# Approximate forward-backward algorithm for a switching linear Gaussian model 

Hugo Hammer ${ }^{1, *}$<br>Department of Mathematical Sciences, Norwegian University of Science and Technology,<br>Trondheim, Norway<br>Håkon Tjelmeland<br>Department of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway


#### Abstract

A hidden Markov model with two hidden layers is considered. The bottom layer is a Markov chain and given this the variables in the second hidden layer are assumed conditionally independent and Gaussian distributed. The observation process is Gaussian with mean values that are linear functions of the second hidden layer. The forward-backward algorithm is not directly feasible for this model as the recursions result in a mixture of Gaussian densities where the number of terms grows exponentially with the length of the Markov chain. By dropping the less important Gaussian terms an approximate forward-backward algorithm is defined. Thereby one gets a computationally feasible algorithm that generates samples from an approximation to the conditional distribution of the unobserved layers given the data. The approximate algorithm is also used as a proposal distribution in a MetropolisHastings setting, and this gives high acceptance rates and good convergence and mixing properties. The model considered is related to what is known as Switching linear dynamical systems. The proposed algorithm can in principle also be used for these models and the potential use of the algorithm is therefore large. In simulation examples the algorithm is used for the problem of seismic inversion. The simulations demonstrate the effectiveness and quality of the proposed approximate algorithm.


Keywords: approximation, forward-backward algorithm, hidden Markov model,

[^0]Metropolis-Hastings algorithm, seismic inversion

## 1. Introduction

In a hidden Markov model (HMM) the observations are incomplete and noisy functions of an underlying unobserved process, where the latent process is assumed to be Markov. The goal is typically to restore the underlying process from the noisy observations and possibly also to estimate unknown parameters in both the latent and observation processes. For some HMM the underlying process consists of two layers, where Switching linear dynamical systems (SLDS) (Bar-Shalom and Li, 1998) is a typical example. In SLDS the bottom layer of the underlying unobserved process is a discrete Markov chain and conditioned on this the second unobserved layer is a Gaussian Markov process. Given the two unobserved layers the observations are assumed independent Gaussian. The mean vector and covariance matrix for the observed value at any time index are functions of the two unobserved states at the same time index. The goal is to restore the unobserved layers. SLDS have been used in many applications, e.g. fault detection in planetary rovers (Dearden and Clancy, 2002), speech recognition (Rosti and Gales, 2004), dancing of bees (Oh et al., 2005), econometrics (Kim, 1994) and machine learning (Lerner et al., 2000; Ghahramani and Hinton, 1998). Larsen et al. (2006) considers the problem of seismic inversion and a model similar to SLDS, but allow the observations to be a function of both past and future values of the hidden Gaussian process. The goal is again to restore the unobserved layers.

Recursive algorithms for HMM have successfully been used in many areas, see the discussions and references in MacDonald and Zucchini (1997), Künsch (2000), Scott (2002) and Cappé et al. (2005). Generalizations to hidden semi-Markov models are discussed in Guédon (2007) and Bulla et al. (2010). When an HMM only has one unobserved layer modelled as a discrete Markov chain, efficient recursive computations known as the forward-backward algorithm can be used. If the observations and the unobserved layer both are Gaussian, then the forward-backward algorithm corresponds to the famous Kalman filter. Forward-backward algorithms for HMM with two unobserved layers have been considered by Bar-Shalom and Li (1998), Barber (2006) and Zoeter and Heskes (2006). The forward-backward
recursions can also be formulated for these models, but are not computationally feasible as they involve a mixture of Gaussian distributions where the number of terms grows exponentially with the length of the Markov chain. In the first set of references given above, approximate forward recursions are defined by substituting the Gaussian mixture by a single Gaussian term. Larsen et al. (2006) define approximate recursions by approximating the marginal distribution for the unobserved continuous process by a product of Gaussian densities.

We consider a model close to the model in Larsen et al. (2006) and separate from SLDS in that the observations can be a function of both past and future values. To get a computationally feasible algorithm, we construct an approximate forward-backward algorithm. In the forward recursions we propose to drop terms associated with small weights in the Gaussian mixture. Thus, our approximation is less dramatic than previous suggestions, but with a corresponding higher computational cost. Clearly, the quality of the approximation depends on the number and importance of the terms that are dropped. Using the approximate forwardbackward algorithm as a proposal distribution in a Metropolis-Hastings algorithm (Smith and Roberts, 1993; Dellaportas and Roberts, 2003) we correct for the induced approximation. Moreover, we use the Metropolis-Hastings acceptance rate as a measure for the quality of the approximation. The proposed algorithm can in principle be used for SLDS and other similar models with the Markov property. The potential use of the algorithm is therefore large.

An alternative strategy to cope with HMM with two hidden layers are sequential Monte Carlo algorithms. Chen and Liu (2000) define a sequential Monte Carlo algorithm for what they call conditional dynamic linear models (CDLM). Our model defined in Section 2 can be rephrased to a CDLM by redefining the state variables. However, Chen and Liu (2000) only consider the filtering problem, whereas our focus is mainly the smoothing problem. Doucet et al. (2000) also use sequential Monte Carlo for a model similar to the one in Chen and Liu (2000), but the focus is again on filtering. Godsill et al. (2004) use sequential Monte Carlo to solve the smoothing problem for a state-space model. To generalize the procedure in Godsill et al. (2004) to handle a CDLM constitutes an alternative avenue for solving the problem we discuss.

In Hammer et al. (2010) a slightly modified variant of the simulation algorithm we present have successfully been used to invert real seismic data from an oilfield offshore Norway. Ulvmoen and Hammer (2009) also focus on the seismic inversion problem and use the procedure proposed in the present article to evaluate the quality of a much faster, but rougher approximation strategy. We also consider the problem of seismic inversion, but the focus here is to evaluate the efficiency of our proposed algorithm. In the seismic inversion setting, the Markov chain represents lithology-fluid classes along a vertical trace through the underground, the intermediate Gaussian layer represents elastic parameters of the rock along the same trace, and the observations are seismic data. The focus is to restore the unobserved Markov chain. Parameter estimation is clearly also of interest, but not considered here.

The paper is organized as follows. Section 2 introduces necessary notation for our hidden Markov model. In Section 3 we develop the approximate forward-backward algorithm. Section 4 gives a brief introduction to the seismic inversion application and explains how the hidden Markov model is the core part of the resulting model. We also evaluate the algorithm in simulation examples. Finally, Section 5 provides conclusions.

## 2. The switching linear Gaussian model

We represent (multivariate) Gaussian distributions in its canonical form, as this simplifies the forward-backward recursions. A Gaussian distribution with mean vector $\mu \in \mathbb{R}^{r}$ and covariance matrix $\Sigma \in \mathbb{R}^{r \times r}$ is then parameterized by the precision matrix $Q=\Sigma^{-1}$ and the vector $q=Q \mu$, and we use $\mathrm{N}(q, Q)$ to denote this distribution. The corresponding density we denote by $\mathrm{N}(u \mid q, Q)$, which reads

$$
\begin{equation*}
\mathrm{N}(u \mid q, Q)=\frac{\sqrt{|Q|}}{(2 \pi)^{\frac{r}{2}}} \exp \left\{-\frac{1}{2} q^{T} Q^{-1} q\right\} \exp \left\{-\frac{1}{2}\left[u^{T} Q u-2 q^{T} u\right]\right\} . \tag{1}
\end{equation*}
$$

Consider a three layer hidden Markov model $\left\{\left(x_{i}, y_{i}, z_{i}\right)\right\}_{i=1}^{n}$ as visualized in Figure 1, where $x_{i} \in\{1, \ldots, L\}, y_{i}=\left(y_{i 1}, \ldots, y_{i r}\right)^{T} \in \mathbb{R}^{r}$ and $z_{i}=\left(z_{i 1}, \ldots, z_{i s}\right)^{T} \in$ $\mathbb{R}^{s}$ for $i=1, \ldots, n$. We call this a switching linear Gaussian model. We require the number of possible values for $x_{i}, L$, to be small. In the seismic data example in Section 4 we have $L=4$. We let $x_{1: n}=\left(x_{1}, \ldots, x_{n}\right)^{T}$ be a stationary, aperiodic


Figure 1: Directed acyclic graph (DAG) representation of the hidden Markov model discussed in Sections 2 and 3.
and ergodic Markov chain with transition matrix

$$
\begin{equation*}
\mathrm{P}=\left[p\left(x_{i} \mid x_{i-1}\right)\right]_{x_{i-1}, x_{i}=1}^{L} \tag{2}
\end{equation*}
$$

Thus, the marginal distribution of $x_{1}$, which we denote by $p\left(x_{1}\right)$, equals the limiting distribution induced by P. Conditioned on $x_{1: n}$ we assume the elements of $y_{1: n}=$ $\left(y_{1}, \ldots, y_{n}\right)^{T}$ to be independent and Gaussian distributed, where the mean vector and precision matrix of $y_{i}$ are (known) functions of $x_{i}$ and denoted by $\mu\left(x_{i}\right)$ and $Q\left(x_{i}\right)$, respectively. Thus,

$$
\begin{equation*}
y_{i} \mid x_{1: n} \sim \mathrm{~N}\left(q\left(x_{i}\right), Q\left(x_{i}\right)\right), \text { where } q\left(x_{i}\right)=Q\left(x_{i}\right) \mu\left(x_{i}\right) \tag{3}
\end{equation*}
$$

Given $y_{1: n}$ we assume the elements of $z_{1: n}=\left(z_{1}, \ldots, z_{n}\right)^{T}$ to be independent and Gaussian, and the mean vector and precision matrix of $z_{i}$ are $a_{i}^{T} y_{i-1}+b_{i}^{T} y_{i}+c_{i}^{T} y_{i+1}$ and $R_{i}$, respectively, i.e.

$$
\begin{equation*}
z_{i} \mid y_{1: n} \sim \mathrm{~N}\left(A_{i}^{T} y_{i-1}+B_{i}^{T} y_{i}+C_{i}^{T} y_{i+1}, R_{i}\right) \tag{4}
\end{equation*}
$$

where $A_{i}=a_{i} R_{i}, B_{i}=b_{i} R_{i}$ and $C_{i}=c_{i} R_{i}$. Note that we allow the coefficient matrices $a_{i}, b_{i}$ and $c_{i}$ to vary with $i$, and in particular we require $a_{1}=c_{n}=0$.

## 3. The forward-backward algorithm

In this section we define an approximate forward-backward algorithm for the model described in Section 2. We first derive the exact forward recursions. Starting with $\pi\left(x_{1: n}, y_{1: n} \mid z_{1: n}\right)$ we integrate out $y_{i}$ and $x_{i}$ for $i=1, \ldots, n$ in turn to get
the distributions $\pi\left(x_{i: n}, y_{i+1: n} \mid z_{1: n}\right)$ and $\pi\left(x_{i: n}, y_{i: n} \mid z_{1: n}\right)$ for $i=1, \ldots, n$. This is the basis for the backward simulation part of the algorithm, which sequentially for $i=n, \ldots, 1$ generates $x_{i}$ from $\pi\left(x_{i} \mid x_{i+1: n}, y_{i+1: n}, z_{1: n}\right) \propto \pi\left(x_{i: n}, y_{i+1: n} \mid z_{1: n}\right)$ and $y_{i}$ from $\pi\left(y_{i} \mid x_{i: n}, y_{i+1: n}, z_{1: n}\right) \propto \pi\left(x_{i: n}, y_{i: n} \mid z_{1: n}\right)$.

### 3.1. Forward integration

The conditional distribution of interest is $\pi\left(x_{1: n}, y_{1: n} \mid z_{1: n}\right)$. However, to avoid notationally cumbersome special cases for $i=n-1$ and $n$ we also introduce $x_{n+1}$, $y_{n+1}, y_{n+2}$ and $z_{n+1}$. We make these auxiliary variables independent of the variables of interest by setting $p\left(x_{n+1} \mid x_{n}\right)=1 / L$ and also adopting (3) for $i=n+1$ and $n+2$ and (4) for $i=n+1$ with $A_{n+1}=B_{n+1}=C_{n+1}=0$ and $R_{n+1}=I$. As a function of $x_{1: n}$ and $y_{1: n}$ we then have

$$
\begin{gather*}
\pi\left(x_{1: n}, y_{1: n} \mid z_{1: n}\right) \propto \pi\left(x_{1: n+1}, y_{1: n}, z_{1: n+1} \mid y_{n+1: n+2}\right) \propto p\left(x_{1}\right) \prod_{i=2}^{n+1} p\left(x_{i} \mid x_{i-1}\right) \\
\prod_{i=1}^{n} \mathrm{~N}\left(y_{i} \mid q\left(x_{i}\right), Q\left(x_{i}\right)\right) \cdot \prod_{i=1}^{n+1} \mathrm{~N}\left(z_{i} \mid A_{i}^{T} y_{i-1}+B_{i}^{T} y_{i}+C_{i}^{T} y_{i+1}, R_{i}\right) \tag{5}
\end{gather*}
$$

To get a more compact notation in the development of the forward recursions we define

$$
\begin{equation*}
T_{0}\left(y_{1: 2}\right)=\mathrm{N}\left(z_{1} \mid B_{1}^{T} y_{1}+C_{1}^{T} y_{2}, R_{1}\right) \tag{6}
\end{equation*}
$$

and

$$
\begin{equation*}
T_{i}\left(x_{i}, y_{i: i+2}\right)=\mathrm{N}\left(y_{i} \mid q\left(x_{i}\right), Q\left(x_{i}\right)\right) \mathrm{N}\left(z_{i+1} \mid A_{i+1}^{T} y_{i}+B_{i+1}^{T} y_{i+1}+C_{i+1}^{T} y_{i+2}, R_{i+1}\right) \tag{7}
\end{equation*}
$$

for $i=1, \ldots, n$, so that

$$
\begin{equation*}
\pi\left(x_{1: n}, y_{1: n} \mid z_{1: n}\right) \propto p\left(x_{1}\right) T_{0}\left(y_{1: 2}\right) \prod_{i=1}^{n} p\left(x_{i+1} \mid x_{i}\right) T_{i}\left(x_{i}, y_{i: i+2}\right) \tag{8}
\end{equation*}
$$

Starting with this expression we integrate and sum out $y_{i}$ and $x_{i}$ for $i=1, \ldots, n$ in turn and define $U_{i}\left(x_{i}, y_{i: i+1}\right)$ and $V_{i}\left(x_{i}, y_{i+1: i+2}\right)$ for $i=1, \ldots, n$ so that the result becomes

$$
\begin{equation*}
\pi\left(x_{i: n}, y_{i: n} \mid z_{1: n}\right) \propto U_{i}\left(x_{i}, y_{i: i+1}\right) \prod_{j=i}^{n} p\left(x_{j+1} \mid x_{j}\right) T_{j}\left(x_{j}, y_{j: j+2}\right) \tag{9}
\end{equation*}
$$

and

$$
\begin{equation*}
\pi\left(x_{i: n}, y_{i+1: n} \mid z_{1: n}\right) \propto V_{i}\left(x_{i}, y_{i+1: i+2}\right) p\left(x_{i+1} \mid x_{i}\right) \prod_{j=i+1}^{n} p\left(x_{j+1} \mid x_{j}\right) T_{j}\left(x_{j}, y_{j: j+2}\right) \tag{10}
\end{equation*}
$$

The following theorem gives the relation between the $U_{i}$ and $V_{i}$ functions.

Theorem 1. Consider the hidden Markov model defined in Section 2 and the notation introduced above. We then have

$$
\begin{equation*}
U_{1}\left(x_{1}, y_{1: 2}\right)=p\left(x_{1}\right) T_{0}\left(y_{1: 2}\right) \tag{11}
\end{equation*}
$$

and the recursions

$$
\begin{equation*}
U_{i+1}\left(x_{i+1}, y_{i+1: i+2}\right)=\sum_{x_{i}=1}^{L} V_{i}\left(x_{i}, y_{i+1: i+2}\right) p\left(x_{i+1} \mid x_{i}\right), i=1, \ldots, n-1 \tag{12}
\end{equation*}
$$

and

$$
\begin{equation*}
V_{i}\left(x_{i}, y_{i+1: i+2}\right)=\int U_{i}\left(x_{i}, y_{i: i+1}\right) T_{i}\left(x_{i}, y_{i: i+2}\right) d y_{i}, i=1, \ldots, n \tag{13}
\end{equation*}
$$

The theorem is proven as follows. By comparing (8) and (9) for $i=1$ we get (11). Next, by summing out $x_{i}$ in (10) and comparing with (9) we get (12). Finally integrating out $y_{i}$ in (9) and comparing with (10) one gets (13).

In the following we use the notation

$$
D_{1}=\left[\begin{array}{l}
\mathbf{I}  \tag{14}\\
\mathbf{0} \\
\mathbf{0}
\end{array}\right], D_{3}=\left[\begin{array}{l}
\mathbf{0} \\
\mathbf{0} \\
\mathbf{I}
\end{array}\right], D_{12}=\left[\begin{array}{cc}
\mathbf{I} & \mathbf{0} \\
\mathbf{0} & \mathbf{I} \\
\mathbf{0} & \mathbf{0}
\end{array}\right] \text { and } D_{23}=\left[\begin{array}{cc}
\mathbf{0} & \mathbf{0} \\
\mathbf{I} & \mathbf{0} \\
\mathbf{0} & \mathbf{I}
\end{array}\right]
$$

where $\mathbf{0}$ and $\mathbf{I}$ are an $r \times r$ matrix with all elements equal to zero and the $r$ dimensional identity matrix, respectively. Then the following theorem gives how to compute the $U_{i}$ and $V_{i}$ functions recursively.

Theorem 2. Consider the hidden Markov model defined in Section 2 and the notation introduced above. We then have

$$
\begin{equation*}
U_{i}\left(x_{i}, y_{i: i+1}\right) \propto \sum_{j=1}^{N_{i}} \gamma_{i j}\left(x_{i}\right) \mathrm{N}\left(y_{i: i+1} \mid g_{i j}, G_{i j}\right) \tag{15}
\end{equation*}
$$

and

$$
\begin{equation*}
V_{i}\left(x_{i}, y_{i+1: i+2}\right) \propto \sum_{j=1}^{N_{i}} \kappa_{i j}\left(x_{i}\right) \mathrm{N}\left(y_{i+1: i+2} \mid k_{i j}\left(x_{i}\right), K_{i j}\left(x_{i}\right)\right) \tag{16}
\end{equation*}
$$

for $i=1, \ldots, n$, where $N_{i}=L^{i-1}$ and $\gamma_{i j}\left(x_{i}\right) \in \mathbb{R}, g_{i j} \in \mathbb{R}^{2 r \times 1}, G_{i j} \in \mathbb{R}^{2 r \times 2 r}$, $\kappa_{i j}\left(x_{i}\right) \in \mathbb{R}, k_{i j}\left(x_{i}\right) \in \mathbb{R}^{2 r \times 1}$ and $K_{i j}\left(x_{i}\right) \in \mathbb{R}^{2 r \times 2 r}$ can be computed recursively. Initial values $\gamma_{11}\left(x_{1}\right), g_{11}$ and $G_{11}$ are

$$
\gamma_{11}\left(x_{1}\right)=p\left(x_{1}\right), \quad g_{11}=\left[\begin{array}{c}
B_{1}  \tag{17}\\
C_{1}
\end{array}\right] z_{1}, \quad G_{11}=\left[\begin{array}{c}
B_{1} \\
C_{1}
\end{array}\right] R_{1}^{-1}\left[B_{1}^{T} C_{1}^{T}\right]
$$

For $i=1, \ldots, n$ and $j=1, \ldots, N_{i}$ the $\kappa_{i j}\left(x_{i}\right), k_{i j}\left(x_{i}\right)$ and $K_{i j}\left(x_{i}\right)$ can be obtained from $\gamma_{i j}\left(x_{i}\right), g_{i j}$ and $G_{i j}$ by

$$
\begin{align*}
& \kappa_{i j}\left(x_{i}\right)=\gamma_{i j}\left(x_{i}\right) \sqrt{\frac{\left|G_{i j}\right| \cdot\left|Q\left(x_{i}\right)\right|}{\left|H_{i j}\left(x_{i}\right)\right|}}  \tag{18}\\
& \exp \left\{-\frac{1}{2}\left[g_{i j}^{T} G_{i j}^{-1} g_{i j}+q\left(x_{i}\right)^{T} Q\left(x_{i}\right)^{-1} q\left(x_{i}\right)-h_{i j}\left(x_{i}\right)^{T} H_{i j}\left(x_{i}\right)^{-1} h_{i j}\left(x_{i}\right)\right]\right\} \\
& k_{i j}\left(x_{i}\right)=K_{i j}\left(x_{i}\right) D_{23}^{T} H_{i j}\left(x_{i}\right)^{-1} h_{i j}\left(x_{i}\right) \tag{19}
\end{align*}
$$

and

$$
\begin{equation*}
K_{i j}\left(x_{i}\right)=\left(D_{23}^{T} H_{i j}\left(x_{i}\right)^{-1} D_{23}\right)^{-1} \tag{20}
\end{equation*}
$$

where

$$
h_{i j}\left(x_{i}\right)=D_{1} q\left(x_{i}\right)+D_{12} g_{i j}+\left[\begin{array}{ccc}
A_{i+1}^{T} & B_{i+1}^{T} & C_{i+1}^{T} \tag{21}
\end{array}\right]^{T} z_{i+1}
$$

and

$$
\begin{align*}
H_{i j}\left(x_{i}\right)= & D_{1} Q\left(x_{i}\right) D_{1}^{T}+D_{12} G_{i j} D_{12}^{T}+\delta(i \geq n-1) D_{3} D_{3}^{T} \\
& +\left[\begin{array}{lll}
A_{i+1}^{T} & B_{i+1}^{T} & C_{i+1}^{T}
\end{array}\right]^{T} R_{i+1}^{-1}\left[\begin{array}{lll}
A_{i+1}^{T} & B_{i+1}^{T} & C_{i+1}^{T}
\end{array}\right] \tag{22}
\end{align*}
$$

where $\delta(\cdot)$ is the indicator function, i.e. $\delta(E)=1$ when $E$ is true and $\delta(E)=0$ otherwise. Finally, for $i=2, \ldots, n, j=1, \ldots, N_{i-1}$ and $l=1, \ldots, L$

$$
\begin{align*}
\gamma_{i, j+(l-1) N_{i-1}}\left(x_{i}\right) & =p\left(x_{i} \mid l\right) \kappa_{i-1, j}(l),  \tag{23}\\
g_{i, j+(l-1) N_{i-1}} & =k_{i-1, j}(l) \tag{24}
\end{align*}
$$

and

$$
\begin{equation*}
G_{i, j+(l-1) N_{i-1}}=K_{i-1, j}(l) \tag{25}
\end{equation*}
$$

The theorem is proved by induction. Reordering terms in $T_{0}\left(y_{1: 2}\right)$ straightforwardly gives (15) for $i=1$ and initial values (17). Starting with (13), rearranging terms and using well known properties of the multivariate Gaussian distribution gives (16) and (18) through to (22). We represent Gaussian distributions in the canonical form and not by the mean vector and covariance matrix and this causes the somewhat unfamiliar expressions in (19) and (20).

The $D_{3} D_{3}^{T}$ term that appears in (22) for $i=n-1$ and $n$ ensures that the $H_{n-1, j}$ and $H_{n j}$ matrices are invertible. The $D_{3} D_{3}^{T}$ term does not influence the variables of interest, $x_{1: N}$ and $y_{1: n}$, only the auxiliary variables $y_{n+1}$ and $y_{n+2}$. Finally, (12) straightforwardly gives (15) and (23) through to (25) by a reordering of the terms.

As the number of terms in (15) and (16) grows exponentially with $i$ the recursive algorithm is computationally feasible only for small values of $n$. In the next section we propose to approximate $U_{i}\left(x_{i}, y_{i: i+1}\right)$ and $V_{i}\left(x_{i}, y_{i+1: i+2}\right)$ by ignoring less important terms.

### 3.2. Approximate forward integration algorithm

In this section we propose an approximate, but computationally feasible version of the recursions developed above. We first compute the (exact) representations of $U_{1}\left(x_{1}, y_{1: 2}\right)$ and $V_{1}\left(x_{1}, y_{2: 3}\right)$ as given in Theorem 2. The starting point for finding an approximation for $U_{i}\left(x_{i}, y_{i: i+1}\right)$ and $V_{i}\left(x_{i}, y_{i+1: i+2}\right)$ is an approximation of $V_{i-1}\left(x_{i-1}, y_{i: i+1}\right)$ on the form

$$
\begin{equation*}
\tilde{V}_{i-1}\left(x_{i-1}, y_{i: i+1}\right) \propto \sum_{j=1}^{\tilde{N}_{i-1}\left(x_{i-1}\right)} \widetilde{\kappa}_{i-1, j}\left(x_{i-1}\right) \mathrm{N}\left(y_{i: i+1} \mid \widetilde{k}_{i-1, j}\left(x_{i-1}\right), \widetilde{K}_{i-1, j}\left(x_{i-1}\right)\right) \tag{26}
\end{equation*}
$$

where we use tilde to distinguish approximate quantities from exact ones. The approximate representation is of the same form as (16), except that in (26) the number of terms may depend on the value of $x_{i-1}$. Of course, for $i=2$ we use $\widetilde{V}_{i-1}\left(x_{i-1}, y_{i: i+1}\right)=V_{i-1}\left(x_{i-1}, y_{i: i+1}\right)$. For $i>2$ we define $\widetilde{V}_{i}\left(x_{i}, y_{i+1: i+2}\right)$ from $\widetilde{V}_{i-1}\left(x_{i-1}, y_{i: i+1}\right)$ in to steps. First we use the recursions in Theorem 2 to find an approximation $U_{i}^{\star}\left(x_{i}, y_{i: i+1}\right)$ to $U_{i}\left(x_{i}, y_{i: i+1}\right)$ and a first approximation $V_{i}^{\star}\left(x_{i}, y_{i+1: i+2}\right)$ to $V_{i}\left(x_{i}, y_{i+1: i+2}\right)$. Thereafter we drop the less important terms in $V_{i}^{\star}\left(x_{i}, y_{i+1: i+2}\right)$ to get a final approximation $\widetilde{V}_{i}\left(x_{i}, y_{i+1: i+2}\right)$. More precisely, we set

$$
\begin{equation*}
U_{i}^{\star}\left(x_{i}, y_{i: i+1}\right) \propto \sum_{j=1}^{N_{i}^{\star}} \gamma_{i j}^{\star}\left(x_{i}\right) \mathrm{N}\left(y_{i: i+1} \mid g_{i j}^{\star}, G_{i j}^{\star}\right) \tag{27}
\end{equation*}
$$

and

$$
\begin{equation*}
V_{i}^{\star}\left(x_{i}, y_{i+1: i+2}\right) \propto \sum_{j=1}^{N_{i}^{\star}} \kappa_{i j}^{\star}\left(x_{i}\right) \mathrm{N}\left(y_{i+1: i+2} \mid k_{i j}^{\star}\left(x_{i}\right), K_{i j}^{\star}\left(x_{i}\right)\right), \tag{28}
\end{equation*}
$$

where $N_{i}^{\star}=\sum_{l=1}^{L} \tilde{N}_{i-1}(l)$. Corresponding to (23) through to (25), $\gamma_{i j}^{\star}\left(x_{i}\right), g_{i j}^{\star}$ and $G_{i j}^{\star}$ are defined by

$$
\begin{align*}
\gamma_{i, j+\sum_{t=1}^{l-1} \widetilde{N}_{i-1}(t)}\left(x_{i}\right) & =p\left(x_{i} \mid l\right) \widetilde{\kappa}_{i-1, j}(l),  \tag{29}\\
g_{i, j+\sum_{t=1}^{\star-1} \tilde{N}_{i-1}(t)} & =\widetilde{k}_{i-1, j}(l) \tag{30}
\end{align*}
$$

and

$$
\begin{equation*}
G_{i, j+\sum_{t=1}^{l-1} \widetilde{N}_{i-1}(t)}^{\star}=\widetilde{K}_{i-1, j}(l) \tag{31}
\end{equation*}
$$

for $i=2, \ldots, n, j=1, \ldots, \widetilde{N}_{i-1}(l)$ and $l=1, \ldots, L$. Finally, $\kappa_{i j}^{\star}\left(x_{i}\right), k_{i j}^{\star}\left(x_{i}\right)$ and $K_{i j}^{\star}\left(x_{i}\right)$ are defined by replacing $\kappa_{i j}\left(x_{i}\right), k_{i j}\left(x_{i}\right), K_{i j}\left(x_{i}\right), \gamma_{i j}\left(x_{i}\right), g_{i j}, G_{i j}$ and $N_{i}$ with corresponding starred quantities in (18) through to (22).

Which terms in (28) that are of less importance is not obvious as the terms are functions of $y_{i+1: i+2}$ which is still unspecified when the decision about what terms to drop has to be made. Natural strategies are either to maximize over or to integrate out $y_{i+1: i+2}$ before comparing the terms. Maximizing over $y_{i+1: i+2}$ is obtained by evaluating the Gaussian densities in (28) at their mean values. Thus, for a threshold value $\varepsilon$ this gives that we should drop terms in (28) that have

$$
\begin{equation*}
\frac{\kappa_{i j}^{\star}\left(x_{i}\right) \mathrm{N}\left(\mu_{i j}^{\star}\left(x_{i}\right) \mid k_{i j}^{\star}\left(x_{i}\right), K_{i j}^{\star}\left(x_{i}\right)\right)}{\max _{k=1, \ldots, N_{i}^{\star}\left(x_{i}\right)}\left\{\kappa_{i k}^{\star}\left(x_{i}\right) \mathrm{N}\left(\mu_{i k}^{\star}\left(x_{i}\right) \mid k_{i k}^{\star}\left(x_{i}\right), K_{i k}^{\star}\left(x_{i}\right)\right)\right\}}<\varepsilon, \tag{32}
\end{equation*}
$$

where $\mu_{i j}^{\star}\left(x_{i}\right)=K_{i j}^{\star}\left(x_{i}\right)^{-1} k_{i j}^{\star}\left(x_{i}\right)$. With the second strategy, integrating out $y_{i+1}$ and $y_{i+2}$, only $\kappa_{i j}^{\star}\left(x_{i}\right)$ remains to compare. Thus, again for a given threshold $\varepsilon$, we drop all terms that corresponds to a $\kappa_{i j}^{\star}\left(x_{i}\right)$ that have

$$
\begin{equation*}
\frac{\kappa_{i j}^{\star}\left(x_{i}\right)}{\max _{k=1, \ldots, N_{i}^{\star}\left(x_{i}\right)}\left\{\kappa_{i k}^{\star}\left(x_{i}\right)\right\}}<\varepsilon \tag{33}
\end{equation*}
$$

In the simulation examples in Section 4 we adopt the first strategy, but we do not expect the second strategy to behave much differently. One should note that we decide what terms to drop separately for each possible value of $x_{i}$, and as a result the number of remaining terms, $\tilde{N}_{i}\left(x_{i}\right)$, becomes a function of $x_{i}$.

Clearly, alternative term dropping strategies may be defined. First, one may use the term dropping step for $U_{i}^{\star}\left(x_{i}, y_{i: i+1}\right)$ instead, but we do not expect this to make much difference. Second, instead of choosing a specific threshold value $\varepsilon$, one may fix the number of terms we want to keep and drop the necessary number of small terms. Thereby the memory requirements for running the algorithm will be known in advance, but the quality of the approximation may be more variable than with the strategy we have chosen.

### 3.3. Backward simulation

When the (exact or approximate) forward integration is done and necessary quantities stored in memory, backward simulation is straight forward. Here we
give the necessary equations for the approximate, computational feasible algorithm. We initiate auxiliary variables $x_{n+1}, y_{n+1: n+2}$ and $z_{n+1}$ with arbitrary values and sequentially for $i=n, \ldots, 1$ first simulate $x_{i}$ from

$$
\begin{equation*}
\pi^{\star}\left(x_{i} \mid x_{i+1: n}, y_{i+1: n}, z_{1: n}\right) \propto V_{i}^{\star}\left(x_{i}, y_{i+1: i+2}\right) p\left(x_{i+1} \mid x_{i}\right) \tag{34}
\end{equation*}
$$

and then $y_{i}$ from

$$
\begin{equation*}
\pi^{\star}\left(y_{i} \mid x_{i: n}, y_{i+1: n}, z_{1: n}\right) \propto U_{i}^{\star}\left(x_{i}, y_{i: i+1}\right) T_{i}\left(x_{i}, y_{i: i+2}\right) \tag{35}
\end{equation*}
$$

The first is a discrete distribution and the second a mixture of $r$-variate Gaussian densities, so both are easy to sample from. The resulting realization is thereby simulated from an approximation to the conditional distribution $\pi\left(x_{1: n}, y_{1: n} \mid z_{1: n}\right)$,

$$
\begin{align*}
& \pi^{\star}\left(x_{1: n}, y_{1: n} \mid z_{1: n}\right)= \\
& \quad \prod_{i=1}^{n}\left[\pi^{\star}\left(x_{i} \mid x_{i+1: n+1}, y_{i+1: n+2}, z_{1: n+1}\right) \pi^{\star}\left(y_{i} \mid x_{i: n+1}, y_{i+1: n+2}, z_{1: n+1}\right)\right] . \tag{36}
\end{align*}
$$

One should note that evaluating $\pi^{\star}\left(x_{1: n}, y_{1: n} \mid z_{1: n}\right)$ is straight forward for a generated sample $\left(x_{1: n}, y_{1: n}\right)$, but to do this correctly one must of course remember to include the normalizing constants in the two conditional distributions $\pi^{\star}\left(x_{i} \mid x_{i+1: n+1}, y_{i+1: n+2}, z_{1: n+1}\right)$ and $\pi^{\star}\left(y_{i} \mid x_{i: n+1}, y_{i+1: n+2}, z_{1: n+1}\right)$.

### 3.4. Simulation from the hidden Markov model

The error introduced by the approximation discussed above may be corrected for by adopting $\pi^{\star}\left(x_{1: n}, y_{1: n} \mid z_{1: n}\right)$ as a proposal distribution in an independent proposal Metropolis-Hastings scheme. The resulting acceptance rate can then also be used as a measure for the quality of the approximation.

## 4. Simulation examples

We study the approximate forward-backward algorithm in a number of simulation exercises. We implement the algorithm in $\mathrm{C}++$, where a list is used to store the Gaussian mixture. Each element in the list is a term in the Gaussian mixture. Using a list makes it is easy to remove terms with small weight following Section 3.2.

In the simulation examples we focus on the problem of seismic inversion from the petroleum industry. Our objective here is to demonstrate that our approximate algorithm is able to solve a problem of significant practical importance. Here we consider simulated data only, but in Hammer et al. (2010) we also apply it on real seismic data. Seismic inversion is the discipline of predicting lithology-fluid characteristics in a reservoir from seismic data. Numerous introductory books to seismic terminology and inversion exist, see for example Sheriff and Geldart (1995) and references therein.

Seismic data is created by an explosion which sends sound waves into the ground. Parts of the waves are reflected, returned upwards and observed by microphones (geo- or hydrophones). These observations are the basis for the seismic data. A forward model, describing what we observe for given lithology-fluid characteristics, is known from physics theory. In seismic inversion we are interested in the corresponding inverse problem.

The simulation example is organized as follows. In Section 4.1 we present the seismic model and in Section 4.2 we present our choices of parameters in the seismic model. Further in Section 4.3 we link the seismic model to the switching model in Section 2 and explain how we simulate efficiently by taking advantage of the approximate forward-backward algorithm in Section 3. Finally in Sections 4.4 and 4.5 we present simulation results.

### 4.1. Seismic model

Our forward model is similar to the ones in Buland et al. (2003) and Larsen et al. (2006). When dealing with seismic data depth is typically not referenced by distance, but time used by the sound wave from the surface to a location in the underground and back, called two way travel time. An important problem not considered here is how to convert travel times to depths. Following Buland et al. (2003) and Larsen et al. (2006) we discretize the travel time and formulate the problem in a Bayesian setting. Let $i=1, \ldots, n$ denote $n$ two way travel times along a vertical profile and let $x_{i}$ denote the lithology-fluid class in location $i$. As prior for $x_{1: n}=\left(x_{1}, \ldots, x_{n}\right)^{T}$ we adopt a Markov chain as specified by (2). Assuming an isotropic and elastic medium, the material properties at a location $i$ are uniquely
defined by the P -wave velocity $\left(\alpha_{i}\right)$, S-wave velocity $\left(\beta_{i}\right)$ and density $\left(\rho_{i}\right)$ at that location. Let $y_{i}=\left(\ln \alpha_{i}, \ln \beta_{i}, \ln \rho_{i}\right)^{T}$. The distribution of $y_{1: n}=\left(y_{1}, \ldots, y_{n}\right)^{T}$ given $x_{1: n}$ is based on a rock physics model (Avseth et al., 2005) and we assume a Gaussian distribution as specified by (3).

We consider seismic data for $s$ offset values, or angles, $\theta_{1}, \ldots, \theta_{s}$. For each depth location $i$ and offset value $\theta_{j}$ a reflection coefficient $r_{i j}$ results from $y_{1: n}$. For this we use what is known as a weak contrast approximation to the Zoeppritz equations (Aki and Richards, 1980; Buland and Omre, 2003) and get for $r_{i}=\left(r_{i 1}, \ldots, r_{i s}\right)^{T}$,

$$
\begin{equation*}
r_{i}=\Gamma \frac{y_{i+1}-y_{i-1}}{2} \text { for } i=2, \ldots, n-1 \tag{37}
\end{equation*}
$$

where

$$
\Gamma=\left[\begin{array}{llll}
\gamma_{\alpha}\left(\theta_{1}\right) & \gamma_{\alpha}\left(\theta_{2}\right) & \cdots & \gamma_{\alpha}\left(\theta_{s}\right)  \tag{38}\\
\gamma_{\beta}\left(\theta_{1}\right) & \gamma_{\beta}\left(\theta_{2}\right) & \cdots & \gamma_{\beta}\left(\theta_{s}\right) \\
\gamma_{\rho}\left(\theta_{1}\right) & \gamma_{\rho}\left(\theta_{2}\right) & \cdots & \gamma_{\rho}\left(\theta_{s}\right)
\end{array}\right], \begin{aligned}
& \gamma_{\alpha}(\theta)=\frac{1}{2}\left(1+\tan ^{2}(\theta)\right) \\
& \gamma_{\beta}(\theta)=-4 \overline{\beta / \alpha}^{2} \sin ^{2}(\theta) \\
& \\
& \gamma_{\rho}(\theta)=\frac{1}{2}\left(1-4 \overline{\beta / \alpha}^{2} \sin ^{2}(\theta)\right)
\end{aligned}
$$

and one has assumed the ratio $\beta_{i} / \alpha_{i}$ to have an approximately constant value $\overline{\beta / \alpha}$ in the reservoir. The difference in (37) is an approximation to a derivative in the corresponding continuous model. For $i=1$ and $n$ we correspondingly use forward and backward differences, respectively. Finally, seismic observation $d_{i j}$ is obtained for each location $i$ and offset $\theta_{j}$ through a convolution of the reflection coefficients,

$$
\begin{equation*}
d_{i j}=\sum_{u=-k}^{k} \omega_{u j} r_{i-u, j}+\varepsilon_{i j} \tag{39}
\end{equation*}
$$

where $\left\{\omega_{u j}\right\}_{u=-k}^{k}$ defines a wavelet for each offset $\theta_{j}$ and $\varepsilon_{i j}$ is Gaussian observation noise. Similar to Buland and Omre (2003) we assume the main part of the noise to have a correlation structure corresponding to the wavelet. The argument for this is that both the signal and noise parts are the results of sound waves going through the (same) underground. More precisely, we set

$$
\begin{equation*}
\varepsilon_{i j}=\sum_{u=-k}^{k} w_{u j} \varepsilon_{i-u, j}^{1}+\varepsilon_{i j}^{2}, \tag{40}
\end{equation*}
$$

where $\varepsilon_{i j}^{1}$ and $\varepsilon_{i j}^{2}$ are independent Gaussian white noise with $\operatorname{Var}\left(\varepsilon_{i j}^{1}\right)=\sigma_{1}^{2}$ and $\operatorname{Var}\left(\varepsilon_{i j}^{2}\right)=\sigma_{2}^{2}$.

### 4.2. Parameter values

Our base case parameter values are chosen to be realistic for the seismic inversion application and are based on the values adopted in Larsen et al. (2006). We have $L=4$ classes for $x_{i}$, where $x_{i}=1,2$ and 3 represent gas-, oil- and brine (water) saturated sandstone, respectively, and $x_{i}=4$ represents shale. Sandstone is porous and allows flow of gas, oil and water, whereas the shale porosity is negligible and thereby acts as a barrier to fluid flow. Our choice of transition matrix P is based on values used in Larsen et al. (2006), but we consider a coarser seismic resolution than done there. Numbering the nodes from bottom to top, we use

$$
\mathrm{P}=\left[\begin{array}{cccc}
0.9441 & 0 & 0 & 0.0559  \tag{41}\\
0.0431 & 0.9146 & 0 & 0.0424 \\
0.0063 & 0.0230 & 0.9422 & 0.0284 \\
0.0201 & 0.0202 & 0.1006 & 0.8591
\end{array}\right]
$$

The zero elements are important in the seismic application as these represent known physical properties. Water has a higher density than oil, which again has a higher density than gas. Thus, water can not be above gas or oil and oil can not be above gas, unless separated by a non-porous shale layer. The corresponding marginal probabilities for $x_{i}$ are $[0.24,0.16,0.38,0.22]$.

As discussed in Section 4.1 we use $y_{i} \in \mathbb{R}^{3}$, where the three elements represent logarithms of P - and S-wave velocities and density, respectively. In Larsen et al. (2006) the distribution of $y_{i} \mid x_{i}$ is represented as empirical distributions given by a set of corresponding $x_{i}$ and $y_{i}$ values. We use the same set of $\left(x_{i}, y_{i}\right)$ values to estimate mean vectors and covariance matrices for the four assumed Gaussian distributions. The resulting mean vectors are $\mu(1)=[8.052,7.492,7.688]^{T}, \mu(2)=$ $[8.071,7.472,7.730]^{T}, \mu(3)=[8.121,7.467,7.746]^{T}$ and $\mu(4)=[8.166,7.546,7.846]^{T}$, and for each value of $x_{i}$, the diagonal and off-diagonal entries in the following matrices give corresponding standard deviations and correlations

$$
\left[\begin{array}{lll}
0.031 & 0.876 & 0.322  \tag{42}\\
0.876 & 0.033 & 0.271 \\
0.322 & 0.271 & 0.012
\end{array}\right], \quad\left[\begin{array}{lll}
0.027 & 0.891 & 0.384 \\
0.891 & 0.032 & 0.295 \\
0.384 & 0.295 & 0.009
\end{array}\right]
$$



Figure 2: Scatter plots of P - and S -wave velocities and density of samples from the distribution adopted for $y_{i} \mid x_{i}$ in BC. Red, green, blue and black is used for gas-, oil- and brine-saturated sandstone and shale, respectively.

$$
\left[\begin{array}{lll}
0.022 & 0.912 & 0.453  \tag{43}\\
0.912 & 0.032 & 0.317 \\
0.453 & 0.315 & 0.008
\end{array}\right], \quad\left[\begin{array}{lll}
0.044 & 0.982 & 0.935 \\
0.982 & 0.068 & 0.917 \\
0.935 & 0.917 & 0.015
\end{array}\right]
$$

Figure 2 shows scatter plots of simulated P - and S-wave velocities and density according to the specified distributions. Here and in all the following we use red, green, blue and black for gas-, oil- and brine-saturated sandstone and shale, respectively. We observe that shale is well separated from the other classes and that gas-saturated sandstone is reasonably well separated from oil- and brine-saturated sandstone, whereas there is more overlap between oil- and brine-saturated sandstone.

To specify the model for $d_{1: n} \mid y_{1: n}$ we must give what offsets to use, the wavelet for each offset, and the variances $\sigma_{1}^{2}$ and $\sigma_{2}^{2}$. Still following Larsen et al. (2006) we use $s=5$ offsets $\theta=0^{\circ}, 10^{\circ}, 20^{\circ}, 30^{\circ}$ and $40^{\circ}$ and adopt an offset independent Ricker wavelet,

$$
\begin{equation*}
w(u, \theta)=\left\{1-2(\pi \phi u)^{2}\right\} \exp \left\{-(\pi \phi u)^{2}\right\}, u=-k, \ldots, k \tag{44}
\end{equation*}
$$

with $\phi=0.11$ and $k=10$. For the error variances we use $\sigma_{1}^{2}=0.015^{2}$ and $\sigma_{2}^{2}=\sigma_{1}^{2} / 10^{4}$, which also corresponds to values used in Larsen et al. (2006).

The above defines our base case parameter set, which we refer to as BC. We define four more parameter sets, which are small modifications of BC. First the variances of the noise terms $\varepsilon_{i j}^{1}$ and $\varepsilon_{i j}^{2}$ are scaled to give higher and lower signal-tonoise ratios, and we denote these by LN and MN , respectively. We use $\sigma_{1}^{2}=0.0085^{2}$ and $0.026^{2}$ for LN and MN , respectively, still keeping $\sigma_{2}^{2}=\sigma_{1}^{2} / 10^{4}$. The next two

| Parameter set | BC | LN | MN | RL | RM |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Signal-to-noise-ratio | 1.34 | 2.22 | 0.53 | 1.30 | 1.36 |
| $\varepsilon$ | $2.5 \cdot 10^{-3}$ | $2.5 \cdot 10^{-3}$ | $1.8 \cdot 10^{-3}$ | $1.7 \cdot 10^{-3}$ | $1.7 \cdot 10^{-3}$ |
| acceptance rate | 0.44 | 0.43 | 0.37 | 0.46 | 0.42 |

Table 1: Signal-to-noise ratios, value used for the tuning parameter $\varepsilon$, and the resulting MetropolisHastings acceptance rate for the various parameter sets.
cases are obtained from BC by a scaling of the covariance matrices in the rock physics model, $y_{i} \mid x_{i}$. We define rock physics models with less variance (RL) and more variance (RM) by multiplying the covariance matrices defined by (42) and (43) by $1 / 2$ and 2 , respectively. We want the signal-to-noise ratios for RL and RM to be about the same as for BC and obtain this by modifying the noise variance $\sigma_{1}^{2}$, still keeping $\sigma_{2}^{2}=\sigma_{1}^{2} / 10^{4}$. When defining the signal-to-noise ratio we consider variability in $d_{1: n}$ originating from $x_{1: n}$ as signal and the remaining variability in $d_{1: n}$ as noise, see Hammer (2008) for the precise definition. This gave $\sigma_{1}^{2}=1.65 \cdot 10^{-2}$ and $1.10 \cdot 10^{-2}$ for RL and RM, respectively. Table 1 gives the resulting signal-to-noise ratios for all five parameter sets.

### 4.3. Simulating from the seismic model

We want to simulate $x_{1: n}$ and $y_{1: n}$ conditioned on $d_{1: n}$ in the seismic model defined above. A key point in the construction of an effective simulation algorithm is to take advantage of the approximate forward-backward algorithm defined in Section 3. We achieve this by introducing the additional variable $z_{1: n}=\left(z_{1}, \ldots, z_{n}\right)^{T}$, where $z_{i}=\left(z_{i 1}, \ldots, z_{i s}\right)^{T} \in \mathbb{R}^{s}$ for $i=1, \ldots, n$,

$$
\begin{equation*}
z_{i}=r_{i}+\varepsilon_{i}^{1} \tag{45}
\end{equation*}
$$

and $\varepsilon_{i}^{1}=\left(\varepsilon_{i 1}^{1}, \ldots, \varepsilon_{i s}^{1}\right)^{T}$. The distributions for $x_{1: n}, y_{1: n}$ and $z_{1: n}$ is then as specified in Section 2. In (4) we have $A_{i}=-\Gamma R_{i} / 2, B_{i}=\mathbf{0}$ and $C_{i}=\Gamma R_{i} / 2$ for $i=$ $2, \ldots, n-1$, and using forward and backward difference at the boundaries $B_{1}=$ $-\Gamma R_{1}, C_{1}=\Gamma R_{1}, A_{n}=-\Gamma R_{n}$ and $B_{n}=\Gamma R_{n}$. Finally we have $R_{i}=\sigma_{1}^{-2} I$ for $i=1, \ldots, n$, where $I$ is the identity matrix. Combining (39), (40) and (45) we get
the relation between $z_{1: n}$ and $d_{1: n}$,

$$
\begin{equation*}
d_{i j} \mid z \sim \mathrm{~N}\left(\sigma_{2}^{-2} \sum_{u=-k}^{k} w\left(u, \theta_{j}\right) \cdot z_{i-u, j}, \sigma_{2}^{-2} I\right) \tag{46}
\end{equation*}
$$

We construct a Metropolis-Hastings algorithm (Smith and Roberts, 1993; Dellaportas and Roberts, 2003) consisting of two updates in each iteration. The first update is a block Gibbs update for $y_{1: n}$ and $z_{1: n}$. The joint full conditional for these are Gaussian and therefore easy to sample from. The second update in each iteration is a joint Metropolis-Hastings update for $x_{1: n}$ and $y_{1: n}$ by using the approximate forward-backward algorithm as the proposal distribution.

### 4.4. Evaluation of the approximate forward-backward algorithm

In this section we report the results for one Metropolis-Hastings run for each of the five parameter sets defined in Section 4.2 with $n=100$ and use this to evaluate the performance of the proposed approximate forward integration algorithm. In each case we first simulate $x_{1: n}, y_{1: n}, z_{1: n}$ and $d_{1: n}$ according to the model specified in Sections 2 and 4.1 and thereafter use the algorithm proposed in Section 4.3 to sample from the resulting posterior distribution $\pi\left(x_{1: n}, y_{1: n}, z_{1: n} \mid d_{1: n}\right)$. We evaluate the quality of the approximate algorithm by the acceptance rates and the convergence and mixing properties of the simulated Markov chains.

Following the optimal strategies for choice of Metropolis-Hastings tuning parameters found in Roberts et al. (1997) and Roberts and Rosenthal (1998) we find a value for our tuning parameter $\varepsilon$ for each of the five parameter sets by aiming at a Metropolis-Hastings acceptance rate of about 0.4. It should be noted that our situation differs from what is discussed in the two references, so it is not clear that this is an optimal strategy in our situation. However, we found it to be a reasonable first try and it has worked satisfactory in all our runs. Table 1 reports both the $\varepsilon$ values used and the resulting acceptance rates.

Figures 3 to 7 present simulation results for each of the five parameter sets. The upper rows show the simulated "true" values. Note that we use the same realization of $x_{1: n}$ in all cases to make comparison easier. The lower rows consist of three parts. To the left the "true" $x_{1: n}$ is replotted for easier comparison, in the middle each state of the Metropolis-Hastings run is plotted side by side, and the plots to the


Figure 3: Simulation results for parameter set BC: The upper row shows, from left to right, the simulated "true" $x_{1: n}$, elastic parameters $\exp \left(y_{1: n}\right), z_{1: n}$ and $d_{1: n}$. The lower row contains posterior simulation results. From left to right, the lower row shows the true $x_{1: n}$ (replotted for easier comparison), each state of the Metropolis-Hastings run plotted side by side, and the resulting estimated marginal posterior probabilities.
right show the resulting estimated marginal probabilities for each node $i$. The runs shown are all initiated by setting all $x_{i}=1$ and drawing $y_{1: n}$ and $z_{1: n}$ values from the corresponding full conditional. In all the runs the initial state is left within very few iterations and the burn-in phases are not even visible in the figures. We have also tried starting with all $x_{i}=4$ and other initial values, but without experiencing any burn-in problems. The results clearly show that the approximate forwardbackward algorithm gives a good approximation to the distribution of interest and produces very good mixing properties when used as a proposal distribution in a Metropolis-Hastings setting.

For the data shown in Figures 3 to 7 we have also tried the algorithm described in Section 3.4 for simulating from $\pi\left(x_{1: n}, y_{1: n} \mid z_{1: n}\right)$. Again we tuned the value of $\varepsilon$ as described above. The results indicated quite good convergence and mixing properties, but here some stickiness in the runs could be observed. As the algorithm is an independent proposal procedure, the latter should come as no surprise. Figure 8 shows the total number of Gaussian terms stored for each node $i$. Comparing the


Figure 4: Simulation results for parameter set LN: See Figure 3 for an explanation of the different parts of the figure.


Figure 5: Simulation results for parameter set MN: See Figure 3 for an explanation of the different parts of the figure.


Figure 6: Simulation results for parameter set RL: See Figure 3 for an explanation of the different parts of the figure.


Figure 7: Simulation results for parameter set RM: See Figure 3 for an explanation of the different parts of the figure.


Figure 8: Simulation results for each of the five parameter sets: Number of terms stored in the approximate forward integration algorithm, $\sum_{x_{i-1}} \tilde{N}_{i}\left(x_{i-1}\right)$, as a function of node number $i$ when simulating $x_{1: n}, y_{1: n} \mid z_{1: n}$. Note the different horizontal scales.
five models we observe that more terms are required for noisy models. This is also as one should expect, in the extreme case when the noise level goes to infinity the importance of the terms are only decided by the prior. Studying the five graphs separately we observe a growing trend in the number of terms, but it grows much slower than the exponential increase of all terms. We have also tried runs for $n>100$ and the results there supports this conclusion. We also note the large fluctuations in the number of terms and that many of the abrupt changes happens close to where the true $x_{i}$ changes value.

### 4.5. Inversion results

For each of our five parameter sets we repeated the simulation exercise shown in Figures 3 through to 7 ten times, now also randomizing over the value of the true $x_{1: n}$. For each run we used the MCMC results to estimate a confusion matrix $\left[c_{i j}\right]$ where $c_{i j}$ is the average posterior probability for class $j$ in nodes with true class $i$. Table 2 shows the resulting confusion matrices. As one would expect we observe a tendency that higher signal-to-noise ratios gives better predictions. However, LN and RM have comparable performances, demonstrating that the model noise in $y_{1: n} \mid x_{1: n}$ is informative. We also tried an alternative definition of the signal-to-noise ratio by considering variability originating from both $x_{1: n}$ and $y_{1: n}$ as signal, see the discussion in Hammer (2008), but this did not produce a better explanation of the figures in Table 2. Studying the individual confusion matrices we see that shale is

## BC

|  | gas | oil | brine | shale |
| :---: | :---: | :---: | :---: | :---: |
| gas | 0.668 | 0.324 | 0.004 | 0.004 |
| oil | 0.353 | 0.581 | 0.063 | 0.003 |
| brine | 0.005 | 0.064 | 0.891 | 0.040 |
| shale | 0.001 | 0.005 | 0.028 | 0.967 |


|  | LN |  |  |  | MN |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | gas | oil | brine | shale |  | gas | oil | brine | shale |
| gas | 0.837 | 0.163 | 0.000 | 0.000 | gas | 0.611 | 0.279 | 0.102 | 0.009 |
| oil | 0.173 | 0.824 | 0.002 | 0.001 | oil | 0.285 | 0.323 | 0.367 | 0.025 |
| brine | 0.000 | 0.136 | 0.833 | 0.093 | brine | 0.047 | 0.115 | 0.740 | 0.098 |
| shale | 0.007 | 0.003 | 0.013 | 0.977 | shale | 0.015 | 0.014 | 0.130 | 0.842 |


|  | RL |  |  |  |  |  | RM |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | gas | oil | brine | shale |  | gas | oil | brine | shale |  |
| gas | 0.787 | 0.197 | 0.016 | 0.001 | gas | 0.881 | 0.112 | 0.005 | 0.001 |  |
| oil | 0.156 | 0.530 | 0.297 | 0.022 | oil | 0.104 | 0.768 | 0.122 | 0.006 |  |
| brine | 0.006 | 0.059 | 0.918 | 0.017 | brine | 0.001 | 0.033 | 0.924 | 0.042 |  |
| shale | 0.000 | 0.001 | 0.035 | 0.964 | shale | 0.001 | 0.006 | 0.063 | 0.930 |  |

Table 2: Confusion matrices for the various parameter sets. In the tables element $(i, j)$ is the estimated average posterior probability for class $j$ for nodes where the true class is $i$.
most frequently classified correctly in all five cases. Considering Figure 2 this comes as no surprise. The misclassification between gas- and oil-saturated sandstone is significant in all cases. Oil-saturated sandstone is frequently misclassified to brinesaturated sandstone in MN, RL and RM, whereas such a misclassification is less frequent in BC and rare in LN .

## 5. Closing remarks

We have revisited the seismic inversion problem as a hidden Markov model with both continuous and discrete hidden variables. We split the model into a switching linear Gaussian model and a Gaussian linear model. To handle the first part computationally we propose an approximate forward-backward algorithm. In a number of simulation exercises we demonstrate the effectiveness of the approximation and how this makes inversion of the seismic model computationally feasible. The approximate algorithm includes a tuning parameter $\varepsilon$. To choose a value for $\varepsilon$ one must compromise between memory usage and computation time on one side and approximation accuracy on the other. We have found no automatic way to set the value of $\varepsilon$, but our experience is that it is relatively easy to find a reasonable value by trial and error. What makes the choice of $\varepsilon$ non-trivial is that it is used to decide what terms to drop in the forward recursions when information from the data is available from one side only. The importance of the various terms becomes available first when the following backward recursions have been done.

We think the inversion problem in the switching linear Gaussian model for seismic inversion is harder than the problems previously considered for switching linear dynamical systems (Zoeter and Heskes, 2006; Bar-Shalom and Li, 1998) and switching state space models (Barber, 2006). Within an interval with the same value for $x_{1: n}$, the seismic data does not depend on the mean value of the continuous variables. By the difference taken in (37) the mean value of the continuous variables influence the data only when the value of $x_{1: n}$ is changing. This induces larger posterior uncertainty in $x_{1: n}$ and it becomes correspondingly more important to have an approximate forward-backward algorithm that realistically represents this uncertainty. Thus, we think the importance of including more Gaussian terms in the forward recursion is larger for the seismic model than for the cases previously
considered in Zoeter and Heskes (2006), Bar-Shalom and Li (1998) and Barber (2006).

We define an approximate forward recursion by dropping Gaussian terms with small weights. In the references mentioned above an approximation is obtained by taking a single Gaussian density that (approximately) represents the whole Gaussian mixture. It is clearly also possible to define an approximate forward recursion by following an intermediate strategy, finding groups of terms in the Gaussian mixture that have similar mean and covariance and approximate these by a single Gaussian term. However, the computational cost of finding what terms to merge is quadratic in the number of terms, whereas the cost of finding what terms to drop grows linearly with the number of terms. Thus, unless the number of Gaussian terms necessary to obtain a sufficiently good approximation is dramatically reduced when using the merging strategy, our simple dropping strategy is preferable. We have done a little experimentation with the merging strategy for our seismic inversion model, but without success. However, we think the merging strategy may have a potential if the continuous variable $y_{i}$ is univariate.

The focus of the simulation examples of this paper is the computational problem associated with the hidden Markov seismic model. We have not considered inversion of real seismic data. To answer a real inversion problem one must also solve the associated parameter estimation problem. Preliminary experimentation with maximum likelihood estimation from simulated data indicates that it is not possible to estimate all the model parameters only from seismic data. Either one must adopt a Bayesian view with informative priors, or information about (at least some of) the parameters must be obtained from other data sources.

## References

Aki, K. and Richards, P. G. (1980). Quantitative seismology: Theory and methods, W. H. Freeman and Company.

Avseth, P., Mukerji, T. and Mavko, G. (2005). Quantitative seismic interpretation -Applying rock physics tools to reduce interpretation risk, Cambridge University Press.

Bar-Shalom, Y. and Li, X.-R. (1998). Estimation and Tracking: Principles, Techniques and Software, Artech House, Norwood, MA.

Barber, D. (2006). Expectation correction for smoothed inference in switching linear dynamical systems, Journal of machine learning research 7: 2515-2540.

Buland, A., Kolbjørnsen, O. and Omre, H. (2003). Rapid spatially coupled AVO inversion in the Fourier domain, Geophysics 68: 824-836.

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

Bulla, J., Bulla, I. and Nenadić, O. (2010). hsmm - an R package for analyzing hidden semi-Markov models, Computational Statistics $\&$ Data Analysis 54: 611619.

Cappé, O., Moulines, E. and Rydén, T. (2005). Inference in Hidden Markov Models, Springer.

Chen, R. and Liu, J. S. (2000). Mixture Kalman filters, Journal of the Royal Statistical Society, Series B 62: 493-508.

Dearden, R. and Clancy, D. (2002). Particle filters for real-time fault detection in planetary rovers, Proceedings of the 13th International Workshop on Principles of Diagnosis (DX02), pp. 1-6.

Dellaportas, P. and Roberts, G. O. (2003). An introduction to MCMC, in J. Møller (ed.), Spatial Statistics and Computational Methods, number 173 in Lecture Notes in Statistics, Springer, Berlin, pp. 1-41.

Doucet, A., Godsill, S. and Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering, Statistics and Computing 10: 197-208.

Ghahramani, Z. and Hinton, G. E. (1998). Variational learning for switching statespace models, Neural Computation 12: 963-996.

Godsill, S. J., Doucet, A. and West, M. (2004). Monte Carlo smoothing for nonlinear time series, Journal of the American Statistical Association 99: 156-168.

Guédon, Y. (2007). Exploring the state sequence space for hidden Markov and semi-Markov chains, Computational Statistics \& Data Analysis 51: 2379-2409.

Hammer, H., Kolbjørnsen, O., Tjelmeland, H. and Buland, A. (2010). Lithology and fluid prediction from prestack seismic data using a Bayesian model with Markov process prior, Technical report, Norwegian Computing Center. In preparation.

Hammer, H. L. (2008). Topics in stochastic simulation, with an application to seismic inversion, PhD thesis, Norwegian University of Science and Technology. Thesis number 73:2008.

Kim, C. J. (1994). Dynamic linear models with Markovswitching, Journal of Econometrics 60.

Künsch, H. (2000). State space models and hidden Markov models, in O. BarndorffNielsen, D. Cox and C. Klüppelberg (eds), Complex Stochastic Systems, number 87 in Monographs on Statsitics and Applied Probability, Chapman \& Hall/CRC.

Larsen, A. L., Ulvmoen, M., Omre, H. and Buland, A. (2006). Bayesian lithology/fluid prediction and simulation on the basis of a Markov-chain prior model, Geophysics 71 issue 5: R69-R78.

Lerner, U., Parr, R., Koller, D. and Biswas, G. (2000). Bayesian fault detection and diagnosis in dynamic systems, In Proc. AAAI, pp. 531-537.

MacDonald, I. and Zucchini, W. (1997). Hidden Markov and Other Models for Discrete-Valued Time Series, Chapman and Hall.

Oh, S. M., Rehg, J. M., Balch, T. and Dallaert, F. (2005). Data-driven memc for learning and inference in switching linear dynamic systems, Proc. 20th National Conference on Artificial Intelligence (AAAI-2005), Pittsburgh, PA.

Roberts, G. O., Gelman, A. and Gilks, W. R. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms, Annals of Applied Probability 7: 110-120.

Roberts, G. O. and Rosenthal, J. S. (1998). Optimal scaling of discrete approximations to Langevin diffusions, Journal of the Royal Statistical Society. Series B 60: 255-268.

Rosti, A.-V. and Gales, M. (2004). Rao-blackwellised Gibbs sampling for switching linear dynamical systems, In Intl. Conf. Acoust., Speech, and Signal Proc. (ICASSP) 1: 809-812.

Scott, A. L. (2002). Bayesian methods for hidden Markov models: Recursive compution in the 21st century, Journal of the American Statistical Association 97: 337-351.

Sheriff, R. E. and Geldart, L. P. (1995). Exploration Seismology, Cambridge University Press.

Smith, A. and Roberts, G. O. (1993). Bayesian computation via the Gibbs sampler and related Markov chain Monte Carlo methods (with discussion), Journal of the Royal Statistical Society, Series B 55: 2-24.

Ulvmoen, M. and Hammer, H. (2009). Bayesian lithology/fluid inversion - comparison of two algorithms, Computational Geosciences 14: 357-367.

Zoeter, O. and Heskes, T. (2006). Deterministic approximate inference techniques for conditionally Gaussian state space models, Statistics and Computing 16: 279292.


[^0]:    * Corresponding author
    ${ }^{1}$ Present address: Faculty of Engineering, Oslo University College, PO box 4 St. Olavs plass, 0130 Oslo, Norway. Phone number: +4745210 184, Fax: $(+47) 22453205$ (Att: Hugo Hammer)

