

Exploiting locality in high-dimensional Factorial hidden Markov models

Lorenzo Rimella

*Department of Mathematics and Statistics
Lancaster University
Lancaster, LA1 4YF, UK*

L.RIMELLA@LANCASTER.AC.UK

Nick Whiteley

*Institute for Statistical Science
School of Mathematics
University of Bristol
Bristol, BS8 1TW, UK
and the Alan Turing Institute, UK*

NICK.WHITELEY@BRISTOL.AC.UK

Editor: Barbara Engelhardt

Abstract

We propose algorithms for approximate filtering and smoothing in high-dimensional Factorial hidden Markov models. The approximation involves discarding, in a principled way, likelihood factors according to a notion of locality in a factor graph associated with the emission distribution. This allows the exponential-in-dimension cost of exact filtering and smoothing to be avoided. We prove that the approximation accuracy, measured in a local total variation norm, is “dimension-free” in the sense that as the overall dimension of the model increases the error bounds we derive do not necessarily degrade. A key step in the analysis is to quantify the error introduced by localizing the likelihood function in a Bayes’ rule update. The factorial structure of the likelihood function which we exploit arises naturally when data have known spatial or network structure. We demonstrate the new algorithms on synthetic examples and a London Underground passenger flow problem, where the factor graph is effectively given by the train network.

Keywords: Factorial hidden Markov models, filtering, smoothing, EM algorithm, high-dimensions

1. Introduction

Since early appearance in the statistical literature (Baum and Petrie, 1966; Baum et al., 1970) and popularization in speech recognition (Rabiner, 1989), Hidden Markov models (HMMs) have been used to solve a broad range of problems ranging from texture recognition (Bose and Kuo, 1994), to gene prediction (Stanke and Waack, 2003), and weather forecasting (Hughes et al., 1999). The influential paper of Ghahramani and Jordan (1997) introduced the class of Factorial hidden Markov models (FHMMs), in which the hidden Markov chain is a multivariate process, with a-priori independent coordinates. This structure provides a rich modelling framework to capture complex statistical patterns in data sequences and it is particularly suited to applications with straightforward local dependencies: traffic modelling, weather forecasting, etc.

The main objective of the present paper is to develop rigorous insight into how inference in FHMMs can be scaled up to high-dimensional problems. The problem one immediately faces when scaling up exact inference techniques for FHMM's is the computational cost of the matrix-vector operations underlying the well known forward-backward algorithm, which computes conditional distributions over hidden states given data, and in turn the Baum-Welch algorithm to perform maximum-likelihood estimation of static parameters. This cost typically grows exponentially with the dimension of the underlying state-space. Ghahramani and Jordan (1997) derived a variant of the forward-backward algorithm which achieves a degree of efficiency by exploiting the structure of FHMMs, but the exponential-in-dimension scaling of cost cannot be avoided.

As an alternative to exact inference, Ghahramani and Jordan (1997) also proposed two families of variational methods for FHMMs, which allow approximate solution of the smoothing problem: computing conditional distributions over hidden states given past and future data. In turn, this allows approximate maximum likelihood parameter estimation via an expectation-maximization (EM) algorithm. The attractive feature of these variational approximations is, in typical FHMMs, that their computational complexity is a low-order polynomial in state dimension. However, little seems to be known about their performance in the context of FHMMs from a theoretical point of view, other than the fact that by construction they minimize a Kullback-Liebler divergence criterion. As surveyed recently by Blei et al. (2017)(Section 5), there are a number of strands of research into theoretical properties of variational methods for some classes of statistical models, such as convergence analysis for mixture models (Wang and Titterton, 2006), and consistency studies for stochastic block models (Celisse et al., 2012; Bickel et al., 2013). However analysis specifically for FHMMs appears to be lacking, and to the authors' knowledge there are currently no detailed mathematical studies of how variational approximation errors for FHMMs scale with dimension, data record length, model parameters, etc.

Boyan and Koller (1998, 1999), proposed and studied inference methods in Dynamic Bayesian Networks which involve recursively approximating belief-state distributions by the product of their marginals and then propagating the result to the next time step. Particle filtering algorithms in the same vein appeared in Ng et al. (2002), Brandao et al. (2006) and Besada-Portas et al. (2009). Ground-breaking theoretical work of Rebeschini and Van Handel (2015) proved that similar algorithms can be used to conduct particle filtering efficiently in high-dimensions, using techniques based on the Dobrushin Comparison Theorem (see for instance Georgii, 2011). Subsequently, Rebeschini and van Handel (2014) refined their analysis through generalized Dobrushin comparison theorems. Finke and Singh (2017) extended these ideas from particle filtering to particle smoothing and studied dimension-independence of the asymptotic Monte Carlo variance. In the present work we do not consider sampling-based approximate inference techniques because we focus on FHMMs with discrete hidden states. If one were to consider FHMMs with continuous hidden states then following the blueprint of the aforementioned works one could devise particle filtering and smoothing versions of the approximate inference algorithms we propose, but this is beyond the scope of the present work. Another key difference between the aforementioned works and ours is that they do not considered FHMMs, but rather models in which posterior dependence across dimensions arises from the prior. For FHMMs, this dependence arises only from the emission distribution likelihood function and we use not

only factorization but also the key notion of *localization*, introduced below, to deal with this.

1.1 Setting and Contributions

In this paper we propose and study approximate inference algorithms for FHMMs called the *Graph Filter* and *Graph Smoother*. In some ways, the Graph Filter and Smoother are similar in spirit to variational methods: they involve constructing approximate posterior distributions over hidden states which factorize across dimension. However, unlike variational methods, they do not involve explicit minimization of a Kullback-Liebler divergence criterion and therefore avoid the fixed-point iterations or other numerical optimization procedures which variational methods typically involve. Another contrast between the variational methods of Ghahramani and Jordan (1997) and the Graph Filter and Smoother is that the latter share the recursive-in-time structure of the forward-backward algorithm. The forward pass, which conducts the task of *filtering*, can therefore be used for prediction in online settings. Variational methods for FHMMs work in a batch setting, suitable for offline data analysis.

In taking first steps towards rigorously understanding the quality of factorization-based approximations it is natural and desirable to impose on the FHMM:

1. structure which allows one to express connections between the scaling properties of the approximation errors, computational cost and dimensionality, and
2. assumptions that allow the weakness of dependence across dimensions to be quantified.

In a FHMM, by definition, hidden state variables are a-priori independent across dimensions, so any dependence across dimensions under the corresponding posterior distributions must arise from the emission distribution. Therefore the emission distribution is where we impose structure and assumptions as per 1. and 2. above. As made precise in equation (2) below, we assume the emission distribution likelihood function has a factorial structure, expressed in terms of a factor graph. Whilst any FHMM can be expressed in this general form (there could be just one factor, which is a function of all dimensions of the hidden state) the Graph Filter and Smoother allow the exponential-in-dimension cost of inference to be significantly reduced when there are several factors, each depending on some “small” subset of the hidden state dimensions. Part of the novelty of our theoretical results is that in this setting it turns out that graph distance on the factor graph can be used to quantify the strength of statistical dependence across dimension, and this allows us to express approximation error in terms of quantities such as the cardinalities of neighborhoods on the factor graph, which do not necessarily grow with overall dimension of the FHMM. The main message conveyed by our theoretical results is then that the Graph Filter and Smoother have “dimension-free” error performance, in the sense that as the dimension of the underlying state space increases, the assumptions we make do not necessarily become more demanding, and the error bounds we derive do not necessarily degrade. The proofs are heavily influenced by the Dobrushin Comparison techniques used by Rebeschini and Van Handel (2015).

One situation in which this factorial emission likelihood structure naturally arises is when, at each time step, the observed data consist of a set of spatially-indexed values, and the factor graphs captures this spatial structure, or at least some topological aspects thereof.

This is the case, for example, in models of transport flow data (Hofleitner et al., 2012; Woodard et al., 2017), where observations measure the numbers of vehicles or passengers at various interconnected locations on a rail, road or flight network, and the hidden variables indicate levels of congestion in spatially localized regions. Here the factor graph arises naturally from the transport network. In this paper we provide numerical results for such a model analyzing data from the London Tube underground rail network. Another domain in which this structure arises is spatio-temporal epidemiology, where the data are noisy, incomplete measurement of disease states across a collection of locations with pairwise dependence (Park and Ionides, 2020; Asfaw et al., 2021). Spatial localization of observations is a key element of geophysical models, often treated in terms of continuous hidden variables, with inference performed e.g. using particle filters (Van Leeuwen, 2009). The FHMM structure we consider encompasses discrete-state counterparts of these models. The same is true of high-dimensional multivariate stochastic volatility models with cross-leverage effects (Xu and Jasra, 2019), where it is common to use discrete states to capture switching between volatility regimes. Discrete-state models also arise when detecting change-points in spatio-temporal data such as temperatures and air-pollution levels (Knoblauch and Damoulas, 2018). Such models could be cast as FHMMs with factorial emission likelihood structure.

Section 2 introduces the class of FHMMs we consider and background on exact inference by filtering and smoothing. Section 3 introduces the Graph Filter and Smoother, it discusses their computational complexity and our main results, Theorems 2 and 3. An important component in our analysis is a preliminary result concerning approximate Bayes’ rule updates using localized factorial likelihoods, Proposition 1, which may be of independent interest.

Numerical experiments are reported in Section 4. We illustrate that the scaling of errors with algorithm parameters and model attributes indicated by our theoretical results can have a substantial impact in practice by avoiding exponential-in-dimension computational cost. Considering synthetic data so that true parameter values are known, we show how the Graph Smoother can be used within an approximate EM algorithm to fit the parameters of an FHMM and we compare accuracy to the approximate EM approach of Ghahramani and Jordan (1997) involving variational approximations.

Finally, we illustrate the utility of the Graph Filter and Smoother analysing Oyster card “tap” data on the London Underground train network, estimating parameters and performing prediction of passenger inflow-outflow through stations, showing that Graph Filter-Smoother can be used to model complex spatial interactions over time. Some possible extensions and future research directions are described in Section 5. All proofs and a collection of supporting results are in the online appendix (**Online appendix**). All the codes are implemented in Python 3 and available on Github (**Online code**).

2. Factorial hidden Markov models, filtering and smoothing

Over a time horizon of length $T \in \mathbb{N}$, a HMM is a pair of processes $(X_t)_{t \in \{0, \dots, T\}}$, $(Y_t)_{t \in \{1, \dots, T\}}$ where the unobserved process $(X_t)_{t \in \{0, \dots, T\}}$ is a Markov chain, and each observation Y_t is conditionally independent of all other variables given X_t .

We consider the case where the state-space of $(X_t)_{t \in \{0, \dots, T\}}$ is of product form, $X_t = (X_t^v)_{v \in V} \in \mathbb{X}^V$ where \mathbb{X} and V are finite sets, and write $L := \mathbf{card}(\mathbb{X})$ and $M := \mathbf{card}(V)$.

Each $(Y_t)_{t \in \{1, \dots, T\}}$ is valued in a set \mathbb{Y} which could be a discrete set, \mathbb{R}^d or some subset thereof.

We write $\mu_0(x)$ for the probability mass function of X_0 , $p(x, z)$ for the transition probability of $(X_t)_{t \in \{0, \dots, T\}}$ from x to z , and $g(x, y)$ for the conditional probability mass or density function of Y_t given X_t . In the literature, $g(x, \cdot)$ is often called the emission distribution.

For any $U \subseteq V$ and $x = (x^v)_{v \in V} \in \mathbb{X}^V$ we shall use the shorthand $x^U = (x^v)_{v \in U}$, and similarly for any probability mass function μ on \mathbb{X}^V we shall denote its marginal associated with U by μ^U . When \mathcal{K} is any partition of the set V , we shall say that μ factorizes with respect to \mathcal{K} if:

$$\mu(x) = \prod_{K \in \mathcal{K}} \mu^K(x^K), \quad \forall x \in \mathbb{X}^V,$$

and in this situation we will use the shorthand:

$$\mu = \bigotimes_{K \in \mathcal{K}} \mu^K.$$

The total variation distance between probability mass functions on \mathbb{X}^U , say μ and ν , is denoted by:

$$\|\mu - \nu\| := \sup_{A \in \sigma(\mathbb{X}^U)} |\mu(A) - \nu(A)|,$$

with the obvious overloading of notation $\mu(A) = \sum_{x^U \in A} \mu(x^U)$ and where $\sigma(\mathbb{X}^U)$ is the power set of \mathbb{X}^U . When μ, ν are probability mass functions on \mathbb{X}^V it will be convenient to denote the local total variation (LTV) distance associated with $U \subseteq V$,

$$\|\mu - \nu\|_U := \sup_{A \in \sigma(\mathbb{X}^U)} |\mu^U(A) - \nu^U(A)|.$$

2.1 Factorial hidden Markov Models

In FHMMs (Ghahramani and Jordan, 1997) the transition probabilities of the hidden Markov chain are assumed to factorize in the following manner:

$$p(x, z) = \prod_{v \in V} p^v(x^v, z^v), \tag{1}$$

where each $p^v(x^v, z^v)$ is a transition probability on \mathbb{X} . Directed acyclic graphs showing the conditional independence structures of HMMs and FHMMs are shown in Figure 1. As motivated in section 1.1, in this paper we consider a factorial structure for the likelihood function $x \mapsto g(x, y)$. Let F be a finite set and let $\mathcal{G} = (V, F, E)$ be a factor graph associated with $x \mapsto g(x, y)$, that is a bi-partite graph with vertex sets V, F and edge set E such that $g(x, y)$ can be written in terms of factors:

$$g(x, y) = \prod_{f \in F} g^f(x^{N(f)}, y), \tag{2}$$

where $N(\cdot)$ is the neighbourhood function,

$$N(w) := \{w' \in V \cup F : (w, w') \in E\}, \quad w \in V \cup F. \tag{3}$$




Figure 1: A conditional independence structure of a HMM (a) and an FHMM (b)

In applications each observation y will typically be multivariate and each likelihood factor $g^f(x^{N(f)}, y)$ may depend on y only through some subset of its constituent variates, but the details will be model specific, so we do not introduce them at this stage. We note that equations 1 and 2 are the mathematical formalization of “local structures”, indeed equation 1 says that each component of the hidden variable is evolving independently, while equation 2 stands that the emission distribution can be decomposed in pieces that can be built locally.




Figure 2: An example of factor graph for a fixed time step t of an FHMM where $V = \{1, 2, 3, 4, 5\}$ and $F = \{f_1, f_2, f_3, f_4\}$.

2.2 Filtering and Smoothing

The tasks of filtering and smoothing at time t are to compute, respectively, the conditional distributions of X_t given the realized observations (y_1, \dots, y_t) and (y_1, \dots, y_T) . We shall denote the corresponding probability mass functions by π_t and $\pi_{t|T}$. As well as facilitating inference about hidden states, computing these distributions is often a key step in estimating static parameters of the HMM.

Filtering can be conducted in a forward pass through the data:

$$\pi_0 := \mu_0, \quad \pi_t := F_t \pi_{t-1}, \quad F_t := C_t P, \quad t \in \{1, \dots, T\}, \quad (4)$$

where the “prediction” operator \mathbf{P} and the “correction” operator \mathbf{C}_t act on probability mass functions μ as:

$$\mathbf{P}\mu(x) := \sum_{z \in \mathbb{X}^V} p(z, x)\mu(z), \quad \mathbf{C}_t\mu(x) := \frac{g(x, y_t)\mu(x)}{\sum_{z \in \mathbb{X}^V} g(z, y_t)\mu(z)}, \quad x \in \mathbb{X}^V. \quad (5)$$

If μ is considered a prior distribution, the operator \mathbf{C}_t can be understood as applying a Bayes’ rule update using the likelihood function $g(x, y_t)$.

Amongst various smoothing algorithms, we focus on the forward-filtering, backward-smoothing method presented by Kitagawa (1987), which involves a backward in time recursion performed after filtering:

$$\pi_{T|T} := \pi_T, \quad \pi_{t|T} := \mathbf{R}_{\pi_t}\pi_{t+1|T}, \quad t \in \{T-1, \dots, 0\}, \quad (6)$$

where for probability mass functions μ, ν the operator $\mathbf{R}_\nu\mu$ is defined as:

$$\mathbf{R}_\nu\mu(x) := \sum_{z \in \mathbb{X}^V} \frac{p(x, z)\nu(x)}{\sum_{\tilde{x} \in \mathbb{X}^V} p(\tilde{x}, z)\nu(\tilde{x})}\mu(z), \quad x \in \mathbb{X}^V.$$

Ghahramani and Jordan (1997) showed that for FHMMs the complexity of (4) and (6) together is $\mathcal{O}(TML^{M+1})$. The L^M part of this complexity makes implementation prohibitively costly as the dimension M grows.

3. Approximate filtering and smoothing

To introduce our approximate filtering and smoothing techniques — called the Graph Filter and Graph Smoother — consider the filtering and smoothing recursions, (4) and (6), and fix any partition \mathcal{K} of V . Suppose that one has already obtained an approximation to π_{t-1} , call it $\tilde{\pi}_{t-1}$, which factorizes with respect to \mathcal{K} . Then due to (1), $\mathbf{P}\tilde{\pi}_{t-1}$, also factorizes with respect to \mathcal{K} . However $\mathbf{C}_t\mathbf{P}\tilde{\pi}_{t-1}$ does not factorize with respect to \mathcal{K} in general. In Section 3.1 we shall define an approximation to the Bayes update operator \mathbf{C}_t , denoted $\tilde{\mathbf{C}}_t^m$, where m is a parameter, such that $\tilde{\pi}_t := \tilde{\mathbf{C}}_t^m\mathbf{P}\tilde{\pi}_{t-1}$ does factorize with respect \mathcal{K} .

Once $(\tilde{\pi}_t)_{t \in \{0, \dots, T\}}$ have been computed in this manner, a sequence of approximate smoothing distributions $(\tilde{\pi}_{t|T})_{t \in \{0, \dots, T\}}$ will be obtained by setting $\tilde{\pi}_{T|T} := \tilde{\pi}_T$ and then $\tilde{\pi}_{t-1|T} := \mathbf{R}_{\tilde{\pi}_t}\tilde{\pi}_{t|T}$, for $t = T, T-1, \dots, 0$. Due to (1) and the fact that $(\tilde{\pi}_t)_{t \in \{0, \dots, T\}}$ each factorizes with respect to \mathcal{K} , it follows that $(\tilde{\pi}_{t|T})_{t \in \{0, \dots, T\}}$ also factorize with respect to \mathcal{K} .

The key ingredient in all of this is finding a way to approximate the action of the Bayes update operator \mathbf{C}_t in an accurate but computationally inexpensive manner. Our next objective is to introduce the details of how we do so.

3.1 Approximate Bayes updates via localization and factorization

Let $d : (V \cup F)^2 \rightarrow \mathbb{N}$ be the graph distance on \mathcal{G} , that is $d(w, w')$ is the number of edges in a shortest path between w, w' . Augmenting the definition of the neighborhood function (3), define, for any $J \subseteq V$,

$$\begin{aligned} N_v^r(J) &:= \{v' \in V \text{ such that } \exists v \in J \text{ with } d(v, v') \leq 2r + 2\}, \\ N_f^r(J) &:= \{f \in F \text{ such that } \exists v \in J \text{ with } d(v, f) \leq 2r + 1\}. \end{aligned}$$

The sets $N_v^r(J)$ and $N_f^r(J)$ will play an important role in our theoretical results, we need these two definitions due to the bi-partite nature of the graph \mathcal{G} .

For a given probability mass function μ on \mathbb{X}^V , a partition of V denoted \mathcal{K} and $m \geq 0$ define:

$$\tilde{\mathcal{C}}_t^{m,K} \mu(x^K) := \frac{\sum_{z \in \mathbb{X}^V : z^K = x^K} \prod_{f \in N_f^m(K)} g^f(z^{N(f)}, y_t) \mu(z)}{\sum_{z \in \mathbb{X}^V} \prod_{f \in N_f^m(K)} g^f(z^{N(f)}, y_t) \mu(z)}, \quad x^K \in \mathbb{X}^K, K \in \mathcal{K}, \quad (7)$$

$$\tilde{\mathcal{C}}_t^m \mu := \bigotimes_{K \in \mathcal{K}} \tilde{\mathcal{C}}_t^{m,K} \mu. \quad (8)$$

Note the dependence of $\tilde{\mathcal{C}}_t^m$ on \mathcal{K} is not shown in the notation. Here m is key parameter which, as we shall see in more detail later, influences both the accuracy and computational cost of our approximate inference algorithms.

To see the motivation for (7)-(8), observe from the definition of the Bayes update operator (5) and factorial likelihood function (2) that the marginal distribution of $\mathcal{C}_t \mu$ associated with some $K \in \mathcal{K}$ is given by:

$$(\mathcal{C}_t \mu)^K(x^K) := \frac{\sum_{z \in \mathbb{X}^V : z^K = x^K} \prod_{f \in F} g^f(z^{N(f)}, y_t) \mu(z)}{\sum_{z \in \mathbb{X}^V} \prod_{f \in F} g^f(z^{N(f)}, y_t) \mu(z)}. \quad (9)$$

The definition (7)-(8) thus embodies two ideas: *localization*, in that $\tilde{\mathcal{C}}_t^{m,K} \mu$ is an approximation to the exact marginal $(\mathcal{C}_t \mu)^K$ obtained by replacing the likelihood function $\prod_{f \in F} g^f(z^{N(f)}, y_t)$ in (9) by the “local-to- K ” product $\prod_{f \in N_f^m(K)} g^f(z^{N(f)}, y_t)$; and *factorization*, in that $\tilde{\mathcal{C}}_t^m \mu$ factorizes with respect to \mathcal{K} by construction. Figure 3 illustrates sub-graphs of \mathcal{G} associated with each of the neighborhoods $N_f^m(K)$, $K \in \mathcal{K}$.

It is important to note that the factorization idea alone is not enough for our purposes: computing the marginal distribution $(\mathcal{C}_t \mu)^K$ has a cost which is exponential in M in general for likelihoods of the form (2), even when μ factorizes with respect to \mathcal{K} . So taking $\bigotimes_{K \in \mathcal{K}} (\mathcal{C}_t \mu)^K$ as an approximation to $\mathcal{C}_t \mu$ would offer no computational advantage. This distinguishes our setup from that of (Rebeschini and Van Handel, 2015; Finke and Singh, 2017) as discussed in Section 1, and is the reason we introduce localization through the parameter m .

More detailed consideration of the complexity of computing $\tilde{\mathcal{C}}_t^m \mu$ is given after our first main result, Proposition 1, which quantifies the approximation error associated with $\tilde{\mathcal{C}}_t^m$ and is one of the building blocks in the overall analysis of our approximate filtering and smoothing method.




Figure 3: An example of sub-graphs associated with the neighbourhoods $N_f^m(K)$, $K \in \mathcal{K}$, when $V = \{1, 2, 3, 4, 5\}$, $F = \{f_1, f_2, f_3, f_4\}$, $\mathcal{K} = \{\{1, 2\}, \{3\}, \{4\}, \{5\}\}$ and $m = 0$.

In order to state Proposition 1 we need some further definitions. Firstly let us introduce the following attributes of the factor graph \mathcal{G} .

$$d(J, J') := \min_{w \in J} \min_{w' \in J'} d(w, w'), \quad J, J' \subseteq V \cup F, \quad (10)$$

$$n_K := \frac{1}{2} \max_{v \in V} d(K, v), \quad K \in \mathcal{K},$$

$$n := \max_{K \in \mathcal{K}} n_K,$$

$$\Upsilon := \max_{v \in V} \text{card}(N(v)),$$

$$\Upsilon^{(2)} := \max_{v \in V} \text{card}(N_v^0(v)),$$

$$\tilde{\Upsilon} := \max_{v, v' \in V} \text{card}(N(v) \cap N(v')). \quad (11)$$

Note the dependence of n on \mathcal{K} is not shown in the notation.

Secondly, given a probability mass function μ on \mathbb{X}^V and a random variable $X \sim \mu$, we shall denote by μ_x^v the conditional distribution of X^v given $\{X^{V \setminus v} = x^{V \setminus v}\}$, and define

$$C_{v, v'}^\mu := \frac{1}{2} \sup_{x, z \in \mathbb{X}^V: x^{V \setminus v'} = z^{V \setminus v'}} \|\mu_x^v - \mu_z^v\|, \quad v, v' \in V,$$

$$\text{Corr}(\mu, \beta) := \max_{v \in V} \sum_{v' \in V} e^{\beta d(v, v')} C_{v, v'}^\mu,$$

where $\beta > 0$ is a given constant.

Proposition 1 Fix any partition \mathcal{K} of V and any $t \in \{1, \dots, T\}$. Suppose there exists $\kappa \in (0, 1)$ such that:

$$\kappa \leq g^f(x^{N(f)}, y_t) \leq \frac{1}{\kappa}, \quad \forall x \in \mathbb{X}^V, f \in F. \quad (12)$$

Assume that for a given probability mass function μ on \mathbb{X}^V there exists $\beta > 0$ such that:

$$2\kappa^{-2\Upsilon} \text{Corr}(\mu, \beta) + e^{2\beta\Upsilon(2)} (1 - \kappa^{2\tilde{\Upsilon}}) \leq \frac{1}{2}. \quad (13)$$

Then for any $K \in \mathcal{K}$, $J \subseteq K$ and $m \in \{0, \dots, n\}$,

$$\left\| \mathbf{C}_t \mu - \tilde{\mathbf{C}}_t^m \mu \right\|_J \leq 4e^{-\beta} (1 - \kappa^{b(m, \mathcal{K})}) \mathbf{card}(J) e^{-\beta m}, \quad (14)$$

where $b(m, \mathcal{K}) := 2 \max_{K \in \mathcal{K}} \max_{v \notin N_v^{m-1}(K)} \{\mathbf{card}(N(v))\}$, with the convention that the maximum over an empty set is zero.

The proof of Proposition 1 is given in appendix A. The term $\text{Corr}(\mu, \beta)$ quantifies the strength of dependence across the coordinates of $X = (X^v)_{v \in V} \sim \mu$. We say that μ satisfies the decay of correlation if it exists a β such that $\text{Corr}(\mu, \beta)$ is bounded above. The hypothesis (13) places a combined constraint on this dependence, the constant κ which in (12) controls the oscillation of the likelihood function factors $g^f(x^{N(f)}, y_t)$, and the graph attributes Υ , $\Upsilon^{(2)}$ and $\tilde{\Upsilon}$.

Turning to the bound (14), let us examine its dependence on the partition \mathcal{K} and the parameter m . The quantity $b(m, \mathcal{K})$ is non-increasing with m . In practice, $b(m, \mathcal{K})$ will often be decreasing with m , and is always zero when $m = n$ since then $N_v^{m-1}(K)$ is V . Also $b(m, \mathcal{K})$ will often decrease as \mathcal{K} becomes more coarse and in the extreme case of the trivial partition $\mathcal{K} = \{V\}$, the constant $b(m, \mathcal{K})$ is always zero, because $N_v^{m-1}(K)$ is again V , and so $v \notin V$ is equivalent to $v \in \emptyset$. Combined with the $e^{-\beta m}$ term, this means $\left\| \mathbf{C}_t \mu - \tilde{\mathbf{C}}_t^m \mu \right\|_J$ can be made small by choosing the partition \mathcal{K} to be suitably coarse and m to be suitably large.

It is important to note that $\text{Corr}(\mu, \beta)$, κ and Υ appearing in the hypotheses (12) and (13), and the quantities on the right hand side of (14) do not necessarily have any dependence on M , the overall dimension of the state-space. For instance, when $\mu = \otimes_{v \in V} \mu^v$ then $\text{Corr}(\mu, \beta) = 0$ and one can easily construct families of FHMMs of increasing dimension in which κ , Υ , $\Upsilon^{(2)}$, $\tilde{\Upsilon}$, $b(m, \mathcal{K})$ are independent of M : consider the simple case where the factor graph is a chain as shown in Figure 4, and the dimension of the model is increased by adding f_5 and $X^{(6)}$ then f_6 and $X^{(7)}$ as shown by the dashed lines. In this situation, for any $v \in V$ the cardinality of $N(v)$ and $N_v^0(v)$ remain unchanged as the dimension of the model increases.

Algorithm 1 shows the steps involved in computing $\tilde{\mathbf{C}}_t^m \mu$ in the case that μ factorizes with respect to \mathcal{K} . To simplify considerations of the computational cost of Algorithm 1, let us suppose that for each $K \in \mathcal{K}$ there exists a collection of elements in \mathcal{K} that is a partition of $N_v^m(K)$. This is a typical feature of regular graphs such as lattices. In this case the complexity of Algorithm 1 is readily found to be:

$$\mathcal{O} \left(\mathbf{card}(\mathcal{K}) \max_{K \in \mathcal{K}} \mathbf{card}(N_f^m(K)) L_{K \in \mathcal{K}}^{\max \mathbf{card}(N_v^m(K))} \right).$$




Figure 4: Solid lines indicate a chain factor graph with 4 likelihood factors and $V = \{1, 2, 3, 4, 5\}$, hence $M = 5$. Dashed lines indicate extension to $M = 6, 7$ by adding f_5 and $X^{(6)}$ then f_6 and $X^{(7)}$.

Algorithm 1 Approximate Bayes update

Require: $\mathcal{K}, (N_f^m(K))_{K \in \mathcal{K}}, (N_v^m(K))_{K \in \mathcal{K}}, (\mu^K)_{K \in \mathcal{K}}, (g^f(\cdot, y))_{f \in F}$

- 1: **for** $K \in \mathcal{K}$ **do**
 - 2: $\hat{K} \leftarrow \{K' \in \mathcal{K} : K' \cap N_v^m(K) \neq \emptyset\}$
 - 3: **for** $x \in \mathbb{X}^{\hat{K}}$ **do**
 - 4: $\hat{\mu}(x) \leftarrow \prod_{K' \in \hat{K}} \mu^{K'}(x^{K'})$
 - 5: **for** $f \in N_f^m(K)$ **do**
 - 6: $\hat{\mu}(x) \leftarrow \hat{\mu}(x) \cdot g^f(x^{N(f)}, y)$
 - 7: **Normalize** $\hat{\mu}$ **to a probability mass function on** $\mathbb{X}^{\hat{K}}$
 - 8: **Marginalize out components** $\hat{K} \setminus K$: $\tilde{\mu}^K \leftarrow \hat{\mu}^K$
- return** $(\tilde{\mu}^K)_{K \in \mathcal{K}}$
-

Crucially the exponent of L in this cost, which is proportional to $\max_{K \in \mathcal{K}} \text{card}(N_v^m(K))$, does not necessarily grow with the overall dimension M ; recall that $N_v^m(K)$ is, in words, the set of vertices belonging to V which are within $2m + 2$ graph distance of the set K . Thus $\text{card}(N_v^m(K))$ captures the density of edges in the graph in the neighborhood (defined by m) of K . Clearly this is a local rather than global characteristic of the graph. Taking the $\max_{K \in \mathcal{K}}$ of these cardinalities is a simple upper-bound on these local quantities across the graph. In the case of regular graphs such as lattices and in particular the chain example in Figure 4, the term $\max_{K \in \mathcal{K}} \text{card}(N_v^m(K))$ is independent of the total number of vertices in the graph. These complexity considerations for Algorithm 1 suggest that the overall cost of a filtering and smoothing method built around the approximate Bayes update operator $\tilde{\mathcal{C}}_t^m$ may avoid the exponential-in- M factor in the cost $\mathcal{O}(TML^{M+1})$ of exact filtering and smoothing for FHMMs.

3.2 Graph Filter

As an approximation to the operator F_t introduced in Section 2.2 we now define:

$$\tilde{F}_t^m := \tilde{\mathcal{C}}_t^m \mathbf{P},$$

where the dependence of \tilde{F}_t^m on \mathcal{K} , inherited from $\tilde{\mathcal{C}}_t^m$, is not shown in the notation. The approximate filtering distributions are then defined by the recursion:

$$\tilde{\pi}_0 := \mu_0, \quad \tilde{\pi}_t := \tilde{F}_t^m \tilde{\pi}_{t-1}, \quad t \in \{1, \dots, T\} \quad (15)$$

Our next result, Theorem 2, builds from Proposition 1 and quantifies the approximation error associated with $(\tilde{\pi}_t)_{t \in \{0, \dots, T\}}$. In order to state it we need to introduce, further to (10)-(11), the definitions for $J \subseteq V$,

$$\begin{aligned} \tilde{J} &:= \{v \in J : \forall f \in N(v), N(f) \subseteq J\}, \\ \partial J &:= J \setminus \tilde{J}, \\ \partial N(J) &:= \{f \in N(J) : N(f) \cap V \setminus J \neq \emptyset\}. \end{aligned} \tag{16}$$

Moreover, given a probability mass function μ on \mathbb{X}^V , transition probabilities p on \mathbb{X}^V and two random variables such that $X \sim \mu$ and $Z|X = x \sim p(x, \bullet)$ for $x \in \mathbb{X}^V$, we shall denote by $\mu_{x,z}^v$ the conditional distribution of X^v given $\{X^{V \setminus v} = x^{V \setminus v}, Z = z\}$, and define

$$\begin{aligned} \tilde{C}_{v,v'}^\mu &:= \frac{1}{2} \sup_{z \in \mathbb{X}^V} \sup_{\substack{x, \hat{x} \in \mathbb{X}^V: \\ x^{V \setminus v'} = \hat{x}^{V \setminus v'}}} \|\mu_{x,z}^v - \mu_{\hat{x},z}^v\|, \quad v, v' \in V \\ \widetilde{\text{Corr}}(\mu, \beta) &:= \max_{v \in V} \sum_{v' \in V} e^{\beta d(v,v')} C_{v,v'}^\mu, \end{aligned}$$

where $\beta > 0$ is a given constant. We note that if the components of X are independent, i.e. $\mu = \bigotimes_{v \in V} \mu^v$, then $\tilde{C}_{v,v'}^\mu = 0$ for any $v \neq v'$ and so $\widetilde{\text{Corr}}(\mu, \beta) = 0$ for any $\beta > 0$.

Theorem 2 *Fix any collection of observations $\{y_1, \dots, y_T\}$ and any partition \mathcal{K} of V . There exists a region $\mathcal{R}_0 \subseteq (0, 1)^3$ depending only on $\tilde{\Upsilon}, \Upsilon$ and $\Upsilon^{(2)}$, such that if, for given $(\epsilon_-, \epsilon_+, \kappa) \in \mathcal{R}_0$,*

$$\epsilon_- \leq p^v(x^v, z^v) \leq \epsilon_+, \quad \text{and} \quad \kappa \leq g^f(x^{N(f)}, y_t) \leq \frac{1}{\kappa},$$

for all $x, z \in \mathbb{X}^V, v \in V, f \in F, t \in \{1, \dots, T\}$, then for $\beta > 0$ small enough depending only on $\tilde{\Upsilon}, \Upsilon, \Upsilon^{(2)}, \epsilon_-, \epsilon_+$ and κ , we have that for any μ_0 satisfying:

$$\widetilde{\text{Corr}}(\mu_0, \beta) \leq 2e^{-\beta} \left(1 - \frac{\epsilon_-}{\epsilon_+}\right) + 2e^{2\beta} \Upsilon^{(2)} \left(1 - \kappa^{2\tilde{\Upsilon}}\right)$$

and for any $K \in \mathcal{K}, J \subseteq K$ and $m \in \{0, \dots, n\}$:

$$\|\pi_t - \tilde{\pi}_t\|_J \leq \alpha_1(\beta) \left(1 - \kappa^{a(\mathcal{K})}\right) \mathbf{card}(J) + \gamma_1(\beta) \left(1 - \kappa^{b(\mathcal{K}, m)}\right) \mathbf{card}(J) e^{-\beta m}, \quad \forall t \in \{1, \dots, T\},$$

where $\pi_t, \tilde{\pi}_t$ are given by (4) and (15) with initial state distribution μ_0 ; $\alpha_1(\beta), \gamma_1(\beta)$ are constants depending only on β , and

$$\begin{aligned} a(\mathcal{K}) &:= 2 \max_{K \in \mathcal{K}} \max_{v \in \partial K} \mathbf{card}(N(v) \cap \partial N(K)), \\ b(m, \mathcal{K}) &:= 2 \max_{K \in \mathcal{K}} \max_{v \notin N_v^{m-1}(K)} \mathbf{card}(N(v)), \end{aligned}$$

with the convention that the maximum over an empty set is zero.

The proof of Theorem 2 is in appendix A. Explicit expressions for \mathcal{R}_0 , β , $\alpha_1(\beta)$ and $\gamma_1(\beta)$ are given in the proof of the theorem and its supporting results. The full assumption on the initial distribution μ_0 can be found in appendix A.

The second term on the right hand side of the bound on $\|\pi_t - \tilde{\pi}_t\|_J$ given in Theorem 2 is, up to a numerical constant, equal to the upper bound obtained in Proposition 1, see discussion there of its dependence on m and \mathcal{K} . The first term on the right hand side of the bound on $\|\pi_t - \tilde{\pi}_t\|_J$ depends on the neighborhood structure of the factor graph \mathcal{G} . Loosely speaking, the constant $a(\mathcal{K})$ is small when the graph is sparsely connected and the partition is coarse, and in the extreme case of the trivial partition $\mathcal{K} = \{V\}$, the constant $a(\mathcal{K})$ is zero because, in the notation of (16), $\tilde{V} = V$ and so ∂V is empty. The quantities in the hypotheses and bound of Theorem 2 exhibit the same dimension-free qualities as discussed after Proposition 1.

Algorithm 2 Graph Filter

Require: $\mathcal{K}, (N_f^m(K))_{K \in \mathcal{K}}, (N_v^m(K))_{K \in \mathcal{K}}, (\mu_0^K)_{K \in \mathcal{K}}, (p^v(\cdot, \cdot))_{v \in V}, (g^f(\cdot, \cdot))_{f \in F}, (y_t)_{t=\{1, \dots, T\}}$

- 1: **for** $K \in \mathcal{K}$ **do**
- 2: $\tilde{\pi}_0^K \leftarrow \mu_0^K$
- 3: **Compute** $p^K(\cdot, \cdot) \leftarrow \prod_{v \in K} p^v(\cdot, \cdot)$
- 4: **for** $t \in \{1, \dots, T\}$ **do**
- 5: **for** $K \in \mathcal{K}$ **do**
- 6: **for** $z^K \in \mathbb{X}^K$ **do**
- 7: $\hat{\pi}^K(z^K) \leftarrow \sum_{x^K \in \mathbb{X}^K} p^K(x^K, z^K) \cdot \tilde{\pi}_{t-1}^K(x^K)$
- 8: $(\tilde{\pi}_t^K)_{K \in \mathcal{K}} \leftarrow$ **Algorithm 1** $\left(\mathcal{K}, (N_f^m(K))_{K \in \mathcal{K}}, (N_v^m(K))_{K \in \mathcal{K}}, (\hat{\pi}^K)_{K \in \mathcal{K}}, (g^f(\cdot, y_t))_{f \in F} \right)$
- 9: **return** $((\tilde{\pi}_t^K)_{K \in \mathcal{K}})_{t=\{0, \dots, T\}}$

Implementation of the approximate filtering method is shown in Algorithm 2, which we shall refer to from now on as the Graph Filter. Noting that the complexity of computing $\hat{\pi}^K \leftarrow (\mathbb{P}\mu)^K$, which is $\mathcal{O}(TL \max_{K \in \mathcal{K}} \text{card}(K)^2)$, is dominated by the cost of Algorithm 1, with an additional 2 at the exponent of L . The overall complexity of Algorithm 2 is then:

$$\mathcal{O} \left(T \max_{K \in \mathcal{K}} \text{card}(K) \max_{K \in \mathcal{K}} \text{card}(N_f^m(K)) L^{2 \max_{K \in \mathcal{K}} \text{card}(N_v^m(K))} \right).$$

Remark that when implementing the correction step the Bayes update is performed in the for loop over the partition. Even though such an implementation is faster, it requires a single loop over the partition, the theoretical computational cost does not change.

The reader is referred back to the end of section 3.1 for detailed discussion of the exponent of L appearing here. We note that for the chain graph example in Figure 4 with $\mathcal{K} = \{\{1\}, \{2\}, \dots\}$, the complexity is:

$$\mathcal{O} \left(T \max_{K \in \mathcal{K}} \text{card}(K) \min\{2(m+1), M-1\} L^{2 \min\{2(m+1)+1, M\}} \right).$$

Thus the exponent of L has the dimension free property of becoming independent of M depending on m when M is large enough.

3.3 Graph Smoother

The approximate smoothing distributions are defined by simply substituting the approximate filtering distributions into (6):

$$\tilde{\pi}_{T|T} := \tilde{\pi}_T, \quad \tilde{\pi}_{t|T} := \mathbf{R}_{\tilde{\pi}_t} \tilde{\pi}_{t+1|T}, \quad t \in \{T-1, \dots, 0\}. \quad (17)$$

Theorem 3 *Fix any collection of observations $\{y_1, \dots, y_T\}$ and any partition \mathcal{K} of V . There exists a region $\tilde{\mathcal{R}}_0 \subseteq (0, 1)^3$ depending only on $\tilde{\Upsilon}, \Upsilon$ and $\Upsilon^{(2)}$, such that if, for given $(\epsilon_-, \epsilon_+, \kappa) \in \tilde{\mathcal{R}}_0$,*

$$\epsilon_- \leq p^v(x^v, z^v) \leq \epsilon_+, \quad \text{and} \quad \kappa \leq g^f(x^{N(f)}, y_t) \leq \frac{1}{\kappa},$$

for all $x, z \in \mathbb{X}^V, v \in V, f \in F, t \in \{1, \dots, T\}$, then for $\beta > \log(2)$ small enough depending only on $\tilde{\Upsilon}, \Upsilon, \Upsilon^{(2)}, \epsilon_-, \epsilon_+$ and κ , we have that for any μ_0 satisfying:

$$\widetilde{\text{Corr}}(\mu_0, \beta) \leq 2e^{-\beta} \left(1 - \frac{\epsilon_-}{\epsilon_+}\right) + 2e^{2\beta} \Upsilon^{(2)} \left(1 - \kappa^{2\tilde{\Upsilon}}\right)$$

and for any $K \in \mathcal{K}, J \subseteq K$ and $m \in \{0, \dots, n\}$:

$$\|\tilde{\pi}_{t|T} - \pi_{t|T}\|_J \leq \alpha_2(\beta, \epsilon_-, \epsilon_+) \left(1 - \kappa^{a(\mathcal{K})}\right) \mathbf{card}(J) + \gamma_2(\beta, \epsilon_-, \epsilon_+) \left(1 - \kappa^{b(\mathcal{K}, m)}\right) \mathbf{card}(J) e^{-\beta m},$$

where $\pi_{t|T}, \tilde{\pi}_{t|T}$ are given by (6) and (17) with initial distribution μ_0 , $\alpha_2(\beta, \epsilon_-, \epsilon_+)$ and $\gamma_2(\beta, \epsilon_-, \epsilon_+)$ are constants depending on $\epsilon_-, \epsilon_+, \beta$ and

$$a(\mathcal{K}) := 2 \max_{K \in \mathcal{K}} \max_{v \in \partial K} \mathbf{card}(N(v) \cap \partial N(K)),$$

$$b(m, \mathcal{K}) := 2 \max_{K \in \mathcal{K}} \max_{v \notin N_v^{m-1}(K)} \mathbf{card}(N(v)),$$

with the convention that the maximum over an empty set is zero.

The proof of Theorem 3 is in appendix A. The only difference between the bound in this theorem and that in Theorem 2 are the constants $\alpha_2(\beta, \epsilon_-, \epsilon_+)$ and $\gamma_2(\beta, \epsilon_-, \epsilon_+)$, explicit expressions can be deduced from the proof. The full assumption on the initial distribution μ_0 can be found in appendix A.

The approximate smoothing method, shown in Algorithm 3 has complexity

$$\mathcal{O} \left(T \mathbf{card}(\mathcal{K}) L^{3 \max_{K \in \mathcal{K}} \mathbf{card}(K)} \right).$$

Assuming $3 \max_{K \in \mathcal{K}} \mathbf{card}(K)$ is smaller than $2 \max_{K \in \mathcal{K}} \mathbf{card}(N_v^m(K))$, which is typically the case in practice, the overall complexity of Algorithm 2 combined with Algorithm 3 is:

$$\mathcal{O} \left(T \mathbf{card}(\mathcal{K}) \max_{K \in \mathcal{K}} \mathbf{card}(N_f^m(K)) L^{2 \max_{K \in \mathcal{K}} \mathbf{card}(N_v^m(K))} \right).$$

The reader is referred back to the end of section 3.1 for detailed discussion of the exponent of L appearing here.

Algorithm 3 Graph Smoother

Require: $\mathcal{K}, (p^v(\cdot, \cdot))_{v \in V}, (\tilde{\pi}_t^K)_{K \in \mathcal{K}, t = \{0, \dots, T\}}$
 1: **for** $K \in \mathcal{K}$ **do**
 2: $\tilde{\pi}_{T|T}^K \leftarrow \tilde{\pi}_T^K$
 3: **Compute** $p^K(\cdot, \cdot) := \prod_{v \in K} p^v(\cdot, \cdot)$
 4: **for** $t \in \{T-1, \dots, 0\}$ **do**
 5: **for** $K \in \mathcal{K}$ **do**
 6: **for** $z^K \in \mathbb{X}^K$ **do**
 7: **for** $x^K \in \mathbb{X}^K$ **do**
 8: $\overleftarrow{p}^K(z^K, x^K) \leftarrow p^K(x^K, z^K) \tilde{\pi}_t^K(x^K)$
 9: **Normalize** $\overleftarrow{p}^K(z^K, \cdot)$ **to a probability mass function on** \mathbb{X}^K
 10: **for** $x^K \in \mathbb{X}^K$ **do**
 11: $\tilde{\pi}_{t|T}^K(x^K) \leftarrow \sum_{z^K \in \mathbb{X}^K} \overleftarrow{p}^K(z^K, x^K) \tilde{\pi}_{t+1|T}^K(z^K)$
 12: **return** $((\tilde{\pi}_{t|T}^K)_{K \in \mathcal{K}})_{t = \{1, \dots, T\}}$

4. Numerical results

Section 4.1 describes a class of FHMMs with conditionally Gaussian observations used as a running example in Ghahramani and Jordan (1997) and which we shall use in our numerical experiments. The purpose of the first set of experiments, in Section 4.1, is to illustrate the practical implications of our theoretical results, assessing the performance of the Graph Filter and Smoother methods against exact filtering and smoothing, both in terms of accuracy and computational speed. In Section 4.1 we compare the performance of EM algorithms for parameter estimation built around the Graph Smoother and variational approximations presented in Ghahramani and Jordan (1997). In Section 4.2 we outline a model of traffic flow on the London Underground and illustrate parameter estimation and prediction using the Graph Filter and Smoother.

We used the University of Bristol’s BlueCrystal High Performance Computing machine. The experiments were run on either one or two standard compute nodes each with 2 x 2.6GHz 8-CORE INTEL E5-2670 (SANDYBRIDGE) chips and 4GB of RAM per core.

4.1 Synthetic data

With \mathbb{X} a finite subset of \mathbb{Z} , $V = \{1, \dots, M\}$ and $\mathbb{Y} = \mathbb{R}^{d_y}$, consider the Gaussian emission model from Ghahramani and Jordan (1997):

$$g(x, y) = |\Sigma|^{-\frac{1}{2}} (2\pi)^{-\frac{d_y}{2}} \exp \left\{ -\frac{1}{2} [y - a(x)]^T \Sigma^{-1} [y - a(x)] \right\},$$

where $a(x)$ is a vector whose entries may depend on x .

We specialize to the case $d_y = M - 1$ and the specific forms of Σ and $a(x)$:

$$\Sigma = \sigma^2 I \quad \text{and} \quad a(x) = \left(a^f(x) \right)_{f \in \{1 \dots M-1\}} \quad \text{with} \quad a^f(x) := c \left(x^f + x^{f+1} \right), \quad (18)$$




Figure 5: Factor graph for the model (18) in the case $F = \{f^1, f^2, f^3, f^4\}$ and $V = \{1, 2, 3, 4, 5\}$.

where $c > 0$ is a constant. Under these assumptions, $N(f) = \{f, f + 1\}$ and

$$g^f(x^{N(f)}, y) = \frac{1}{\sqrt{2\pi\sigma}} \exp \left\{ -\frac{[y^f - c(x^f + x^{f+1})]^2}{2\sigma^2} \right\}.$$

The corresponding factor graph \mathcal{G} is a chain, illustrated in Figure 5. We also assume that the transition probabilities and initial distribution have identical components across V ,

$$p(x, z) = \prod_{v \in V} p^v(x^v, z^v), \quad \mu_0(x) = \prod_{v \in V} \mu_0^v(x^v), \quad p^v = \hat{p}, \quad \mu_0^v = \hat{\mu}_0, \quad \forall v \in V.$$

Throughout the experiments we fix the partition \mathcal{K} as:

$$\mathcal{K} = \{\{1\}, \dots, \{M\}\}.$$

Accuracy and speed performance for filtering and smoothing We took $\mathbb{X} = \{0, 1\}$ and simulated three data sets of length $T = 500$ from the model with parameters:

$$\hat{\mu}_0(x^v) = 1, \quad x^v = 1, \quad \forall v \in V, \quad \{\hat{p}(x^v, z^v)\}_{x^v, z^v \in \mathbb{X}} = \begin{pmatrix} 0.6 & 0.4 \\ 0.2 & 0.8 \end{pmatrix}, \quad c = 1, \quad \sigma^2 = 1. \quad (19)$$

First consider the execution time of the approximate filtering and smoothing method, i.e., the combination of Algorithm 2 and Algorithm 3, as a function of the parameters m and M . Figure 6 shows execution time as m and M vary. The execution time of exact filtering and smoothing using the algorithm of Ghahramani and Jordan (1997), henceforth “GJ”, is included for reference.

It is apparent from the top row of plots that with m fixed, the execution time of the Graph Filter and Smoother initially increases super-linearly with M up to some point which depends on m , and from then on it is linear in M . This is most visually evident for the large values of m and is consistent with the complexity of the combined Graph Filter and Smoother method discussed in Section 3.3, which for the model considered here is:

$$\mathcal{O} \left(TM \min\{2(m+1), M-1\} L^{2 \min\{2(m+1)+1, M\}} \right). \quad (20)$$

By contrast, the execution time of GJ increases exponentially with M , making its implementation extremely expensive in high-dimensional cases.

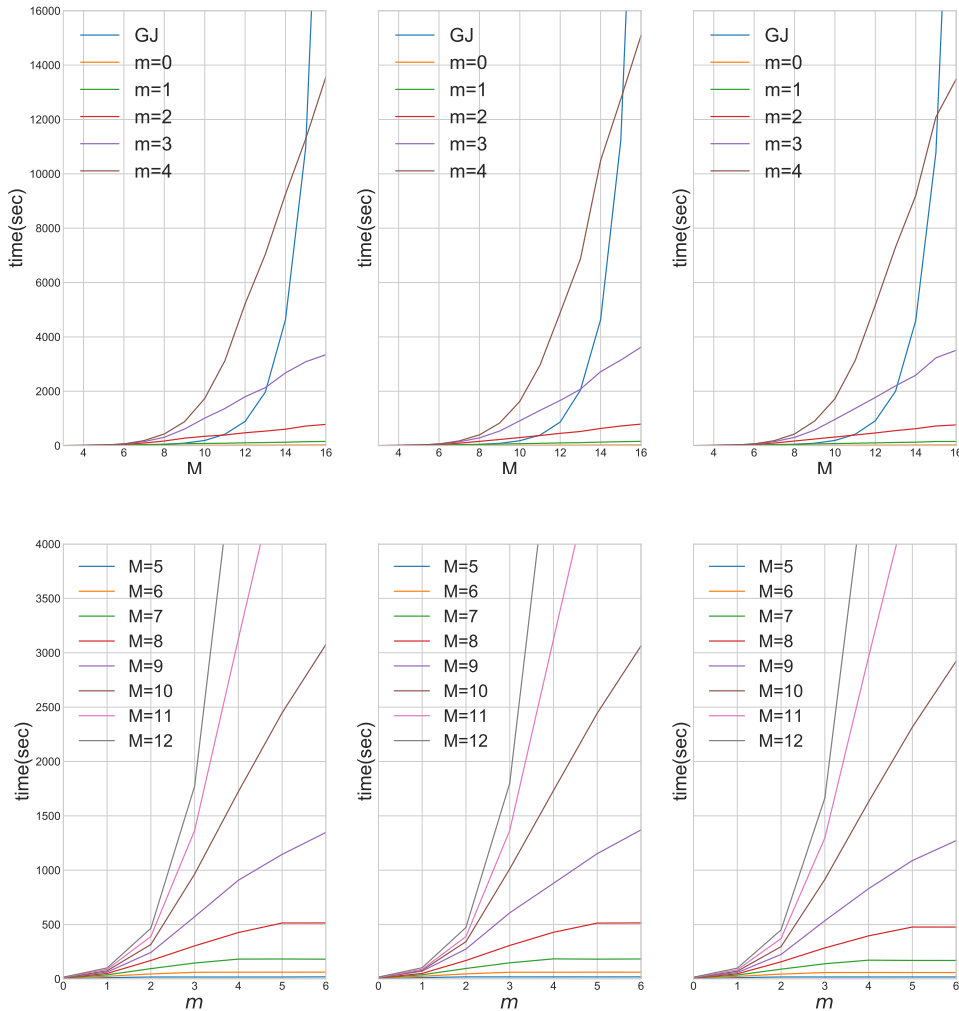


Figure 6: Execution time for the combined filtering and smoothing algorithms as a function of m and M . Each vertical pair of plots corresponds to one of three simulated data sets. GJ is the exact filtering and smoothing algorithm of Ghahramani and Jordan (1997).

When M is fixed, it is clear from the bottom row of plots in Figure 6 that the execution time of the Graph Filter and Smoother is super-linear in m up to some point which depends on M , and then is constant in m . Again this is consistent with (20). The phenomenon of the cost becoming constant in m arises because as m grows, eventually all factors are included in the products in $\tilde{C}_t^{m,K}$, see (7).

We now examine accuracy. Recall two important characteristics of the bound of Theorem 3: the bound does not depend on the overall dimension, M , and decays exponentially with m . The region $\tilde{\mathcal{R}}_0$ in Theorem 3 is non-empty, but for the specific parameter settings in (19) there does not exist $(\epsilon_-, \epsilon_+, \kappa) \in \tilde{\mathcal{R}}_0$ such that the assumptions of the theorem on p^v




Figure 7: The LTV distance between the approximate and exact marginal smoothing distributions averaged over both the components X_t^i, \dots, X_t^{i+4} , with $i = 1, \dots, 8$, and the $T = 500$ time steps. With m fixed the average LTV is constant in M for M large enough. The three plots correspond to the three simulated data sets.

and g^f hold. Thus technically Theorem 3 does not hold in this example. However, Figure 7 and Figure 8 encouragingly show that the LTV between the exact and approximate smoothing distributions exhibits the characteristics of not depending on the overall dimension, M , and decaying exponentially with m .

Comparison to variational inference within EM for parameter estimation Our next objective is to illustrate the accuracy of parameter estimation using the Graph Smoother within an approximate EM algorithm. We shall compare performance to the approximate EM approach of Ghahramani and Jordan (1997) in which variational approximations to the smoothing distributions are employed. For background on EM see Dempster et al. (1977) and Ghahramani and Jordan (1997) (Section 3.1).

Ghahramani and Jordan (1997) (Sections 3.4 and 3.5) describe two families of variational distributions for FHMM which can be used to compute the E-step in EM approximately: a “fully-factorized” scheme in which the variational distribution is chosen to statistically decouple all state variables, $(X_t^v)_{v \in V}$, $t = 0, \dots, T$, in the HMM, and a “structured” approximation, in which the variational distribution is Markovian in time but statistically decouples state variables across V . We shall refer to the former as *completely decoupled* and the latter as *spatially decoupled*.

The time complexity of computing the approximate smoothing distributions using either the completely decoupled or spatially decoupled schemes is:

$$\mathcal{O}(ITL^2M^2(M-1)^4), \quad (21)$$




Figure 8: The LTV distance between the approximate and exact marginal smoothing distributions of certain components of X_t^1, \dots, X_t^M and the $T = 500$ time steps. The LTV decreases exponentially with m . The three plots correspond to the three simulated data sets.

where I is the number of iterations of the fixed-point equations needed to find the variational approximation. In our experiments we found that $I = 20$ was sufficient for convergence, indeed Ghahramani and Jordan (1997) (pag. 254) suggest 2-10 iterations is typically sufficient. Recall that for M large enough, (20) is exponential in m , but linear in M , while (21) scales no faster than M^6 . Whether or not the variational approximations can be computed more quickly than the Graph Smoother is dependent on the model in question. In our experiments we did not find a substantial difference in speed.

Details of the EM updates using the Graph Smoother are given in appendix B. The only difference between these updates and those using the variational approximations is in the E-step, where the expectation is simply taken with respect to the corresponding approximate smoothing distribution.

Method	$\mu_0(0), \mu_0(1)$	c	σ^2	$p(0, 0), p(0, 1), p(1, 0), p(1, 1)$
True values	0.000, 1.000	2.000	4.000	0.600, 0.400, 0.200, 0.800
Graph Smoother $m = 0$	0.137, 0.863	1.753	4.642	0.072, 0.928, 0.518, 0.482
Graph Smoother $m = 1$	0.001, 0.999	1.779	4.542	0.075, 0.925, 0.549, 0.451
Completely decoupled Variational Bayes	0.384, 0.616	1.498	6.562	0.075, 0.925, 0.084, 0.916
Spatially decoupled Variational Bayes	0.393, 0.607	1.960	6.223	0.362, 0.638, 0.385, 0.615

Table 1: Parameters estimates for the case $M = 3$ with Graph Filter-Smoother and variational Bayes at the end of the EM algorithm. The estimates are found by taking the mean over the different initializations.




Figure 9: $M = 3$. Estimation of c and σ^2 using approximate EM based on the Graph Smoother and the completely and spatially decoupled variational approximations by Ghahramani and Jordan (1997). Horizontal axes correspond to EM iterations. 20 different EM initializations shown for each algorithm setting.

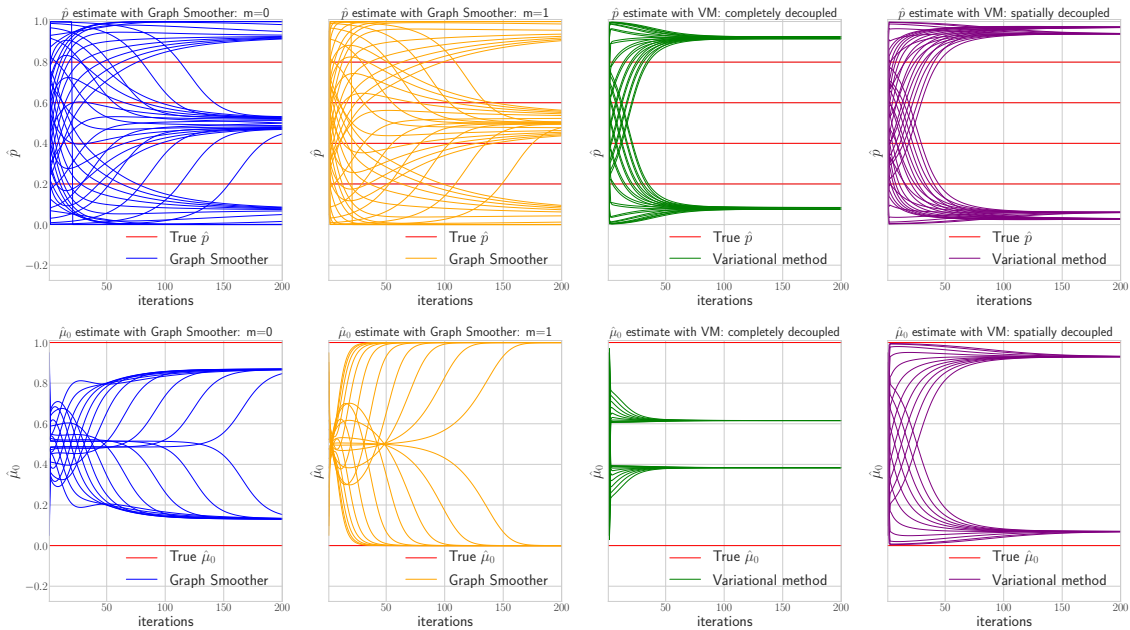


Figure 10: $M = 3$. Estimation of $\hat{\mu}_0$ and \hat{p} using approximate EM based on the Graph Smoother and the completely and spatially decoupled variational approximations by Ghahramani and Jordan (1997). Horizontal axes correspond to EM iterations. 20 different EM initializations shown for each algorithm setting. Traces corresponding to the 2 elements of the initial distribution $\hat{\mu}_0$ and 4 elements of the transition matrix \hat{p} are superimposed on each plot.

Method	$\mu_0(0), \mu_0(1)$	c	σ^2	$p(0, 0), p(0, 1), p(1, 0), p(1, 1)$
True values	0.000, 1.000	2.000	4.000	0.600, 0.400, 0.200, 0.800
Graph Smoother $m = 0$	0.206, 0.794	1.740	4.790	0.492, 0.508, 0.094, 0.906
Graph Smoother $m = 1$	0.496, 0.504	1.765	4.651	0.556, 0.444, 0.096, 0.904
Completely decoupled Variational Bayes	0.469, 0.531	1.833	6.260	0.343, 0.657, 0.336, 0.664
Spatially decoupled Variational Bayes	0.026, 0.974	1.423	5.804	0.044, 0.956, 0.040, 0.960

Table 2: Parameters estimates for the case $M = 10$ with Graph Filter-Smoother and variational Bayes at the end of the EM algorithm. The estimates are found by taking the mean over the different initializations.




Figure 11: $M = 10$. Estimation of c and σ^2 using approximate EM based on the Graph Smoother and the completely and spatially decoupled variational approximations by Ghahramani and Jordan (1997). Horizontal axes correspond to EM iterations. 20 different EM initializations shown for each algorithm setting.

In our experiments we considered the model described in Section 4.1, with $\mathbb{X} = \{0, 1\}$ and $T = 200$, and generated a data set with true parameter values:

$$\hat{\mu}_0(x^v) = 1, x^v = 1, \forall v \in V, \quad \{\hat{p}(x^v, z^v)\}_{x^v, z^v \in \mathbb{X}} = \begin{pmatrix} 0.6 & 0.4 \\ 0.2 & 0.8 \end{pmatrix}, \quad c = 2 \quad \text{and} \quad \sigma^2 = 4.$$

The EM algorithms based on the Graph Smoother and the fully and spatially decoupled variational approximations were run for 20 different EM initializations. Tables 1 and 2 show the mean of the final estimates over the different EM initializations for $M = 3$ and $M = 10$, respectively. Graphical results are presented in figures 9 and 10 for $M = 3$ and figures 11 and 12 for $M = 10$.

In Figure 9 the EM algorithms associated with the Graph Smoother converge to points closer to the true parameter values than those using the variational approximations. Using $m = 1$ rather than $m = 0$ in the former yields a slight increase in accuracy. In Figure 10 the results using the Graph Smoother are not more accurate in all cases, but the completely decoupled variational method generally performs badly. In Figure 11 the results for the Graph Smoother are again more accurate. In Figure 12 it is notable that the estimates of the transition probabilities are a little more accurate with $m = 0$ rather than $m = 1$, but substantially more accurate than with either of the variational schemes. Again the completely decoupled variational approximation performs poorly. These observations are numerically evident also from tables 1 and 2.

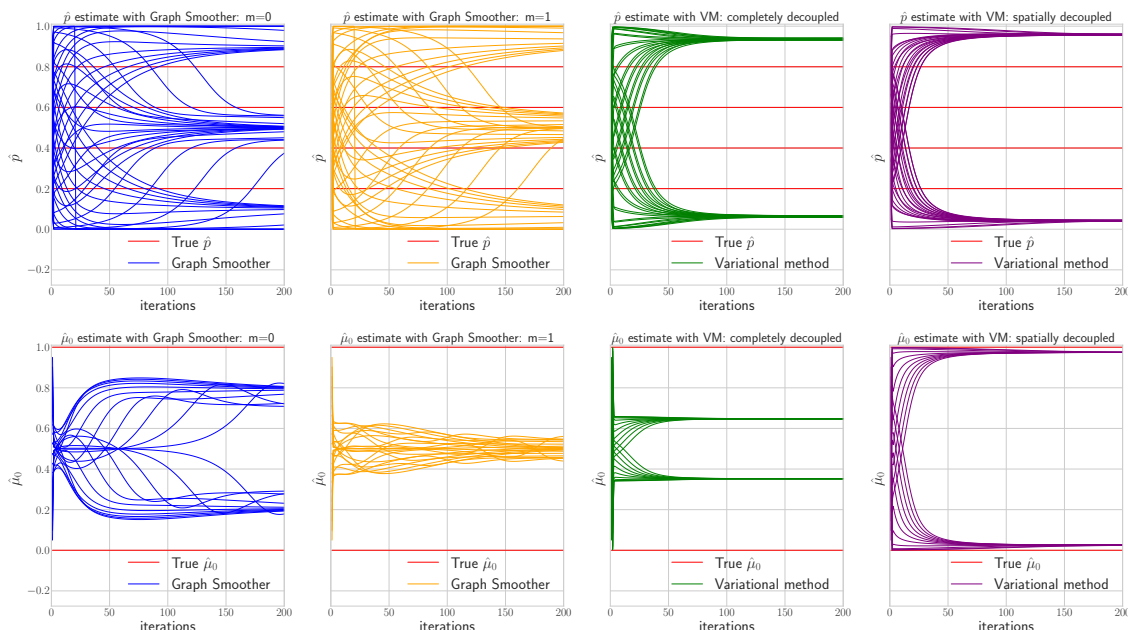


Figure 12: $M = 10$. Estimation of $\hat{\mu}_0$ and \hat{p} using approximate EM based on the Graph Smoother and the completely and spatially decoupled variational approximations by Ghahramani and Jordan (1997). Horizontal axes correspond to EM iterations. 20 different EM initializations shown for each algorithm setting. Traces corresponding to the 2 elements of the initial distribution $\hat{\mu}_0$ and 4 elements of the transition matrix \hat{p} are superimposed on each plot.

4.2 Analyzing traffic flows on the London Underground

Transport For London, the operator of the London Underground, has made publicly available “tap” data, consisting of a 5% sample of all Oyster card journeys in a week during November 2009 (Transport for London, 2018). The data consist of the locations and times of entry to and exit from the transport network for each trip.

Similar Transport for London data have been analyzed by Silva et al. (2015), who developed models of numbers of trips between pairs of stations in the Underground network in order to quantify the effects of shocks such as line and station closures, and to predict traffic volumes. The modelling approach described by Silva et al. (2015) (Supporting Information) is very sophisticated, including several components such as regression of the numbers of passengers entering stations onto time, a cascade of nonparametric binomial models for the numbers of passengers inside the transport system who entered at each station and a Bayesian probabilistic flow model. One of many attractive features of this approach is that it avoids the computational cost of network tomography models for traffic data (Guimera et al., 2005; Colizza et al., 2006; Newman et al., 2011) which is prohibitive in the context of large transport systems due to the exponential growth of problem size in the number of network links. Similar computational difficulties are encountered with some dynamic Bayesian network models of flow on transport networks. For example, Hofleitner et al. (2012) propose a dynamic mixture model for travel times where the mixture component represents a time-

varying congestion state associated with each link in the transport network. In principle, inference in this model can be performed using a particle filter, but as noted by Woodard et al. (2017), due to high-dimensionality the cost of doing so accurately (with respect to Monte Carlo error) is very demanding.

It is not our objective to conduct as detailed modelling exercise as in these works, but rather to establish a proof of principle that the Graph Filter and Smoother are naturally suited to the topological structure of transport networks and show promise calibrating the model by estimating parameters, and for prediction. This leaves potential for a deeper investigation of traffic modelling using FHMMs in future work.




Figure 13: Locations in longitude and latitude of 20 stations from the the Central line and the Jubilee line of the London underground. Stations are represented as yellow nodes, while train lines are green edges.

4.2.1 DATA AND MODEL

We consider 20 stations on a portion of the Central line and the Jubilee line. Stations' names and geographical locations are shown in figure 13. The dataset consists of integer counts of the inflow and the outflow of passengers from Monday to Friday per station every 10 minutes from 00:00 am to 00:00 am of the next day. The data are split into training given by Monday, Tuesday and Wednesday and test consisting of Thursday and Friday.

In constructing our model we consider two factors per station, one for the inflow, the other for the outflow: for each $f \in \{1, \dots, 20\}$, each corresponding to a station, we denote by $y_t^{f,\text{in}}$ and $y_t^{f,\text{out}}$ respectively the counts of passenger inflow and outflow at time t . We consider a hidden state variable associated with each direction of travel on each tube line segment connecting a pair of stations. These state variables are written $x_t^{i,j}$ to indicate the state of the line segment between stations i and j at time t in the direction from i to j .

Thus the index set V of the FHMM has one element corresponding to each direction of travel between each pair of stations which are connected by a single line segment on the tube network. The state-space of each variable $x_t^{i,j}$ is defined to be $\mathbb{X} = \{0, 1, 2, 3\}$, with the interpretation of increasing levels of congestion on the line-segment and in the direction of travel corresponding to (i, j) .

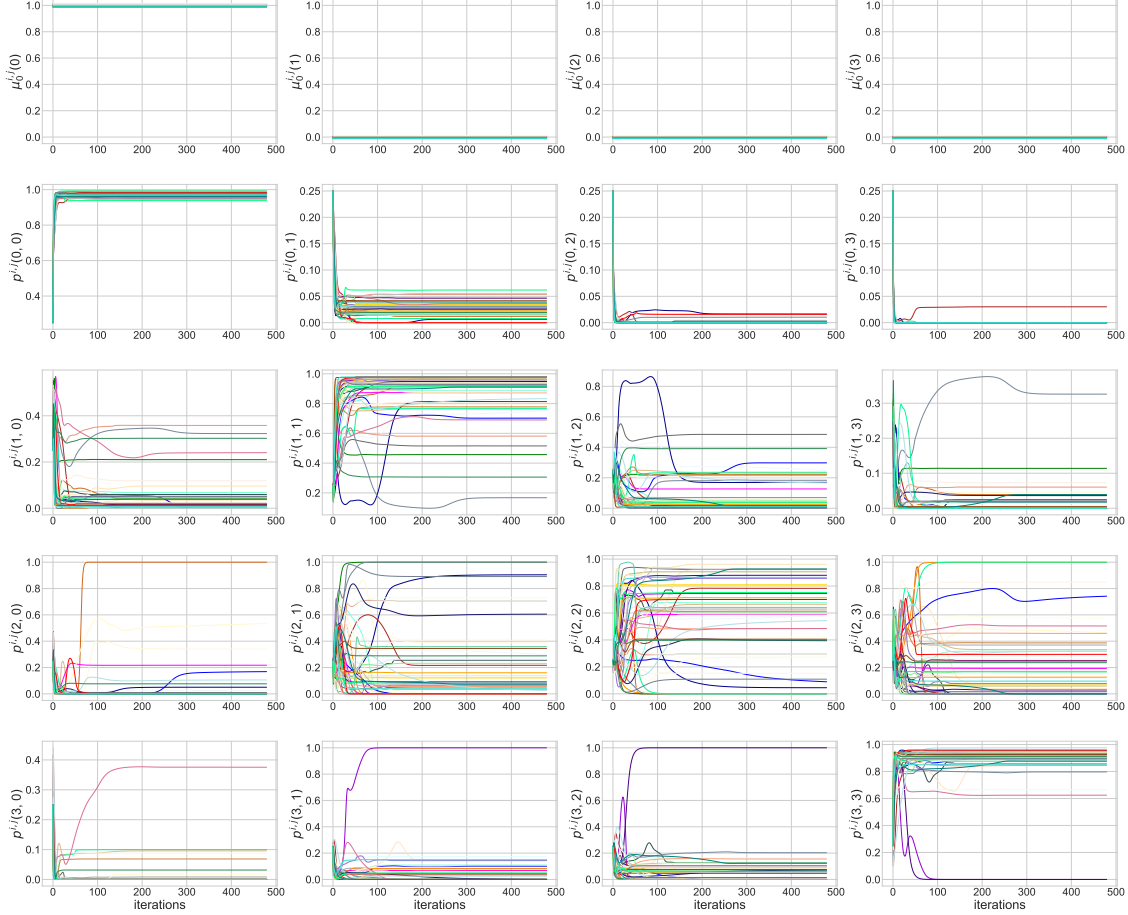


Figure 14: Estimation of the of initial distribution and transition matrix for each tube line using the approximate EM algorithm built around the Graph Filter-Smoother algorithm. See appendix B for algorithm details. Each coloured line corresponds to a different line and direction in the tube’s network.

The emission distribution is given by:

$$g(x, y) = \prod_{f=1}^{20} g^{f,\text{in}} \left(x^{f,N(f)}, y^{f,\text{in}} \right) g^{f,\text{out}} \left(x^{N(f),f}, y^{f,\text{out}} \right), \quad x \in \mathbb{X}^V,$$




Figure 15: EM estimates using the Graph Filter-Smoother algorithm, of $\lambda^{f,\text{in}}$ and $\lambda^{f,\text{out}}$. Each coloured line correspond to a different station flow (inflow or outflow) in the tube's network.

where $g^{f,\text{in}}(x^{f,N(f)}, y^{f,\text{in}})$ and $g^{f,\text{out}}(x^{N(f),f}, y^{f,\text{out}})$ are Poisson distributions:

$$g^{f,\text{in}}(x^{f,N(f)}, y^{f,\text{in}}) = \frac{\left(\lambda^{f,\text{in}} \sum_{j \in N(f)} x^{f,j}\right)^{y^{f,\text{in}}}}{(y^{f,\text{in}})!} e^{-\left(\lambda^{f,\text{in}} \sum_{j \in N(f)} x^{f,j}\right)}$$

$$g^{f,\text{out}}(x^{N(f),f}, y^{f,\text{out}}) = \frac{\left(\lambda^{f,\text{out}} \sum_{i \in N(f)} x^{i,f}\right)^{y^{f,\text{out}}}}{(y^{f,\text{out}})!} e^{-\left(\lambda^{f,\text{out}} \sum_{i \in N(f)} x^{i,f}\right)}$$

where $x_t^{f,N(f)} = (x_t^{f,j})_{j \in N(f)}$ and $x_t^{N(f),f} = (x_t^{i,f})_{i \in N(f)}$. Each pair of parameters $\lambda^{f,\text{in}}, \lambda^{f,\text{out}}$ has the interpretation as the overall intensity of inflow and outflow at the station corresponding to factor f . The initial distribution and the transition matrix are written:

$$\mu_0(x) = \prod_{i=1}^{20} \prod_{j \in N(i)} \mu_0^{i,j}(x^{i,j}), \quad x \in \mathbb{X}^V$$

$$p(x, z) = \prod_{i=1}^{20} \prod_{j \in N(i)} p^{i,j}(x^{i,j}, z^{i,j}), \quad x, z \in \mathbb{X}^V.$$

For this model we have $M = 48$ (the number of line segments between stations multiplied by two for the directions of travel on each line segment) and $L = 4$, hence the cardinality of the overall state-space \mathbb{X}^V is 4^{48} . The factor graph has $\max_{K \in \mathcal{K}} \text{card}(N_v^m(K)) = 7$ and $\max_{K \in \mathcal{K}} \text{card}(N_f^m(K)) = 4$ hence the computational cost of the Graph Filter-Smoother, when $m = 0$ and $\mathcal{K} = \{\{1\}, \dots, \{48\}\}$, is $\mathcal{O}[(48 \cdot 4)TL^{2 \cdot 7}]$, compared to $\mathcal{O}(48TL^{48+1})$ for the FHMM forward-backward algorithm proposed by Ghahramani and Jordan (1997).

The total number of parameters of the model 1000: 40 from the flow intensity (2 for each of the 20 stations), 768 from the transition matrices, and 192 from the initial distribution.

We now describe four sets of experiments. In the first set we demonstrate parameter estimation. In the second and third sets of experiments we illustrate predictive capabilities and comparison to a Long short-term Memory (LSTM) recurrent neural network, without and with missing data. In the fourth set of experiments we consider prediction in the presence of structural changes corresponding to hypothetical disruption and modification of the tube network.

Parameter estimation. The parameters are estimated using an approximate EM algorithm, where the approximate smoothing distributions from the Graph Filter-Smoother are used in the E-step, more details are available in appendix B. We took $m = 0$ and \mathcal{K} as in subsection 4.1 (partition of singleton over V). The algorithm was run on the training set with random EM initializations for the estimated parameters. The estimates are shown in figure 14 and figure 15. The estimates are found to be robust with respect to the variability in the EM initializations.

The estimate of the initial distribution for each $x_0^{i,j}$ puts very high probability on state 0, which is to be expected since the tube lines are closed in the very early hours of the morning.

The interpretation of the estimated transition probabilities is less straightforward, but we can discern interesting structure: upon inspecting the results we found that the state variables $x_t^{i,j}, i \neq j$ could be roughly partitioned into four groups according to their estimated transition probabilities. The first group, which we shall heuristically refer to as “stable”, are those for which the transition probabilities put most weight on maintaining a constant state, i.e. high probabilities on the diagonal of the transition matrix. Most state variables belong to this group. The second group, which we shall refer to as “quiet”, consists of those variables for which there is little probability of transitioning to state $\{3\}$, which corresponds to the highest level of flow. State variables in the “quiet” group typically correspond to line segments connecting hubs and low-flow stations (e.g Liverpool Street inflow with Aldgate outflow); high-flow stations and low-flow stations (e.g Canary Wharf inflow with North Greenwich outflow); low-flow stations and low-flow stations (e.g. Stepney Green inflow with Whitechapel outflow). We call the third group “busy”, distinguished by transition probabilities in which transitions to states $\{2, 3\}$ are somewhat likely. These estimates are typical for lines connecting high-flow stations with high-flow stations, e.g. Liverpool Street inflow with Bank outflow.

The interpretation of figure 15 is straightforward, indeed $\lambda^{f,\cdot}$ scales the intensity of the flow, hence bigger estimates of $\lambda^{f,\cdot}$ are referred to stations with higher flows.

Passenger flow prediction without missing data. Prediction over the test data is performed through the posterior predictive of Graph Filter-Smoother, i.e. per each time step $t \geq 0$ a posterior predictive sample y_{t+1} is obtained as:

$$x_t \sim \tilde{\pi}_{t|t}, \quad x_{t+1} \sim p(x_t, \cdot), \quad y_{t+1} \sim g(x_{t+1}, \cdot), \quad (22)$$

where g, p are computed through the EM algorithm run on the training set (Monday, Tuesday, Wednesday) and $\tilde{\pi}_{\cdot|\cdot}$ is obtained by running recursively the Graph Filter on the test data (Thursday, Friday). Note that (22) provides a sample from the posterior predictive, whose mean is then used as prediction for the test set. For ease of presentation we report the plots on four stations, plots on all the other stations are available in appendix C. The

results indicate that the model is able to track the peaks of the inflow and outflow that occur during the morning and afternoon rush hours, which vary in magnitude from station to station. Moreover, credible intervals show satisfying coverage of the true data.

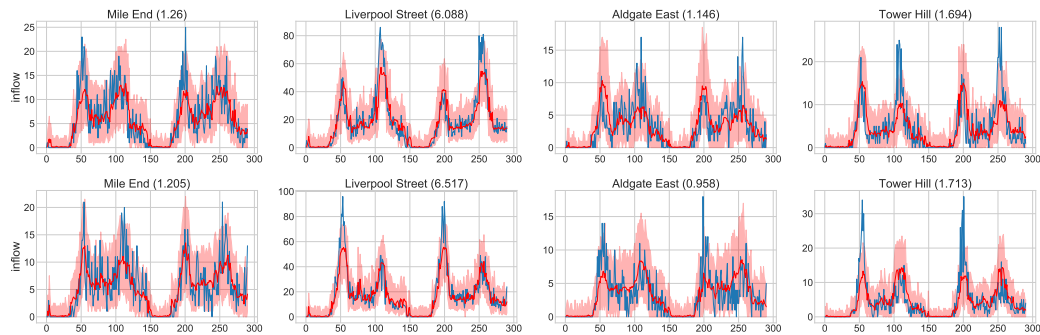


Figure 16: One step-ahead posterior predictive mean (solid red line) and 95% credible intervals (red bands) using the Graph Filter-Smoother on four stations. Blue solid lines stand for the observed data from Thursday to Friday. The first row shows the inflow, the second row shows the outflow. The name of the station is reported at the top of each plot, along with the estimate of the corresponding λ^f .

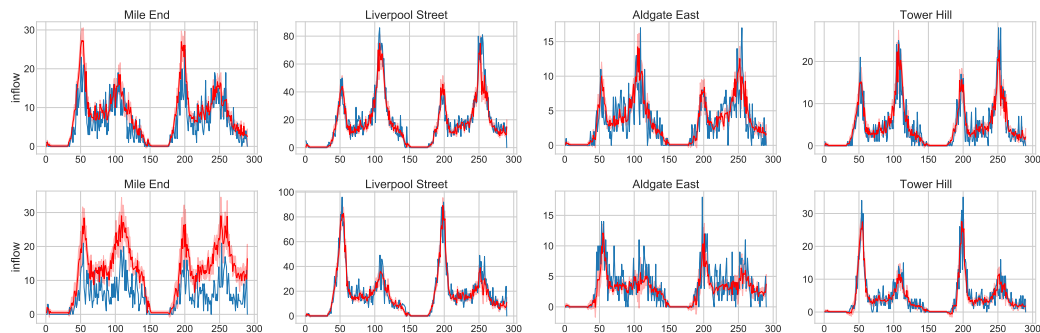


Figure 17: One step-ahead prediction with the LSTM. Per time step, a sample of size 100 is built over initializations and random number seeds used in training the LSTM, where solid red lines show the mean and red bands show the region between the 0.025 and the 0.975 quantiles. Blue solid lines stand for the observed data from Thursday to Friday. The first row reports the inflow, the second row reports the outflow. The name of the station is written at the top of each plot.

We compare the proposed method with an LSTM trained on one-step-ahead prediction over the inflow-outflow, more details about the architecture and training are available in appendix C. The LSTM takes as input inflow-outflow data over all the stations at time t and output predictions for time $t + 1$. 100 LSTMs (different initializations and seeds) are trained on Monday, Tuesday and Wednesday. As for the Graph Filter-Smoother method, testing is performed on Thursday and Friday. We report in figure 17 the LSTMs predictions on four

stations only, more details are available in the appendix C. An LSTM does not itself provide any uncertainty quantification associated with its predictions. We employ a commonly used heuristic of considering the variability of predictions across different initializations and random number seeds of the stochastic gradient algorithm used to train the LSTM. Figure 17 shows that different trainings of the LSTM lead to similar performances with bands that are narrower than figure 16.

Method	No missing	Missing without peak	Missing with peak
Graph Filter-Smoother	4.796 ± 0.011	5.399 ± 0.014	5.493 ± 0.019
LSTM	4.639 ± 0.138	7.427 ± 2.080	6.178 ± 1.425

Table 3: RMSE comparison between the posterior predictive mean of Graph Filter-Smoother and LSTM. “No missing” refers to the performance on the full test set. “Missing without peak” refers to the test set performance when data from $t = 130$ (around 9 pm on Thursday) to $t = 170$ (around 4 am on Friday) are missing. “Missing with peak” refers to the test set performance when data from $t = 230$ (around 2 pm on Friday) to $t = 270$ (around 9 pm on Friday) are missing.

Table 3 reports RMSEs for the posterior predictive mean of Graph Filter-Smoother and the LSTM. The mean and standard deviation of the RMSE for the posterior predictive mean of Graph Filter-Smoother are computed over 100 samples of the posterior predictive mean. The mean and standard deviation of the RMSE for LSTM are computed over 100 LSTMs optimization (i.e. different initializations and seeds). The proposed algorithm has performances that are comparable with the LSTM and we note here that the LSTM has been explicitly trained to minimize the RMSE on the one-step-ahead prediction.

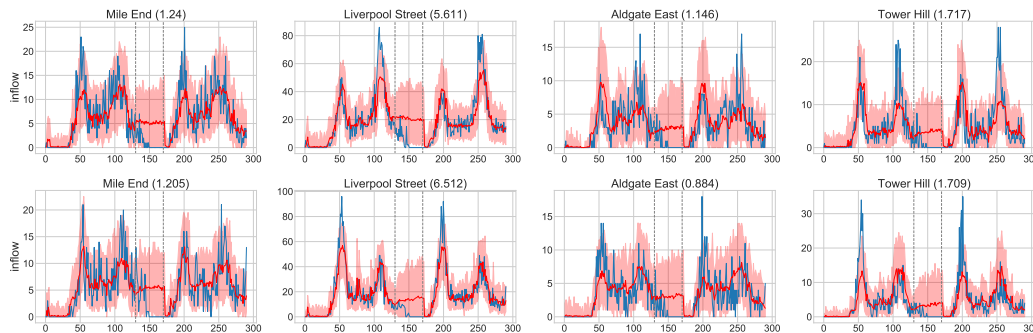


Figure 18: Multi-step-ahead posterior predictive mean (solid red line) and 95% credible intervals (red bands) using the Graph Filter-Smoother on four stations with missing data in a quiet period (without peak). Blue solid lines stand for the observed data from Thursday to Friday (the missing data are included). Grey dashed lines show the start and the end of the missing data window. The name of the station is reported at the top of each plot, along with the estimate of the corresponding λ^f .

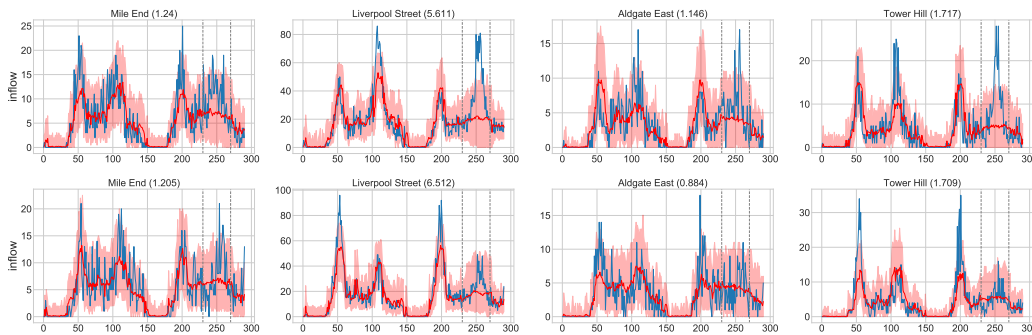


Figure 19: Multi-step-ahead posterior predictive mean (solid red line) and 95% credible intervals (red bands) using the Graph Filter-Smoother on four stations with missing data in a busy period (with peak). Blue solid lines stand for the observed data from Thursday to Friday (the missing data are included). Grey dashed lines show the start and the end of the missing data window. The name of the station is reported at the top of each plot, along with the estimate of the corresponding λ^f .

Passenger flow prediction with missing data. Suppose now that data cannot be collected in the period from t to $t + h$. In practice this could correspond, for example, to hardware malfunctions at stations such that people come in and out without tapping their Oyster cards. The Graph Filter can easily be used to impute missing data by applying the “prediction” operator without the “correction” operator: we can sample from the posterior predictive from t to $t + h$ as follows:

$$\begin{aligned}
 x_t &\sim \tilde{\pi}_t|t, \\
 x_{t+1} &\sim p(x_t, \cdot), \quad \tilde{y}_{t+1} \sim g(x_{t+1}, \cdot), \\
 x_{t+2} &\sim p(x_{t+1}, \cdot), \quad \tilde{y}_{t+1} \sim g(x_{t+2}, \cdot), \\
 &\dots \\
 x_{t+h} &\sim p(x_{t+h-1}, \cdot), \quad \tilde{y}_{t+h} \sim g(x_{t+h}, \cdot).
 \end{aligned}$$

Similarly, the LSTM can do multi-step-ahead prediction by using the output on the current time step as input for the next time step. We note that the considered LSTM is trained exclusively for one-step-ahead prediction (i.e. LSTM maps y_t onto y_{t+1}). As an alternative, one could train LSTM on multi-step-ahead predictions (i.e. LSTM maps y_t onto y_{t+1}, \dots, y_{t+h}), which requires to know the missing data window in advance. However, that is generally not plausible in practice, e.g. missing data in traffic applications are caused by broken sensors and the time they are out of order cannot be predicted in advance.

RMSE performances between LSTM and the proposed method in a missing data scenario are compared in the second and the third column of Table 3. Two cases are distinguished: missing data in a quiet period (second column) and missing data in a busy period (third column). The LSTM performance is evidently not robust, with substantial variability in estimates corresponding to different initializations. On the contrary, the proposed method

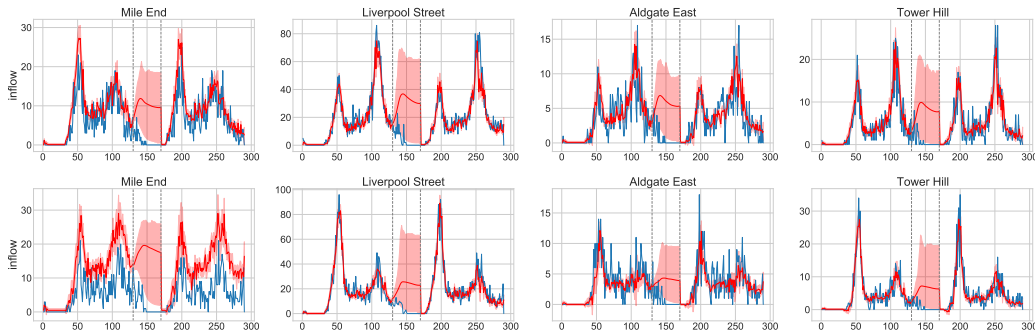


Figure 20: Multi-step-ahead prediction with the LSTM on four stations with missing data in a quiet period (without peak). Per time step, a sample of size 100 is built over different training of the LSTM where solid red lines show the mean and red bands show the region between the 0.025 and the 0.975 quantiles. Blue solid lines stand for the observed data from Thursday to Friday. Grey dashed lines show the start and the end of the missing data window. Stations' names are reported at the top of each plot.

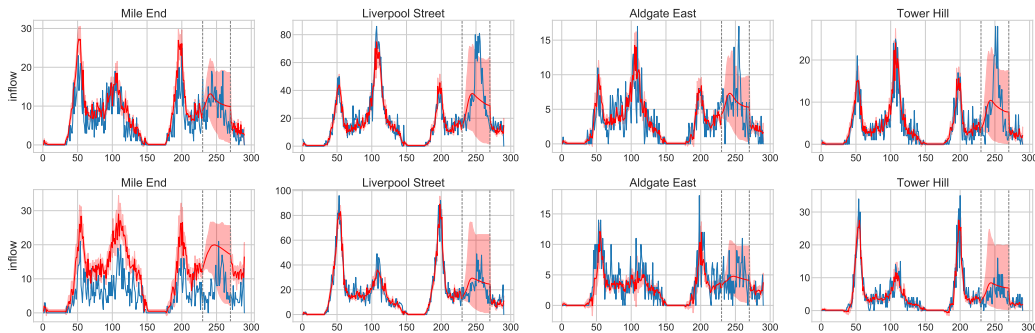


Figure 21: Multi-step-ahead prediction with the LSTM on four stations with missing data in a busy period (with peak). Per time step, a sample of size 100 is built over different training of the LSTM where solid red lines show the mean and red bands show the region between the 0.025 and the 0.975 quantiles. Blue solid lines stand for the observed data from Thursday to Friday. Grey dashed lines show the start and the end of the missing data window. Stations' names are reported at the top of each plot.

is more stable, with a standard deviation that is 100 times lower than LSTM and an RMSE that is significantly lower in mean. This appears in the experiments for missing data in both busy and quiet periods. Graphical illustrations on selected stations can be found in figures 18, 19, 20, 21, more figures can be found in appendix C.

Flow prediction under structural change - transport network disruptions and modifications. In this final set of experiments we demonstrate multi-step ahead prediction in a scenario where a change in the model occurs at a known point in time. We consider two types of change, which have the interpretation of i) hypothetical disruptions (or line segment closures) on the tube network, and ii) hypothetical modifications to the tube network by adding lines which do not currently exist. Making predictions in these scenarios is

achievable because the hidden state variables have a straight-forward interpretation. This is in contrast, for example, to an LSTM, whose parameters are not interpretable in terms of the lines of the tube network and their respective levels of congestion.

For the case of a hypothetical disruption, the procedure is as follows. We first estimate model parameters using the Graph Filter-Smoother and EM as described above, using the data up to time point 200. Beyond time step 200 we then perform multi-step ahead prediction, again as described above, but with the state variable $x_t^{i,j}$ corresponding to the disrupted line conditioned to zero. For the case of a hypothetical addition of a line to the tube network, we again estimate model parameters up to time point 200. For time steps beyond 200 we introduce additional state variables $x_t^{i,j}$ corresponding the new lines and (for simplicity) set the prior Markov transition probability matrix for these individual state variables to the element-wise average of the corresponding matrices for all other lines in the network. We then perform multi-step ahead prediction for the augmented network.

Figure 22 shows how the predicted passenger flows behave in two scenarios: one when a hypothetical disruption on the line between Bank and London Bridge is enforced, and one when the tube network is hypothetically modified by adding a line between Mile End and North Greenwich. These results illustrate that the disruption leads to decreased inflow-outflow at the stations at either end of the affected line segmented, while adding a line segment between two stations results in increases inflow-outflow of the connected stations.

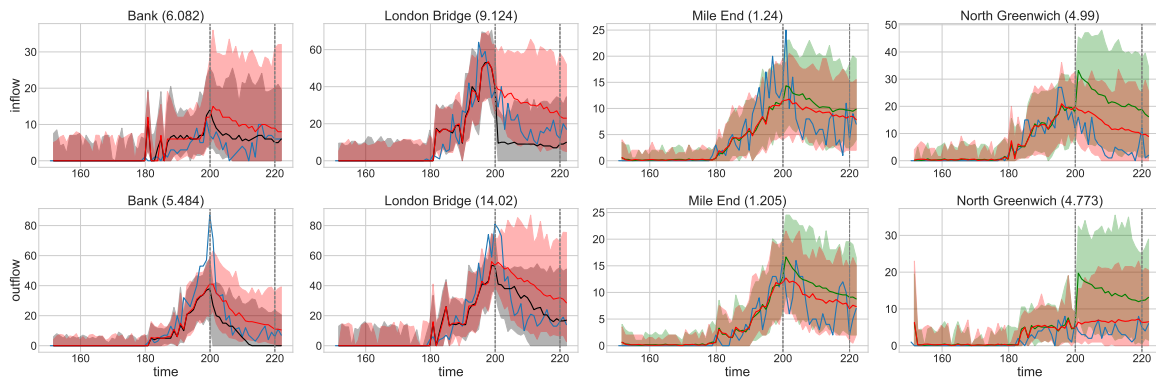


Figure 22: Predicted flows for Bank and London Bridge under a hypothetical disruption at time 200 on the line between Bank and London Bridge (first and second columns), and for Mile End and North Greenwich when adding a line between these two stations at time 200 (third and fourth columns). Top row: in-flow, bottom row: outflow. Black solid line and grey shading: posterior predictive mean and 95% credible interval in the presence of the hypothetical disruption. Green solid line and green shading: posterior predictive mean and 95% credible interval in the case when a hypothetical line is added. Red solid line and light pink shading: posterior predictive mean and 95% credible interval without disruption or addition of a line. Darker pink/brown shading indicates intersection with respectively grey/green. Blue solid lines are the observed real data, in absence of disruption or addition of a line. Each subfigure title gives the station name and estimated $\lambda^{f,\cdot}$ value based on the data up to time 200.

5. Conclusion

There are number of ways in which this work could be extended or generalized. From an algorithmic point of view, it seems natural to explore hybrids between variational methods of the sort proposed by Ghahramani and Jordan (1997) and the Graph Filter-Smoother. In the spirit of variational methods, could a layer of optimization in Kullback-Leibler divergence be somehow combined with the approximations in Graph Filter-Smoother, with the aim of yielding even more accurate approximations?

From a theoretical point of view, it would be interesting to investigate whether the kind of mathematical tools we have used to study the Graph Filter and Smoother could also help rigorously quantify the approximation error in variational methods for FHMMs, which appears to be an unsolved problem. Another interesting question beyond the scope of the present work is whether the generalized Dobrushin Comparison Theorems developed by Rebeschini and van Handel (2014) might help loosen some of the technical assumptions on which our results rely.

From a modelling point of view, there are many directions in which the London Underground example could be further investigated. The model we have considered is a simple prototype, with straightforward spacial interactions among lanes and stations. It would be interesting to see how it performs with spacial interactions over time, i.e. lanes do not evolve independently from each other, and with a richer state space. Even the emission distribution could be chosen more carefully, indeed a zero-inflated Poisson could be used to model more precisely the quiet periods.

6. Acknowledgments

Nick Whiteley’s research is supported by a fellowship from the Alan Turing Institute.

References

- Kidus Asfaw, Joonha Park, Allister Ho, Aaron A King, and Edward Ionides. Statistical inference for spatiotemporal partially observed markov processes via the r package spatpomp. *arXiv preprint arXiv:2101.01157*, 2021.
- Leonard E Baum and Ted Petrie. Statistical inference for probabilistic functions of finite state Markov chains. *The Annals of Mathematical Statistics*, 37(6):1554–1563, 1966.
- Leonard E Baum, Ted Petrie, George Soules, and Norman Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *The Annals of Mathematical Statistics*, 41(1):164–171, 1970.
- Eva Besada-Portas, Sergey M Plis, M Jesus, and Terran Lane. Parallel subspace sampling for particle filtering in dynamic bayesian networks. In *Machine Learning and Knowledge Discovery in Databases*, pages 131–146. Springer, 2009.
- Peter Bickel, David Choi, Xiangyu Chang, and Hai Zhang. Asymptotic normality of maximum likelihood and its variational approximation for stochastic blockmodels. *The Annals of Statistics*, 41(4):1922–1943, 2013.

- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Chinmoy B Bose and Shyh-Shiaw Kuo. Connected and degraded text recognition using hidden Markov model. *Pattern Recognition*, 27(10):1345–1363, 1994.
- Xavier Boyen and Daphne Koller. Tractable inference for complex stochastic processes. In *Proc. 14th Conf. Uncertainty Artif. Intell.*, pages 33–42. Morgan Kaufmann Publishers Inc., 1998.
- Xavier Boyen and Daphne Koller. Exploiting the architecture of dynamic systems. In *Proc. 16th Nat. Conf. Artif. Intell.*, pages 313–320, 1999.
- Bruno Cedraz Brandao, Jacques Wainer, and Siome Klein Goldenstein. Subspace hierarchical particle filter. In *Proc. 19th Brazilian Symp. Comput. Graph. Process*, pages 194–204. IEEE, 2006.
- Alain Celisse, Jean-Jacques Daudin, and Laurent. Pierre. Consistency of maximum-likelihood and variational estimators in the stochastic block model. *Electronic Journal of Statistics*, 6:1847–1899, 2012.
- Vittoria Colizza, Alain Barrat, Marc Barthélemy, and Alessandro Vespignani. The role of the airline transportation network in the prediction and predictability of global epidemics. *Proceedings of the National Academy of Sciences*, 103(7):2015–2020, 2006.
- Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38, 1977.
- Axel Finke and Sumeetpal S Singh. Approximate smoothing and parameter estimation in high-dimensional state-space models. *IEEE Transactions on Signal Processing*, 65(22):5982–5994, 2017.
- Hans-Otto Georgii. *Gibbs measures and phase transitions*, volume 9. Walter de Gruyter, 2011.
- Zoubin Ghahramani and Michael I. Jordan. Factorial hidden Markov models. *Machine Learning*, 29(2):245–273, Nov 1997. ISSN 1573-0565. doi: 10.1023/A:1007425814087. URL <https://doi.org/10.1023/A:1007425814087>.
- Roger Guimera, Stefano Mossa, Adrian Turtschi, and LA Nunes Amaral. The worldwide air transportation network: Anomalous centrality, community structure, and cities’ global roles. *Proceedings of the National Academy of Sciences*, 102(22):7794–7799, 2005.
- Aude Hofleitner, Ryan Herring, Pieter Abbeel, and Alexandre Bayen. Learning the dynamics of arterial traffic from probe data using a dynamic Bayesian network. *IEEE Transactions on Intelligent Transportation Systems*, 13(4):1679–1693, 2012.
- James P Hughes, Peter Guttorp, and Stephen P Charles. A non-homogeneous hidden Markov model for precipitation occurrence. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 48(1):15–30, 1999.

- Genshiro Kitagawa. Non-gaussian state-space modeling of nonstationary time series. *Journal of the American statistical association*, 82(400):1032–1041, 1987.
- Jeremias Knoblauch and Theodoros Damoulas. Spatio-temporal bayesian on-line change-point detection with model selection. In *International Conference on Machine Learning*, pages 2718–2727. PMLR, 2018.
- Mark Newman, Albert-Laszlo Barabasi, and Duncan J Watts. *The structure and dynamics of networks*, volume 19. Princeton University Press, 2011.
- Brenda Ng, Leonid Peshkin, and Avi Pfeffer. Factored particles for scalable monitoring. In *Proc. 18th Conf. Uncertainty Artif. Intell.*, pages 370–377. Morgan Kaufmann Publishers Inc., 2002.
- Joonha Park and Edward L Ionides. Inference on high-dimensional implicit dynamic models using a guided intermediate resampling filter. *Statistics and Computing*, 30(5):1497–1522, 2020.
- Lawrence R Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- Patrick Rebeschini and Ramon van Handel. Comparison theorems for gibbs measures. *Journal of Statistical Physics*, 157(2):234–281, 2014.
- Patrick Rebeschini and Ramon Van Handel. Can local particle filters beat the curse of dimensionality? *The Annals of Applied Probability*, 25(5):2809–2866, 2015.
- Ricardo Silva, Soong Moon Kang, and Edoardo M Airoldi. Predicting traffic volumes and estimating the effects of shocks in massive transportation systems. *Proceedings of the National Academy of Sciences*, page 201412908, 2015.
- Mario Stanke and Stephan Waack. Gene prediction with a hidden Markov model and a new intron submodel. *Bioinformatics*, 19(suppl_2):ii215–ii225, 2003.
- Transport for London. Oyster card journey information, 2018. <https://api-portal.tfl.gov.uk/docs>, section: “Oyster card data”.
- Peter Jan Van Leeuwen. Particle filtering in geophysical systems. *Monthly Weather Review*, 137(12):4089–4114, 2009.
- Bo Wang and DM Titterton. Convergence properties of a general algorithm for calculating variational bayesian estimates for a normal mixture model. *Bayesian Analysis*, 1(3):625–650, 2006.
- Dawn Woodard, Galina Nogin, Paul Koch, David Racz, Moises Goldszmidt, and Eric Horvitz. Predicting travel time reliability using mobile phone GPS data. *Transportation Research Part C: Emerging Technologies*, 75:30–44, 2017.
- Yaxian Xu and Ajay Jasra. Particle filters for inference of high-dimensional multivariate stochastic volatility models with cross-leverage effects. *Foundations of Data Science*, 1(1):61, 2019.