Society of Exploration Geophysicists

Zhou H, Li L, Franssen HJ, Gomez-Hernandez J (2012) Pattern recognition in a bimodal aquifer using the normal-score ensemble Kalman filter. Math Geosci 44:1–17

Zoeppritz K (1919) On the reflection and propagation of seismic waves. Gottinger Nachrichten 1(5):66–84

An overview of ensemble Kalman filtering methods applied to groundwater flow

Maxime Conjard, Emilio Sanchéz-Léon, Olaf Cirpka and Henning Omre

Technical Report

An overview of ensemble Kalman filtering methods applied to groundwater flow

Maxime Conjard [*1], Emilio Sanchéz-Léon [2], Olaf Cirpka[2], and Henning Omre[1]

[1]Department of Mathematical Sciences, NTNU
[2]Center for Applied Geoscience, University of Tübingen

March 2021

Abstract

Assimilation of spatio-temporal data is challenging especially when the variables are suspected to display non-Gaussian spatial histograms. For these cases, the selection ensemble Kalman filter has shown to produce encouraging results on synthetic tests cases. The objective of this study is therefore to situate its performance on a real data application when compared to established methods. We first present a review of existing ensemble Kalman filtering methods. We then present a synthetic and real data study where hydraulic conductivity is predicted using pumping test data. The synthetic study confirms the suitability of the selection ensemble Kalman filter when the conductivity field displays a bimodal spatial histogram. The real data study shows that the selection ensemble Kalman filter provides robust results, but its suitability is questioned when considering the added computational cost relative to the ensemble Kalman filter with multiple data assimilation.

1   Introduction

Data assimilation for spatio-temporal phenomena is of interest in many scientific fields such as geosciences, numerical weather prediction and oceanography. When cast in a Bayesian setting, we define a prior distribution for the initial state, a forward and a likelihood model and a posterior distribution. The forward model represents the evolution of the spatio-temporal variables at play while the likelihood model describes the data acquisition procedure. When the prior distribution for the initial state is Gaussian and the forward and likelihood models are linear with additive Gaussian noise (Gauss-linear), the posterior distribution is analytically tractable and can be assessed using the Kalman filter (Kalman, 1960). When the forward and likelihood models are non-linear, the extended Kalman filter (McElhoe, 1966) can be used to approximate the posterior distribution. It however requires repeated evaluations of sensitivty matrices which can be costly for high-dimensional problems. The ensemble Kalman filter (EnKF) (Evensen, 1994), a Monte-Carlo implementation of the Bayesian update, provides an alternative approach. Ensemble members are simulated from the initial distribution and serve as inputs to the forward model. The ensemble members are then sequentially conditioned on the data using the covariance matrix estimated from the ensemble itself. In addition to allowing for non-linearity in the

---

*Corresponding author: maxime.conjard@ntnu.no

forward and likelihood models, the ensemble can in principle be simulated from any initial distribution. The EnKF has successfully been applied to numerical weather prediction (Houtekamer et al., 2005), oceanography (Bertino et al., 2003), reservoir simulation (Aanonsen et al., 2009) and groundwater flow (Hendricks Franssen and Kinzelbach, 2008). Many extensions to the method have been developed to handle its intrinsic weaknesses. Methods (Anderson and Anderson, 1999; Sætrom and Omre, 2012) such as covariance inflation, can be used to counteract the systematic underestimation of the spread of the posterior ensemble. Covariance localization (Hamill et al., 2001) and hierarchical approaches (Myrseth and Omre, 2010) are used to improve covariance matrix estimates which come under scrutiny when the ensemble size is smaller than the dimension of the state space vector, which is usually the case. Gaussian mixture models (GMM) (Dovera and Della Rossa, 2010) and Gaussian anamorphosis (GA) (Simon and Bertino, 2009) are used to conserve non-Gaussianity in the posterior distribution. Iterative methods (Li and Reynolds, 2009) and so called multiple data assimilation (Emerick and Reynolds, 2013) are used to improve data match in cases where the non linearity in the forward and likelihood models is too important thereby making data match difficult.

We present a review of ensemble Kalman filtering methods that have been successfully used across a broad range of applications. The goal of this study is to evaluate the newly developed Selection Ensemble Kalman filter (SEnKF) (Conjard and Omre, 2021, 2020) and establish its applicability to real data experiments in regard to established methods. The SEnKF is an extension to the EnKF that can model multimodality, skewness and peakedness in the prior and posterior distributions. The prior model is defined to be a selection-Gaussian distribution (Arellano-Valle et al., 2006; Arellano-Valle and del Pino, 2004). The selection-Gaussian distribution is a conjugate prior to Gauss-linear forward and likelihood models and under these assumptions, Bayesian inversion (Omre and Rimstad, 2021; Forberg et al., 2021; Forberg et al., 2021) and sequential data assimilation (Conjard and Omre, 2021) produce analytically tractable posterior distributions. The motivation behind the development of the SEnKF is to extend the methodology to non-linear forward and likelihood models.

The relevance of the SEnKF is investigated through two case studies in which pumping test data is used to estimate the variables that govern groundwater flow in an aquifer. The synthetic case study determines the applicability of the SEnKF to groundwater flow in which the variables of interest exhibit spatial bimodality. The real case study illustrates the pertinence of the SEnKF in real data study where the variables of interest are suspected to exhibit spatial bimodality, by comparing the results with established methods.

In this paper $\boldsymbol{y} \sim f(\boldsymbol{y})$ denotes a random variable $\boldsymbol{y}$ distributed according to the probability density function (pdf) $f(\boldsymbol{y})$, or alternatively according to the corresponding cumulative distribution function (cdf) $F(\boldsymbol{y})$. Moreover, $\varphi_n(\boldsymbol{y}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the pdf of the Gaussian n-vector $\boldsymbol{y}$ with expectation n-vector $\boldsymbol{\mu}$ and covariance $(n \times n)$-matrix $\boldsymbol{\Sigma}$. Further $\Phi_n(A; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the probability of the aforementioned Gaussian n-vector $\boldsymbol{y}$ to be in $A \subset \mathbb{R}^n$. We also use $\boldsymbol{i}_n$ to denote the all-ones n-vector and $\boldsymbol{I}_n$ to denote the identity $(n \times n)$-matrix.

In Section 2, we introduce the two hidden Markov model, the Kalman model and selection Kalman model, that provides the theoretical foundations for the rest of the study. In Section 3, a review of the different ensemble filtering methods is given. In Section 4, the results from the synthetic and real data case studies are presented. In Section 5, conclusions are forwarded.

2 Hidden Markov model

Consider the real valued state n-vector $\boldsymbol{r}_t$ for $t \in \mathcal{T}_r : \{0, 1, \ldots, T, T+1\}$. Let $\boldsymbol{r}$ denote $\{\boldsymbol{r}_0, \boldsymbol{r}_1, \ldots, \boldsymbol{r}_T, \boldsymbol{r}_{T+1}\}$ and $\boldsymbol{r}_{i:j}$ denote $\{\boldsymbol{r}_i, \boldsymbol{r}_{i+1}, \ldots, \boldsymbol{r}_j\}, \forall (i, j) \in \mathcal{T}_r^2, i \le j$. Assume that the temporal m-vectors of observations $\boldsymbol{d}_t$ for $t \in \mathcal{T}_d : \{0, 1, \ldots, T\}$ are available, and define $\boldsymbol{d} = \{\boldsymbol{d}_0, \boldsymbol{d}_1, \ldots, \boldsymbol{d}_T\}$ and $\boldsymbol{d}_{i:j} = (\boldsymbol{d}_i, \ldots, \boldsymbol{d}_j)$ accordingly. The model specified thereafter defines a hidden Markov (HM) model (Cappé et al., 2005) as displayed in Figure 1.


Figure 1: Graph of the hidden Markov model.

Prior model: The prior model on $\boldsymbol{r}$ consists of an initial and a forward model,

$$f(\boldsymbol{r}) = f(\boldsymbol{r}_0) f(\boldsymbol{r}_{1:T+1} | \boldsymbol{r}_0), \tag{1}$$

where $f(\boldsymbol{r}_0)$ is the pdf of the initial state and $f(\boldsymbol{r}_{1:T+1} | \boldsymbol{r}_0)$ defines the forward model.

The forward model given the initial state $[\boldsymbol{r}_{1:T+1} | \boldsymbol{r}_0]$ is defined as

$$f(\boldsymbol{r}_{1:T+1} | \boldsymbol{r}_0) = \prod_{t=0}^{T} f(\boldsymbol{r}_{t+1} | \boldsymbol{r}_t), \tag{2}$$

with

$$[\boldsymbol{r}_{t+1} | \boldsymbol{r}_t] = \omega_t(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^r) \sim f(\boldsymbol{r}_{t+1} | \boldsymbol{r}_t), \tag{3}$$

where $\omega_t(\cdot, \cdot) \in \mathbb{R}^n$ is the forward function with random n-vector $\boldsymbol{\epsilon}_t^r$. This forward model only involves the variable at the previous time step $\boldsymbol{r}_t$, hence it defines a first-order Markov chain.

Likelihood model: The likelihood model for $[\boldsymbol{d} | \boldsymbol{r}]$ is defined as conditional independent with single-site response,

$$f(\boldsymbol{d} | \boldsymbol{r}) = \prod_{t=0}^{T} f(\boldsymbol{d}_t | \boldsymbol{r}_t), \tag{4}$$

with

$$[\boldsymbol{d}_t | \boldsymbol{r}_t] = \psi_t(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^d) \sim f(\boldsymbol{d}_t | \boldsymbol{r}_t), \tag{5}$$

where $\psi_t(\cdot, \cdot) \in \mathbb{R}^m$ is the likelihood function with random m-vector $\boldsymbol{\epsilon}_t^d$.

Posterior model: The posterior model for the HM model in Figure 1 is also a Markov chain (Cappé et al., 2005; Moja et al., 2018) and is given by

$$
\begin{aligned}
[\boldsymbol{r} | \boldsymbol{d}] \sim f(\boldsymbol{r} | \boldsymbol{d}) &= const \times f(\boldsymbol{d} | \boldsymbol{r}) f(\boldsymbol{r}) \\
&= const \times f(\boldsymbol{d}_0 | \boldsymbol{r}_0) f(\boldsymbol{r}_0) \prod_{t=1}^{T} f(\boldsymbol{d}_t | \boldsymbol{r}_t) f(\boldsymbol{r}_t | \boldsymbol{r}_{t-1}) f(\boldsymbol{r}_{T+1} | \boldsymbol{r}_T) \\
&= f(\boldsymbol{r}_0 | \boldsymbol{d}) \prod_{t=1}^{T} f(\boldsymbol{r}_t | \boldsymbol{r}_{t-1}, \boldsymbol{d}_{t:T}) f(\boldsymbol{r}_{T+1} | \boldsymbol{r}_T),
\end{aligned}
\tag{6}
$$

The HM model defined in this section gives a very general framework for the evaluation of the posterior distribution defined in Equation 6. Its analytical tractability depends however on the assumptions that are made on the prior and likelihood models. If the initial state in the prior model is Gaussian and both forward and likelihood models are Gauss-linear, the posterior distribution is analytically tractable, and it can be assessed using the Kalman recursions (Kalman, 1960). If the initial state is selection-Gaussian and both forward and likelihood models are Gauss-linear, the posterior distribution is also analytically tractable, and it can be assessed with the Selection Kalman recursions (Conjard and Omre, 2021).

Focus in the current study is on filtering/forecasting, hence on assessing $f(\boldsymbol{r}_{T+1}|\boldsymbol{d}_{0:T})$, which usually can be done by filtering recursions.

2.1 Traditional Kalman model

The traditional Kalman model assumes that the initial state is distributed according to a Gaussian distribution,

$$f(\boldsymbol{r}_0) = \varphi_n(\boldsymbol{r}_0; \boldsymbol{\mu}_{0|\cdot}^r, \boldsymbol{\Sigma}_{0|\cdot}^r) \tag{7}$$

with parametrization $\Theta^G = (\boldsymbol{\mu}_{0|\cdot}^r, \boldsymbol{\Sigma}_{0|\cdot}^r)$. It also assumes that the forward and likelihood models are Gauss-linear:

$$[\boldsymbol{r}_{t+1}|\boldsymbol{r}_t] = \omega_t(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^r) = \boldsymbol{A}_t \boldsymbol{r}_t + \boldsymbol{\epsilon}_t^r \rightarrow \varphi_n(\boldsymbol{r}; \boldsymbol{A}_t \boldsymbol{r}_t, \boldsymbol{\Sigma}_t^r) \tag{8}$$

$$[\boldsymbol{d}_t|\boldsymbol{r}_t] = \psi_t(\boldsymbol{r}_t, \boldsymbol{\epsilon}_t^d) = \boldsymbol{H}_t \boldsymbol{r}_t + \boldsymbol{\epsilon}_t^d \rightarrow \varphi_m(\boldsymbol{d}; \boldsymbol{H}_t \boldsymbol{r}_t, \boldsymbol{\Sigma}_t^d) \tag{9}$$

where $\boldsymbol{A}$ is a $(n \times n)$-matrix and $\boldsymbol{H}$ is a $(m \times n)$-matrix, $\boldsymbol{\epsilon}_t^r$ is a centered Gaussian n-vector with co-variance $(n \times n)$-matrix $\boldsymbol{\Sigma}_t^r$ and $\boldsymbol{\epsilon}_t^d$ a centered Gaussian m-vector with covariance $(m \times m)$-matrix $\boldsymbol{\Sigma}_t^d$. Under these assumptions, the posterior model in Equation 6 is Gaussian and analytically tractable. The recursion,

$$f(\boldsymbol{r}_{t+1}|\boldsymbol{d}_{0:t}) = \int f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t) f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t}) \mathrm{d}\boldsymbol{r}_t,$$
$$= \int f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t) [f(\boldsymbol{d}_t|\boldsymbol{d}_{0:t-1})]^{-1} f(\boldsymbol{d}_t|\boldsymbol{r}_t) f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t-1}) \mathrm{d}\boldsymbol{r}_t, \tag{10}$$

for $t = 0, \ldots, T$, can be evaluated using the Kalman Filter (Kalman, 1960) with the following recursion formulas,

$$
\begin{aligned}
f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t-1}) &= \varphi_n(\boldsymbol{r}_t; \boldsymbol{\mu}_{t|t-1}^r, \boldsymbol{\Sigma}_{t|t-1}^r), \\
f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t}) &= \varphi_n(\boldsymbol{r}_t; \boldsymbol{\mu}_{t|t}^r, \boldsymbol{\Sigma}_{t|t}^r), \\
\boldsymbol{\mu}_{t|t}^r &= \boldsymbol{\mu}_{t|t-1}^r + \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T (\boldsymbol{H}_t \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T + \boldsymbol{\Sigma}_t^d)^{-1} (\boldsymbol{d}_t - \boldsymbol{H}\boldsymbol{\mu}_{t|t-1}^r), \\
\boldsymbol{\Sigma}_{t|t}^r &= \boldsymbol{\Sigma}_{t|t-1}^r - \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T (\boldsymbol{H}_t \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T + \boldsymbol{\Sigma}_t^d)^{-1} \boldsymbol{H}_t \boldsymbol{\Sigma}_{t|t-1}^r, \\
f(\boldsymbol{r}_{t+1}|\boldsymbol{d}_{0:t}) &= \varphi_n(\boldsymbol{r}_{t+1}; \boldsymbol{\mu}_{t+1|t}^r, \boldsymbol{\Sigma}_{t+1|t}^r), \\
\boldsymbol{\mu}_{t+1|t}^r &= \boldsymbol{A}_t \boldsymbol{\mu}_{t|t}^r, \\
\boldsymbol{\Sigma}_{t+1|t}^r &= \boldsymbol{A}_t \boldsymbol{\Sigma}_{t|t}^r \boldsymbol{A}_t^T + \boldsymbol{\Sigma}_t^r.
\end{aligned}
\tag{11}
$$

The $(n \times m)$-matrix $\boldsymbol{K}_t = \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T (\boldsymbol{H}_t \boldsymbol{\Sigma}_{t|t-1}^r \boldsymbol{H}_t^T + \boldsymbol{\Sigma}_t^d)^{-1}$ is denoted the Kalman weights. Hence the Kalman filter can be determined analytically by the recursions above, and it is closed in the class of Gaussian pdfs. The posterior model is Gaussian and the filtering/forecasting distribution is of course also Gaussian and it is assessable from the Kalman recursions $f(\boldsymbol{r}_{T+1}|\boldsymbol{d}_{0:T}) = \varphi_n(\boldsymbol{r}; \boldsymbol{\mu}_{T+1|T}^r, \boldsymbol{\Sigma}_{T+1|T}^r)$.

2.2 Selection Kalman model

In the selection Kalman model, the distribution for the initial state $f(\boldsymbol{r}_0)$ is assumed to be in the class of selection-Gaussian distributions (Arellano-Valle et al., 2006; Omre and Rimstad, 2021). Consider an auxiliary Gaussian $(n + n)$-vector $[\tilde{\boldsymbol{r}}_0, \boldsymbol{v}]$,

$$\begin{bmatrix} \tilde{\boldsymbol{r}}_0 \\ \boldsymbol{v} \end{bmatrix} \sim \varphi_{2n} \left(\begin{bmatrix} \tilde{\boldsymbol{r}}_0 \\ \boldsymbol{v} \end{bmatrix}; \begin{bmatrix} \boldsymbol{\mu}_{0|\cdot}^{\tilde{r}} \\ \boldsymbol{\mu}_v \end{bmatrix}, \begin{bmatrix} \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}} & \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}} \boldsymbol{\Gamma}_{v|\tilde{r}}^T \\ \boldsymbol{\Gamma}_{v|\tilde{r}} \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}} & \boldsymbol{\Sigma}_v \end{bmatrix} \right), \tag{12}$$

with expectation n-vectors $\boldsymbol{\mu}_{0|\cdot}^{\tilde{r}}$ and $\boldsymbol{\mu}_v$, correlation $(n \times n)$-matrix $\boldsymbol{\Gamma}_{v|\tilde{r}}$, and where $\boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}}$, $\boldsymbol{\Sigma}_v$, and $\boldsymbol{\Sigma}_{v|\tilde{r}}$ are covariance $(n \times n)$-matrices with $\boldsymbol{\Sigma}_v = \boldsymbol{\Gamma}_{v|\tilde{r}} \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}} \boldsymbol{\Gamma}_{v|\tilde{r}}^T + \boldsymbol{\Sigma}_{v|\tilde{r}}$. Define a selection set $A \subset \mathbb{R}^n$ of dimension n and let $\boldsymbol{r}_0 = [\tilde{\boldsymbol{r}}_0 | \boldsymbol{v} \in A]$; then, $\boldsymbol{r}_0$ is in the class of selection-Gaussian distribution and its pdf is,

$$\begin{aligned} f(\boldsymbol{r}_0) &= \left[\Phi_n(A; \boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v) \right]^{-1} \\ &\times \Phi_n(A; \boldsymbol{\mu}_v + \boldsymbol{\Gamma}_{v|\tilde{r}}(\boldsymbol{r}_0 - \boldsymbol{\mu}_{0|\cdot}^{\tilde{r}}), \boldsymbol{\Sigma}_{v|\tilde{r}}) \times \varphi_n(\boldsymbol{r}_0; \boldsymbol{\mu}_{0|\cdot}^{\tilde{r}}, \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}}). \end{aligned} \tag{13}$$

It can therefore be parametrized with $\Theta^{SG} = (\boldsymbol{\mu}_{0|\cdot}^{\tilde{r}}, \boldsymbol{\mu}_v, \boldsymbol{\Sigma}_{0|\cdot}^{\tilde{r}}, \boldsymbol{\Sigma}_{v|\tilde{r}}, \boldsymbol{\Gamma}_{v|\tilde{r}}, A)$. Note that the class of Gaussian distributions constitutes a subset of the class of selection-Gaussian distributions with $\boldsymbol{\Gamma}_{v|\tilde{r}} = 0 \times \boldsymbol{I}_n$. The dependence in $[\tilde{r}, v]$ represented by $\boldsymbol{\Gamma}_{v|\tilde{r}}$ and the selection subset A are crucial user-defined parameters. The selection-Gaussian model may represent multimodal, skewed, and/or peaked marginal distributions (Omre and Rimstad, 2021). As in the traditional Kalman model, see Section 2.1, the selection Kalman model also assumes the forward and likelihood models to be Gauss-linear, see Equation 8. Under these assumptions, the posterior model in Equation 6 is selection-Gaussian and analytically tractable (Conjard and Omre, 2021). The selection Kalman model is defined by the Gaussian augmented state $2n$-vector $\boldsymbol{u}_t = [\tilde{\boldsymbol{r}}_t, \boldsymbol{v}]$ and the selection set A. The forward model in Equation 8 is adjusted to account for the augmented state vector,

$$\boldsymbol{u}_{t+1} = \begin{bmatrix} \boldsymbol{A}_t & 0 \\ 0 & \boldsymbol{I}_n \end{bmatrix} \boldsymbol{u}_t + \begin{bmatrix} \boldsymbol{\epsilon}_t^r \\ 0 \end{bmatrix}, \tag{14}$$

and so is the likelihood model in Equation 9,

$$\boldsymbol{d}_t = \begin{bmatrix} H & 0 \\ 0 & 0 \end{bmatrix} \boldsymbol{u}_t + \begin{bmatrix} \boldsymbol{\epsilon}_t^d \\ 0 \end{bmatrix} \tag{15}$$

These forward and likelihood models are Gauss linear, and the recursion for the augmented state vector is ,

$$f(\boldsymbol{u}_{t+1}|\boldsymbol{d}_{0:t}) = \int f(\boldsymbol{u}_{t+1}|\boldsymbol{u}_t)[f(\boldsymbol{d}_t|\boldsymbol{d}_{0:t-1})]^{-1} f(\boldsymbol{d}_t|\boldsymbol{u}_t) f(\boldsymbol{u}_t|\boldsymbol{d}_{0:t-1}) \mathrm{d}\boldsymbol{u}_t. \tag{16}$$

This recursion is identical to Equation 10, and hence it can be evaluated using the Kalman recursions in Equation 11. The filtering /forecasting distribution

$$f(\boldsymbol{u}_{T+1}|\boldsymbol{d}_{0:T}) = \varphi_{2n}(\boldsymbol{u}; \boldsymbol{\mu}_{T+1|T}^u, \boldsymbol{\Sigma}_{T+1|T}^u), \tag{17}$$

yielded by the Kalman filter is a Gaussian $2n$-vector with mean $2n$-vector $\boldsymbol{\mu}_{T+1|T}^u$ and covariance $(2n \times 2n)$-matrix $\boldsymbol{\Sigma}_{T+1|T}^u$. In order to obtain the distribution as evaluated by the selection Kalman filter, one needs to assess $f(\tilde{\boldsymbol{r}}_{T+1}|\boldsymbol{v} \in A, \boldsymbol{d}_{0:T})$ which can be done using McMC simulation.

The McMC algorithm used in Conjard and Omre (2021) and Conjard and Omre (2020) to sample realizations from the selection-Gaussian distribution worked well on small state vectors (≤ 500) but had to be adapted to handle high-dimensional problems. In order to sample $\boldsymbol{r}_0 = [\tilde{\boldsymbol{r}}_0 | \boldsymbol{v}_0 \in A]$, it is necessary to sample a truncated Gaussian n-vector $\boldsymbol{v} \in A$. The selection set is typically $[(-\infty, -b] \cup [b, +\infty)])^n$, where n is the dimension of the state vector. This means that the selection set is composed of 2^n disjoint subsets. To ensure satisfactory mixing, it is necessary that the samples come from different subsets of A which becomes difficult to ensure as n increases.

3   Ensemble Filter Algorithm

In the general HM model, see Section 2, the filtering/forecasting distribution $f(\boldsymbol{r}_{T+1}|\boldsymbol{d}_{0:T})$ can be evaluated using sequential Monte-Carlo techniques. Similar techniques for evaluating the smoothing distribution also exist but are not discussed in greater detail here. There are primarily two sequential Monte-Carlo algorithms and their variants that are in frequent use: the particle filter (PF) (Gordon et al., 1993; Van Leeuwen, 2009) and the EnKF (Evensen, 1994). The former has favorable asymptotic characteristics when the number of particles/members tends towards infinity, while the latter is only asymptotically exact for the Gauss-linear case. In high-dimensional models, actually for $n \geq 10$, and a fairly limited number of particles/members, the EnKF has proven to be clearly favorable (Katzfuss et al., 2020; Li et al., 2016). We focus on the properties of the EnKF and its variants in the current study.

3.1   Traditional Ensemble Kalman filter

When the forward and/or likelihood models are nonlinear, Equations 3 and 5 cannot be written in closed form. The analytical tractability of the posterior model in Equation 6 is lost. Hence the filtering/forecasting distribution must be approximated. One way to approximate the forecast recursion, see Equation 10, is the EnKF. Instead of tracking the first two-moments of the forecast recursion as is done in the Kalman Filter, the EnKF uses a Monte-Carlo approach and generates n_e realizations from the initial distribution $f(\boldsymbol{r}_0)$ that form an ensemble $\boldsymbol{e}_0$. The realizations are referred to as ensemble members, and n_e as ensemble size. Each iteration of the EnKF algorithm consists of two steps, conditioning and forwarding, represented by the following recursions,

$$f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t}) = [f(\boldsymbol{d}_t|\boldsymbol{d}_{0:t-1})]^{-1} f(\boldsymbol{d}_t|\boldsymbol{r}_t) f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t-1}), \tag{18}$$

$$f(\boldsymbol{r}_{t+1}|\boldsymbol{d}_{0:t}) = \int f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t) f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t}) \mathrm{d}\boldsymbol{r}_t,$$

that in general are not analytically tractable. The initial ensemble is defined as n_e iid realizations $\boldsymbol{e}_0^r = \{\boldsymbol{r}_0^{u(i)}, i = 1, \dots, n_e\}$ from $f(\boldsymbol{r}_0)$. The distribution from which the prior ensemble is sampled is not limited to Gaussian distributions, it is therefore possible to represent non-Gaussian features in the prior. The prior ensemble can actually be generated by any stochastic algorithm representing prior experience with the variable under study. At time t, the ensemble $\boldsymbol{e}_t^r = \{\boldsymbol{r}_t^{u(i)}, i = 1, \dots, n_e\}$ is assumed to represent $f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t-1})$. The conditioning step is initiated by generating pseudo observations using the likelihood model $\boldsymbol{d}_t^{(i)} = \psi_t(\boldsymbol{r}_t^{u(i)}, \boldsymbol{e}_t^d)$, and defining $\boldsymbol{e}_t = \{(\boldsymbol{r}_t^{u(i)}, \boldsymbol{d}_t^{(i)}), i = 1, \dots, n_e\}$. The conditioning step is the challenging step, and it is performed by linearizing the updating,

$$\text{Assess } \boldsymbol{\Sigma}_{rd} \text{ from } \boldsymbol{e}_t \rightarrow \hat{\boldsymbol{\Sigma}}_{rd} \rightarrow \hat{\boldsymbol{K}}_t = \hat{\boldsymbol{\Gamma}}_{rd}[\hat{\boldsymbol{\Sigma}}_d]^{-1}$$

$$\boldsymbol{r}_t^{c(i)} = \boldsymbol{r}_t^{u(i)} + \hat{\boldsymbol{K}}_t(\boldsymbol{d}_t - \boldsymbol{d}_t^{(i)}); i = 1, \ldots, n_e. \tag{19}$$

The ensemble $\{\boldsymbol{r}^{c(i)}; i = 1, \ldots, n_e\}$ is assumed to represent $f(\boldsymbol{r}_t|\boldsymbol{d}_{0:t})$. The Kalman weights $\boldsymbol{K}_t$ must be inferred from the ensemble $\boldsymbol{e}_t$. Both the linearization and this inference are critical factors for the reliability of the EnKF algorithm, which will be discussed in greater detail later.

The forward step can be made correctly for each ensemble member,

$$\boldsymbol{r}_{t+1}^{u(i)} = \omega_t(\boldsymbol{r}_t^{c(i)}, \boldsymbol{\epsilon}_t^r); i = 1, \ldots, n_e, \tag{20}$$

and the ensemble $\boldsymbol{e}_{t+1}^r = \{\boldsymbol{r}_{t+1}^{u(i)}, i = 1, \ldots, n_e\}$ representing $f(\boldsymbol{r}_{t+1}|\boldsymbol{d}_{0:t})$ is defined.

The traditional EnKF algorithm is presented in Algorithm 1. The outcome of the traditioal EnKF algorithm $\boldsymbol{e}_{T+1}^r = \{\boldsymbol{r}_{T+1}^{u(i)}, i = 1, \ldots, n_e\}$ represents $f(\boldsymbol{r}_{T+1}|\boldsymbol{d}_{0:T})$. Based on $\boldsymbol{e}_{T+1}^r$, both predictions and prediction intervals for $[\boldsymbol{r}_{T+1}|\boldsymbol{d}_{0:T}]$ can be inferred using the ensemble average and ensemble marginal variance respectively. In terms of computer demand, the forwarding step is usually the bottleneck in the EnKF algorithm, as n_e evaluations of the forward model repeated at each iteration can be extremely costly, usually limiting the ensemble size to the order of 100.

Algorithm 1 Traditional Ensemble Kalman Filter (EnKF)

The $(n + m)$-vector $[\boldsymbol{r}, \boldsymbol{d}]$ has covariance matrix:

$$\boldsymbol{\Sigma}_{rd} = \begin{bmatrix} \boldsymbol{\Sigma}_r & \boldsymbol{\Gamma}_{rd} \\ \boldsymbol{\Gamma}_{dr} & \boldsymbol{\Sigma}_d \end{bmatrix}$$

Initiate:

 n_e = No. of ensemble members

 Generate: $\boldsymbol{r}_0^{u(i)} \sim f(\boldsymbol{r}_0)$, $i = 1, ..., n_e$

 Generate: $\boldsymbol{\epsilon}_0^{d(i)}$, $i = 1, ..., n_e$

 $\boldsymbol{d}_0^{(i)} = \psi_0(\boldsymbol{r}_0^{u(i)}, \boldsymbol{\epsilon}_0^{d(i)})$, $i = 1, ..., n_e$

 $\boldsymbol{e}_0 = \{(\boldsymbol{r}_0^{u(i)}, \boldsymbol{d}_0^{(i)}), i = 1, \ldots, n_e\}$

Iterate $t = 0, ..., T$:

 Conditioning:

 Estimate $\boldsymbol{\Sigma}_{rd}$ from $\boldsymbol{e}_t \longrightarrow \hat{\boldsymbol{\Sigma}}_{rd} \rightarrow \hat{\boldsymbol{K}}_t = \hat{\boldsymbol{\Gamma}}_{rd}[\hat{\boldsymbol{\Sigma}}_d]^{-1}$

 $\boldsymbol{r}_t^{c(i)} = \boldsymbol{r}_t^{u(i)} + \hat{\boldsymbol{K}}_t(\boldsymbol{d}_t - \boldsymbol{d}_t^{(i)})$, $i = 1, ..., n_e$

 Forwarding:

 Generate: $\boldsymbol{\epsilon}_t^{r(i)}$, $i = 1, ..., n_e$

 $\boldsymbol{r}_{t+1}^{u(i)} = \omega_t(\boldsymbol{r}_t^{c(i)}, \boldsymbol{\epsilon}_t^{r(i)})$, $i = 1, ..., n_e$

 Generate: $\boldsymbol{\epsilon}_{t+1}^{d(i)}$, $i = 1, ..., n_e$

 $\boldsymbol{d}_{t+1}^{(i)} = \psi_{t+1}(\boldsymbol{r}_{t+1}^{u(i)}, \boldsymbol{\epsilon}_{t+1}^{d(i)})$, $i = 1, ..., n_e$

 $\boldsymbol{e}_{t+1} = \{(\boldsymbol{r}_{t+1}^{u(i)}, \boldsymbol{d}_{t+1}^{(i)}), i = 1, \ldots, n_e\}$

End iterate

The ensemble $\boldsymbol{e}_{T+1}^r = \{\boldsymbol{r}_{T+1}^{u(i)}, i = 1, \ldots, n_e\}$ represents $f(\boldsymbol{r}_{T+1} | \boldsymbol{d}_{0:T})$

3.2 Special case: Traditional Kalman model

The traditional Kalman model is a Gauss-linear HM model, and the model assumptions are specified in Section 2.1. The EnKF can be used on this model to assess the filtering/forecasting distribution and

the results can be compared to the exact analytical results from the Kalman recursions. If the model parameters are known and the Kalman weights K_t can be calculated correctly, then the conditioning steps in the EnKF are exactly correct for all t, and the ensemble e_t contains n_e iid elements from the correct Gaussian pdf.

The EnKF algorithm is assumed to be initiated by the ensemble members in e_0^r being iid from a multivariate Gaussian n-vector with unknown parameters. Hence, the approximation in the EnKF, under the traditional Kalman model assumptions, is only the assessment of the Kalman weights K_t from the ensemble e_t.

In order to evaluate the effect of the Kalman weights inference, assume that the ensemble e_t contains iid ensemble members from the correct Gaussian pdf. The conditioning step can then be written,

$$e_t : \{(r_t^{u(i)}, d_t^{(i)}), i = 1, \ldots, n_e\}, \quad iid \quad \varphi_{n+m}\left(\begin{bmatrix} r \\ d \end{bmatrix}; \mu_{rd}, \Sigma_{rd}\right),$$

$$\text{Assess } \Sigma_{rd} \text{ from } e_t \rightarrow \hat{\Sigma}_{rd} \rightarrow \hat{K}_t = \hat{\Gamma}_{rd}[\hat{\Sigma}_d]^{-1},$$

$$r_t^{c(i)} = r_t^{u(i)} + \hat{K}_t(d_t - d_t^{(i)}), i = 1, \ldots, n_e, \tag{21}$$

where d_t is the observed data. The reliability of the EnKF is dependent on the Kalman weights inference. Under Gaussian assumptions, the traditional minimum variance unbiased estimator for Σ_{rd} should be used (Mardia et al., 1979). Note that if the estimator for Σ_{rd} is consistent as $n_e \rightarrow \infty$, so is the estimator for K_t. Then the EnKF is asymptotically correct as $n_e \rightarrow \infty$ under the traditional Kalman model assumptions. This characteristic is not very impressive since the filter is analytically tractable under these assumptions. Unfortunately, there are not more general asymptotic results for the EnKF. In spite of this, the EnKF has proven to be very reliable for high-dimensional problems with weakly non-linear forward and likelihood models. Challenges arising from having only a finite number of ensemble members will be discussed in the following paragraphs.

3.2.1 Ensemble shrinkage

The inference of the Kalman weights $(n \times m)$-matrix K_t from the ensemble e_t of size n_e is the challenge under traditional Kalman model assumptions. For the estimator $\hat{K}_t$ to have full rank, one needs $n_e > n + m$ (Sætrom and Omre, 2013) and then $\hat{K}_t$ is know to be matrix-T distributed. If $n_e \leq n + m$, additional rank related inference problems occur. In this section, we assume that $n_e > n + m$ and return to rank challenges in Section 3.3.1. The conditioning in the EnKF appears as Gaussian ensemble conditioning using plug-in estimate $\hat{K}_t$ for the Kalman weight K_t. This standard statistical procedure is known to slightly underestimate the variability, but may be justifiable for one assimilation step. However, if the conditioning step is repeated numerous times, the ensemble variability may be severely underestimated and the ensemble can collapse. Thereafter data assimilation is impossible since the Kalman weights cannot be estimated. It is therefore important to perform each conditioning step precisely in a way that also captures model parameter uncertainty. We will evaluate two major effects in the conditioning step of the EnKF:

1. Ensemble center drift

2. Ensemble variability reduction

The ensemble members in e_t are assumed to be iid $(n + m)$-dimensional Gaussian vectors with expectation μ_{rd} and covariance Σ_{rd}, which is convenient when evaluating one single conditioning step from t to $t + 1$.

3.2.1.1 Ensemble center drift   We evaluate the ensemble average by studying the expected value of the conditioning step of one ensemble member, which should be the conditional expectation $\boldsymbol{\mu}_{r|d}$,

$$
\begin{aligned}
\mathbb{E}(\boldsymbol{r}_t^{c(i)}) &= \mathbb{E}(\boldsymbol{r}_t^{u(i)}) + \mathbb{E}(\hat{\boldsymbol{K}}_t(\boldsymbol{d}_t - \boldsymbol{d}_t^{(i)})), \\
&= \boldsymbol{\mu}_r + \mathbb{E}(\hat{\boldsymbol{K}}_t)(\boldsymbol{d}_t - \boldsymbol{H}\boldsymbol{\mu}_r) - Cov(\hat{\boldsymbol{K}}_t, \boldsymbol{d}_t^{(i)}), \\
&= \boldsymbol{\mu}_{r|d} + [\mathbb{E}(\hat{\boldsymbol{K}}_t) - \boldsymbol{K}_t][\boldsymbol{d}_t - \boldsymbol{H}\boldsymbol{\mu}_r] - Cov(\hat{\boldsymbol{K}}_t, \boldsymbol{d}_t^{(i)}).
\end{aligned}
\tag{22}
$$

The conditioning step is therefore not automatically centred at $\boldsymbol{\mu}_{r|d}$ as wanted. However, if $\hat{\boldsymbol{K}}_t$ is an unbiased estimator for $\boldsymbol{K}_t$ and $(\hat{\boldsymbol{K}}_t, \boldsymbol{d}_t^{(i)})$ are independent then correct centering for each updated ensemble member is ensured.

Define the sub-ensemble $\boldsymbol{e}_t^{-i} = \boldsymbol{e}_t \setminus \{\boldsymbol{r}_t^{u(i)}, \boldsymbol{d}_t^{(i)}\}$ and let the estimator for $\boldsymbol{\Sigma}_{rd}$ based on $\boldsymbol{e}_t^{-i}$ be the regular unbiased covariance estimator, defining $\hat{\boldsymbol{\Sigma}}_{rd}^{-i}$, according to the suggestion in Loe and Tjelmeland (2021). Then the estimator for $\boldsymbol{K}_t$, $\hat{\boldsymbol{K}}_t^{-i} = \hat{\boldsymbol{\Gamma}}_{rd}^{-i}[\hat{\boldsymbol{\Sigma}}_d^{-i}]^{-1}$, will be unbiased and independent of $\boldsymbol{d}_t^{(i)}$ as required above.

If the conditioning of each ensemble member is correctly centred, then the ensemble average will also be unbiased for $\boldsymbol{\mu}_{r|d}$. In Section 3.2.1.3, an algorithm preventing ensemble center drift is defined.

3.2.1.2 Ensemble variability reduction   We evaluate the ensemble variance, which is determined by the variance in each member and the covariance between members. It is important to note that the objective is not to minimize the variance, but rather to represent uncertainty in the EnKF algorithm realistically.

Assume that we ensure correct centering of the conditioning by using the conditioning weights $\hat{\boldsymbol{K}}_t^{-i}$ for each ensemble member. The variance of one ensemble member after conditioning is,

$$
\begin{aligned}
\text{Var}(\boldsymbol{r}_t^{c(i)}) &= \text{Var}(\boldsymbol{r}_t^{u(i)}) + \text{Var}(\hat{\boldsymbol{K}}_t^{-i}[\boldsymbol{d}_t - \boldsymbol{d}_t^{(i)}]) + 2\,\text{Cov}(\boldsymbol{r}_t^{u(i)}, \hat{\boldsymbol{K}}_t^{-i}[\boldsymbol{d}_t - \boldsymbol{d}_t^{(i)}]), \\
&= \boldsymbol{\Sigma}_{r|d} + \mathbb{E}([\hat{\boldsymbol{K}}_t^{-i} - \boldsymbol{K}_t]\boldsymbol{Q}_1[\hat{\boldsymbol{K}}_t^{-i} - \boldsymbol{K}_t]^T),
\end{aligned}
\tag{23}
$$

where $\boldsymbol{Q}_1$ is a $(m \times m)$- matrix related to the likelihood model (Sætrom and Omre, 2013).

The variance is the conditional variance $\boldsymbol{\Sigma}_{r|d}$, plus a term caused by the lack of knowledge about the model parameters, represented by $\boldsymbol{K}_t$. The EnKF is based on plugging in the estimator $\hat{\boldsymbol{K}}_t^{-i}$ for $\boldsymbol{K}_t$ and do not capture the latter variability. Consequently, there is a downward bias in the variability in each updated ensemble member.

The ensemble variance will also be influenced by the coupling of the updated ensemble members. Positive coupling causes a downward bias in the ensemble variance relative to the individual member variance.

The covariance between two updated ensemble members caused by the conditioning step is,

$$
\begin{aligned}
\text{Cov}(\boldsymbol{r}_t^{c(i)}, \boldsymbol{r}_t^{c(j)}) &= \text{Cov}(\boldsymbol{r}_t^{u(i)}, \boldsymbol{r}_t^{u(j)}) + \text{Cov}(\hat{\boldsymbol{K}}_t^{-i}[\boldsymbol{d}_t - \boldsymbol{d}_t^{(i)}], \hat{\boldsymbol{K}}_t^{-j}[\boldsymbol{d}_t - \boldsymbol{d}_t^{(j)}]) \\
&\quad + \text{Cov}(\boldsymbol{r}_t^{u(i)}, \hat{\boldsymbol{K}}_t^{-j}[\boldsymbol{d}_t - \boldsymbol{d}_t^{(j)}]) + \text{Cov}(\hat{\boldsymbol{K}}_t^{-i}[\boldsymbol{d}_t - \boldsymbol{d}_t^{(i)}], \boldsymbol{r}_t^{u(j)}), \\
&= \mathbb{E}([\hat{\boldsymbol{K}}_t^{-i} - \boldsymbol{K}_t]\boldsymbol{Q}_2[\hat{\boldsymbol{K}}_t^{-j} - \boldsymbol{K}_t]^T) + 2\,\text{Cov}(\boldsymbol{r}_t^{u(i)}, [\hat{\boldsymbol{K}}_t^{-j} - \boldsymbol{K}_t]\boldsymbol{q}_2),
\end{aligned}
\tag{24}
$$

where $\boldsymbol{q}_2$ is an m-vector and $\boldsymbol{Q}_2$ a $(m \times m)$- matrix both defined by the likelihood model (Sætrom and Omre, 2013). The covariance is non-zero, since $[\hat{\boldsymbol{K}}_t^{-i}, \hat{\boldsymbol{K}}_t^{-j}]$ are correlated through many common ensemble members and so are $[\boldsymbol{r}_t^{u(i)}, \hat{\boldsymbol{K}}_t^{-j}]$ since the former is used to calculate the latter. The covariance will actually be positive, and hence the ensemble variability will have a downward bias relative to the individual ensemble member variability.

Consequently, both member variance and coupling will contribute to underestimating the variability in the EnKF. In order to correct for this downward bias, we need to account for the uncertainty in estimating the Kalman weights $\boldsymbol{K}_t$ from the ensemble $\boldsymbol{e}_t$. For the traditional Kalman model with $n_e > n + m$, the estimator $\hat{\boldsymbol{K}}_t^{-i}$ is known to be matrix-T distributed with parameters defined by the model parameters. Let $\tilde{\boldsymbol{K}}_t^i, i = 1, \ldots, n_e$ be n_e samples from this matrix-T distribution with plug-in parameter values $\hat{\boldsymbol{K}}_t^{-i}, i = 1, \ldots, n_e$, respectively (Sætrom and Omre, 2011). If these simulated weights are used in the conditioning step instead of the plug in estimates themselves, the member variance will be approximately correct and coupling will approximately be zero. The interdependence in the weights is not fully removed since the plug-in estimates in the matrix-T model parameters are dependent.

The simulation algorithm for a matrix-T distribution can be performed in three steps. Firstly, generate an ensemble $\boldsymbol{e}_G : \{(\boldsymbol{r}^j, \boldsymbol{d}^j), j = 1, \ldots, n_e\}$ from a centred multi-Gaussian vector with plug-in estimates $\hat{\boldsymbol{\Sigma}}_{rd}$ from the EnKF; secondly assess $\boldsymbol{\Sigma}_{rd}$ from $\boldsymbol{e}_G$ to obtain $\tilde{\boldsymbol{\Sigma}}_{rd}$ and thirdly compute $\tilde{\boldsymbol{K}} = \tilde{\boldsymbol{\Gamma}}_{rd}[\tilde{\boldsymbol{\Sigma}}_d]^{-1}$. For $n_e > n + m$, the resulting $\tilde{\boldsymbol{K}}$ is matrix-T distributed. This approach can also be used when $n_e \le n + m$, but then the distributional characteristics of $\tilde{\boldsymbol{K}}$ remain unspecified.

3.2.1.3   Practical implemetation   Based on the discussion in the two previous sections, we recommend using the approach detailed in Algorithm 2 for the conditioning step.

Algorithm 2 Hierarchical EnKF

$\boldsymbol{e}_t : \{(\boldsymbol{r}_t^{u(i)}, \boldsymbol{d}_t^{(i)}), i = 1, \dots, n_e\}$

Conditioning:

For $i = 1, \dots, n_e$

Define $\boldsymbol{e}_t^{-i} = \boldsymbol{e}_t \backslash \{\boldsymbol{r}_t^{u(i)}, \boldsymbol{d}_t^{(i)}\}$

Assess $\boldsymbol{\Sigma}_{rd}$ from $\boldsymbol{e}_t^{-i} \longrightarrow \hat{\boldsymbol{\Sigma}}_{rd}^{-i}$

Generate $\boldsymbol{e}^{Gi} : \{(\boldsymbol{r}^{(j)}, \boldsymbol{d}^{(j)}), j = 1, \dots, n_e - 1\}$ iid $\varphi_{n+m}\left(\begin{bmatrix} \boldsymbol{r} \\ \boldsymbol{d} \end{bmatrix}; 0\mathbf{i}_{n+m}, \hat{\boldsymbol{\Sigma}}_{rd}^{-i}\right)$

Assess $\boldsymbol{\Sigma}_{rd}$ from $\boldsymbol{e}^{Gi} \longrightarrow \tilde{\boldsymbol{\Sigma}}_{rd}^i \rightarrow \tilde{\boldsymbol{K}}_t^i = \tilde{\boldsymbol{\Gamma}}_{rd}^i [\tilde{\boldsymbol{\Sigma}}_d^i]^{-1}$

$\boldsymbol{r}_t^{c(i)} = \boldsymbol{r}_t^{u(i)} + \tilde{\boldsymbol{K}}_t^i (\boldsymbol{d}_t - \boldsymbol{d}_t^i)$

End for

Forwarding:

$\boldsymbol{r}_{t+1}^{u(i)} = \omega_t(\boldsymbol{r}_t^{c(i)}) + \boldsymbol{\epsilon}_t^r; i = 1, \dots, n_e$

$\boldsymbol{d}_{t+1}^{(i)} = \psi_{t+1}(\boldsymbol{r}_{t+1}^{u(i)}) + \boldsymbol{\epsilon}_t^d; i = 1, \dots, n_e$

$\boldsymbol{e}_{t+1} : \{(\boldsymbol{r}_{t+1}^{u(i)}, \boldsymbol{d}_{t+1}^{(i)}); i = 1, \dots, n_e\}$

This conditioning procedure ensures correct centering by using the ensemble $\boldsymbol{e}_t^{-i}$ and an unbiased estimator for $\boldsymbol{K}_t$, and approximately correct variance by capturing the uncertainty in $\boldsymbol{K}_t$ estimates. The latter correction also reduces the coupling of the ensemble members.

3.2.1.4 Ensemble inflation Ensemble shrinkage, cause by repeated conditioning steps in the EnKF, is frequently observed in practice. Ensemble inflation (Anderson and Anderson, 1999) is a pragmatic technique that treats the problem rather than correcting for the cause,

$$\boldsymbol{r}_I^{u(i)} = \alpha \times (\boldsymbol{r}^{u(i)} - \bar{\boldsymbol{r}}^u) + \bar{\boldsymbol{r}}^u, i = 1, \dots, n_e, \tag{25}$$

where $\bar{\boldsymbol{r}}^u$ is the ensemble average and $\alpha \geq 1$ is the inflation factor and where the index I denotes the inflated ensemble members. Inflation increases the variance of each ensemble member but does not reduce the coupling between the members. The value is often claimed to be assigned from experience although trial and error seems to be somewhat prevalent. Adaptive inflation (Anderson, 2007) can also used to systematically estimate the inflation factor. The subjectivity of the inflation factor in correcting the ensemble variability reduces the confidence one should have in the prediction intervals provided by the EnKF when using inflation.

3.3 General case: Hidden Markov model

In the general case, when the initial distribution is non-Gaussian and/or the forward and likelihood models are non-linear, the EnKF only approximates the filtering/forecasting distribution $f(\boldsymbol{r}_{T+1}|\boldsymbol{d}_{0:T})$. In this section, we consider the HM model defined in Section 2, with additive Gaussian errors. The HM model is therefore defined as,

$$
\begin{aligned}
\boldsymbol{r}_0 &\sim f(\boldsymbol{r}_0). \\
[\boldsymbol{r}_{t+1}|\boldsymbol{r}_t] &= \omega_t(\boldsymbol{r}_t) + \boldsymbol{\epsilon}_t^r \rightarrow f(\boldsymbol{r}_{t+1}|\boldsymbol{r}_t) = \varphi_n(\boldsymbol{r};\omega_t(\boldsymbol{r}_t),\boldsymbol{\Sigma}_t^r). \\
[\boldsymbol{d}_t|\boldsymbol{r}_t] &= \psi_t(\boldsymbol{r}_t) + \boldsymbol{\epsilon}_t^d \rightarrow f(\boldsymbol{d}_t|\boldsymbol{r}_t) = \varphi_m(\boldsymbol{d};\psi_t(\boldsymbol{r}_t),\boldsymbol{\Sigma}_t^d),
\end{aligned}
\tag{26}
$$

with n-vector $\boldsymbol{\epsilon}_t^r$ and m-vector $\boldsymbol{\epsilon}_t^d$ being centred Gaussian with covariance matrices $\boldsymbol{\Sigma}_t^r$ and $\boldsymbol{\Sigma}_t^d$ respectively. In the following paragraph, we discuss three issues arising when the forward and likelihood models are non-linear:

1. Instability in the Kalman weight inference

2. Drift towards Gaussianity

3. Linearized conditioning

3.3.1 Instability in the Kalman weight inference

The ensemble $\boldsymbol{e}_t^{rd}$, representing $f(\boldsymbol{r}_t^u,\boldsymbol{d}_t)$ is in the general case non-Gaussian. The approximate linearized conditioning is performed with the Kalman weights $\boldsymbol{K}_t$ which must be estimated from the ensemble. As default, one estimates the covariance matrix of $[\boldsymbol{r}_t^u,\boldsymbol{d}_t]$ denoted $\hat{\boldsymbol{\Sigma}}_{rd}$ using the traditional unbiased covariance estimator and calculates the weights $\hat{\boldsymbol{K}}_t = \hat{\boldsymbol{\Gamma}}_{rd}[\hat{\boldsymbol{\Sigma}}_d]^{-1}$. In practice, one may only calculate the two submatrices of $\hat{\boldsymbol{\Sigma}}_{rd}$ that define $\hat{\boldsymbol{K}}_t$. There is of course estimator uncertainty but in the traditional EnKF algorithm, one uses directly the plug-in estimate $\hat{\boldsymbol{K}}_t$ and ignores this uncertainty.

Two sources of uncertainty appear in $\hat{\boldsymbol{K}}_t$. Firstly, there is estimation uncertainty since the ensemble size is finite, $n_e < \infty$. Secondly, if the ensemble size is small, $n_e \le n + m$, $\hat{\boldsymbol{K}}_t$ will also lack full rank. Estimation error in $\hat{\boldsymbol{K}}_t$ will distribute the weights somewhat arbitrarily to the m-vector $[\boldsymbol{d}_t - \boldsymbol{d}_t^i]$. The major concern is that observation dimensions which should have no influence on conditioning are by chance assigned influence. Moreover, it is notorious that the traditional unbiased covariance estimator is not robust with respect to deviations from Gaussianity (Huber, 1981). In statistics, robust shrinkage estimators are frequently used (Ledoit and Wolf, 2004) to account for this instability.

Rank deficiency in $\hat{\boldsymbol{K}}_t$ occurs if $n_e \le n + m$ which is often the case. This effect will also cause the weights to be arbitrarily distributed. The only way to correct for this effect is to introduce subjective information in the estimator. We discuss two alternative approaches for improving the Kalman weights estimate:

1. Covariance localization

2. Hierarchical approach

3.3.1.1 Covariance localization Covariance localization (Hamill et al., 2001; Bocquet and Carrassi, 2017) can be seen as a pragmatic inference alternative in the sense that it both reduces the estimation uncertainty in $\boldsymbol{\Sigma}_{rd}$ and adresses the rank defiency issue. It is conceptually simple as it operates

directly on the traditional unbiased covariance estimator by multiplying it element-wise with a damping $(n + m) \times (n + m)$-matrix $\boldsymbol{\Omega}$,

$$\tilde{\boldsymbol{\Sigma}}_{rd} = \boldsymbol{\Omega} \circ \hat{\boldsymbol{\Sigma}}_{rd}, \tag{27}$$

where $\circ$ is the element-wise matrix product and $\hat{\boldsymbol{\Sigma}}_{rd}$ is the traditional unbiased covariance estimator. This localization matrix $\boldsymbol{\Omega}$ must be assigned based on experience with the correlation structure in $[\boldsymbol{r}, \boldsymbol{d}]$ and it must be positive definite. The piecewise rational function defined in Gaspari and Cohn (1999) is often used to define $\boldsymbol{\Omega}$.

The Kalman weight estimator is then $\tilde{\boldsymbol{K}}_t = \tilde{\boldsymbol{\Gamma}}_{rd}[\tilde{\boldsymbol{\Sigma}}_d]^{-1}$ which appears with shrinkage towards the weights implicitly defined by $\boldsymbol{\Omega}$. Moreover, the estimated Kalman weights $\tilde{\boldsymbol{K}}_t$ have improved rank properties.

3.3.1.2 Hierarchical approach   The hierarchical approach Myrseth and Omre (2010) can be considered as Bayesian inference in a Gaussian model. The model parameters $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ are seen as random variables and are assigned prior pdfs, which add subjectivity in the model. By using conjugate prior models, analytical tractability is ensured. Consider the Gauss-linear HM model, and define the model parameters to be random variables in a Bayesian inference setting,

$$[\boldsymbol{r}_t^u, \boldsymbol{d}_t | \boldsymbol{\mu}_{rd}, \boldsymbol{\Sigma}_{rd}] \sim \varphi_{n+m}\left(\begin{bmatrix} \boldsymbol{r} \\ \boldsymbol{d} \end{bmatrix}; \boldsymbol{\mu}_{rd}, \boldsymbol{\Sigma}_{rd}\right), \tag{28}$$

and assume further $\boldsymbol{e}_t^{rd} = \{(\boldsymbol{r}_t^{u(i)}, \boldsymbol{d}_t^i), i = 1, \dots, n_e\}$ with n_e iid samples. Assign a conjugate prior model to the model parameters,

$$f(\boldsymbol{\mu}_{rd} | \boldsymbol{\Sigma}_{rd}) = \varphi_{n+m}(\boldsymbol{\mu}; \boldsymbol{\mu}_p, \tau_p \boldsymbol{\Sigma}_{rd}), \tag{29}$$
$$f(\boldsymbol{\Sigma}_{rd}) = \mathrm{W}^{-1}(\boldsymbol{\Sigma}; \boldsymbol{\Sigma}_p, \nu_p),$$

where W^{-1} denoted the inverse Wishart pdf, with user defined hyper parameters being $(n+m)$-vector $\boldsymbol{\mu}_p$, scale factor $\tau_p \in \mathbb{R}_+$, positive definite covariance $((n+m) \times (n+m))$-matrix $\boldsymbol{\Sigma}_p$ and degree of freedom $\nu_p \in \mathbb{R}_{[n+m+1, +\infty]}$. Note that the prior expectation for $\boldsymbol{\Sigma}_{rd}$ is $\mathbb{E}(\boldsymbol{\Sigma}_{rd}) = [\nu_p - (n+m) - 1]^{-1} \boldsymbol{\Sigma}_p$.

The corresponding posterior distribution for the covariance matrix $\boldsymbol{\Sigma}_{rd}$ is,

$$f(\boldsymbol{\Sigma}_{rd} | \boldsymbol{e}_t) = \mathrm{W}^{-1}(\boldsymbol{\Sigma}; \boldsymbol{\Sigma}_{pp}, \nu_p + n_e), \tag{30}$$

with,

$$\boldsymbol{\Sigma}_{pp} = \boldsymbol{\Sigma}_p + (n_e - 1)\hat{\boldsymbol{\Sigma}}_{rd} + \left[\frac{1}{n_e} + \tau_p\right]^{-1} (\hat{\boldsymbol{\mu}}_{rd} - \boldsymbol{\mu}_p)(\hat{\boldsymbol{\mu}}_{rd} - \boldsymbol{\mu}_p)^T, \tag{31}$$

and hence the hierarchical estimator for $\boldsymbol{\Sigma}_{rd}$,

$$\tilde{\boldsymbol{\Sigma}}_{rd} = \mathbb{E}(\boldsymbol{\Sigma}_{rd} | \boldsymbol{e}_t) = [\nu_p + n_e - (n+m) - 1]^{-1} \boldsymbol{\Sigma}_{pp}. \tag{32}$$

The hierarchical estimator $\tilde{\boldsymbol{\Sigma}}_{rd}$ appears as a weighted average of the prior assignment, the traditional unbiased covariance estimator and the deviations of the expectations. The latter term is often ignored by using plug-in estimates for the expectation in the prior model. Note that in the limit when $n_e \to \infty$, one obtains $\tilde{\boldsymbol{\Sigma}}_{rd} = \hat{\boldsymbol{\Sigma}}_{rd}$, the traditional estimator. Moreover, even for cases where $n_e < n + m$ the hierarchical

estimator $\bar{\Sigma}_{rd}$ has full rank which is ensured by the user-specified prior Σ_p having full rank. The Kalman weight estimator is then $\bar{K}_t = \bar{\Gamma}_{rd}[\bar{\Sigma}_d]^{-1}$ which has full rank and appears with shrinkage towards the weights implicitly defined by the user-specified prior Σ_p. The challenge is to assign a suitable prior model for Σ_{rd} with expectation $[v_p - (n + m) - 1]^{-1}\Sigma_p$ which must reflect experience with the correlation structure in $[r_t^u, d_t]$. The confidence in the prior assignment is reflected in the degree of freedom $v_p > n + m + 1$.

3.3.2 Drift towards Gaussianity

The EnKF algorithm is, as mentioned previously, asymptotically correct, when $n_e \to \infty$, for Gauss-linear model HM model. Consider an initial ensemble e_0 generated from a non-Gaussian initial model $f(r_0)$. Assume further that the forward and likelihood functions are continuous.

The sequential updates in the EnKF will make the ensemble e_t more and more Gaussian - the ensemble drifts towards Gaussianity. This drift is caused by the successive linearized updates when conditioning on the data.

Several variants of the EnKF exist that address this issue:

1. Gaussian anamorphosis (GA) EnKF

2. Gaussian Mixture (GM) EnKF

3. Selection (S) EnKF

3.3.2.1 Gaussian anamorphosis (GA) EnKF The idea is to transform the ensemble to be marginally Gaussian before conditioning, the transformed ensemble will be marginally Gaussian, and to carry out the conditioning with the transformed ensemble. After conditioning, the ensemble is then back-transformed. Applications have shown that Gaussian anamorphosis can successfully prevent the drift towards Gaussianity, see Zhou et al. (2012).

Consider a univariate random variable y with cdf $F_Y(y)$ and a random sample $(y_1, \ldots, y_{n_y})$ iid from $F_Y(y)$. The cdf can be estimated as,

$$\hat{F}_Y(y) = J\left\{ \frac{1}{n_y} \sum_{j=1}^{n_y} \mathbf{1}(y_j \le y) \right\}, \tag{33}$$

where $J\{\cdot\}$ is some semi parametric smoother of the empirical stepwise cdf estimator in the argument. The smoother ensures that the back-transform is real valued.

The univariate Gaussian transform of one sample y_0 from $F_Y(y)$ is defined as,

$$\tilde{y}_0 = \Phi_1(\hat{F}_Y^{-1}(y_0); 0, 1). \tag{34}$$

Note that $\tilde{y}_0$ has an approximate standard Gaussian pdf. Similarly the back-transform of the univariate Gaussian sample $\tilde{u}_0$ is,

$$u_0 = \hat{F}_Y(\Phi_1^{-1}(\tilde{u}_0; 0, 1)). \tag{35}$$

The smoothing of $\hat{F}_Y(y)$ ensures that $u_0 \in \mathbb{R}$ and do not belong to the set $\{y_1, \ldots, y_{n_y}\}$ only.

The univariate transformation of an ensemble $e_t : \{(r_t^{u(i)}, d_t^{(i)}), i = 1, \ldots, n_e\}$ entails that for each ensemble member, each of the $(n + m)$ dimensions must be transformed independently. Hence only approximate univariate Gaussianity is ensured, the multivariate characteristics remaining unspecified. The

latter entails that the linearized conditioning is only approximately correct. Algorithm 3 presents the GA EnKF procedure.

Algorithm 3 Gaussian anamorphosis EnKF

$\boldsymbol{e}_t : \{ (\boldsymbol{r}_t^{u(i)}, \boldsymbol{d}_t^{(i)}), i = 1, \dots, n_e \}$

Conditioning:

Assess all $(n + m)$ dimensions by $F_Y(y)$ from $\boldsymbol{e}_t \rightarrow \hat{F}_Y(y)$

Univariate Gaussian transform of $\boldsymbol{e}_t$ by $\hat{F}_Y(y) \rightarrow \tilde{\boldsymbol{e}}_t$

Univariate Gaussian transform of observations $\boldsymbol{d}_t$ by $\hat{F}_Y(y) \rightarrow \tilde{\boldsymbol{d}}_t$

Estimate $\boldsymbol{\Sigma}_{rd}$ from $\tilde{\boldsymbol{e}}_t \rightarrow \hat{\boldsymbol{\Sigma}}_{rd} \rightarrow \hat{\boldsymbol{K}}_t = \hat{\boldsymbol{\Gamma}}_{rd} [\hat{\boldsymbol{\Sigma}}_d]^{-1}$

$\tilde{\boldsymbol{r}}_t^{c(i)} = \tilde{\boldsymbol{r}}_t^{u(i)} + \hat{\boldsymbol{K}}_t (\tilde{\boldsymbol{d}}_t - \tilde{\boldsymbol{d}}_t^{(i)}), \ i = 1, \dots, n_e$

Univariate back-transform of $\tilde{\boldsymbol{r}}_t^{c(i)}$ by $\hat{F}_Y(y) \rightarrow \boldsymbol{r}_t^{c(i)}$

Forwarding:

$\boldsymbol{r}_{t+1}^{u(i)} = \omega_t (\boldsymbol{r}_t^{c(i)}) + \boldsymbol{\epsilon}_t^r ; i = 1, \dots, n_e$

$\boldsymbol{d}_{t+1}^{(i)} = \psi_{t+1} (\boldsymbol{r}_{t+1}^{u(i)}) + \boldsymbol{\epsilon}_t^d ; i = 1, \dots, n_e$

$\boldsymbol{e}_{t+1} : \{ (\boldsymbol{r}_{t+1}^{u(i)}, \boldsymbol{d}_{t+1}^{(i)}); i = 1, \dots, n_e \}$

3.3.2.2   Gaussian Mixture (GM) EnKF   The basic idea is to specify the prior initial model as a Gaussian mixture (GM) model which can represent multimodal variables. Let the forward and likelihood models of the HM model be Gauss-linear. The posterior pdf will then also be a GM model and be analytically tractable. The conditioning step can be made independently for each component of the GM model, and the associated weight for each component can also be calculated. For general forward and likelihood models, ensemble based filtering algorithms must be used. The difficulty is that each ensemble member must carry a mode indicator assigned to one of the components, and that this indicator may change during the forwarding step. These filtering algorithms have proven robust against the drift towards Gaussianity (Li et al., 2016; Ackerson and Fu, 1970; Chen and Liu, 2000; Smith, 2007; Dovera and Della Rossa, 2010; Bengtsson et al., 2003), at least for low-dimensional models.

Consider a set of n-dimensional Gaussian pdfs, $\varphi_n(\boldsymbol{r}; \boldsymbol{\mu}_r^l, \boldsymbol{\Sigma}_r^l); l = 1, \dots, L$, denoted components, and a set of normalized mixture weights $\pi : \{\pi_1, \dots, \pi_L\}$. The prior initial model is specified to be a GM model :

$$f(\boldsymbol{r}_0) = \sum_{l=1}^{L} \pi_l \times \varphi_n(\boldsymbol{r}_0; \boldsymbol{\mu}_r^l, \boldsymbol{\Sigma}_r^l). \tag{36}$$

Note that a particular realization $\boldsymbol{r}_s$ will belong to the k-th component of the mixture with probability:

$$\lambda_k(\boldsymbol{r}_s) = \left[\sum_{l=1}^{L} \pi_l \times \varphi_n(\boldsymbol{r}_s; \boldsymbol{\mu}_r^l, \boldsymbol{\Sigma}_r^l) \right]^{-1} \times \varphi_n(\boldsymbol{r}_s; \boldsymbol{\mu}_r^k, \boldsymbol{\Sigma}_r^k). \tag{37}$$

Consider further the Gauss-linear likelihood model,

$$f(\boldsymbol{d}|\boldsymbol{r}) = \varphi_m(\boldsymbol{d}; \boldsymbol{H}\boldsymbol{r}, \boldsymbol{\Sigma}_{d|r}). \tag{38}$$

The associated posterior model will also be a Gaussian mixture model (Grana et al., 2017),

$$f(\boldsymbol{r}|\boldsymbol{d}) = \sum_{l=1}^{L} \pi_{l|d} \times \varphi_n(\boldsymbol{r}; \boldsymbol{\mu}_{r|d}^l, \boldsymbol{\Sigma}_{r|d}^l), \tag{39}$$

where the conditional expectations and covariances are obtained by component wise conditioning on the observations $\boldsymbol{d}$. The posterior mixture weights are defined by

$$\pi_{k|d} = \left[\sum_{l=1}^{L} \pi_l \times \varphi_m(\boldsymbol{d}; \boldsymbol{H}\boldsymbol{\mu}_r^l, \boldsymbol{H}\boldsymbol{\Sigma}_r^l \boldsymbol{H}^T + \boldsymbol{\Sigma}_{d|r}) \right]^{-1} \pi_k \times \varphi_m(\boldsymbol{d}; \boldsymbol{H}\boldsymbol{\mu}_r^k, \boldsymbol{H}\boldsymbol{\Sigma}_r^k \boldsymbol{H}^T + \boldsymbol{\Sigma}_{d|r}). \tag{40}$$

For general forward and likelihood models, the ensemble representation must contain a mode indicator associated with each ensemble member,

$$\boldsymbol{e}_t : \{(\boldsymbol{r}_t^{u(i)}, l_t^{(i)}, \boldsymbol{d}_t^{(i)}, i = 1, \dots, n_e\}, \tag{41}$$

with mode indicator $l_t^{(i)} \in \{1, \dots, L\}$. The dynamic updating of this mode indicator often appears as challenging. The conditioning/forward steps in the GM-EnKF are detailed in Algorithm 4.

Algorithm 4 Gaussian mixture EnKF

$\boldsymbol{e}_t^{-l} : \{(\boldsymbol{r}_t^{u(i)}, \cdot, \boldsymbol{d}_t^{(i)}), i = 1, \dots, n_e\}$

Conditioning:

Assess $f(\boldsymbol{r}_t^u)$ from $\boldsymbol{e}_t^{-l} \rightarrow \hat{f}(\boldsymbol{r}_t^u) = \sum_{l=1}^{L} \hat{\pi}_l \times \varphi_n(\boldsymbol{r}, \hat{\mu}_r^l, \hat{\Sigma}_r^l)$

Assign ensemble member i to mode indicator $l_t^{(i)} \in \{1, \dots, L\}$ with probability $\{\lambda_l(\boldsymbol{r}_t^{u(i)}), l = 1, \dots, L\}$

Define ensemble $\boldsymbol{e}_t : \{(\boldsymbol{r}_t^{u(i)}, l_t^{(i)}, \boldsymbol{d}_t^{(i)}, i = 1, \dots, n_e\}$

Assess $\boldsymbol{\Sigma}_{rd}^l; l = 1, \dots, L$ from $\boldsymbol{e}_t \rightarrow \hat{\boldsymbol{\Sigma}}_{rd}^l \rightarrow \hat{\boldsymbol{K}}_t^l = \hat{\boldsymbol{\Gamma}}_{rd}^l [\hat{\boldsymbol{\Sigma}}_d^l]^{-1}$

$\boldsymbol{r}_t^{c(i)} = \boldsymbol{r}_t^{u(i)} + \hat{\boldsymbol{K}}_t^{l_t^{(i)}} (\boldsymbol{d}_t - \boldsymbol{d}_t^{(i)}), i = 1, \dots, n_e$

Forwarding:

$\boldsymbol{r}_{t+1}^{u(i)} = \omega_t(\boldsymbol{r}_t^{c(i)}) + \boldsymbol{\epsilon}_t^r, i = 1, \dots, n_e$

$\boldsymbol{d}_{t+1}^{(i)} = \psi_{t+1}(\boldsymbol{r}_{t+1}^{u(i)}) + \boldsymbol{\epsilon}_t^d, i = 1, \dots, n_e$

$\boldsymbol{e}_{t+1}^{-l} : \{(\boldsymbol{r}_{t+1}^{u(i)}, \boldsymbol{d}_{t+1}^{(i)}), i = 1, \dots, n_e\}$

The challenging part of the algorithm is to assess the GM model $f(\boldsymbol{r}_t^u)$. Other versions of the (GM) EnKF algorithm use the EM-algorithm, particle filters or clustering techniques. If the dimension of $\boldsymbol{r}$ is of some size, estimating a suitable GM model will be complicated.

3.3.2.3 Selection (S)EnKF The basic idea is to specify the prior initial model as a selection-Gaussian pdf, see Section 2.2, which can represent multimodal, skewed and/or peaked variables. For general forward and likelihood models, one must rely on ensemble based filtering algorithms. These algorithms are inspired by the analytically tractable model discussed in Section 2.2, and have proven to be robust with regard to drift towards Gaussianity (Conjard and Omre, 2020). Let the prior initial distribution $f(\boldsymbol{r}_0)$ be a selection-Gaussian pdf with parameters $\Theta^{SG} = (\boldsymbol{\mu}_{\tilde{r}}, \boldsymbol{\Sigma}_{\tilde{r}}, \boldsymbol{\mu}_v, \boldsymbol{\Sigma}_{v|\tilde{r}}, \boldsymbol{\Gamma}_{v|\tilde{r}}, A)$, see Equations 18 and 19. The auxiliary variables $(\tilde{\boldsymbol{r}}_0, \boldsymbol{v})$ are then jointly Gaussian and the variable of interest is $\boldsymbol{r}_0 = [\tilde{\boldsymbol{r}}_0 | \boldsymbol{v} \in A]$ which is selection-Gaussian.

The initial ensemble $\boldsymbol{e}_0$ of the EnKF algorithm contains realizations of the auxiliary variables $[\tilde{\boldsymbol{r}}_0, \boldsymbol{v}]$ which are jointly Gaussian. The SEnKF algorithm is identical to the EnKF algorithm defined on these auxiliary variables. The forward model is given by

$$\begin{bmatrix} \tilde{\boldsymbol{r}}_{t+1} \\ \boldsymbol{v}_{t+1} \end{bmatrix} = \begin{bmatrix} \psi_t(\tilde{\boldsymbol{r}}_t, \boldsymbol{\epsilon}_t^r) \\ \boldsymbol{v}_t \end{bmatrix}, \tag{42}$$

while the likelihood model is given by

$$\boldsymbol{d}_t = \omega_t(\tilde{\boldsymbol{r}}_t, \boldsymbol{\epsilon}_t^d). \tag{43}$$

Based on these models, an algorithm identical to the traditional EnKF algorithm is activated, to obtain the ensemble $\boldsymbol{e}_{T+1} = \{(\tilde{\boldsymbol{r}}_{T+1}^{u(i)}, \boldsymbol{v}_{T+1}^{u(i)}), i = 1, \dots, n_e\}$. Note that a time index is added to $\boldsymbol{v}$ to account for the data assimilation up to time T. The expectation vector $\boldsymbol{\mu}_{\tilde{r}v}$ and covariance matrix $\boldsymbol{\Sigma}_{\tilde{r}v}$ are estimated from $\boldsymbol{e}_{T+1}$, and based on the jointly Gaussian $\varphi_{2n}((\tilde{\boldsymbol{r}}, \boldsymbol{v}); \hat{\boldsymbol{\mu}}_{\tilde{r}v}, \hat{\boldsymbol{\Sigma}}_{\tilde{r}v})$, the filter variable of interest $[\boldsymbol{r}_{T+1} | \boldsymbol{d}_{0:T}] = [\tilde{\boldsymbol{r}}_{T+1} | \boldsymbol{v} \in A, \boldsymbol{d}_{0:T}]$ is assessed by McMC simulation.

The conditioning and forwarding steps of the SEnKF algorithm are specified in Algorithm 5.

Algorithm 5 Selection EnKF

$\boldsymbol{e}_t = \{(\tilde{\boldsymbol{r}}_t^{u(i)}, \boldsymbol{v}_t^{u(i)}, \boldsymbol{d}_t^i), i = 1, \dots, n_e\}$

Conditioning:

Estimate $\boldsymbol{\Sigma}_{\tilde{r}vd}$ from $\boldsymbol{e}_t \longrightarrow \hat{\boldsymbol{\Sigma}}_{\tilde{r}vd} \rightarrow \hat{\boldsymbol{K}}_t = \hat{\boldsymbol{\Gamma}}_{rvd}[\hat{\boldsymbol{\Sigma}}_d]^{-1}$

$\begin{bmatrix} \tilde{\boldsymbol{r}}_t^{c(i)} \\ \boldsymbol{v}_t^{c(i)} \end{bmatrix} = \begin{bmatrix} \tilde{\boldsymbol{r}}_t^{u(i)} \\ \boldsymbol{v}_t^{u(i)} \end{bmatrix} + \hat{\boldsymbol{K}}_t(\boldsymbol{d}_t - \boldsymbol{d}_t^{(i)}), i = 1, \dots, n_e$

Forwarding:

Generate $\boldsymbol{\epsilon}_t^{\tilde{r}(i)}, i = 1, \dots, n_e$

$\begin{bmatrix} \tilde{\boldsymbol{r}}_{t+1}^{u(i)} \\ \boldsymbol{v}_{t+1}^{u(i)} \end{bmatrix} = \begin{bmatrix} \omega_t(\tilde{\boldsymbol{r}}_t^{c(i)}, \boldsymbol{\epsilon}_t^{\tilde{r}(i)}) \\ \boldsymbol{v}_t^{c(i)} \end{bmatrix}, i = 1, \dots, n_e$

Generate $\boldsymbol{\epsilon}_t^{d(i)}, i = 1, \dots, n_e$

$\boldsymbol{d}_{t+1}^{(i)} = \psi_{t+1}(\tilde{\boldsymbol{r}}_{t+1}^{u(i)}, \boldsymbol{\epsilon}_{t+1}^{d(i)}), i = 1, \dots, n_e$

$\boldsymbol{e}_{t+1} = \{(\tilde{\boldsymbol{r}}_{t+1}^{u(i)}, \boldsymbol{v}_{t+1}^{u(i)}, \boldsymbol{d}_{t+1}^i), i = 1, \dots, n_e\}$

3.3.3 Linearized conditioning

The ensemble $\boldsymbol{e}_t$ representing $f(\boldsymbol{r}_t^u, \boldsymbol{d}_t)$ is in the general case non-Gaussian. The conditioning step is initiated by assessing $\boldsymbol{\Sigma}_{rd}$ from the ensemble to obtain $\hat{\boldsymbol{\Sigma}}_{rd}$. The cross covariance matrix $\boldsymbol{\Gamma}_{rd}$ only captures the linear relationship between $[\boldsymbol{r}_t^u, \boldsymbol{d}_t]$, not the non-linear features of the observation function $\psi_t(\cdot, \cdot)$. Since the linearization of $\psi_t(\cdot, \cdot)$ is based on the ensemble $\boldsymbol{e}_t$, it is made over the span of the ensemble members. The observation error will also influence the impact of extreme pseudo-observations. Narrow span ensembles with larger observation errors tend to give good linear approximation. We discuss two alternative approaches for reducing the effect of this approximate linearization in the conditioning step,

1. Multiple data assimilation.

2. Iterative data assimilation.

3.3.3.1 Multiple data assimilation (MDA) EnKF The basic idea is to split the conditioning step into a sequence of conditioning sub-steps (Emerick and Reynolds, 2013). Each sub-step uses a larger variance in the likelihood model, and these variances are balanced such that the conditioning is correct for the Gauss-linear HM model. The empirical linearization is over the span of the ensemble, which is narrowed by sub-conditioning. Moreover, the likelihood variance is increased for each sub-conditioning step. These two effects make the MDA EnKF algorithm efficient in practice (Emerick and Reynolds, 2013).

The likelihood model in the general HM model, $f(\boldsymbol{d}_t | \boldsymbol{r}_t)$, is defined to have a non-linear likelihood function with additive Gaussian error term, see Equation 38. This likelihood model can be decomposed as,

$$
\begin{aligned}
f(\boldsymbol{d}_t | \boldsymbol{r}_t^u) &= \prod_{i=1}^{L} [(f(\boldsymbol{d}_t | \boldsymbol{r}_t^u))]^{w_i}, \\
&= \prod_{i=1}^{L} [(\varphi_m(\boldsymbol{d}; \psi_t(\boldsymbol{r}_t^u), \boldsymbol{\Sigma}_t^d))]^{w_i}, \\
&= const \times \prod_{i=1}^{L} exp\left(-\frac{1}{2}(\boldsymbol{d}_t - \psi_t(\boldsymbol{r}_t^u))\left[\frac{1}{w_i}\boldsymbol{\Sigma}_t^d\right]^{-1}(\boldsymbol{d}_t - \psi_t(\boldsymbol{r}_t^u))^T\right),
\end{aligned}
\tag{44}
$$

with positive decomposition factors $\boldsymbol{w} : \{w_i, i = 1, \dots, L\}$ such that $\sum_{i=1}^{L} w_i = 1$, ordered in decreased order. The posterior model in the conditioning step can therefore be expressed by

$$
\begin{aligned}
f(\boldsymbol{r}_t^u | \boldsymbol{d}_t) &= const \times f(\boldsymbol{d}_t | \boldsymbol{r}_t^u) f(\boldsymbol{r}_t^u), \\
&= const \times exp\left(-\frac{1}{2}(\boldsymbol{d}_t - \psi_t(\boldsymbol{r}_t^u))\left[\frac{1}{w_1}\boldsymbol{\Sigma}_t^d\right]^{-1}(\boldsymbol{d}_t - \psi_t(\boldsymbol{r}_t^u))^T\right) \\
&\quad \times \dots \\
&\quad \times exp\left(-\frac{1}{2}(\boldsymbol{d}_t - \psi_t(\boldsymbol{r}_t^u))\left[\frac{1}{w_L}\boldsymbol{\Sigma}_t^d\right]^{-1}(\boldsymbol{d}_t - \psi_t(\boldsymbol{r}_t^u))^T\right) \\
&\quad \times f(\boldsymbol{r}_t^u).
\end{aligned}
\tag{45}
$$

This decomposition justifies the sequential conditioning procedure. The sub-conditioning steps starts from the right hand side, by conditioning on $\boldsymbol{d}_t$ associated with variance $\frac{1}{w_L}\boldsymbol{\Sigma}_t^d$. Then the resulting

updated r_t^u is conditioned on d_t with variance $\frac{1}{w_{L-1}}\Sigma_t^d$, etc. Finally conditioning on d_t associated with variance $\frac{1}{w_1}\Sigma_t^d$ provides the correct posterior distribution $f(r_t^u|d_t)$.

The span of the ensemble members narrows and the likelihood variance increases in the sub-conditioning steps. In practice, this makes the conditioning piece-wise linear, which of course is beneficial for non-linear likelihood functions.

The conditioning and forwarding steps of the MDA EnKF algorithm are specified in Algorithm 6.

Algorithm 6 Multiple data assimilation EnKF

Conditioning:

Set $r_{t,1}^{u(i)} = r_t^{u(i)}$, $i = 1,...,n_e$

For $k = 1,\dots,L$

$\boldsymbol{\epsilon}_{t,k}^{d(i)} \sim \varphi_m(\boldsymbol{\epsilon};0,\frac{1}{\sqrt{w_k}}\Sigma_t^d)$, $i = 1,...,n_e$

$d_{t,k}^{(i)} = \psi_t(r_{t,k}^{u(i)},\boldsymbol{\epsilon}_{t,k}^{d(i)})$, $i = 1,...,n_e$

$e_{t,k} = \{(r_{t,k}^{u(i)},d_{t,k}^i), i = 1,...,n_e\}$

Estimate $\boldsymbol{\Sigma}_{rd}$ from $e_{t,k} \longrightarrow \hat{\boldsymbol{\Sigma}}_{rd,k}$

$r_{t,k+1}^{u(i)} = r_{t,k}^{u(i)} + \hat{\boldsymbol{\Gamma}}_{r,d}\hat{\boldsymbol{\Sigma}}_{d,k}^{-1}(d_t - d_{t,k}^i)$, $i = 1,...,n_e$

End For

$r_t^{c(i)} = r_{t,L+1}^{u(i)}$, $i = 1,...,n_e$

Forwarding:

$r_{t+1}^{u(i)} = \omega_t(r_t^{c(i)}) + \boldsymbol{\epsilon}_t^r$, $i = 1,\dots,n_e$

$d_{t+1}^{(i)} = \psi_{t+1}(r_{t+1}^{u(i)}) + \boldsymbol{\epsilon}_t^d$, $i = 1,\dots,n_e$

$e_{t+1} : \{(r_{t+1}^{u(i)},d_{t+1}^{(i)}), i = 1,\dots,n_e\}$

3.3.3.2 Iterative(I) EnKF The basic idea is to rephrase the conditioning step as an optimization problem (Li and Reynolds, 2009; Sakov et al., 2018), which can be done exactly for a Gauss-linear HMM. If the likelihood is non-linear, it is intuitive to replace the conditioning step by an iterative method such as Gauss-Newton or Levenberg-Marquardt where the gradients are either calculated by solving the adjoint problem or estimated using the ensemble.

Consider the conditioning step in a Gauss-linear HM model, the posterior pdf $f(r|d)$ is also Gaussian

and the posterior expectation therefore coincides with the posterior mode given by,

$$\underset{\boldsymbol{r}}{\operatorname{argmax}}\{f(\boldsymbol{r}|\boldsymbol{d})\} = \underset{\boldsymbol{r}}{\operatorname{argmax}}\{\log f(\boldsymbol{r}|\boldsymbol{d})\}, \tag{46}$$

$$= \underset{\boldsymbol{r}}{\operatorname{argmax}}\{(\boldsymbol{r} - \boldsymbol{\mu}_r)^T \boldsymbol{\Sigma}_r^{-1} (\boldsymbol{r} - \boldsymbol{\mu}_r) + (\boldsymbol{Hr} - \boldsymbol{d})^T \boldsymbol{\Sigma}_d^{-1} (\boldsymbol{Hr} - \boldsymbol{d})\}, \tag{47}$$

$$= \boldsymbol{\mu}_r - \boldsymbol{\Sigma}_r \boldsymbol{H}^T (\boldsymbol{H}^T \boldsymbol{\Sigma}_r \boldsymbol{H} + \boldsymbol{\Sigma}_d)^{-1} (\boldsymbol{d} - \boldsymbol{H}\boldsymbol{\mu}_r). \tag{48}$$

Note that if the substitutions $\boldsymbol{\mu}_r = \boldsymbol{r}_t^u$ and $\boldsymbol{d} = \boldsymbol{d}_t - \boldsymbol{\epsilon}_t^d$ are made in the expression above, it coincides with the EnKF conditioning step. Consequently, the conditioning optimization expression for Gauss-linear HM models is,

$$\boldsymbol{r}_t^c = \underset{\boldsymbol{r}}{\operatorname{argmax}}\{(\boldsymbol{r}^u - \boldsymbol{r})^T \hat{\boldsymbol{\Sigma}}_r^{-1} (\boldsymbol{r}^u - \boldsymbol{r}) + (\boldsymbol{Hr} - \boldsymbol{d}_s)^T \hat{\boldsymbol{\Sigma}}_d^{-1} (\boldsymbol{Hr} - \boldsymbol{d}_s)\}. \tag{49}$$

Note that $\boldsymbol{d}^s = \boldsymbol{d} - \boldsymbol{\epsilon}_t^d$ is a random pseudo truth since $\boldsymbol{d}_t = \boldsymbol{Hr}_t^u + \boldsymbol{\epsilon}_t^d$ while $\boldsymbol{r}_t^u$ is a randomized expectation in the prior. In the Gauss-linear HM model, the objective function is quadratic and the solution can be calculated analytically. This solution coincides with the EnKF conditioning step.

For a general HM model where the likelihood model is defined as,

$$[\boldsymbol{d}_t|\boldsymbol{r}_t^u] = \psi_t(\boldsymbol{r}_t^u) + \boldsymbol{\epsilon}_t^d, \tag{50}$$

one will intuitively define the conditioning optimization step as,

$$\boldsymbol{r}_t^c = \underset{\boldsymbol{r}}{\operatorname{argmax}}\{(\boldsymbol{r}^u - \boldsymbol{r})^T \boldsymbol{\Sigma}_r^{-1} (\boldsymbol{r}^u - \boldsymbol{r}) + (\psi_t(\boldsymbol{r}) - \boldsymbol{d}^s)^T \boldsymbol{\Sigma}_d^{-1} (\psi_t(\boldsymbol{r}) - \boldsymbol{d}^s)\}. \tag{51}$$

The maximization is solved using an iterative method such as Gauss-Newton or Levenberg-Marquardt where the Jacobian matrices are assessed either from the ensemble or evaluated numerically by for instance solving the adjoint problem. The conditioning and forwarding steps of the iterative EnKF are detailed in Algorithm 7

Algorithm 7 Iterative EnKF

$\boldsymbol{e}_t : \{(\boldsymbol{r}_t^{u(i)}, \boldsymbol{d}_t^{(i)}), i = 1, \dots, n_e\}$

Conditioning:

Estimate $\boldsymbol{\Sigma}_{rd}$ from $\boldsymbol{e}_t \longrightarrow \hat{\boldsymbol{\Sigma}}_{rd}$

$$\boldsymbol{r}_t^{c(i)} = \underset{\boldsymbol{r}}{\operatorname{argmax}}\{(\boldsymbol{r}_t^{u(i)} - \boldsymbol{r})^T \hat{\boldsymbol{\Sigma}}_r^{-1} (\boldsymbol{r}_t^{u(i)} - \boldsymbol{r}) + (\psi_t(\boldsymbol{r}) - \boldsymbol{d}_t)^T \hat{\boldsymbol{\Sigma}}_d^{-1} (\psi_t(\boldsymbol{r}) - \boldsymbol{d}_t)\},$$
$$i = 1, \dots, n_e.$$

Forwarding:

$\boldsymbol{r}_{t+1}^{u(i)} = \omega_t(\boldsymbol{r}_t^{c(i)}) + \boldsymbol{\epsilon}_t^r, i = 1, \dots, n_e$

$\boldsymbol{d}_{t+1}^{(i)} = \psi_{t+1}(\boldsymbol{r}_{t+1}^{u(i)}) + \boldsymbol{\epsilon}_t^d, i = 1, \dots, n_e$

$\boldsymbol{e}_{t+1} : \{(\boldsymbol{r}_{t+1}^{u(i)}, \boldsymbol{d}_{t+1}^{(i)}), i = 1, \dots, n_e\}$

---

4 Ensemble methods applied to groundwater flow

In hydrogeological applications, it is common practice to use physics-based numerical models to simulate and predict groundwater flow. These models describe groundwater flow by solving numerically the groundwater flow equation on a discretized domain. To produce reasonable simulations, the model parameters (e.g., the hydraulic conductivity) are usually estimated from observations of model states (e.g., hydraulic heads and soil water content) through model inversion.

In recent years, data assimilation and ensemble Kalman-based methods have become popular methods in hydrogeology, mainly because the explicit calculation of the gradient matrices needed for updating parameters is not needed. Instead, they are evaluated by ensemble averaging within a Monte Carlo framework.

In the current study, we investigate the performance of the ensemble methods described in previous sections when estimating spatially distributed aquifer parameters of a groundwater-flow model subject to pumping tests. In the following, we present the governing equation describing groundwater flow, and introduce the main characteristics of the research field site Lauswiesen, near Tübingen, Germany. The objective is to situate the performance of the SEnKF with respect to established methods. To that end, we then present a synthetic study based on the real field site in which the reference log-conductivity field displays a two-layer structure with a bimodal spatial histogram that is similar to the assumed characteristics of the study area. The design of the experiment is purposely favorable to the SEnKF in order to show-case its potential and motivate its use in real data applications in which a bimodal distribution might better represent real material properties. We finally present the results of the data assimilation experiments performed using real pumping tests implemented as the Lauswiesen field site. Earlier direct-push injection-logging (DPIL) measurements from the field site suggest, as previously mentioned, a two-layer structure in the log-conductivity field. It provides a good scenario to test the SEnKF, whose performances are evaluated in regards to five established methods.

Groundwater Flow Equation

We consider transient groundwater flow in a porous aquifer described by:

$$
\begin{aligned}
&S_o(\boldsymbol{x})\frac{\partial h(\boldsymbol{x},t)}{\partial t} - \nabla \cdot (K(\boldsymbol{x})\nabla h(\boldsymbol{x},t)) = W_0, \forall \boldsymbol{x}, \forall t \\
&h(\boldsymbol{x},t_0) = h_0, \forall \boldsymbol{x} \\
&h(\boldsymbol{x} \in \Gamma_{in}, t) = h_{in}, \forall t \\
&h(\boldsymbol{x} \in \Gamma_{out}, t) = h_{out}, \forall t \\
&-(K(\boldsymbol{x} \in \Gamma_{in})\nabla h(\boldsymbol{x} \in \Gamma_{in}, t)) \cdot \boldsymbol{n} = q_0, \forall t,
\end{aligned}
\tag{52}
$$

where $\boldsymbol{x}$ denotes the spatial coordinates, $S_o(\cdot)$ the specific storativity, $K(\cdot)$ is the spatially distributed hydraulic conductivity, $h(\cdot,\cdot)$ is the hydraulic head, t is time, and t_o the initial time. The parameter W_o represents volumetric sources or sinks (e.g., injection/extraction wells). The in- and outflow boundaries of the domain Γ_{in} and Γ_{out}, respectively, have Dirichlet boundary conditions. The rest of the domain's boundaries Γ_n have Neumann boundary conditions. The parameters h_{in}, h_{out} and h_o are the fixed hydraulic heads along Γ_{in} and Γ_{out} and initial hydraulic head, respectively. The vector $\boldsymbol{n}$ is the outer normal unit vector to the boundary. The specific discharge $q(\cdot,\cdot)$ follows Darcy's law $q(\boldsymbol{x},t) = -(K(\boldsymbol{x})\nabla h(\boldsymbol{x},t))$.
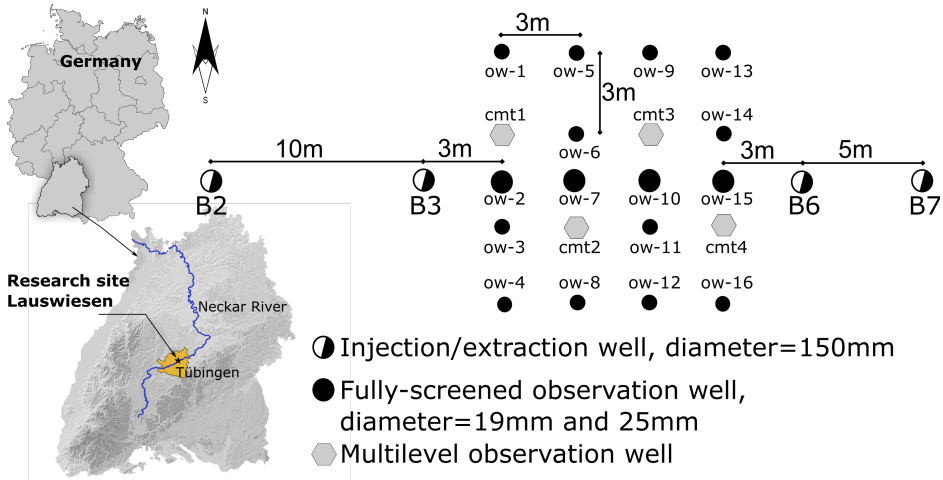
Figure 2: Hydrogeological research site Lauswiesen: Location and distribution of the injection/ extraction and observation wells in the well field used during the field experiments. From Sánchez-León et al. (2020)

Hydrogeological Research Site Lauswiesen and Experimental Design

The research site Lauswiesen is located in the Neckar Valley, Tübingen, Germany (Figure 2). The aquifer consists of 8m to 9m thick gravel with small amounts of sand (~ 10%) and fines (≤ 10%). It is overlain by about 2m of alluvial fines, and underlain by low-conductive claystones. The depth to groundwater is ~ 4m. Previous field investigations report a mean hydraulic conductivity value of the aquifer of $3 \times 10^{-3} \mathrm{m\,s}^{-1}$ (Doro et al., 2015; Sánchez-León et al., 2020), which is a typical value for deposits dominated by gravel-sheet sediments. In Figure 3, we show a schematic representation of the main subsurface units, the wells available at the field site, and the general setting of the pumping tests.

The wells available at the field site cover an area of about 30m × 10m and include four fully screened wells (B-wells), aligned with the main natural groundwater flow direction (Figure 2). Between wells B3 and B6, a total of 20 observation wells are distributed in a 5m × 4m regular grid. From the 20 observation wells, four of them (cmt-wells) have a multi-channel tube with six depth-discrete observation ports each, while the rest (ow-wells) are fully screened and vary between 1.9cm and 2.5cm in diameter (see also Figure 3).

To assess the performance of the methods presented in this work, we adopt the experiment presented in Sánchez-León et al. (2020), which considers the physical subsurface properties at the field site, as well as the available equipment. To demonstrate the applicability of the methods to real aquifers, we use a subset of the data collected during a field hydraulic tomography experiment (Sánchez-León et al., 2020). Hydraulic tomography is a method in which the aquifer is sequentially stressed at specific isolated sections, and the hydraulic responses are measured at many observation points. To minimize the propagation of the hydraulic pressure changes of the Neckar river into the investigated area, a nested-cell flow field with two injection and two extraction wells is proposed (Sánchez-León et al., 2020), producing a sta-

Figure 3: Pumping test design. Adapted from Sánchez-León et al. (2020)

ble local artificial ambient flow field with high extraction and injection rates applied at the outer wells, and lower injection/extraction rates at the inner wells. We selected data from three of six pumping tests, with water injection in well B3 at three different isolated depths and rates, and extracted from wells B6 and B7 over the entire aquifer thickness. The pumping rates values are summarized in Table 1.

Injection/extraction rates (ls^{-1})						
Test	B2	B3top	B3mid	B3bot	B6	B7
3a	5.2	1.8	0.9	0.6	-2.5	-9
3b	4.9	2.2	0.82	0.6	-2.0	-8.7
1b	4.8	2.0	0.87	0.6	-2.4	-9

Table 1: Pumping tests setup: top, mid and bot refer to the top, middle and bottom sections generated in the inner injection well B3. Positive flow rates represent injection of water, and negative flow rates refer to extraction rates

With an extensive DPIL campaign, two major aquifer layers are identified (Lessoff et al., 2010), a highly permeable and relatively homogeneous upper layer and a less conductive and more heterogeneous lower layer (Figure 4). DPIL measurements only provide relative values of hydraulic conductivity, and while it is possible to transform them to absolute values, e.g. by defining the dependence between relative and absolute hydraulic conductivities, the combined uncertainty from the measurements and the inversion method are a strong deterrent to using inverted measurements to fit a prior model. We

(a) 3D scatter plot   (b) Spatial histogram
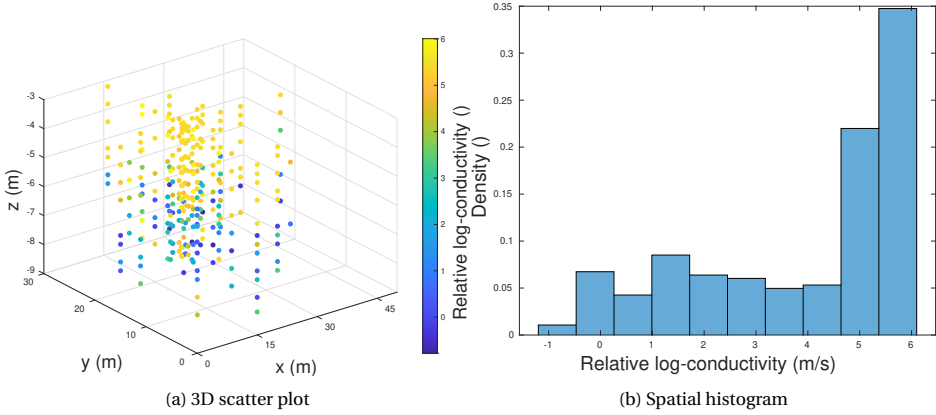
Figure 4: Relative conductivity measurements from the DPIL measurement campaign (Lessoff et al., 2010)

rather refer to the evaluation of the aquifer made in Sánchez-León et al. (2020) to specify a mean and correlation structure for the hydraulic conductivity fields. The DPIL measurements however, serve as a motivation for using a selection-Gaussian prior that describes the two-layer system.

4.1   Groundwater model and data assimilation setups

We consider a discretized spatio-temporal random field, $\{h(\boldsymbol{x}, t), \boldsymbol{x} \in \mathcal{L}_r \subset \mathbb{R}^3\}$ where $\boldsymbol{x}$ is the spatial reference on the regular spatial grid $\mathcal{L}_r$, $t \in \mathcal{L}_t : \{0, t_1, t_2, \ldots, T\}$ are discretized time points, and $h(\cdot, \cdot) \in \mathbb{R}$ represents the hydraulic head (m). We also consider two discretized spatial random fields, $\{\log(K(\boldsymbol{x})), \boldsymbol{x} \in \mathcal{L}_r \subset \mathbb{R}^3\}$ and $\{S_0(\boldsymbol{x}), \boldsymbol{x} \in \mathcal{L}_r \subset \mathbb{R}^3\}$; with $K(\cdot) \in \mathbb{R}_\oplus$ representing the hydraulic conductivity ($\mathrm{m\,s^{-1}}$), and $S_0(\cdot) \in \mathbb{R}_\oplus$ representing aquifer storativity ($\mathrm{m^{-1}}$).

To simulate three dimensional groundwater flow, we use a finite volume numerical model implemented in Matlab. Figure 5 shows the model domain in yellow, the area of interest where the injection, extraction and observation wells are located is shown in blue. The grid covers the whole domain rather than only the area of interest such that pumping effects at the boundaries are minimized. While boundary conditions are assigned at the outer border of the model domain, parameters are updated only for the area of interest. The total number of spatial grid nodes is $n = 14007$, with finer grid elements in the area of interest than outside, reducing the total number of elements in the model and hence, model run times. The area of interest contains 6256 grid nodes. The pumping test times are discretized in 27 steps, from 0 to 7,000s when it is assumed that steady state is achieved. The discretized hydraulic head field at time t may be represented by the n-vector $\boldsymbol{h}_t$ and the log-conductivity and storativity fields by the n-vectors $\log(\boldsymbol{K})$ and $\boldsymbol{S}_0$, respectively. The storativity is assumed to be constant on the whole field such that $\boldsymbol{S}_0 = S_0 \boldsymbol{i}_n$, where S_0 is considered unknown. The state-space vector which contains the quantities of interest, is therefore composed of $\log(\boldsymbol{K})$ and S_0. The data consists of hydraulic head measurements collected at the active observation wells during water extraction. The forward model describes the evolution of the hydraulic head field in time for a given set of parameters (state-space vector). The likelihood
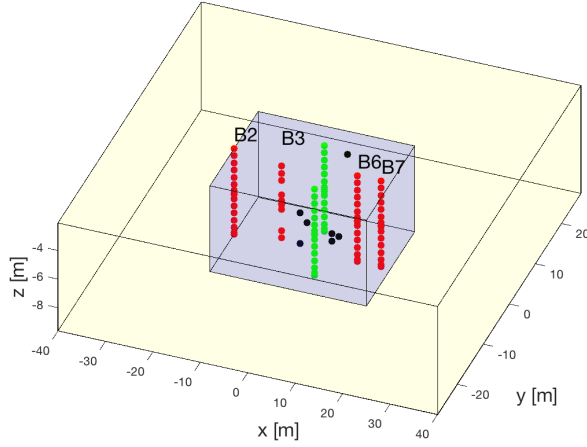
Figure 5: Three dimensional model used for the groundwater flow simulations. The whole model domain is shown in yellow, and the area of interest in blue. The injection and extraction wells are in red (B2, B3, B6, B7), the fully screened measurement wells are in green, and the active multi-layer wells in black. The observation wells shown correspond to those active during pumping test 1b

model selects hydraulic head values only at locations where it is measured. The forward model is obtained by discretizing the groundwater flow equation (Equation 52) with a three-dimensional numerical model that solves either transient flow or directly steady-state by the finite volume method on rectangles, and is defined as:

$$\boldsymbol{h}_{t+1} = \omega_t(\log(\boldsymbol{K}), S_0, \boldsymbol{h}_t; h_{in}, h_{out}, q_0, W_0) \tag{53}$$

Constant-head values are defined at the left ($h_{in} = 0.01$m) and right ($h_{out} = 0$m) boundaries of the model domain, and no-flux boundaries (q_0) at the front, back, bottom, and top faces. The initial head field is defined by $\boldsymbol{h}_0 = h_0 \boldsymbol{i}_n$, such that $h_0 = 0$ m. The injection and extraction rates are defined in Table 1. The likelihood model is defined as follow,

$$\boldsymbol{d}_t = \boldsymbol{H}\boldsymbol{h}_t + \boldsymbol{\epsilon}_t^d \tag{54}$$

The matrix $\boldsymbol{H}$ is the observation ($m \times n$)-matrix. The wells $w_s, s = 1,\dots,m$ contain $l_s, s = 1,\dots,m$ grid cells respectively. The observation matrix is therefore defined as,

$$[H]_{i,j} = \frac{1}{l_i}, j \in w_i, i = 1,\dots,m. \tag{55}$$

Measurements from the cmt-wells are constrained to a specific point in depth, leading to a single non-zero entry for the corresponding rows in the observation matrix $\boldsymbol{H}$. For the ow-wells, which are fully screened along the entire aquifer thickness, non-zero entries in $\boldsymbol{H}$ cover a number of grid elements denoted l. The simulated drawdown at these wells is averaged over the number of grid-cells representing the well. The corresponding rows in the observation matrix have therefore l non-zero entries equal to

Figure 6: Reference log-conductivity field

$\frac{1}{l}$. The random m-vector $\boldsymbol{\epsilon}_t^d$ is centered Gaussian with covariance $\boldsymbol{\Sigma}_d^t = \sigma_d^{t\,2} \boldsymbol{I}_m$. The measurement error level σ_d^t will be specified later. We consider drawdowns rather than the absolute hydraulic heads as observations, where drawdown is defined as the depletion in piezometric head due to groundwater injection and/or extraction. All the data assimilation experiments will be conducted with $n_e = 500$ ensemble members. Note that only the elements of the state space vector $[\log(\boldsymbol{K}), S_0]$ are of interest in this study.

4.2 Synthetic Case Study

For the synthetic pumping test we generate a reference log-conductivity field $\log(\boldsymbol{K})^{ref}$ from a selection-Gaussian distribution with parameters $\boldsymbol{\Theta}_{ref}^{SG} = (\mu_{\tilde{F}}^{ref}, \mu_v^{ref}, \sigma_{\tilde{F}}^{ref}, \boldsymbol{\Sigma}_{\tilde{F}}^{\rho}, \gamma^{ref}, A^{ref})$. The parameters values are listed in Table 5. The spatial structure of the generated log-hydraulic conductivity field is in agreement with the two-layer system of hydraulic conductivity distributions recognized at the field site (see Figure 6). We assume an effective aquifer storativity S_0 of $8.6 \times 10^{-3} \mathrm{m}^{-1}$ which is typical for confined systems.

For the SEnKF, see Section 3.3.2.3, the prior ensemble for the log-conductivity field is generated from a selection-Gaussian random field $\log(\boldsymbol{K})$ with parameters $\boldsymbol{\Theta}_K^{SG} = (\mu_{\tilde{F}}^k, \mu_v^k, \sigma_{\tilde{F}}^k, \boldsymbol{\Sigma}_{\tilde{F}}^{\rho}, \gamma^k, A^k)$. The parameters (γ^k, A^k) are chosen such that the prior marginal distributions be bimodal, with modes close to those of the spatial histogram of the reference log conductivity field $log(\boldsymbol{K})^{ref}$ (Figure 7). For the EnKF, see Section 3.1, the prior ensemble for the log-conductivity field is generated from a Gaussian random field $\log(\boldsymbol{K})$ defined by $\varphi_n(\log(\boldsymbol{K}); \mu_r \boldsymbol{i}_n, \sigma_r^2 \boldsymbol{\Sigma}_r^{\rho})$ where $\mu_r = -7$ and $\sigma_r = \sqrt{1.6}$ are chosen such that the marginal prior distribution covers the spatial histogram of the reference log-conductivity field (Figure 7). The spatial correlation $(n \times n)$-matrix $\boldsymbol{\Sigma}_r^{\rho}$ is defined by the second order exponential spatial correla-

Table 2: Parameter values for reference and prior model

Ref	$\mu_{\bar{F}}^{ref}()$	$\mu_v^{ref}()$	$\sigma_{\bar{F}}^{ref}()$	$\gamma^{ref}()$	$A^{ref}()$
	-7	0	0.6	0.99	$[(-\infty, -0.1] \cup [0.1, +\infty)]^n$
Prior SEnKF	$\mu_{\bar{F}}^k()$	$\mu_v^k()$	$\sigma_{\bar{F}}^k()$	$\gamma^k()$	$A^k()$
	-7	0	1.6	0.95	$[[-3, -0.2] \cup [0.3, 3.3]]^n$

tion function $\rho(\tau; d_x, d_y, d_z) = \exp[-(\tau_x^2/d_x^2 + \tau_y^2/d_y^2 + \tau_z^2/d_z^2)]$. This spatial correlation function is used to generate the reference field and for both prior models with parameters values $d_x = 8$m, $d_y = 4$m and $d_z = 1$m. While the prior for the SEnKF has a larger marginal variance, it has a comparable spread
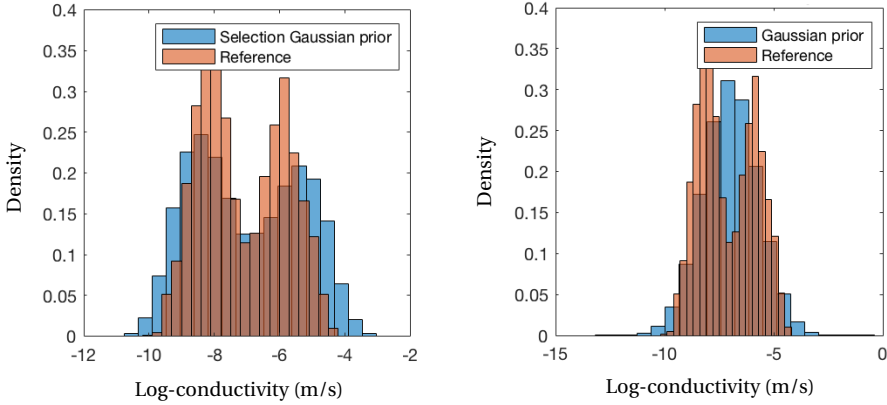
Figure 7: Prior marginal distribution for the SEnKF (left) and the EnKF (right) compared to the spatial histogram of the reference log-conductivity field

to the prior for the EnKF (Figure 7). The prior distribution for the effective storativity S_0 is defined by $\varphi_1(S_0; 8.6 \times 10^{-3}, 1 \times 10^{-6})$. During the data assimilation, covariance localization, see Section 3.3.1.1, is used in both methods.

4.2.1   Results

Figures 8 and 9 show the predicted drawdown recorded at the nine observation wells defined in the model for the SEnKF and EnKF, respectively. The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval. The red lines are the true unperturbed synthetic drawdown, and the red crosses represent drawdown perturbed by observation error with standard deviation $\sigma_d^t = 0.05$. The value of σ_d^t is representative of the expected accuracy of modern sensors for hydraulic head.
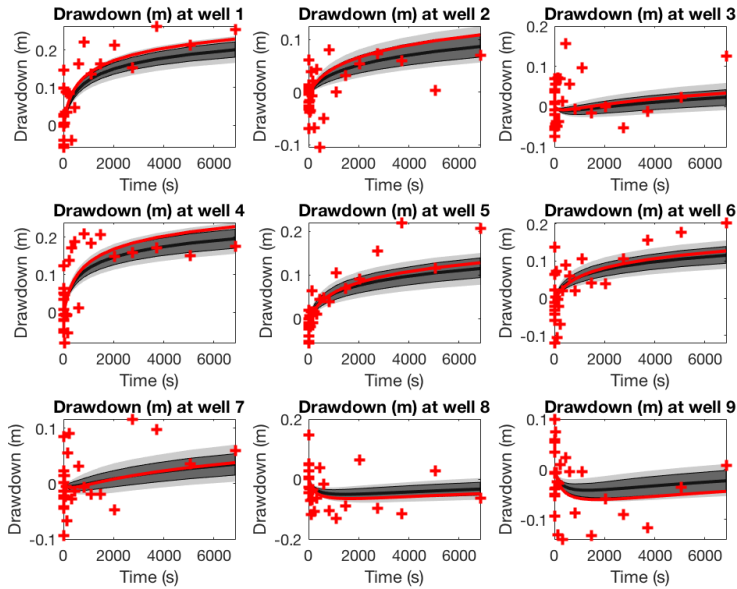
Figure 8: Predicted drawdown from the SEnKF at the observation wells for pumping test 1b compared to the true (red line) and measured drawdowns (red crosses). The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval.

The SEnKF and the EnKF appear to predict convincingly the true unperturbed drawdown. The spread of the predicted drawdown is larger for the SEnKF than for the EnKF, and it seems more commensurate with the observed drawdown for the former than the latter. This suggests that the EnKF might be underestimating the uncertainty despite the good predictions of the true unperturbed drawdowns.

Figure 10 displays the 3D reference log-conductivity field together with the 3D predicted log-conductivity field from the SEnKF and the EnKF. Recall that the predicted spatial variables provided by the SEnKF and the EnKF will be smoother than the truth. The SEnKF prediction represents to some extent the two-layer structure observed in the reference field while the EnKF prediction exhibits the aforementioned regression towards the mean.

Figure 11 shows the predicted log-conductivity along well 2 for the SEnKF and the EnKF. The SEnKF appears to better predict the log-conductivity along this well as it is able to identify the aforementioned layered structure of the reference field, at least at this measurement well, while the EnKF does not. These observations are consistent with the analysis of Figure 10. Recall again that spatial predictions are smoother than the truth.

Figure 12 shows the coverage probability of the reference log-conductivity field for the SEnKF and the EnKF prediction intervals. Visual inspection of the curves indicates that the SEnKF slightly overestimates the prediction intervals while the EnKF clearly underestimates them. Conservatively overestimating the prediction intervals appear preferable to the opposite.
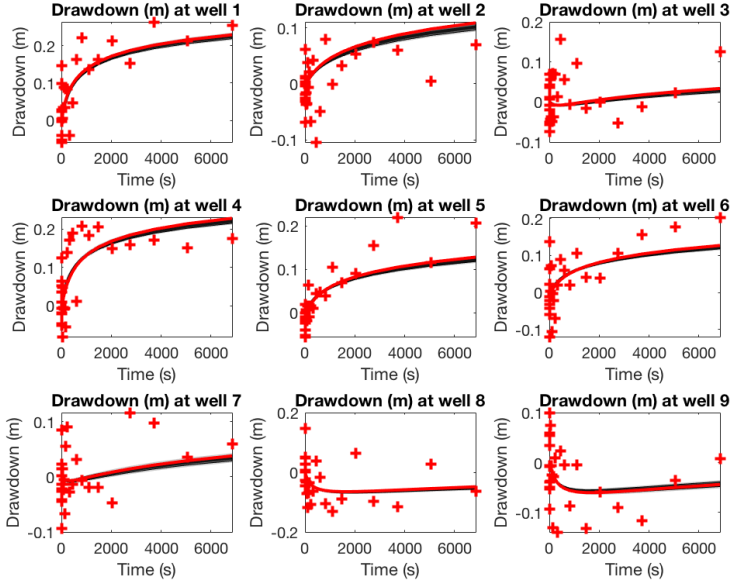
Figure 9: Predicted drawdown from the EnKF at the observation wells for pumping test 1b compared to the true (red line) and measured drawdowns (red crosses). The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval.

Table 3: RMSE and SSIM and coverage comparing the predictions of the different methods to the true log diffusivity field

		SEnKF	EnKF
Initial	MAE	1.26	1.15
Posterior	MAE	1.05	1.10
	$SSIM$	0.0704	0.0464

To further quantify the performance of the data assimilation experiments, we estimate the mean absolute error (MAE) between the predicted log-conductivity fields and the reference field. We want the MAE to be as small as possible. We also compare the structural similarity index (SSIM) (e.g. Zhou et al. (2004)), a metric used in image analysis to measure the degradation of the structural information of an image. We want the SSIM to be as large as possible. The results are detailed in Table 3, both metrics are favorable to the SEnKF. In particular the SEnKF offers a 17% relative reduction in MAE while the EnKF only gives 4%.

Figure 10: Reference log-conductivity field (left), mean log-conductivity prediction from the SEnKF (center) and mean log-conductivity prediction from the EnKF (right)

4.2.2   Closing remarks

The synthetic case study demonstrate that the SEnKF provides more accurate predictions of the reference log-conductivity field than the EnKF. Moreover, the prediction intervals for the drawdowns provided by the SEnKF appear as more realistic than the ones provided by the EnKF. To further test the performance of the SEnKF for real applications, where information about the true aquifer properties is limited, we apply the presented methodology to data collected from real pumping tests.

4.3   Real Case Study

For the study with real data, we consider drawdown data from three different pumping tests performed at the Lauswiesen field site. These tests are part of a larger experiment implemented at the field site as described in Sánchez-León et al. (2020). The pumping tests are labeled as pumping tests 3b, 3a and 1b. Slightly different injection/extraction rates are applied for each test (see Table 1). Additionally, different observation wells are monitored during each test. The collected dataset is shown in Figure 13. When assimilating each pumping test separately, data match is easily achieved for pumping tests 3b and 3a for all methods (not shown). This is however not the case for pumping test 1b. We therefore focus on assimilating pumping test 1b and evaluate whether the different methods detailed in Section 3 influence the results. Table 4 presents the six methods that are considered in this study.

Table 4: Methods overview with different features (✓: included, ✗: excluded)

	Prior	Localization	Inflation	MDA
EnKF	Gaussian	✓	✗	✗
EnKF/Infl	Gaussian	✓	✓	✗
GA EnKF	Gaussian	✓	✗	✗
IEnKF	Gaussian	✓	✗	✗
SEnKF	Selection-Gaussian	✓	✗	✓
MDA EnKF	Gaussian	✓	✗	✓

For the SEnKF, the prior ensemble for the log-conductivity field $\log(\boldsymbol{K})$ is generated from a selection-

Figure 11: Predicted log($\mathbf{K}$) along observation well 2 in pumping test 1b using the SEnKF (left) and the EnKF (right). The prediction is in blue, the true log conductivity in red and the area in gray denotes the 80% prediction interval

Gaussian random field, the parameter values are listed in Table 5. Figure 14 shows the prior marginal distribution overlain with relative log-conductivity measurements collected from available DPIL measurements. The prior distribution roughly reproduces the shape of the relative log($\mathbf{K}$) values of the DPIL measurements. The parameters (γ, A) are chosen so that the prior marginal is bimodal, with modes allowing for high conductivity pockets as suggested by the DPIL measurements.

For the other five methods, the prior ensemble of the log-conductivity field is generated from a Gaussian random field log($\mathbf{K}$) defined by $\varphi_n(\log(\mathbf{K}); \mu_r \mathbf{i}_n, \sigma_r^2 \mathbf{\Sigma}_r^\rho)$ where $\mu_r = -6$ and $\sigma_r = \sqrt{1.6}$.

For both prior models, the spatial correlation $(n \times n)$-matrix $\mathbf{\Sigma}^\rho$ is defined by the second order exponential spatial correlation function $\rho(\tau; d_x, d_y, d_z) = \exp\left[-(\tau_x^2/d_x^2 + \tau_y^2/d_y^2 + \tau_z^2/d_z^2)\right]$. The parameters values are listed in Table 5. The parameter values for both prior models are chosen in accordance with Sánchez-León et al. (2020), and adjusted to ensure that the drawdowns simulated with the initial ensemble covers the measured drawdowns for all three pumping tests. The prior distribution for the effective

Table 5: Parameter values for the selection-Gaussian prior

$\mu_{\bar{r}}()$	$\mu_v()$	$\sigma_{\bar{r}}^2()$	$\gamma()$	A()	$d_x(m)$	$d_y(m)$	$d_z(m)$
-6	0	1.6	0.99	$[(-\infty, 0] \cup [0.2, +\infty)]^n$	8	4	1

Figure 12: Reference conductivity field coverage for the SEnKF and the EnKF. Prediction α intervals (horizontal), coverage of the reference conductivity field in estimated $1 - \alpha$ prediction intervals (vertical)

storativity S_0 is defined by $\varphi_1(S_0; 8.6 \times 10^{-3}, 1 \times 10^{-6})$. During the data assimilation, covariance localization, see Section 3.3.1.1, is used in all methods. When inflation is used, see Section 3.2.1.4, an inflation factor $\alpha = \sqrt{1.02}$ is chosen. When MDA is used, see Section 3.3.3.1, eight substeps are considered. The IEnKF, see Section 3.3.3.2, uses Levenberg-Marquardt iterations and considers sensitivity matrices calculated directly from the ensemble. The method closely follows the procedure presented in Luo et al. (2015), and our implementation is based on the code made available by the authors. The standard deviation of the observation error is set to $\sigma_d^t = 0.05$ in line with the expected accuracy of modern sensors for hydraulic head.

4.3.1 Results

The methods are compared considering the following:

1. **Data match:** Predicted drawdowns at the observation wells of pumping test 1b conditioned on the drawdown measurements collected during pumping test 1b.

2. **Validation I:** Predicted drawdowns at the observation wells of pumping test 3a conditioned on the drawdown measurements collected during pumping test 1b.

3. **Validation II:** Predicted drawdowns at the observation wells of an additional pumping test with different injection, extraction and observation wells conditioned on the drawdown measurements collected during pumping test 1b.

4. **Validation III:** Predicted log-conductivity along well B6 and B7 conditioned on the drawdown measurements collected during pumping test 1b. The results are compared to flowmeter measurements along these two wells.

(a) Pumping test 3b

(b) Pumping test 3a
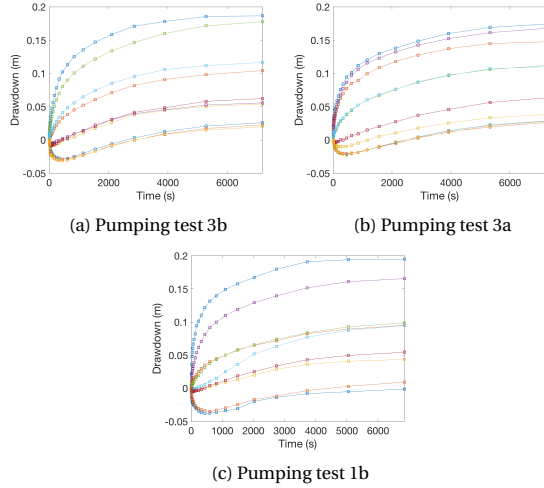
(c) Pumping test 1b

Figure 13: Drawdown measurements $\boldsymbol{d}_{3b}$, $\boldsymbol{d}_{3a}$ and $\boldsymbol{d}_{1b}$ from the three pumping tests 3b, 3a and 1b respectively

We first investigate the Data match results. Figure 15 displays the predicted drawdowns from the EnKF at the observation wells for pumping test 1b. The data match is poor and the ensemble spread exceedingly narrow, especially at wells where the data match is the worst. This could be indicative of ensemble collapse, or of a non-linearity that is too important for the EnKF to accommodate. However, the log-conductivity fields shows no sign of collapse and conserve an acceptable amount of spread. The non-linearity can either have its origins in the forward model, or from missing information in the prior distribution of parameters. Considering that data match is easily achieved in the synthetic study, it would stand to reason that the source of the non-linearity lies in the complexity of the subsurface that is not modelled by the prior distribution.

Upon inspection of Figures 16-19 containing predicted drawdowns at the observation wells for pumping test 1b from the five other methods under consideration, it is determined that two methods display good data match, the IEnKF and the SEnKF, while one, the MDA EnKF, gives an acceptable data match. For the others two methods, data match is poor. Note that the SEnKF uses MDA which has a substantial and positive effect on the results.

In the rest of the study, only the three methods with the best data match (IEnKF, SEnKF and MDA EnKF) are further investigated. To validate the updated ensemble of parameters, we test their performance with two additional datasets that are not used during data assimilation.

The first dataset is collected during pumping test 3a (Validation I). Figures 21, 22, 23 shows the predicted drawdowns at the measurement wells for pumping test 3a for the IEnKF, SEnKF and MDA EnKF respectively. The best prediction is given by the MDA EnKF, closely followed by the SEnKF. The IEnKF predictions fail to give any coverage to 3 of the 10 wells.

The second dataset (Validation II) is obtained in a different field campaign, where the settings of

Figure 14: Prior selection-Gaussian distribution overlain with relative log-conductivity measurements collected from the DPIL measurements centered about the assumed mean log-conductivity of the field.

the pumping tests are completely different to those of the tests presented in this work. Validation using data from tests with different settings, and therefore completely different groundwater flow regimes, is a strong assessment of the robustness of the data assimilation experiments. In this additional pumping test, water is extracted at a rate of 3.54 L/s, from an isolated section at a depth of 5.27 m in well B6 . No additional injection of water is implemented. The aquifer responses to pumping are monitored at different depths using the multilevel cmt-wells (see Figure 3). Figures 24, 25, 26 shows the predicted drawdowns at the measurement wells for the additional pumping test. The data measured at well 2 exhibit an odd behaviour, we do not consider it when analyzing the results. The SEnKF and the MDA EnKF give good predictions while the IEnKF fails to give predictions that cover the measured data. These observations are consistent with the results from pumping test 3a.

As a final validation and to supplement the pumping test validations, we consider the results from Validation III. We compare the predicted log-conductivity field to flowmeter data collected along wells B6 and B7. Figure 27 displays the predicted log-conductivity for the IEnKF, SEnKF and MDA EnKF compared to the log-conductivity measured using a flowmeter. All three methods provide an acceptable coverage to the flowmeter data.

4.3.2   Closing remarks

The real case study demonstrates that even though the IEnKF gives the best data match for 1b and provide an acceptable coverage to the flowmeter data, the prediction of the drawdowns for both pumping tests used as validation are underwhelming. The data match for 1b is comparatively less convincing for the MDA EnKF, but the predictions for both pumping tests are satisfactory, and provides a suitable coverage to the flowmeter ata. The SEnKF never provides the best of out three, but is consistently reliable
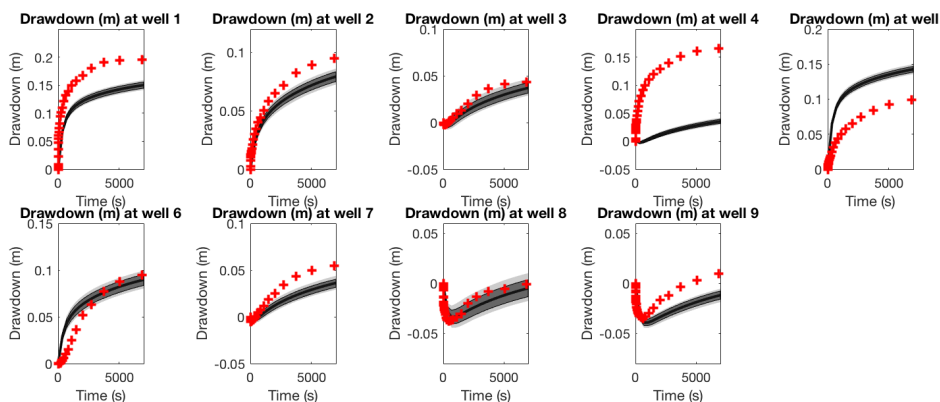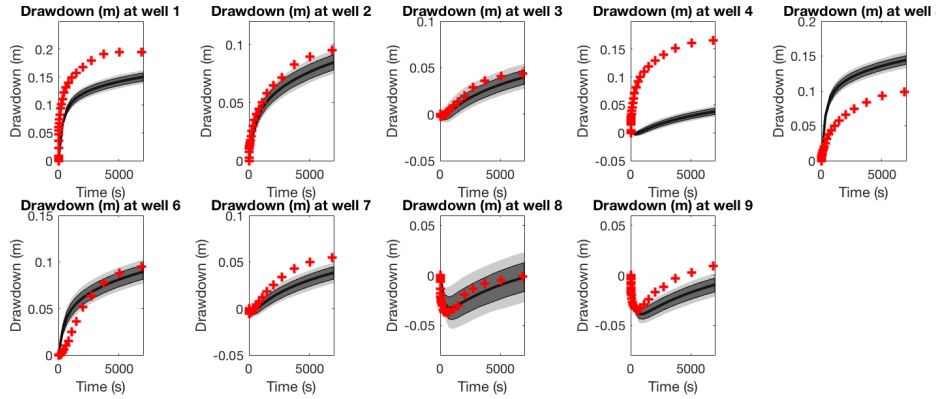
Figure 15: Data Match, EnKF: Predicted drawdown at observations wells for pumping test $1b$ given $\boldsymbol{d}_{1b}$ compared to the measured drawdowns (red line). The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval

throughout.

5   Conclusion

This study presented a review of existing ensemble Kalman filtering methods. The objective is to situate the newly developed SEnKF in the existing literature. The synthetic case study is consistent with previous studies in that it confirms the aptitude of the SEnKF when the quantity of interest displays a bimodal spatial histogram. The real case study shows that in a situation where there is legitimate reasons to suspect the log-conductivity field exhibits spatial bimodality, the SEnKF provides robust if not compelling results. However, considering the added computational burden, the MDA EnKF might be the more cost efficient approach. It would be interesting to see if the comparison between these methods yield different results when conditioning on tracer data. It would however require rewriting the MCMC sampler in the SEnKF to allow for larger grid size because of the resolution needed for tracer test simulation. It would also be interesting to investigate what the comparison would yield if the three pumping tests are assimilated sequentially. Both these matters are left for further research.

Figure 16: Data Match, EnKF with inflation: Predicted drawdown at observations wells for pumping test $1b$ given $\boldsymbol{d}_{1b}$ compared to the measured drawdowns (red line).The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval
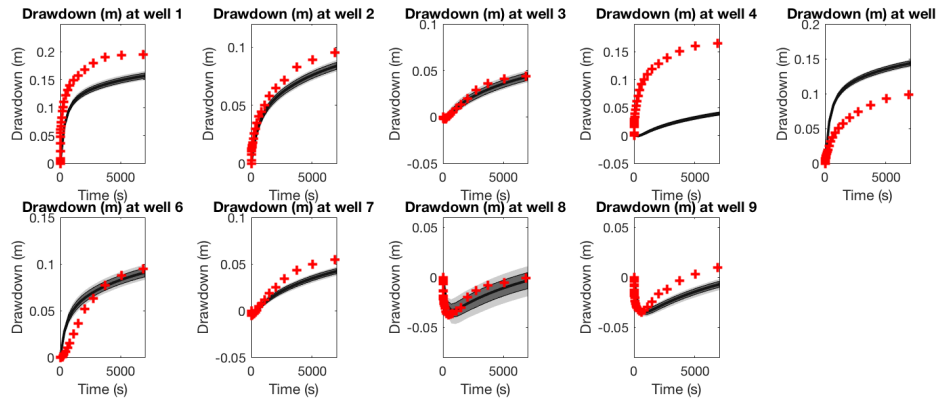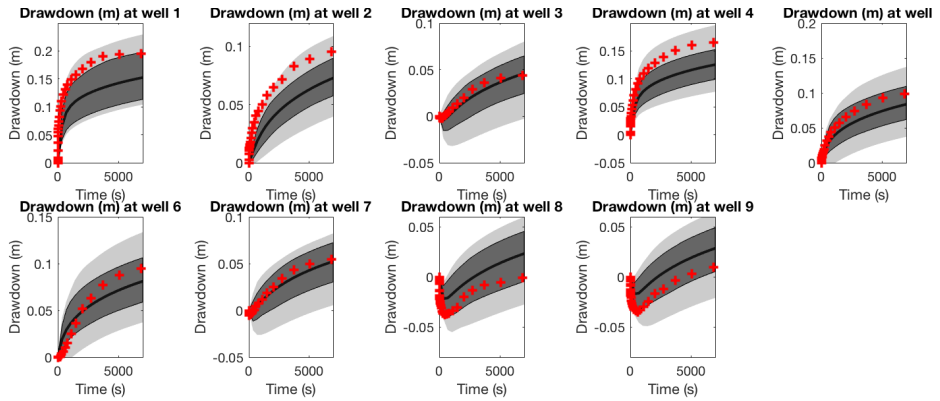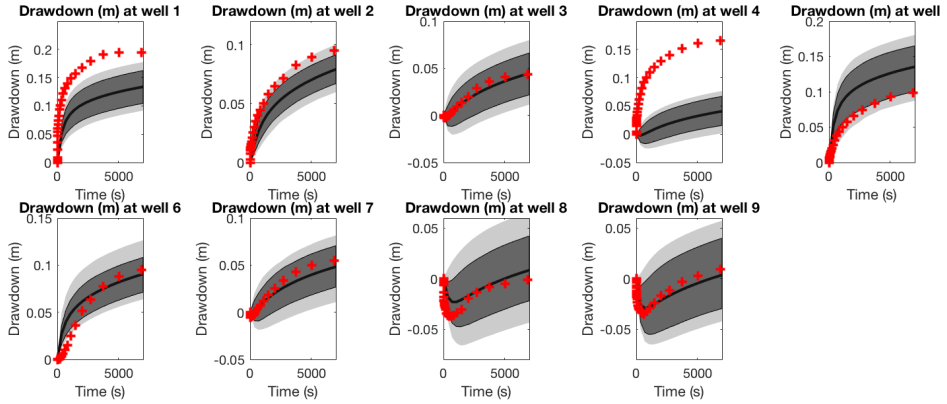

Figure 17: Data Match, GA EnKF: Predicted drawdown at observations wells for pumping test $1b$ given $\boldsymbol{d}_{1b}$ compared to the measured drawdowns (red line). The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval
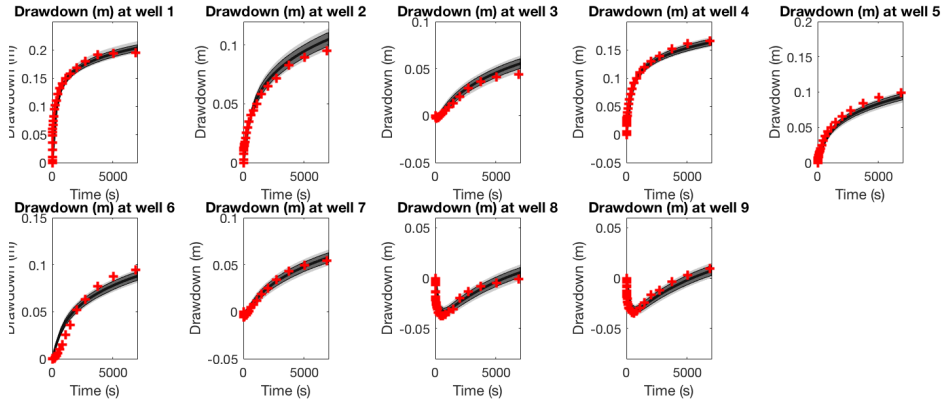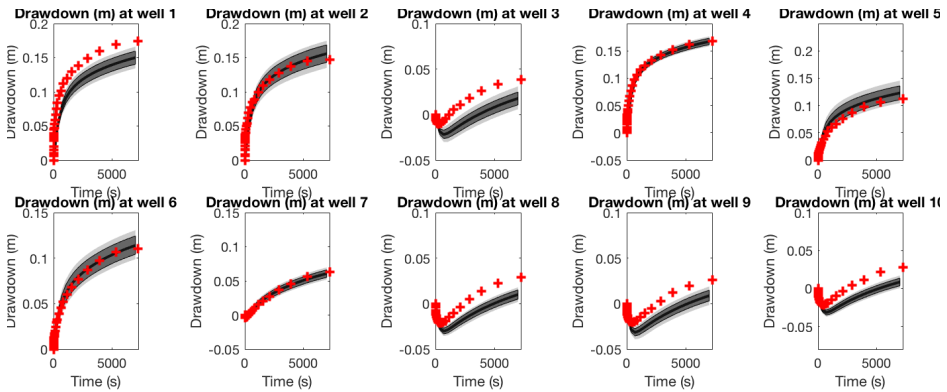
Figure 18: Data Match, IEnKF: Predicted drawdown at observations wells for pumping test $1b$ given $\boldsymbol{d}_{1b}$ compared to the measured drawdowns (red line). The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval

Figure 19: Data Match, SEnKF: Predicted drawdown at observations wells for pumping test $1b$ given $\boldsymbol{d}_{1b}$ compared to the measured drawdowns (red line). The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval

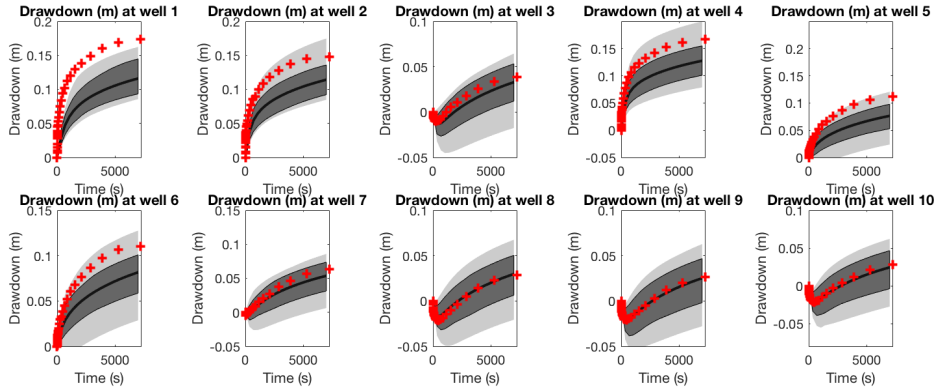Figure 20: Data Match, MDA EnKF: Predicted drawdown at observations wells for pumping test $1b$ given $\boldsymbol{d}_{1b}$ compared to the measured drawdowns (red line). The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval
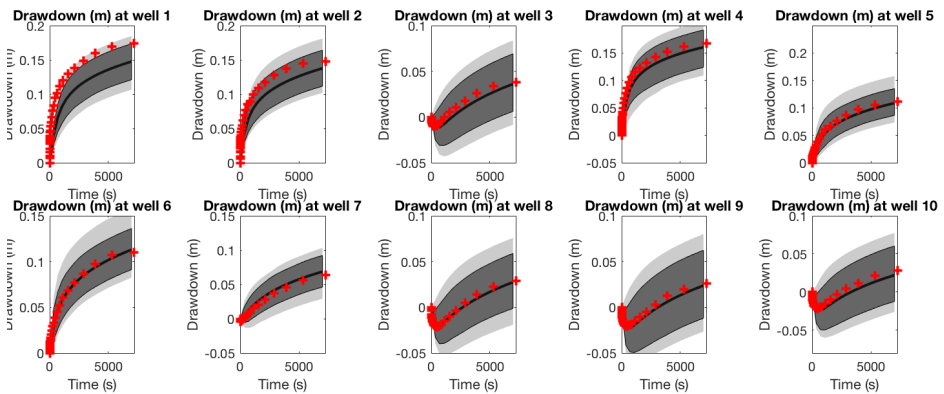

Figure 21: Validation I, IEnKF: Predicted head at observations wells for test $3a$ given $\boldsymbol{d}_{1b}$. The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval

Figure 22: Validation I, SEnKF: Predicted drawdown at observations wells for test $3a$ given $\boldsymbol{d}_{1b}$ compared to the measured drawdowns (red line). The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval

Figure 23: Validation I, MDA EnKF: Predicted head at observations wells for test $3a$ given $\boldsymbol{d}_{1b}$. The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval
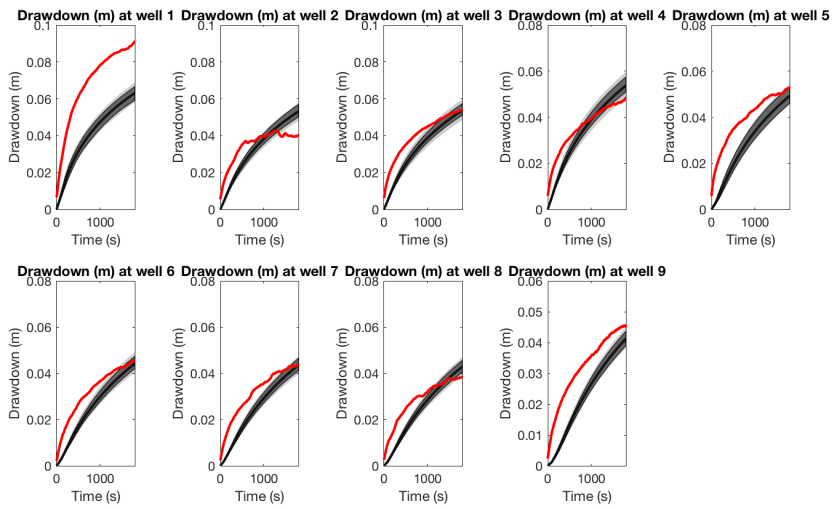
Figure 24: Validation II, IEnKF: Predicted drawdown at observations wells for the additional test given $\boldsymbol{d}_{1b}$ compared to the measured drawdowns (red line) . The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval
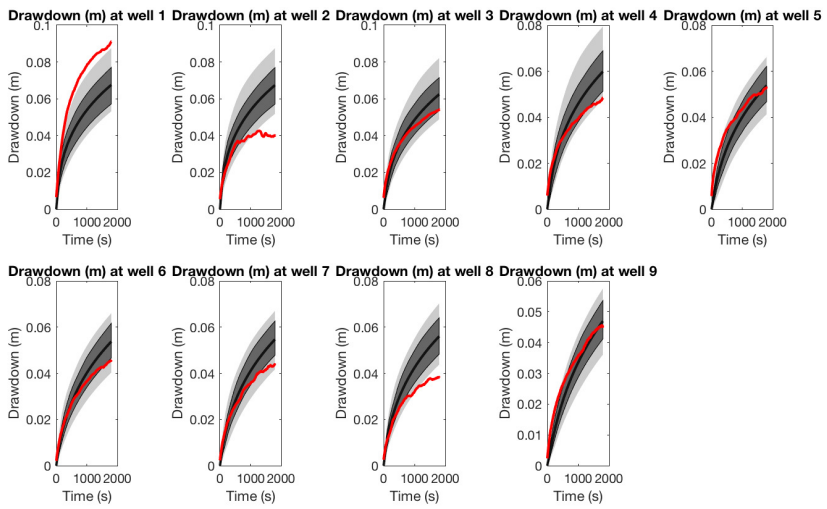
Figure 25: Validation II, SEnKF: Predicted drawdown at observations wells for the additional test given $\boldsymbol{d}_{1b}$ compared to the measured drawdowns (red line). The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval
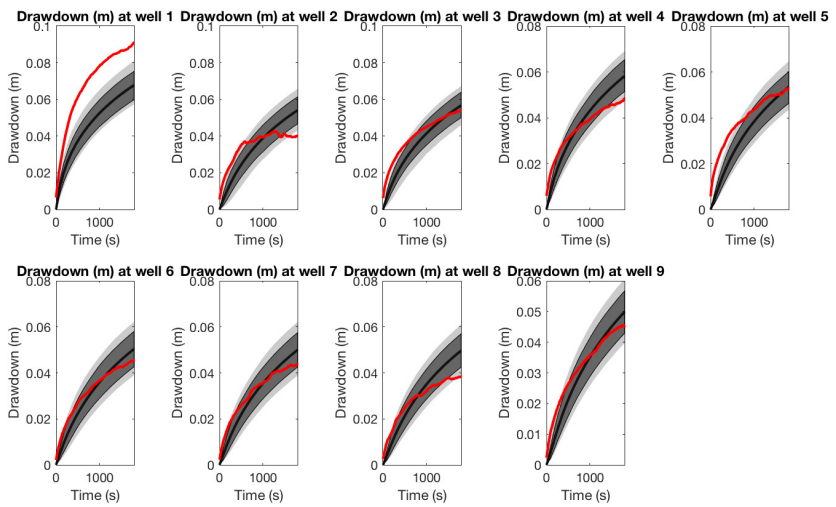
Figure 26: Validation II, MDA EnKF: Predicted drawdown at observations wells for the additional test given $\boldsymbol{d}_{1b}$ compared to the measured drawdowns (red line). The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval
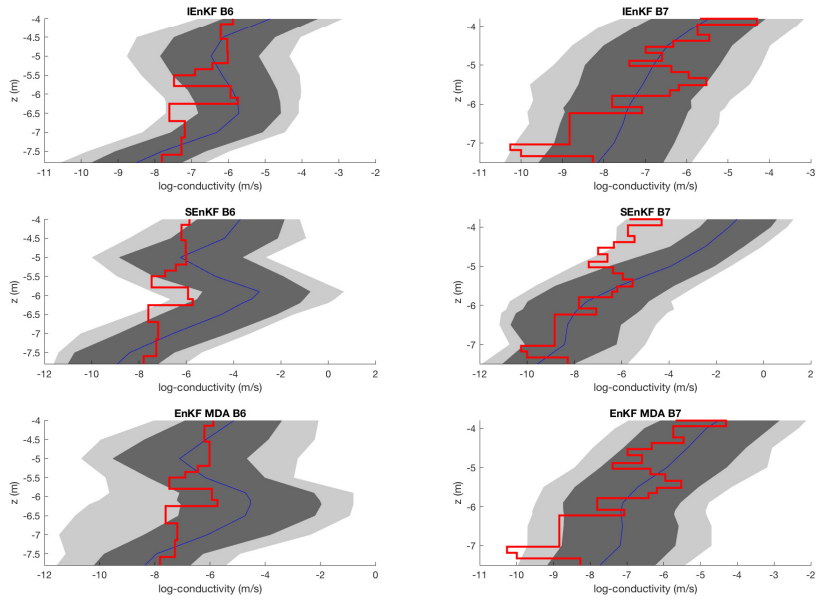
Figure 27: Validation III: predicted log(K) (blue) along wells B6 and B7 using the IEnKF (top), SEnKF (mid) and MDA EnKF (bot) compared to flowmeter measurements(red). The dark gray area represents the 70% prediction interval, the light gray area represents the 90% prediction interval

Acknowledgements

References

Aanonsen, S., Nævdal, G., Oliver, D., Reynolds, A., and Vallès, B. (2009). The ensemble Kalman filter in reservoir engineering–a review. *SPE Journal*, 14:393–412.

Ackerson, G. and Fu, K. (1970). On state estimation in switching environments. *IEEE Transactions on Automatic Control*, 15(1):10–17.

Anderson, J. (2007). An adaptive covariance inflation error correction algorithm for ensemble filters. *Tellus A: Dynamic Meteorology and Oceanography*, 59:210 – 224.

Anderson, J. and Anderson, S. L. (1999). A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review*, 127:2741–2758.

Arellano-Valle, R. and del Pino, G. (2004). *Skew-Elliptical Distributions and Their Applications: A Journey Beyond Normality*, chapter From symmetric to asymmetric distributions: A unified approach, pages 113–133. Chapman & Hall/CRC.

Arellano-Valle, R. B., Branco, M. D., and Genton, M. G. (2006). A unified view on skewed distributions arising from selections. *Canadian Journal of Statistics*, 34(4):581–601.

Bengtsson, T., Snyder, C., and Nychka, D. (2003). Toward a nonlinear ensemble filter for high-dimensional systems. *Journal of Geophysical Research: Atmospheres*, 108(D24).

Bertino, L., Evensen, G., and Wackernagel, H. (2003). Sequential data assimilation techniques in oceanography. *International Statistical Review*, 71(2):223–241.

Bocquet, M. and Carrassi, A. (2017). Four-dimensional ensemble variational data assimilation and the unstable subspace. *Tellus A: Dynamic Meteorology and Oceanography*, 69(1):1304504.

Cappé, O., Moulines, E., and Ryden, T. (2005). *Inference in Hidden Markov Models (Springer Series in Statistics)*. Springer-Verlag, Berlin, Heidelberg.

Chen, R. and Liu, J. S. (2000). Mixture Kalman filters. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(3):493–508.

Conjard, M. and Omre, H. (2020). Data assimilation in spatio-temporal models with non-Gaussian initial states—the selection ensemble Kalman model. *Applied Sciences*, 10(17):5742.

Conjard, M. and Omre, H. (2021). Spatio-temporal inversion using the selection Kalman model. *Frontiers in Applied Mathematics and Statistics*, To appear.

Doro, K., Cirpka, O., and Leven, C. (2015). Tracer tomography: design concepts and field experiments using heat as a tracer. *Ground water*, 53:139–48.

Dovera, L. and Della Rossa, E. (2010). Multimodal ensemble Kalman filtering using Gaussian mixture models. *Computational Geosciences*, 15(2):307–323.

Emerick, A. A. and Reynolds, A. C. (2013). Ensemble smoother with multiple data assimilation. *Computers & Geosciences*, 55:3–15.

Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99(C5):10143– 10162.

Forberg, O., Kjøsnes, Ø., and Omre, H. (2021). Bayesian seismic avo inversion for reservoir variables with bimodal spatial histograms. *Geophysics*, 0:1–141.

Forberg, O. B., Grana, D., and Omre, H. (2021). Bayesian inversion of time-lapse seismic avo data for multimodal reservoir properties. *IEEE Transactions on Geoscience and Remote Sensing*, pages 1–16.

Gaspari, G. and Cohn, S. (1999). Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, 125:723–757.

Gordon, N. J., Salmond, D. J., and Smith, A. F. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)*, 140(2):107–113.

Grana, D., Fjeldstad, T., and Omre, H. (2017). Bayesian Gaussian mixture linear inversion for geophysical inverse problems. *Mathematical Geosciences*, 49:493–515.

Hamill, T., Whitaker, J., and Snyder, C. (2001). Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Monthly Weather Review*, 129:2776–2790.

Hendricks Franssen, H.-J. and Kinzelbach, W. (2008). Real-time groundwater flow modeling with the ensemble Kalman filter: Joint estimation of states and parameters and the filter inbreeding problem. *Water Resources Research*, 44(9).

Houtekamer, P. L., Mitchell, H. L., Pellerin, G., Buehner, M., Charron, M., Spacek, L., and Hansen, B. (2005). Atmospheric data assimilation with an ensemble Kalman filter: Results with real observations. *Monthly Weather Review*, 133(3):604–620.

Huber, P. (1981). *Robust Statistics*. John Wiley & Sons, Ltd.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Trans. ASME-J. Basic Eng.*, 82(1):35–45.

Katzfuss, M., Stroud, J. R., and Wikle, C. K. (2020). Ensemble Kalman methods for high-dimensional hierarchical dynamic space-time models. *Journal of the American Statistical Association*, 115(530):866–885.

Ledoit, O. and Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88:365–411.

Lessoff, S., Schneidewind, U., Leven, C., Blum, P., Dietrich, P., and Dagan, G. (2010). Spatial characterization of the hydraulic conductivity using direct-push injection logging. *Water Resources Research*, 46(12).

Li, G. and Reynolds, A. C. (2009). Iterative ensemble Kalman filters for data assimilation. *Society of Petroleum Engineers*, 14(03):962–982.

Li, R., Prasad, V., and Huang, B. (2016). Gaussian mixture model-based ensemble Kalman filtering for state and parameter estimation for a PMMA process. *Processes*, 4(2):9.

Loe, M. K. and Tjelmeland, H. (2021). A generalised and fully Bayesian framework for ensemble updating. *arXiv*, 2103.14565[stat.ME].

Luo, X., Stordal, A., Lorentzen, R., and Nævdal, G. (2015). Iterative ensemble smoother as an approximate solution to a regularized minimum-average-cost problem: Theory and applications. *SPE Journal*, 20(5):962–982.

Mardia, K., Kent, J., and Bibby, J. (1979). *Multivariate analysis*. Probability and mathematical statistics. Acad. Press, London.

McElhoe, B. A. (1966). An assessment of the navigation and course corrections for a manned flyby of Mars or Venus. *IEEE Transactions on Aerospace and Electronic Systems*, 2(4):613–623.

Moja, S., Asfaw, Z., and Omre, H. (2018). Bayesian inversion in hidden Markov models with varying marginal proportions. *Mathematical Geosciences*, 51:463–484.

Myrseth, I. and Omre, H. (2010). Hierarchical ensemble Kalman filter. *SPE Journal*, 15(02):569–580.

Omre, H. and Rimstad, K. (2021). Bayesian spatial inversion and conjugate selection Gaussian prior models. *Journal of Uncertainty Quantification*, To appear.

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

Sætrom, J. and Omre, H. (2012). Ensemble Kalman filtering in a Bayesian regression framework. *9th Geostats. Oslo*.

Sætrom, J. and Omre, H. (2013). Uncertainty quantification in the ensemble Kalman filter. *Scandinavian Journal of Statistics*, 40(4):868–885.

Sakov, P., Haussaire, J.-M., and Bocquet, M. (2018). An iterative ensemble Kalman filter in the presence of additive model error. *Quarterly Journal of the Royal Meteorological Society*, 144(713):1297–1309.

Simon, E. and Bertino, L. (2009). Application of the Gaussian anamorphosis to assimilation in a 3-d coupled physical-ecosystem model of the north atlantic with the EnKF: a twin experiment. *Ocean Science*, 5(4):495–510.

Smith, K. W. (2007). Cluster ensemble Kalman filter. *Tellus A*, 59(5):749–757.

Sánchez-León, E., Leven, C., Erdal, D., and Cirpka, O. A. (2020). Comparison of two ensemble-Kalman based methods for estimating aquifer parameters from real 3-d hydraulic and tracer tomographic tests. *Geosciences*, 10(11):462.

Van Leeuwen, P. J. (2009). Particle filtering in geophysical systems. *Monthly Weather Review*, 137(12):4089–4114.

Zhou, H., Li, L., Franssen, H.-J., and Gomez-Hernandez, J. (2012). Pattern recognition in a bimodal aquifer using the normal-score ensemble Kalman filter. *Mathematical Geosciences*, 44:1–17.

Zhou, W., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612.