Switching State-Space Models

Zoubin Ghahramani
Geoffrey E. Hinton
Department of Computer Science
University of Toronto
6 King's College Road
Toronto, Canada M5S 3H5

Email: zoubin@cs.toronto.edu

Abstract

We introduce a statistical model for times series data with nonlinear dynamics which iteratively segments the data into regimes with approximately linear dynamics and learns the parameters of each of those regimes. This model combines and generalizes two of the most widely used stochastic time series models—the hidden Markov model and the linear dynamical system—and is related to models that are widely used in the control and econometrics literatures. It can also be derived by extending the mixture of experts neural network model (Jacobs et al., 1991) to its fully dynamical version, in which both expert and gating networks are recurrent. Inferring the posterior probabilities of the hidden states of this model is computationally intractable, and therefore the exact Expectation Maximization (EM) alogithm cannot be applied. However, we present a variational approximation which maximizes a lower bound on the log likelihood and makes use of both the forward–backward recursions for hidden Markov models and the Kalman filter recursions for linear dynamical systems.

1 Introduction

Most commonly used probabilistic models of time series can draw their lineage to either the hidden Markov model (HMM) or to the stochastic linear dynamical system, also known as the state-space model (SSM). Hidden Markov models represent information about the past of a sequence through a single discrete random variable–the hidden state. The probability distribution of this state is a function of the previous state represented by a stochastic transition matrix. Knowing the state at any time renders the past, present and future observations to be statistically independent. This is the *Markov* independence property that gives the model its name.

State-space model represent information about the past though a real-valued hidden state vector. Again, conditioned on this state vector, the past, present, and future observations are rendered independent. The dependency between the present state vector and the previous state vector is specified through the dynamic equations of the system and the noise model. When these equations are linear and the noise model is Gaussian, the state-space model is also known as linear dynamical system or Kalman filter model.

Unfortunately, many real-world processes cannot be characterized by either purely discrete or purely linear–Gaussian dynamics. For example, an industrial plant may have multiple discrete modes of behavior, each of which is appropriately described by linear dynamics. Similarly, the pixel intensities in an image of a translating object vary according to linear dynamics for subpixel translations, but as the image moves over a larger range the dynamics change significantly and nonlinearly.

The goal of this paper is to model complex dynamical phenomena which may be characterized by both discrete and continuous dynamics. To this effect, we introduce a probabilistic model called the switching state-space model inspired by the divide-and-conquer principle underlying the mixture of experts neural network (Jacobs et al., 1991). The switching state-space model is a natural generalization of both the hidden Markov model and the state-space model in which the dynamics can transition in a discrete manner from one linear operating regime to another. There is in face a large literature on models of this kind in econometrics, signal processing, and other fields (Harrison and Stevens, 1976; Chang and Athans, 1978; Hamilton, 1989; Shumway and Stoffer, 1991; Bar-Shalom and Li, 1993). In this paper we extend some of these models to allow for multiple real-valued state vectors, draw connections between these fields and the literature on neural computation, and derive a learning algorithm for all the parameters of the system based on a structured variational approximation which rigorously maximizes a lower bound on the log likelihood of the model.

The paper is organized as follows. In the following section we review the background material on state-space models, hidden Markov models, and recent hybrids of the two. In section 3, we describe the generative model—i.e. the probability distribution defined over the observation sequences—for switching state-space models. In section 4, we describe the learning algorithm for switching state-space models which is based on a structured variational approximation to the EM algorithm. In section 5 we present simulation results both in an artificial domain, to assess the quality of the approximate inference method, and in a natural domain. Finally, we conclude with section 6.

2 Background

2.1 State-space models

In a state-space model, a sequence of D-dimensional observation vectors $\{Y_t\}$ where the discrete time index t ranges from $t = 1, \ldots T$, is modeled by specifying (1) a probabilistic relation between the observations and a hidden state vector X_t, and (2) a probabilistic relation between consecutive hidden state vectors.[1] The hidden state vectors obey the Markov independence property, so the joint probability for the sequences of states X_t and observations Y_t can be factored as:

$$P(\{X_t, Y_t\}) = P(X_1)P(Y_1|X_1)\prod_{t=2}^{T} P(X_t|X_{t-1})P(Y_t|X_t), \tag{1}$$

The conditional independences specified by equation (1) can be expressed graphically in the form of Figure 1a. The simplest and most commonly used models of this kind assume that

[1] A table describing the variables and the notation used throughout the paper is provided in Appendix A.

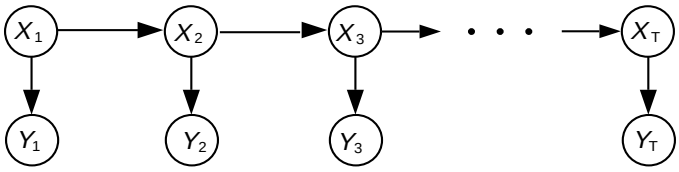

Figure 1: A directed acyclic graph (DAG) specifying conditional independence relations for a state-space model. Each node is conditionally independent from its non-descendents given its parents: The output Y_t is conditionally independent from all other variables given the state X_t; and X_t is conditionally independent from $X_1, \ldots, X_{t-2}$ given X_{t-1}.

the the transition and output functions are linear and time-invariant and the distribution of the state and observation variables is multivariate Gaussian. We will use the term state-space model to refer to this simple form of the model. For such models, the state transition function is written

$$X_t = AX_{t-1} + w_t \tag{2}$$

where A is the state transition matrix and w_t is zero-mean Gaussian state noise. Equation (2) ensures that if $P(X_{t-1})$ is Gaussian, then so is $P(X_t)$. $P(X_1)$ is assumed to be Gaussian. The output function is written

$$Y_t = CX_t + v_t \tag{3}$$

where C is the output matrix and v_t is zero-mean Gaussian output noise with covariance matrix R; $P(Y_t|X_t)$ is therefore also Gaussian:

$$P(Y_t|X_t) = (2\pi)^{-D/2}|R|^{-1/2}\exp\left\{-\frac{1}{2}(Y_t - CX_t)' R^{-1} (Y_t - CX_t)\right\}. \tag{4}$$

Often, the observation vector can be divided into input (or predictor) variables and output (or response) variables. To model the input–output behavior of such a system—i.e. the conditional probability of output sequences given input sequences—the linear Gaussian SSM can be modified to have a state-transition function

$$X_t = AX_{t-1} + BU_t + w_t, \tag{5}$$

where U_t is the input observation vector and B is the (fixed) input matrix.[2]

The problem of *inference* for a state-space model with known parameters consists of estimating the posterior probabilities of the hidden variables given a sequence of observed variables. Since the local likelihood functions for the observations are Gaussian and the priors for the hidden states are Gaussian, the resulting posterior is also Gaussian. The special cases of the inference problem for state-space models play a prominent role in the engineering literature: filtering, smoothing, and prediction (Anderson and Moore, 1979; Goodwin and Sin, 1984). The goal of *filtering* is to compute the probability of the current hidden state X_t given the sequence of inputs and outputs up to time t—$P(X_t|\{Y\}_1^t, \{U\}_1^t)$.[3] The recursive algorithm used to perform this computation is known as the *Kalman filter* (Kalman and

[2]One can also define the state such that $X_{t+1} = AX_t + BU_t + w_t$.

[3]The notation $\{Y\}_1^t$ is short-hand for the sequence $Y_1, \ldots, Y_t$.

Bucy, 1961). The goal of *smoothing* is to compute the probability of X_t given the sequence of inputs and outputs up to time T, where $T > t$. The Kalman filter recursions are used in the forward direction to compute the probability of X_t given $\{Y\}_1^t$ and $\{U\}_1^t$. A similar set of *backward* recursions from T to t complete the computation by accounting for the observations after time t (Rauch, 1963). We will refer to the combined forward and backward recursions for smoothing as the Kalman smoothing recursions (also known as the RTS or Rauch-Tung-Streibel smoother). Finally, the goal of *prediction* is to compute the probability of future states and observations given observations upto time t. Given $P(X_t|\{Y\}_1^t, \{U\}_1^t)$ computed as before, the model is simulated in the forward direction using equations (2) (or (5) if there are inputs) and (3) to compute the probability density of the state or output at future time $t + \tau$.

The problem of *learning* the parameters of a state-space model is known in engineering as the *system identification* problem, and in its most general form assumes access only to sequences of input and output observations. We focus on maximum likelihood learning, in which a single (locally optimal) value of the parameters is estimated, rather than Bayesian approaches which treat the parameters as random variables and compute or approximate the posterior distribution of the parameters given the data. One can also distinguish between on-line and off-line approaches to learning. On-line recursive algorithms, favored in real-time adaptive control applications, can be obtained by computing the gradient or the second derivatives of the log likelihood (Ljung and Söderström, 1983). Similar gradient-based methods can be obtained for off-line methods. An alternative method for off-line learning makes use of the Expectation Maximization (EM) algorithm (Dempster et al., 1977). This procedure iterates between a step that fixes the current parameters and computes posterior probabilities over the hidden states given the observations (the E-step), and a step that uses these probabilities to maximize the expected log likelihood of the parameters (the M-step). For linear Gaussian state-space models, the E-step is exactly the Kalman smoothing problem as defined above, and the M-step simplifies to a linear regression problem (Shumway and Stoffer, 1982; Digalakis et al., 1993). Details on the EM algorithm for state-space models can be found in Ghahramani and Hinton (1996b), as well as in the original Shumway and Stoffer (1982) paper.

It is worth pointing out that the linear Gaussian state-space model is a generalization of a statistical method known as *factor analysis*. Factor analysis models high dimensional data through a smaller number of latent variables or factors (Everitt, 1984). The model relating the factors to the observations is exactly as specified by equation (3): X_t is a Gaussian distributed vector of factor values; Y_t is the observation vector; C is known as the *factor loading matrix*, and v_t is zero-mean Gaussian distributed noise with the further constraint that the elements of the vector v_t are uncorrelated. State-space models are therefore a dynamic generalization of factor analysis which allow the current factor values to depend linearly on the previous factor values.[4]

[4] X must have fewer dimensions that Y for the factor analysis problem to be well-posed. For a state-space model, the equivalent constraint is that the dimensionality of X must be less than the product of the dimension of Y and the length of the observation sequence. This constraint derives from the notion of *observability* in linear system theory (Goodwin and Sin, 1984).

2.2 Hidden Markov models

A hidden Markov models defines a probability distribution over sequences of observations $\{Y_t\}$. This distribution over sequences is obtained by specifying the probability over observations at each time step t given a *discrete* hidden state S_t, and the probability of transitioning from one hidden state to another from one time step to the next. Using the Markov property, the joint probability for the sequences of states S_t and observation Y_t, can be factored in exactly the same manner as equation (1), with S_t taking the place of X_t:

$$P(\{S_t, Y_t\}) = P(S_1)P(Y_1|S_1)\prod_{t=2}^{T} P(S_t|S_{t-1})P(Y_t|S_t). \tag{6}$$

Similarly, the conditional independences in an HMM can be expressed graphically in the same form as Figure 1a. The state is represented by a single multinomial variable that can take one of K discrete values, $S_t \in \{1, \ldots, K\}$. The state transition probabilities, $P(S_t|S_{t-1})$, are specified by a $K \times K$ transition matrix. If the observables are discrete symbols taking on one of L values, the observation probabilities $P(Y_t|S_t)$ can be fully specified as a $K \times L$ observation matrix. For a continuous observation vector, $P(Y_t|S_t)$ can be modeled in many different forms, such as a Gaussian, mixture of Gaussians, or a neural network. HMMs have been applied extensively to problems in speech recognition (Juang and Rabiner, 1991), computational biology (Baldi et al., 1994), and fault detection (Smyth, 1994).

Given an HMM with known parameters and a sequence of observations, two algorithms are commonly used to solve two different forms of the inference problem (Rabiner and Juang, 1986). The first computes the posterior probabilities of the hidden states using a recursive algorithm known as the *forward–backward* algorithm. The computations in the forward pass are exactly analogous to the Kalman filter for SSMs, while the computations in the backward pass are analogous to the backward pass of the Kalman smoothing equations. As noted by Bridle (personal communication, 1985) and Smyth, Heckerman and Jordan (1997), the forward–backward algorithm is a special case of exact inference algorithms for more general graphical probabilistic models (Lauritzen and Spiegelhalter, 1988; Pearl, 1988). The same observation holds true for the Kalman smoothing recursions. The other inference problem commonly posed for HMMs is to compute the single most likely sequence of hidden states. The solution to this problem is given by the *Viterbi* algorithm, which also consists of a forward and backward pass through the model.

To learn maximum likelihood parameters for an HMM given sequences of observations, one can use the well-known *Baum-Welch* algorithm (Baum et al., 1970). This algorithm is a special case of EM that uses the forward–backward algorithm to infer the posterior probabilities of the hidden states in the E-step. The M-step uses expected counts of transitions and observations to re-estimate the transition and output matrices (or linear regression equations in the case where the observations are Gaussian distributed).

The HMM can be augmented to allow for input variables, such that it models the conditional distribution of sequences of output observations given sequences of inputs (Cacciatore and Nowlan, 1994; Bengio and Frasconi, 1995; Meila and Jordan, 1996). The approach used in Bengio and Frasconi's Input Output HMMs (IOHMMs) suggests modeling $P(S_t|S_{t-1}, U_t)$, where U_t is the input, as M separate neural networks, one for each setting of S_{t-1}. This de-

composition ensures that, if a sum-to-one constraint is used on the output of these networks, a valid probability transition matrix is defined at each point in input space.

2.3 Hybrids

For many time series applications, neither the linear Gaussian dynamics of the SSM nor the purely discrete dynamics of the HMM can appropriately model the temporal structure of the data. As a consequence, in fields ranging from econometrics to control engineering, an burgeoning literature has developed on models which combine the discrete transition structure of HMMs with the linear dynamics of SSMs (Harrison and Stevens, 1976; Chang and Athans, 1978; Hamilton, 1989; Shumway and Stoffer, 1991; Bar-Shalom and Li, 1993; Deng, 1993; Kadirkamanathan and Kadirkamanathan, 1996; Chaer et al., 1997). These models are known alternately as hybrid models, state-space models with switching, and jump-linear systems. We briefly review some of the main results in this literature including some recent proposals in the field of neural computation.

The engineering literature on state estimation for state-space models with switching is reviewed in Bar-Shalom and Li (1993). The state estimation problem consists of computing the mean and covariance of the hidden real-valued state vector given the observations (i.e. the filering problem). Shortly after Kalman's results on linear Gaussian state-space models, much attention turned to the problem of state estimation with switching parameters. For example, Ackerson and Fu (1970) consider the problem of state estimation in linear state-space models which receive (unobserved) state and output disturbances coming from Gaussian mixture distributions with Markov transition structure. Chang and Athans (1978) derive the equations for computing the conditional mean and variance of the state when the parameters of a linear state-space model switch according to arbitrary and Markovian dynamics. The prior and transition probabilities of the switching process are assumed to be known. They note that for M models (sets of parameters) and an observation length T, the exact conditional distribution of the state is a Gaussian mixture with M^T components. The conditional mean and variance, which require far less computation, are therefore only summary statistics.

Shumway and Stoffer (1991) consider the problem of learning the parameters of state-space models with a single real-valued hidden state vector and switching output matrices. The probability of chosing a particular output matrix is a pre-specified time-varying function, independent of previous choices. A pseudo-EM algorithm is derived in which the E-step, which in its exact form would require computing a Gaussian mixture with M^T components, is approximated by a single Gaussian at each time step. Kim (1994) extends this to the case where both the state dynamics and the output matrices switch, and where the switching follows Markovian dynamics. Kim uses an approximation in which the exponential Gaussian mixture is collapsed down to M Gaussians at each time step. Other authors have used Markov chain Monte Carlo methods for state and parameter estimation in switching models (Carter and Kohn, 1994; Athaide, 1995) and in other more general dynamic probabilistic networks (Dean and Kanazawa, 1989; Kanazawa et al., 1995).

One can also model nonlinear processes using nonlinear generalizations of the state-space model which do not explicitly representing a switching state. The conditional mean and

variance of the hidden states can then be estimated via the extended Kalman filter (EKF; Goodwin and Sin, 1984). At each time step, the EKF linearizes the system dynamics about the current state estimate and uses the resulting Kalman Filter to estimate the next state. Like the above approaches for hybrid models, this approach therefore also approximates the non-Gaussian state distribution at each time step with a Gaussian. EKF-based approaches have recently been used to address difficult problems in computer vision, such as simultaneous recognition and pose estimation (Rao and Ballard, 1996). The Gaussian approximation can be improved upon by representing non-Gaussian state distributions via a set of samples which are stochastically propagated and reweighted. This approach has been successfully applied to the problem of contour tracking in computer vision (Isard and Blake, 1996; Blake et al., 1995). We have explored elsewhere the use of the EKF in deriving an EM algorithm for general stochastic nonlinear dynamical systems (Ghahramani and Roweis, in preparation). Switching state-space models can be viewed as a nonlinear EKFs in which a fixed number of linearizations are fit to the dynamics. Using the maximum likelihood criterion, the learning algorithm iterates between selecting the points about which such a linearization is least costly and fitting each linear model.

Another related proposal comes from Fraser and Dimitriadis (1993), who combine real-valued and discrete states in a system they call a hidden filter HMM (HFHMM). As a simplification these authors assume that the real-valued state is a known deterministic function of the past observations (i.e. an embedding). The work in this paper departs significantly from Fraser and Dimitriadis' work in several ways. First, the state-space embedding in our paper is learned rather than fixed a-priori. Second, the real-valued state vector is assumed to be a random variable, therefore allowing the propagation of uncertainty in the state. Finally, there are M real-valued state vectors, simultaneously representing M hypothesis that are competing to explain the observations.

With regard to the literature on neural computation, the model presented in this paper is a generalization of the mixtures of experts architecture (Jacobs et al., 1991; Jordan and Jacobs, 1994).[5] Previous dynamical generalizations of the mixture of experts architecture consider the case in which the gating network has Markovian dynamics (Cacciatore and Nowlan, 1994; Kadirkamanathan and Kadirkamanathan, 1996; Meila and Jordan, 1996). One limitation of this generalization is that the entire past sequence is summarized in the value of a single discrete variable (the gating activation), which for a system with M experts can convey on average at most $\log M$ bits of information about the past. In the generalization we consider in this paper both the experts and the gating network have Markovian dynamics. The past is therefore summarized by a state composed of the cross-product of this discrete variable with the combined real-valued state-space of all the experts. This provides a much wider information channel from the past. One advantage of this representation is that the real-valued state can contain componential structure. Thus, attributes such as the position, orientation, and scale of an object in an image, which are most naturally encoded as independent real-valued variables, can be accommodated in the state without the exponential growth required of discrete HMM-like representations.

Before we proceed with the definition of the probabilistic model, it is important to place

[5] It can also be seen as a generalization of mixtures of factor analyzers (Hinton et al., 1996; Ghahramani and Hinton, 1996b).

the work in this paper in the context of the literature we have just reviewed. "Hybrid models", state-space with switching and jump-linear systems all assume that there is a single real-valued state vector. The model considered in this paper generalizes this to multiple real-valued state vectors.[6] We present a learning algorithm for all of the parameters of the model, including the Markov switching parameters. Using a structured variational approximation (Saul and Jordan, 1996), we show that this algorithm maximizes a strict lower bound on the log likelihood of the data, rather than a heuristically motivated pseudo-likelihood. The resulting algorithm has a simple and intuitive flavor: It decouples into forward-backward recursions on a hidden Markov model, and Kalman smoothing recursions on each state-space model. The states of the HMM determine the soft assignment of each observation to a state-space model; the prediction errors of the state-space models determine the observation probabilities for the HMM.

3 The Generative Model

In switching state-space models, the sequence of observations $\{Y_t\}$ is modeled by specifying a probabilistic relation between the observations and a hidden state space comprising M real-valued state vectors, $X_t^{(m)}$, and one discrete state vector S_t. The discrete state, S_t, is modeled as a multinomial variable that can take on M values: $S_t \in \{1, \ldots, M\}$; for reasons that will become obvious we refer to it as the *switch* variable. The joint probability of observations and hidden states can be factored as

$$\begin{aligned} P(\{S_t, X_t^{(1)}, \ldots, X_t^{(M)}, Y_t\}) &= P(S_1) \prod_{t=2}^{T} P(S_t|S_{t-1}) \cdot \prod_{m=1}^{M} P(X_1^{(m)}) \prod_{t=2}^{T} P(X_t^{(m)}|X_{t-1}^{(m)}) \\ &\quad \cdot \prod_{t=1}^{T} P(Y_t|X_t^{(1)}, \ldots, X_t^{(M)}, S_t), \end{aligned} \tag{7}$$

which corresponds graphically to the conditional independences represented by Figure 2. Conditioned on a setting of the switch state, $S_t = m$, the observable is multivariate Gaussian with output equation given by state-space model m. The probability of the observation vector Y_t is therefore

$$P(Y_t|X_t^{(1)}, \ldots, X_t^{(M)}, S_t = m) = (2\pi)^{-\frac{D}{2}} |R|^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}\left(Y_t - C^{(m)}X_t^{(m)}\right)' R^{-1}\left(Y_t - C^{(m)}X_t^{(m)}\right)\right\} \tag{8}$$

where D is the dimension of the observation vector, R is the observation noise covariance matrix, and $C^{(m)}$ is the output matrix for state-space model m (cf. equation (4) for a single linear-Gaussian state-space model). Each real-valued state vector evolves according to the linear Gaussian dynamics of a state-space model with differing initial state, transition matrix, and state noise (equation (2)). The switch state itself evolves according to the discrete Markov transition structure specified by the initial state probabilities $P(S_1)$ and the $M \times M$ state transition matrix $P(S_t|S_{t-1})$.

[6] Note that the state vectors could be concatenated into one large state vector with factorized (block-diagonal) transition matrices (cf. factorial hidden Markov model; Ghahramani and Jordan, 1997). However, this obscures the decoupled structure of the model.

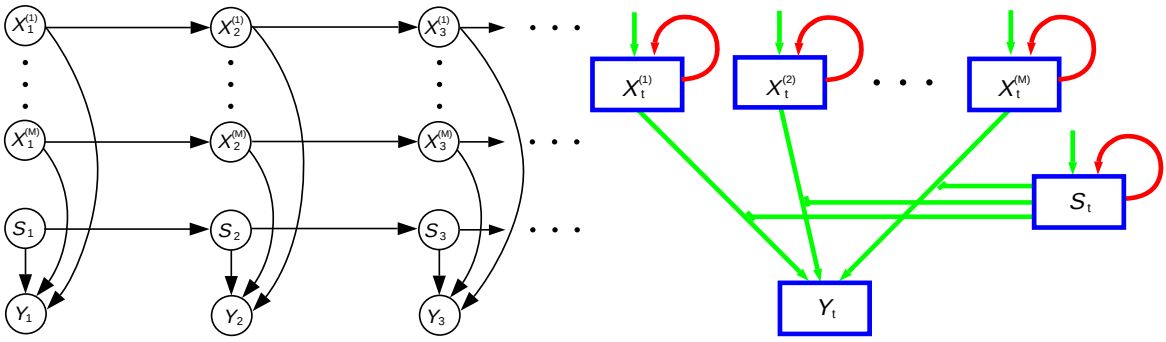

Figure 2: a) Graphical model representation for switching state-space models. S_t is the discete switch variable and $X_t^{(m)}$ are the real-valued state vectors. b) Switching state-space model depicted as a generalization of the mixture of experts. The light arrows correspond to the connections in a mixture of experts. In a switching state-space model, the states of the experts and of the gating network also depend on their previous states (dark arrows).

An exact analogy can be made to the "mixture of experts" architecture for modular learning in neural networks (figure 2b; Jacobs et al, 1991). Each state space model is a linear expert with Gaussian output noise model and linear-Gaussian dynamics. The switch state "gates" the outputs of the M state-space models, and therefore plays the role of a gating network with Markovian dynamics.

We also wish to consider the following three straightforward extensions of the above model:

(Ex1) Differing output covariances, $R^{(m)}$, for each state-space model;

(Ex2) Differing output means, $\mu_Y^{(m)}$, for each state-space model, such that each model is allowed to capture observations in a different operating range;

(Ex3) Conditioning on a sequence of observed input vectors, $\{U_t\}$.

Other extensions are also possible.

4 Learning

An efficient learning algorithm for the parameters of a switching state-space model can be derived by generalizing the Expectation Maximization (EM) algorithm (Baum et al., 1970; Dempster et al., 1977). EM alternates between optimizing a distribution over the hidden states (the E-step) and optimizing the parameters given the distribution over hidden states (the M-step). Any distribution over the hidden states, $Q(\{S_t, X_t\})$, where $X_t = [X_t^{(1)}, \ldots X_t^{(M)}]$ is the combined state of the state-space models, can be used to define a lower bound, $\mathcal{B}$, on the log probability of the observed data:

$$\log P(\{Y_t\}|\theta) \quad = \quad \log \sum_{\{S_t\}} \int P(\{S_t, X_t, Y_t\}|\theta)\, d\{X_t\} \tag{9}$$

$$= \log \sum_{\{S_t\}} \int Q(\{S_t, X_t\}) \left[\frac{P(\{S_t, X_t, Y_t\}|\theta)}{Q(\{S_t, X_t\})} \right] d\{X_t\} \quad (10)$$

$$\geq \sum_{\{S_t\}} \int Q(\{S_t, X_t\}) \log \left[\frac{P(\{S_t, X_t, Y_t\}|\theta)}{Q(\{S_t, X_t\})} \right] d\{X_t\} = \mathcal{B}(Q, \theta), \quad (11)$$

where θ denotes the parameters of the model and we have made use of Jensen's inequality (Cover and Thomas, 1991) to establish (11). Both steps of EM increase the lower bound on the log probability of the observed data. The E-step holds the parameters fixed and sets Q to be the posterior distribution over the hidden states given the parameters,

$$Q(\{S_t, X_t\}) = P(\{S_t, X_t\}|\{Y_t\}, \theta). \quad (12)$$

This maximizes $\mathcal{B}$ with respect to the distribution, turning the lower bound into an equality. The M-step holds the distribution fixed and computes the parameters that maximize $\mathcal{B}$ for that distribution. Since $\mathcal{B} = \log P(\{Y_t\}|\theta)$ at the start of the M-step, and since the E-step does not affect $\log P$, the two steps combined can never decrease $\log P$. Given the change in the parameters produced by the M-step, the distribution produced by the previous E-step is typically no longer optimal, so the whole procedure must be iterated.

Unfortunately, the exact E-step for switching state-space models is intractable. Like the related hybrid models described in section 2.3, the posterior probability of the real-valued states is a Gaussian mixture with M^T terms. This can be seen with reference to the semantics of directed graphs, in particular the d-separation criterion (Pearl, 1988), which implies that the hidden state variables in Figure 2, while marginally independent, become conditionally dependent given the observation sequence. This induced dependency effectively couples all of the real-valued hidden state variables to the discrete switch variable, as a consequence of which the exact posteriors become Gaussian mixtures with an exponential number of terms.[7]

In order to derive an efficient learning algorithm for this system, we relax the EM algorithm by approximating the posterior probability of the hidden states. The basic idea is that, since expectations with respect to P are intractable, rather than setting $Q = P$ in the E-step, a tractable distribution Q is used to *approximate* P. This results in an EM learning algorithm which maximizes a lower bound on the log likelihood. The difference between the bound $\mathcal{B}$ and the log likelihood is given by the Kullback-Liebler (KL) divergence between Q and P (Cover and Thomas, 1991):

$$\mathrm{KL}(Q\|P) = \sum_{\{S_t\}} \int Q(\{S_t, X_t\}) \log \left[\frac{Q(\{S_t, X_t\})}{P(\{S_t, X_t\}|\{Y_t\})} \right] d\{X_t\}. \quad (13)$$

Since the complexity of exact inference in the approximation given by Q is determined by its conditional independence relations, not by its parameters, we can choose Q to have a tractable structure—a graphical representation which eliminates some of the dependencies in P. Given this structure, the parameters of Q are varied to obtain the tightest possible bound by minimizing (13). Therefore, the algorithm alternates between optimizing the parameters

[7] The intractability of the E-step or smoothing problem in the simpler single-state switching model has also been noted in the engineering literature (Chang and Athans, 1978; Bar-Shalom and Li, 1993).

of the distribution Q to minimize (13) (the E-step) and optimizing the parameters of P given the distribution over the hidden states (the M-step). Like in exact EM, both steps increase the lower bound $\mathcal{B}$ on the log likelihood, however equality is not reached in the E-step.

We will refer to the general strategy of using a parameterized approximating distribution as a *variational approximation* and refer to the free parameters of the distribution as *variational parameters*. A completely factorized approximation is often used in statistical physics, where it provides the basis for simple yet powerful *mean field approximations* to statistical mechanical systems (Parisi, 1988). Theoretical arguments motivating approximate E-steps were presented by Neal and Hinton (1993). Saul and Jordan (1996) showed that approximate E-steps could be used to maximize a lower bound on the log likelihood, and proposed the powerful technique of *structured* variational approximations to intractable probabilistic networks. The key insight of Saul and Jordan's work, which the present paper makes use of, is that by judicious use of an approximation Q, exact inference algorithms can be used on the tractable substructures in an intractable network. A general tutorial on variational approximations can be found in Jordan et al. (1998).

The parameters of the switching state-space model are $\theta = \{A^{(m)}, C^{(m)}, \mathcal{Q}^{(m)}, \mu_{X_1}^{(m)}, \mathcal{Q}_1^{(m)}, R, \boldsymbol{\pi}, \Phi\}$, where $A^{(m)}$ is the state dynamics matrix for model m, $C^{(m)}$ is its output matrix, $\mathcal{Q}^{(m)}$ is its state noise covariance, $\mu_{X_1}^{(m)}$ is the mean of the initial state, $\mathcal{Q}_1^{(m)}$ is the covariance of the initial state, R is the (tied) output noise covariance, $\boldsymbol{\pi} = P(S_1)$ is the prior for the discrete Markov process, and $\Phi = P(S_t|S_{t-1})$. Inclusion of extensions (Ex1)–(Ex3) would result in substituting R for $R^{(m)}$, adding means $\mu_Y^{(m)}$ and input matrices $B^{(m)}$.

While there are many possible approximations to the posterior distribution of the hidden variables that one could use for learning and inference in switching state-space models, we focus on the following:

$$Q(\{S_t, X_t\}) = \frac{1}{Z_Q} \left[\psi(S_1) \prod_{t=2}^{T} \psi(S_{t-1}, S_t)\right] \prod_{m=1}^{M} \psi(X_1^{(m)}) \prod_{t=2}^{T} \psi(X_{t-1}^{(m)}, X_t^{(m)}), \tag{14}$$

where the ψ are non-negative *potential* functions which we will define soon, and Z_Q is a normalization constant ensuring that Q integrates to one. Although Q has been written in terms of potential functions rather than conditional probabilities, it corresponds to the simple graphical model shown in Figure 3. The terms involving the switch variables S_t define a discrete Markov chain and the terms involving the state vectors $X_t^{(m)}$ define M *uncoupled* state-space models. Like in mean field approximations we have approximated the stochastically coupled system by removing some of the couplings of the original system. Specifically, we have removed the stochastic coupling between the chains that results from the fact that the observation at time t depends on all the hidden variables at time t. However, we retain the coupling between the hidden variables at successive time steps since these couplings can be handled exactly using the forward–backward and Kalman smoothing recursions. This approximation is therefore structured, in the sense that not all variables are uncoupled.

The discrete switching process is defined by

$$\psi(S_1 = m) = P(S_1 = m)\, q_1^{(m)} \tag{15}$$

$$\psi(S_{t-1}, S_t = m) = P(S_t = m|S_{t-1})\, q_t^{(m)}, \tag{16}$$

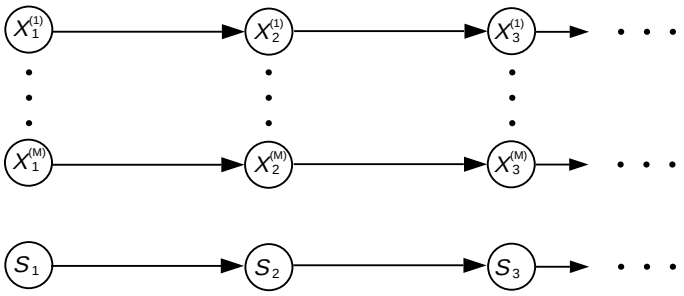

Figure 3: Graphical model representation for the structured variational approximation to the posterior distribution of the hidden states of a switching state-space model.

where the $q_t^{(m)}$ are variational parameters of the Q distribution. These parameters scale the probabilities of each of the states of the switch variable at each time step, so that $q_t^{(m)}$ plays exactly the same role as the observation probability $P(Y_t|S_t = m)$ would play in a regular hidden Markov model. We will soon see that minimizing $KL(Q\|P)$ results in an equation for $q_t^{(m)}$ which supports this intuition.

The uncoupled state-space models in the approximation Q are also defined by potential functions which are related to probabilities in the original system. It is easier to express the log potentials:

$$\log \psi(X_1^{(m)}) = \log P(X_1^{(m)}) + h_1^{(m)} \log P(Y_1|X_1^{(m)}, S_1 = m) \tag{17}$$

$$\log \psi(X_{t-1}^{(m)}, X_t^{(m)}) = \log P(X_t^{(m)}|X_{t-1}^{(m)}) + h_t^{(m)} \log P(Y_t|X_t^{(m)}, S_t = m), \tag{18}$$

where the $h_t^{(m)}$ are also variational parameters of Q. The vector h_t plays a role very similar to the switch variable S_t. Each component $h_t^{(m)}$ can range between 0 and 1. When $h_t^{(m)} = 0$ the posterior probability of $X_t^{(m)}$ under Q does not depend on the observation at time Y_t. When $h_t^{(m)} = 1$, the posterior probability of $X_t^{(m)}$ under Q includes a term which assumes that state-space model m generated Y_t. Therefore, $h_t^{(m)}$ is the *responsibility* assigned to state-space model m for the observation vector Y_t. The difference between $h_t^{(m)}$ and $S_t^{(m)}$ is that $h_t^{(m)}$ is a deterministic parameter, while $S_t^{(m)}$ is a stochastic random variable.

To maximize the lower bound on the log likelihood, $\mathrm{KL}(Q\|P)$ is minimized with respect to the variational parameters $h_t^{(m)}$ and $q_t^{(m)}$ separately for each sequence of observations. Using the definition of P for the switching state-space model (equation (7) and (8)) and the approximating distribution Q, the minimum of KL satisfies the following fixed point equations for the variational parameters (see Appendix B):

$$h_t^{(m)} = Q(S_t = m) \tag{19}$$

$$q_t^{(m)} = \exp\left\{-\frac{1}{2}\left\langle \left(Y_t - C^{(m)}X_t^{(m)}\right)' R^{-1} \left(Y_t - C^{(m)}X_t^{(m)}\right)\right\rangle\right\} \tag{20}$$

where $\langle\cdot\rangle$ denotes expectation over the Q distribution. Intuitively,

- $h_t^{(m)}$ is the *responsibility* assigned to state-space model m for observation vector Y_t

- $q_t^{(m)}$ is a function of the *expected squared error* if state-space model m were to generate Y_t

To compute $h_t^{(m)}$ it is necessary to sum Q over all the S_τ variables not including S_t. This can be done efficiently using the forward–backward algorithm on the switch state variables, with $q_t^{(m)}$ playing exactly the same role as an observation probability associated with each setting of the switch variable. Since $q_t^{(m)}$ is related to the prediction error of model m on data Y_t, this has the intuitive interpretation that the switch state associated with models with smaller expected prediction error on a particular observation will be favored at that time step. However, the forward–backward algorithm ensures that the final responsibilities for the models are obtained after considering the entire sequence of observations.

To compute $q_t^{(m)}$ it is necessary to calculate the expectations of $X_t^{(m)}$ and $X_t^{(m)} X_t^{(m)'}$ under Q:

$$q_t^{(m)} = \exp\left\{ -\frac{1}{2} Y_t' R^{-1} Y_t + Y_t' R^{-1} C^{(m)} \left\langle X_t^{(m)} \right\rangle - \frac{1}{2} tr\left[C^{(m)'} R^{-1} C^{(m)} \left\langle X_t^{(m)} X^{(m)'}{}_t \right\rangle \right] \right\}, \quad (21)$$

where tr is the matrix trace operator and we have used $tr(AB) = tr(BA)$. This can be done efficiently using the Kalman smoothing algorithm on each state-space model, where for model m at time t, the data is weighted by the responsibilities $h_t^{(m)}$.[8] Since the h parameters depend on the q parameters, and vice-versa, the whole process has to be iterated, where each iteration involves calls to the forward–backward and Kalman smoothing algorithms. The learning algorithm for switching state-space models using the above structured variational approximation is summarized in Figure 4.

Deterministic Annealing

The KL divergence minimized in the E step of the variational EM algorithm can have multiple minima in general. One way to visualize these minima is to consider the space of all possible *segmentations* of an observation sequence of length T, where by segmentation we mean a discrete partition of the sequence between the state space models. If there are M SSMs, then there are M^T possible segmentations of the sequence. Given one such segmentation, inferring the optimal distribution for the real-valued states of the SSMs is a convex optimization problem, since these real-valued states are conditionally Gaussian. So the difficulty in the KL minimization lies in trying to find the best (soft) partition of the data.

Like in other combinatorial optimization problems, the possibility of getting trapped in local minima can be reduced by gradually annealing the cost function. We can employ a deterministic variant of the annealing idea by making the following simple modifications to the variational fixed point equations (19) and (20):

$$h_t^{(m)} = \frac{1}{\mathcal{T}} Q(S_t = m) \quad (22)$$

[8] Weighting the data by $h_t^{(m)}$ is equivalent to running the Kalman smoother on the unweighted data using a time-varying observation noise covariance matrix $R_t^{(m)} = R / h_t^{(m)}$.

1. Initialize parameters of the model

2. Repeat until bound on log likelihood has converged:

E step Repeat until convergence of KL$(Q\|P)$:

E.1 Compute $q_t^{(m)}$ from the prediction error of state-space model m on observation Y_t

E.2 Compute $h_t^{(m)}$ using the forward-backward algorithm on the HMM, with observation probabilities $q_t^{(m)}$

E.3 For $m = 1$ to M

Run Kalman smoothing recursions, using the data weighted by $h_t^{(m)}$

M step

M.1 Re-estimate parameters for each state-space model using the data weighted by $h_t^{(m)}$

M.2 Re-estimate parameters for the switching process using Baum-Welch update equations

Figure 4: Learning algorithm for switching state-space models.

$$q_t^{(m)} = \exp\left\{-\frac{1}{2\mathcal{T}}\left\langle\left(Y_t - C^{(m)}X_t^{(m)}\right)' R^{-1}\left(Y_t - C^{(m)}X_t^{(m)}\right)\right\rangle\right\}. \tag{23}$$

Here $\mathcal{T}$ is a *temperature* parameter, which is initialized to a large value and gradually reduced to 1. The above equations maximize a modified form of the bound $\mathcal{B}$ in (11), where the entropy of Q has been multiplied by $\mathcal{T}$ (Ueda and Nakano, 1995).

5 Simulations

5.1 Experiment 1: Variational Segmentation and Deterministic Annealing

The goal of this experiment was to assess the quality of solutions found by the variational inference algorithm, and the effect of using deterministic annealing on these solutions. We generated 200 sequences of length 200 from a simple model which switched between two SSMs. These SSMs and the switching process were defined by:

$$X_t^{(1)} = 0.99\, X_{t-1}^{(1)} + w_t^{(1)} \qquad w_t^{(1)} \sim \mathcal{N}(0,1) \tag{24}$$

$$X_t^{(2)} = 0.9\, X_{t-1}^{(2)} + w_t^{(2)} \qquad w_t^{(2)} \sim \mathcal{N}(0,10) \tag{25}$$

$$Y_t = X_t^{(m)} + v_t \qquad v_t \sim \mathcal{N}(0,0.1) \tag{26}$$

where the switch state m was chosen using priors $\boldsymbol{\pi}^{(1)} = \boldsymbol{\pi}^{(2)} = 1/2$ and transition probabilities $\Phi_{11} = \Phi_{22} = 0.95$; $\Phi_{12} = \Phi_{21} = 0.05$. Five sequences from this data set are shown in in Figure 5, along with the true state of the switch variable.

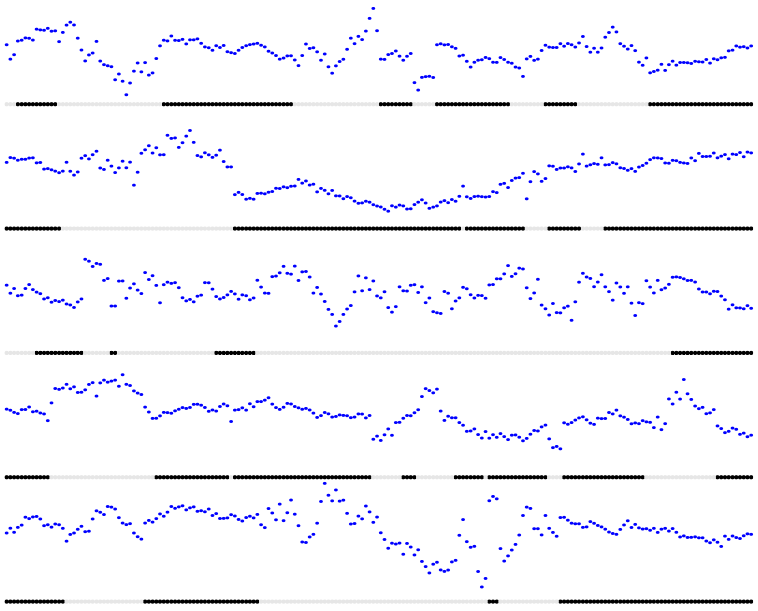

Figure 5: Five data sequences of length 200, with their true segmentations below them. In the segmentations, switch states 1 and 2 are represented with dark and light dots, respectively. Notice that correctly segmenting the sequences based only on the prior knowledge of the dynamics of the two processes is very difficult.

For each sequence, we initialized the inference algorithms with equal responsibilities for the two SSMs and ran the algorithms for 12 iterations. The non-annealed inference algorithm ran at a fixed temperature of $\mathcal{T} = 1$; while the annealed algorithm was initialized to a temperature of $\mathcal{T} = 100$ which was decayed down to 1 over the 12 iteration, using the decay function $\mathcal{T}_{i+1} = \frac{1}{2}\mathcal{T}_i + \frac{1}{2}$. To eliminate the effect of model inaccuracies, we gave both inference algorithms the true parameters of the generative model.

The segmentations found by the non-annealed variational inference algorithm showed little similarity to the true segmentations of the data (Figure 6). Furthermore, the non-annealed algorithm generally underestimated the number of switches, often converging on solutions with no switches at all. The annealed algorithm found segmentations that were more similar to the true segmentations of the data. In fact, annealing significantly increased the mutual information between the estimated and true segmentations (Figure 7).

In light of this, one would expect that the values of the lower bounds $\mathcal{B}$ found with annealing would be higher than those found without annealing. Surprisingly, this was not the case. The mean bounds found with annealing were lower by 13.7 bits over the 200 runs—a statistically significant effect ($p < 0.01$; Wilcoxon signed rank test) although small when compared to the standard deviation of the differences of 83.4 bits. We also computed the value of $\mathcal{B}$ given the true segmentation of the data, which was on average 131.4 bits higher than the bounds found without annealing.

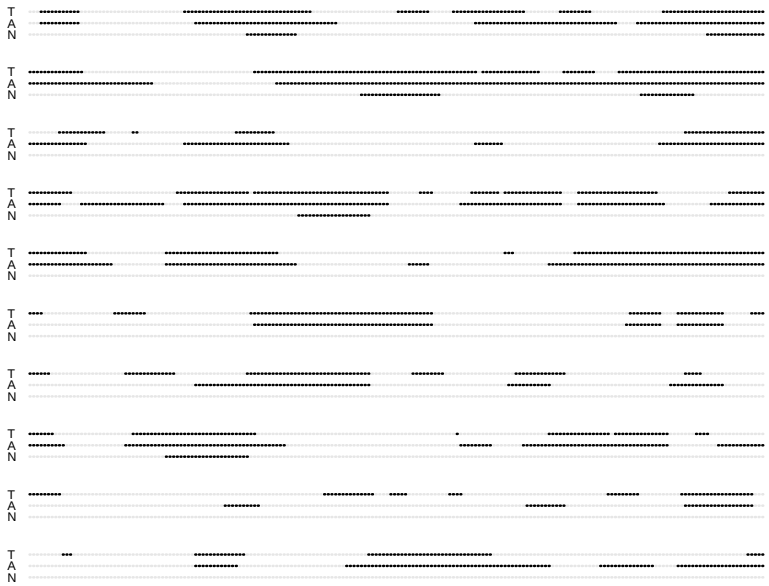

Figure 6: For ten different sequences of length 200, segmentations are shown as sequences of light and dark dots corresponding to the two SSMs generating this data. The rows are the true segmentations (T), the corresponding segmentations found using deterministic annealing (A) and not using annealing (N). These hard segmentations were obtained by thresholding the final $h_t^{(m)}$ values at 0.5. The first five sequences are the ones shown in figure 5.

5.2 Experiment 2: Modelling respiration in a patient with sleep apnea

Switching state-space models should prove useful in modelling time series which have nonlinear dynamics characterized by several different regimes. To illustrate this point we examined a physiological data set from a patient tentatively diagnosed with sleep apnea, which is a medical condition in which patients intermittently stop breathing during sleep, which results in a reflex arousal and gasps of breath. The data was obtained from the repository of time series data sets associated with Santa Fe Time Series Analysis and Prediction Competition (Weigend and Gershenfeld, 1993) and is described in detail in Rigney et al. (1993).[9] We simply wish to highlight the fact that the respiration pattern in sleep apnea is characterized by at least two regimes—no breathing and gasping breathing. Furthermore, in this patient there are also seem to be periods of normal rhythmic breathing (Figure 8).

We trained switching state-space models, varying the random seed, the number of components in the mixture ($M = 2$ to 5) and the dimensionality of the state space in each component ($K = 1$ to 10), on a data set consisting of 1000 consecutive measurements of the chest volume. As controls we also trained simple state-space models (i.e. $M = 1$) varying

[9]The data is available on the web at `http://www.cs.colorado.edu/ ~andreas/ Time-Series/ SantaFe.html#setB`. We used samples 6201–7200 for training and 5201-6200 for testing.

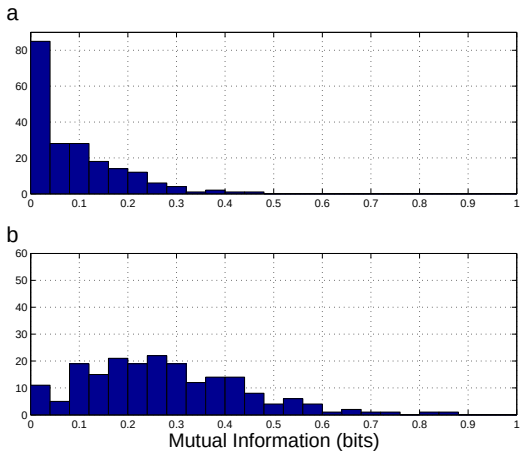

Figure 7: Histogram of the mutual information (in bits per observation) between the segmentations found without annealing (a) and with annealing (b).

the dimension of the state-space from $K = 1$ to 10, and simple hidden Markov models (i.e. $K = 0$) varying the number of discrete hidden states from $M = 2$ to $M = 50$. Simulations were run until convergence or for 200 iterations, whichever came first; convergence was assessed by measuring the change in likelihood (or bound on the likelihood) over consecutive steps of EM.

The likelihood of the simple SSMs and the HMMs was calculated on a test set, which also consisted of 1000 consecutive measurements of the chest volume. For the switching SSMs the likelihood is intractable, so we calculated the lower bound on the likelihood, $\mathcal{B}$. The simple SSMs modeled the data very poorly for $K = 1$, and the performance was flat for values of $K = 2$ to 10 (Figure 9a). The large majority of runs of the switching state-space model resulted in models with higher likelihood than those of the simple SMMs (Figure 9b-e). One consistent exception should be noted: for values of $M = 2$ and $K = 6$ to 10, the switching SSM performed almost identically to the simple SSM. Exploratory experiments suggest that in these cases a single component takes responsibility for all the data, so the model has $M = 1$ effectively. This may be a local minimum problem or a result of poor initialization heuristics. Looking at the learning curves for simple and switching state space models it is easy to see that there are plateaus at the solutions found by the simple one-component SSMs which the switching SSM can get caught in (Figure 10).

The likelihoods for hidden Markov models with around $M = 15$ were comparable to those of the best switching state-space models (Figure 9f). So purely in terms of coding efficiency, there is no advantage to using a switching SSM to model this data. However, it is useful to consider the nature of the solutions learned by the switching SSM.

We illustrate this by showing the segmentation produced by a fairly typical switching SSM with $M = 2$ components of state space dimension $K = 2$. The thick dots at the bottom of the Figures 8a and b indicate the responsibility assigned the one of the two components. This component has clearly specialized to modeling the data during periods of apnea, while the other component models the gasps and periods of rhythmic breathing. These two switching components provide a somewhat more satisfying explanation of the

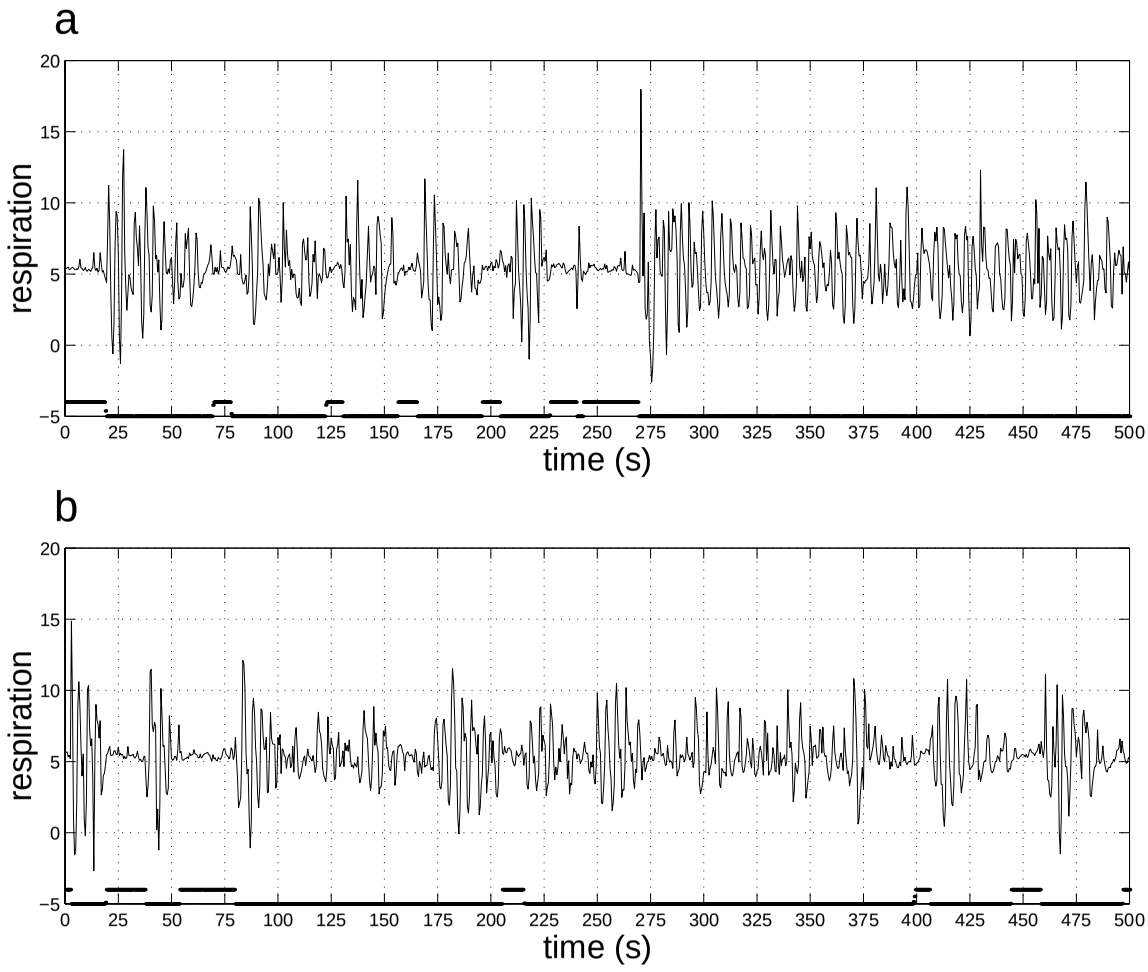

Figure 8: Chest volume (respiration force) of a patient with sleep apnea during two non-continguous time segments of the same night (measurements sampled at 2 Hz). (a) Training data. Apnea is characterized by extended periods of small variability in chest volume, followed by bursts (gasping). Here we see such behaviour around $t = 250$, followed by normal rhythmic breathing. (b) Test data. In this segment we find several instances of apnea and an approximately rhythmic region. (The thick lines at the bottom of each plot are explained in the main text.)

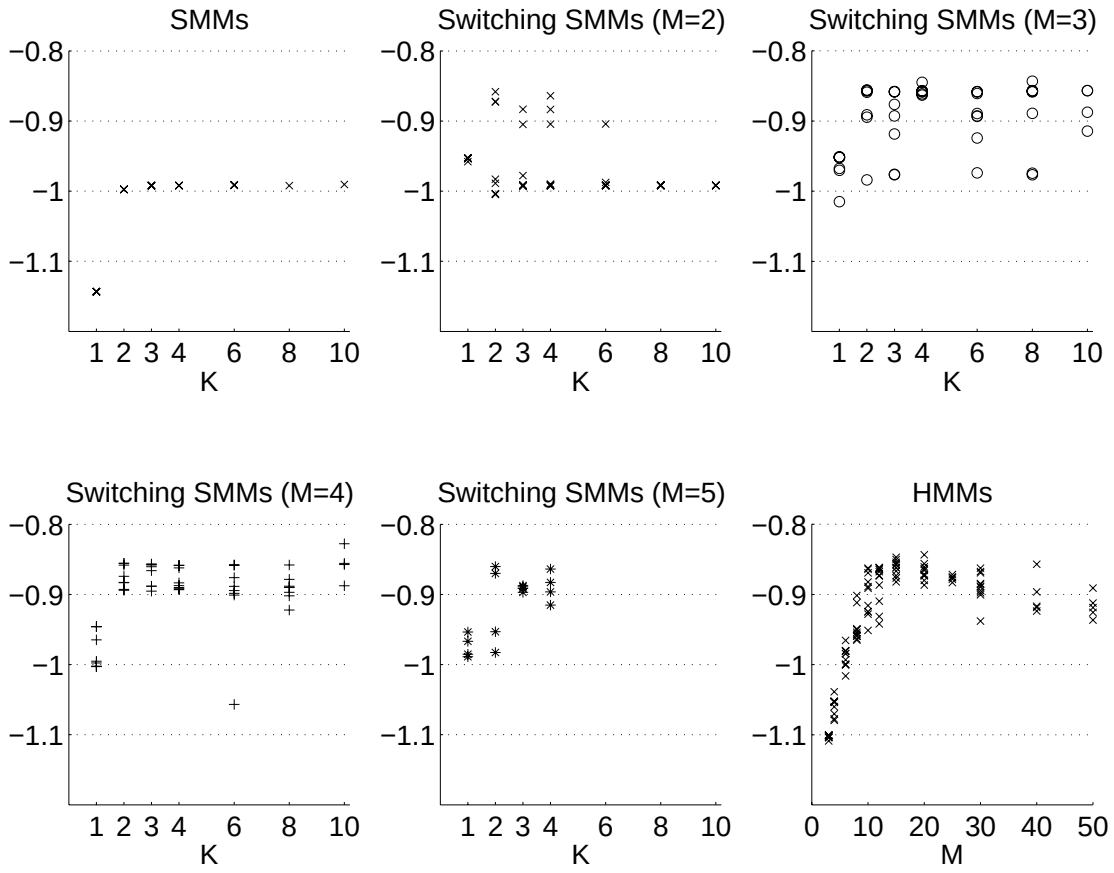

Figure 9: Log likelihood on the test data from a total of almost 400 runs of simple state-space models, switching state-space models with differing numbers of components, and hidden Markov models.

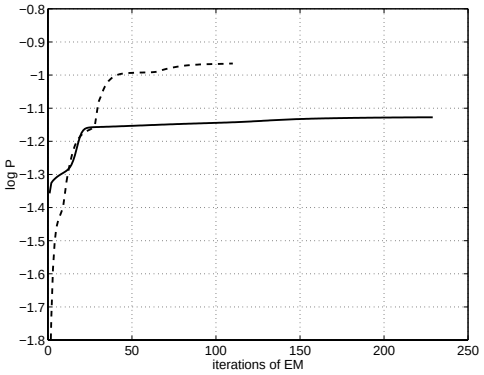

Figure 10: Learning curves for a state space model ($K = 4$) and a switching state-space model ($M = 2, K = 2$).

data than the 10-20 discrete components needed in a comparable HMM.

6 Discussion

The main conclusion we can draw from the first series of experiments is that even when given the correct model parameters, the problem of segmenting a switching time series into its components is difficult. There are combinatorially many alternatives to be considered, and the energy surface suffers from many local minima, so local optimization approaches like the variational method we used are limited by the quality of the initial conditions. Deterministic annealing can be thought of as a sophisticated initialization procedure for the hidden states: the final solution at each temperature provides the initial conditions at the next. We found that annealing substatially improved the quality of the segmentations found.

The second series of experiments suggests that on a real data set believed to have switching dynamics, the switching state-space model can indeed uncover multiple regimes. When it captures these regimes, it generalizes to the test set much better than the simple linear dynamical model. Similar coding efficiency can be obtained by using hidden Markov models, which due to the discrete nature of the state space can model nonlinear dynamics. However, in doing so, the hidden Markov models had to use 10-20 discrete states, which makes their solutions less interpretable.

Variational approximations provide a very powerful tool for inference and learning in complex probabilistic models. We have seen that when applied to the switching state-space model they can incorporate within a single framework well-known exact inference methods like Kalman smoothing and the forward-backward algorithm. However, training more complex models also makes apparent the importance of good methods for model selection and initialization. Bayesian approaches offer a principled way in which a priori knowledge can be used for initialization and models can be selected or weighted using their a posteriori probabilities.

To summarize, the switching state-space model is a fully dynamical extension of the mixture of experts network, which is closely related to well-known models in econometrics and control, and merges ideas underlying hidden Markov models and linear dyanamical systems. While it is clearly not an appropriate model for all time series, in situations where we have some a priori knowledge of multiple approximately linear dynamical regimes, the switching state space model can be used to exploit this knowledge. For example, in the sleep apnea problem it may be possible to train separate models for each of the known regimes in a small labelled portion of the data, and then unleash the full learning algorithm to obtain a better fit using both labelled and unlabelled data. Variational approximations overcome the single most difficult problem in learning switching SSMs, which is that the inference step is intractable. Deterministic annealing further improves on the solutions found by the variational method.

A Notation

Symbol	Size	Description
variables		
Y_t	$D \times 1$	observation vector at time t
$\{Y_t\}$	$D \times T$	sequence of observation vectors $[Y_1, Y_2, \ldots Y_T]$
$X_t^{(m)}$	$K \times 1$	state vector of state-space model (SSM) m at time t
X_t	$KM \times 1$	entire real-valued hidden state at time t: $X_t = [X_t^{(1)}, \ldots, X_t^{(M)}]$
S_t	$M \times 1$	switch state variable (represented either as discrete variable taking on values in $\{1, \ldots M\}$, or as an $M \times 1$ vector $S_t = [S_t^{(1)}, \ldots S_t^{(M)}]'$ where $S_t^{(m)} \in \{0, 1\}$)
model parameters		
$A^{(m)}$	$K \times K$	state dynamics matrix for SSM m
$C^{(m)}$	$D \times K$	output matrix for SSM m
$\mathcal{Q}^{(m)}$	$K \times K$	state noise covariance matrix for SSM m
$\mu_{X_1}^{(m)}$	$K \times 1$	initial state mean for SSM m
$\mathcal{Q}_1^{(m)}$	$K \times K$	initial state noise covariance matrix for SSM m
R	$D \times D$	output noise covariance matrix
$\boldsymbol{\pi}$	$M \times 1$	initial state probabilities for switch state
Φ	$M \times M$	state transition matrix for switch state
variational parameters		
$h_t^{(m)}$	1×1	responsibility of SSM m for Y_t
$q_t^{(m)}$	1×1	related to expected squared error if SSM m generated Y_t
miscellaneous		
X'		matrix transpose of X
$\lvert X \rvert$		matrix determinant of X
$\langle X \rangle$		expected value of X under the Q distribution
dimensions		
D		size of observation vector
T		length of a sequence of observation vectors
M		number of state-space models
K		size of state vector in each state-space model

B Derivation of the variational fixed-point equations

In this appendix we derive the variational fixed-point equations used in the learning algorithm for switching state space models. The plan is the following. First we write out the probability density P defined by a switching state space model. For convenience we will express this probability density in the log domain, through its associated energy function or *hamiltonian*, H. The probability density is related to the hamiltonian through the usual Boltzmann distribution (at a temperature of 1),

$$P(\cdot) = \frac{1}{Z} \exp\{-H(\cdot)\},$$

where Z is a normalization constant required such that $P(\cdot)$ integrates to unity. Expressing the probabilities in the log domain does not affect the resulting algorithm. We then similarly express the approximating distribution Q through its hamiltonian H_Q. Finally, we obtain the variational fixed point equations by setting to zero the derivatives of the KL divergence between Q and P with respect to the variational parameters of Q.

The joint probability of observations and hidden states in a switching state-space model is (equation (7))

$$P(\{S_t, X_t, Y_t\}) = \left[P(S_1)\prod_{t=2}^{T} P(S_t|S_{t-1})\right] \prod_{m=1}^{M} \left[P(X_1^{(m)})\prod_{t=2}^{T} P(X_t^{(m)}|X_{t-1}^{(m)})\right] \prod_{t=1}^{T} P(Y_t|X_t, S_t). \tag{27}$$

We proceed to dissect this expression into its constituent parts. The initial probability of the switch variable at time $t = 1$ is given by

$$P(S_1) = \prod_{m=1}^{M} (\boldsymbol{\pi}^{(m)})^{S_1^{(m)}}, \tag{28}$$

where S_1 is represented by an $M \times 1$ vector $[S_1^{(1)} \ldots S_1^{(M)}]$ where $S_1^{(m)} = 1$ if the switch state is in state m, and 0 otherwise. The probability of transitioning from a switch state at time $t-1$ to a switch state at time t is given by

$$P(S_t|S_{t-1}) = \prod_{m=1}^{M}\prod_{n=1}^{M} (\Phi^{(m,n)})^{S_t^{(m)} S_{t-1}^{(n)}}. \tag{29}$$

The initial distribution for the hidden state variable in state-space model m is Gaussian with mean $\mu_{X_1}^{(m)}$ and covariance matrix $\mathcal{Q}_1^{(m)}$:

$$P(X_1^{(m)}) = (2\pi)^{-K/2}|\mathcal{Q}_1^{(m)}|^{-1/2} \exp\left\{-\frac{1}{2}\left(X_1 - \mu_{X_1}^{(m)}\right)' (\mathcal{Q}_1^{(m)})^{-1} \left(X_1 - \mu_{X_1}^{(m)}\right)\right\}. \tag{30}$$

The probability distribution of the state in state-space model m at time t given the state at time $t-1$ is Gaussian with mean $A^{(m)}X_{t-1}^{(m)}$ and covariance matrix $\mathcal{Q}^{(m)}$:

$$P(X_t^{(m)}|X_{t-1}^{(m)}) = (2\pi)^{-K/2}|\mathcal{Q}^{(m)}|^{-1/2} \exp\left\{-\frac{1}{2}\left(X_t^{(m)} - A^{(m)}X_{t-1}^{(m)}\right)' (\mathcal{Q}^{(m)})^{-1} \left(X_t^{(m)} - A^{(m)}X_{t-1}^{(m)}\right)\right\}. \tag{31}$$

Finally, using (8) we can write:

$$P(Y_t|X_t,S_t) = \prod_{m=1}^{M}\left[(2\pi)^{-D/2}|R|^{-1/2}\exp\left\{-\frac{1}{2}\left(Y_t - C^{(m)}X_t^{(m)}\right)' R^{-1}\left(Y_t - C^{(m)}X_t^{(m)}\right)\right\}\right]^{S_t^{(m)}} \tag{32}$$

since the terms with exponent equal to 0 vanish in the product.

Combining (27)-(32) and taking the negative of the logarithm, we obtain the hamiltonian of a switching state-space model (ignoring constants):

$$\begin{aligned} H &= \frac{1}{2}\sum_{m=1}^{M}\log|\mathcal{Q}_1^{(m)}| + \frac{1}{2}\sum_{m=1}^{M}\left(X_1^{(m)} - \mu_{X_1}^{(m)}\right)'(\mathcal{Q}_1^{(m)})^{-1}\left(X_1^{(m)} - \mu_{X_1}^{(m)}\right) \\ &+ \frac{(T-1)}{2}\sum_{m=1}^{M}\log|\mathcal{Q}^{(m)}| + \frac{1}{2}\sum_{m=1}^{M}\sum_{t=2}^{T}\left(X_t^{(m)} - A^{(m)}X_{t-1}^{(m)}\right)'(\mathcal{Q}^{(m)})^{-1}\left(X_t^{(m)} - A^{(m)}X_{t-1}^{(m)}\right) \\ &+ \frac{T}{2}\log|R| + \frac{1}{2}\sum_{m=1}^{M}\sum_{t=1}^{T}S_t^{(m)}\left(Y_t - C^{(m)}X_t^{(m)}\right)' R^{-1}\left(Y_t - C^{(m)}X_t^{(m)}\right) \\ &- \sum_{m=1}^{M}S_1^{(m)}\log\boldsymbol{\pi}^{(m)} - \sum_{t=2}^{T}\sum_{m=1}^{M}\sum_{n=1}^{M}S_t^{(m)}S_{t-1}^{(n)}\log\Phi^{(m,n)}. \end{aligned} \tag{33}$$

The hamiltonian for the approximating distribution can be analogously derived from the definition of Q (equation (14)):

$$Q(\{S_t,X_t\}) = \frac{1}{Z_Q}\left[\psi(S_1)\prod_{t=2}^{T}\psi(S_{t-1},S_t)\right]\prod_{m=1}^{M}\psi(X_1^{(m)})\prod_{t=2}^{T}\psi(X_{t-1}^{(m)},X_t^{(m)}). \tag{34}$$

The potentials for the initial switch state and switch state transitions are

$$\psi(S_1) = \prod_{m=1}^{M}(\boldsymbol{\pi}^{(m)}q_1^{(m)})^{S_1^{(m)}} \tag{35}$$

$$\psi(S_{t-1},S_t) = \prod_{m=1}^{M}\prod_{n=1}^{M}\left(\Phi^{(m,n)}q_t^{(m)}\right)^{S_t^{(m)}S_{t-1}^{(n)}} \tag{36}$$

The potential for the initial state of state-space model m is

$$\psi(X_1^{(m)}) = P(X_1^{(m)})\left[P(Y_1|X_1^{(m)},S_1=m)\right]^{h_1^{(m)}} \tag{37}$$

and the the potential for the state at time t given the state at time $t-1$ is

$$\psi(X_{t-1}^{(m)},X_t^{(m)}) = P(X_t^{(m)}|X_{t-1}^{(m)})\left[P(Y_t|X_t^{(m)},S_t=m)\right]^{h_t^{(m)}}. \tag{38}$$

The hamiltonian for Q is obtained by combining these terms and taking the negative logarithm:

$$H_Q = \frac{1}{2}\sum_{m=1}^{M}\log|\mathcal{Q}_1^{(m)}| + \frac{1}{2}\sum_{m=1}^{M}\left(X_1^{(m)} - \mu_{X_1}^{(m)}\right)'(\mathcal{Q}_1^{(m)})^{-1}\left(X_1^{(m)} - \mu_{X_1}^{(m)}\right)$$

$$+\frac{(T-1)}{2}\sum_{m=1}^{M}\log|\mathcal{Q}^{(m)}|+\frac{1}{2}\sum_{m=1}^{M}\sum_{t=2}^{T}\left(X_t^{(m)}-A^{(m)}X_{t-1}^{(m)}\right)'(\mathcal{Q}^{(m)})^{-1}\left(X_t^{(m)}-A^{(m)}X_{t-1}^{(m)}\right)$$
$$+\frac{T}{2}\sum_{m=1}^{M}\log|R|+\frac{1}{2}\sum_{m=1}^{M}\sum_{t=1}^{T}h_t^{(m)}\left(Y_t-C^{(m)}X_t^{(m)}\right)'R^{-1}\left(Y_t-C^{(m)}X_t^{(m)}\right)$$
$$-\sum_{m=1}^{M}S_1^{(m)}\log\boldsymbol{\pi}^{(m)}-\sum_{t=2}^{T}\sum_{m=1}^{M}\sum_{n=1}^{M}S_t^{(m)}S_{t-1}^{(n)}\log\Phi^{(m,n)}-\sum_{t=1}^{T}\sum_{m=1}^{M}S_t^{(m)}\log q_t^{(m)}. \quad (39)$$

Comparing H_Q with H we see that the interaction between the $S_t^{(m)}$ and the $X_t^{(m)}$ variables has been eliminated, while introducing two sets of variational parameters: the responsibilities $h_t^{(m)}$ and the bias terms on the discrete Markov chain, $q_t^{(m)}$. In order to obtain the approximation Q which maximizes the lower bound on the log likelihood, we minimize the KL divergence $\mathrm{KL}(Q\|P)$ as a function of these variational parameters

$$\mathrm{KL}(Q\|P) = \sum_{\{S_t\}}\int Q(\{S_t,X_t\})\log\frac{Q(\{S_t,X_t\})}{P(\{S_t,X_t\}|\{Y_t\})}d\{X_t\} \quad (40)$$
$$= \langle H-H_Q\rangle-\log Z_Q+\log Z, \quad (41)$$

where $\langle\cdot\rangle$ denotes expectation over the approximating distribution Q and Z_Q is the normalization constant for Q. Both Q and P define distributions in the exponential family. As a consequence, the zeros of the derivatives of KL with respect to the variational parameters can be obtained simply by equating derivatives of $\langle H\rangle$ and $\langle H_Q\rangle$ with respect to corresponding sufficient statistics (Ghahramani, 1997):

$$\frac{\partial\langle H_Q-H\rangle}{\partial\langle S_t^{(m)}\rangle} = 0 \quad (42)$$
$$\frac{\partial\langle H_Q-H\rangle}{\partial\langle X_t^{(m)}\rangle} = 0 \quad (43)$$
$$\frac{\partial\langle H_Q-H\rangle}{\partial\langle P_t^{(m)}\rangle} = 0 \quad (44)$$

where $P_t^{(m)}=\langle X_t^{(m)}X_t^{(m)\prime}\rangle-\langle X_t^{(m)}\rangle\langle X_t^{(m)}\rangle'$ is the covariance of $X_t^{(m)}$ under Q. Many terms cancel when we subtract the two hamiltonians

$$H_Q-H=\sum_{m=1}^{M}\sum_{t=1}^{T}\frac{1}{2}\left(h_t^{(m)}-S_t^{(m)}\right)\left(Y_t-C^{(m)}X_t^{(m)}\right)'R^{-1}\left(Y_t-C^{(m)}X_t^{(m)}\right)-S_t^{(m)}\log q_t^{(m)} \quad (45)$$

Taking derivatives we obtain

$$\frac{\partial\langle H_Q-H\rangle}{\partial\langle S_t^{(m)}\rangle} = -\log q_t^{(m)}-\frac{1}{2}\left\langle\left(Y_t-C^{(m)}X_t^{(m)}\right)'R^{-1}\left(Y_t-C^{(m)}X_t^{(m)}\right)\right\rangle \quad (46)$$
$$\frac{\partial\langle H_Q-H\rangle}{\partial\langle X_t^{(m)}\rangle} = -\left(h_t^{(m)}-\langle S_t^{(m)}\rangle\right)\left((Y_t-C^{(m)}\langle X_t^{(m)}\rangle)'R^{-1}C^{(m)}\right) \quad (47)$$
$$\frac{\partial\langle H_Q-H\rangle}{\partial P_t^{(m)}} = \frac{1}{2}\left(h_t^{(m)}-\langle S_t^{(m)}\rangle\right)\left(C^{(m)\prime}R^{-1}C^{(m)}\right) \quad (48)$$

From (46) we get the fixed-point equation (20) for $q_t^{(m)}$. Both (47) and (48) are satisfied when $h_t^{(m)} = \langle S_t^{(m)} \rangle$. Using the fact that $\langle S_t^{(m)} \rangle = Q(S_t = m)$ we get (19).

References

Ackerson, G. A. and Fu, K. S. (1970). On state estimation in switching environments. *IEEE Transactions on Automatic Control*, AC-15(1):10–17.

Anderson, B. D. O. and Moore, J. B. (1979). *Optimal Filtering*. Prentice-Hall, Englewood Cliffs, NJ.

Athaide, C. R. (1995). *Likelihood Evaluation and State Estimation for Nonlinear State Space Models*. Ph.D. Thesis, Graduate Group in Managerial Science and Applied Economics, University of Pennsylvania, Philadelphia, PA.

Baldi, P., Chauvin, Y., Hunkapiller, T., and McClure, M. (1994). Hidden Markov models of biological primary sequence information. *Proc. Nat. Acad. Sci. (USA)*, 91(3):1059–1063.

Bar-Shalom, Y. and Li, X.-R. (1993). *Estimation and Tracking*. Artech House, Boston, MA.

Baum, L., Petrie, T., Soules, G., and Weiss, N. (1970). A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *The Annals of Mathematical Statistics*, 41:164–171.

Bengio, Y. and Frasconi, P. (1995). An input–output HMM architecture. In Tesauro, G., Touretzky, D. S., and Leen, T. K., editors, *Advances in Neural Information Processing Systems 7*, pages 427–434. MIT Press, Cambridge, MA.

Blake, A., Isard, M., and Reynard, D. (1995). Learning to track the visual motion of contours. *Artificial Intelligence*, 78:101–133.

Cacciatore, T. W. and Nowlan, S. J. (1994). Mixtures of controllers for jump linear and non-linear plants. In Cowan, J. D., Tesauro, G., and Alspector, J., editors, *Advances in Neural Information Processing Systems 6*, pages 719–726. Morgan Kaufmann Publishers, San Francisco, CA.

Carter, C. K. and Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, 81:541–553.

Chaer, W. S., Bishop, R. H., and Ghosh, J. (1997). A mixture-of-experts framework for adaptive kalman filtering. *IEEE Trans. on Systems, Man and Cybernetics*.

Chang, C. B. and Athans, M. (1978). State estimation for discrete systems with switching parameters. *IEEE Transactions on Aerospace and Electronic Systems*, AES-14(3):418–424.

Cover, T. and Thomas, J. (1991). *Elements of Information Theory*. Wiley, New York.

Dean, T. and Kanazawa, K. (1989). A model for reasoning about persitence and causation. *Computational Intelligence*, 5(3):142–150.

Dempster, A., Laird, N., and Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statistical Society Series B*, 39:1–38.

Deng, L. (1993). A stochastic model of speech incorporating hierarchical nonstationarity. *IEEE Trans. on Speech and Audio Processing*, 1(4):471–474.

Digalakis, V., Rohlicek, J. R., and Ostendorf, M. (1993). ML estimation of a Stochastic Linear System with the EM Algorithm and its Application to Speech Recognition. *IEEE Transactions on Speech and Audio Processing*, 1(4):431–442.

Everitt, B. S. (1984). *An Introduction to Latent Variable Models.* Chapman and Hall, London.

Fraser, A. M. and Dimitriadis, A. (1993). Forecasting probability densities by using hidden Markov models with mixted states. In Wiegand, A. S. and Gershenfeld, N. A., editors, *Time series prediction: Forecasting the future and understanding the past*, SFI Studies in the Sciences of Complexity, Proc. Vol. XV, pages 265–282. Addison-Wesley, Reading, MA.

Ghahramani, Z. (1997). On structured variational approximations. Technical Report CRG-TR-97-1 [ftp://ftp.cs.toronto.edu/pub/ zoubin/struct.ps.gz].

Ghahramani, Z. and Hinton, G. E. (1996a). Parameter estimation for linear dynamical systems. Technical Report CRG-TR-96-2 [ftp://ftp.cs.toronto.edu/pub/zoubin/tr-96-2.ps.gz] .

Ghahramani, Z. and Hinton, G. E. (1996b). The EM algorithm for mixtures of factor analyzers. Technical Report CRG-TR-96-1 [ftp://ftp.cs.toronto.edu/pub/zoubin/tr-96-1.ps.gz] .

Ghahramani, Z. and Jordan, M. I. (1997). Factorial hidden Markov models. *Machine Learning*, 29:245–273.

Goodwin, G. and Sin, K. (1984). *Adaptive filtering prediction and control.* Prentice-Hall.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57:357–384.

Harrison, P. J. and Stevens, C. F. (1976). Bayesian forecasting (with discussion). *J. Royal Statistical Society B*, 38:205–247.

Hinton, G. E., Dayan, P., and Revow, M. (1996). Modeling the manifolds of Images of handwritten digits. *Submitted for Publication.*

Isard, M. and Blake, A. (1996). Contour tracking by stochastic propagation of conditional density. In Buxton, B. and Cipolla, R., editors, *Proc. 4th European Conference on Computer Vision*, volume 1, pages 343–356. Springer Verlag, berlin.

Jacobs, R. A., Jordan, M. I., Nowlan, S. J., and Hinton, G. E. (1991). Adaptive mixture of local experts. *Neural Computation*, 3:79–87.

Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. (1998). An Introduction to variational methods in graphical models. In Jordan, M. I., editor, *Learning in Graphical Models*. Kluwer Academic Publishers.

Jordan, M. I. and Jacobs, R. (1994). Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6:181–214.

Juang, B. H. and Rabiner, L. R. (1991). Hidden Markov models for speech recognition. *Technometrics*, 33:251–272.

Kadirkamanathan, V. and Kadirkamanathan, M. (1996). Recursive estimation of dynamic modular RBF networks. In Touretzky, D., Mozer, M., and Hasselmo, M., editors, *Advances in Neural Information Processing Systems 8*, pages 239–245. MIT Press.

Kalman, R. E. and Bucy, R. S. (1961). New results in linear filtering and prediction. *Journal of Basic Engineering (ASME)*, 83D:95–108.

Kanazawa, K., Koller, D., and Russell, S. J. (1995). Stochastic simulation algorithms for dynamic probabilistic networks. In Besnard, P. and Hanks, S., editors, *Uncertainty in Artificial Intelligence. Proceedings of the Eleventh Conference.*, pages 346–351. Morgan Kaufmann Publishers, San Francisco, CA.

Kim, C.-J. (1994). Dynamic linear models with Markov-switching. *J. Econometrics*, 60:1–22.

Lauritzen, S. L. and Spiegelhalter, D. J. (1988). Local computations with probabilities on graphical structures and their application to expert systems. *J. Royal Statistical Society B*, pages 157–224.

Ljung, L. and Söderström, T. (1983). *Theory and Practice of Recursive Identification*. MIT Press, Cambridge, MA.

Meila, M. and Jordan, M. I. (1996). Learning fine motion by Markov mixtures of experts. In Touretzky, D. S., Mozer, M. C., and Hasselmo, M. E., editors, *Advances in Neural Information Processing Systems 8*. MIT Press.

Neal, R. M. and Hinton, G. E. (1993). A new view of the EM algorithm that justifies incremental and other variants. Technical Report Unpublished manuscript.

Parisi, G. (1988). *Statistical Field Theory*. Addison-Wesley, Redwood City, CA.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA.

Rabiner, L. R. and Juang, B. H. (1986). An Introduction to hidden Markov models. *IEEE Acoustics, Speech & Signal Processing Magazine*, 3:4–16.

Rao, R. P. N. and Ballard, D. H. (1996). A class of stochastic models for invariant recognition, motion, and stereo. *Submitted to Advances in Neural Information Processing Systems 9*.

Rauch, H. E. (1963). Solutions to the linear smoothing problem. *IEEE Transactions on Automatic Control*, 8:371–372.

Rigney, D., Goldberger, A., Ocasio, W., Ichimaru, Y., Moody, G., and Mark, R. (1993). Multi-channel physiological data: Description and analysis. In Weigend, A. and Gershenfeld, N., editors, *Time Series Prediction: Forecasting the future and understanding the past*, SFI Studies in the Sciences of Complexity, Proc. Vol. XV, pages 105–129. Addison-Wesley, Reading, MA.

Saul, L. and Jordan, M. I. (1996). Exploiting tractable substructures in Intractable networks. In Touretzky, D., Mozer, M., and Hasselmo, M., editors, *Advances in Neural Information Processing Systems 8*. MIT Press.

Shumway, R. H. and Stoffer, D. S. (1982). An approach to time series smoothing and forecasting using the EM algorithm. *J. Time Series Analysis*, 3(4):253–264.

Shumway, R. H. and Stoffer, D. S. (1991). Dynamic linear models with switching. *J. Amer. Stat. Assoc.*, 86:763–769.

Smyth, P. (1994). Hidden Markov models for fault detection in dynamic systems. *Pattern Recognition*, 27(1):149–164.

Smyth, P., Heckerman, D., and Jordan, M. I. (1997). Probabilistic independence networks for hidden Markov probability models. *Neural Computation*, 9:227–269.

Ueda, N. and Nakano, R. (1995). Deterministic annealing variant of the EM algorithm. In Tesauro, G., Touretzky, D., and Alspector, J., editors, *Advances in Neural Information Processing Systems 7*, pages 545–552. Morgan Kaufmann.

Weigend, A. and Gershenfeld, N. (1993). *Time Series Prediction: Forecasting the future and understanding the past*. SFI Studies in the Sciences of Complexity, Proc. Vol. XV. Addison-Wesley, Reading, MA.