

Department of Econometrics and Business Statistics

<http://business.monash.edu/econometrics-and-business-statistics/research/publications>

Updating Variational Bayes: Fast sequential posterior inference

Nathaniel Tomasetti, Catherine S. Forbes and
Anastasios Panagiotelis

August 2020

Working Paper 27/20
(Revised working paper 13/19)

Updating Variational Bayes: Fast sequential posterior inference

Nathaniel Tomasetti* Catherine S. Forbes† Anastasios Panagiotelis‡

August 16, 2020

Abstract

Variational Bayesian (VB) methods produce posterior inference in a time frame considerably smaller than traditional Markov Chain Monte Carlo approaches. Although the VB posterior is an approximation, it has been shown to produce good parameter estimates and predicted values when a rich classes of approximating distributions are considered. In this paper we propose the use of recursive algorithms to update a sequence of VB posterior approximations in an on-line, time series setting, with the computation of each posterior update requiring only the data observed since the previous update. We show how importance sampling can be incorporated into online variational inference allowing the user to trade accuracy for a substantial increase in computational speed. The proposed methods and their properties are detailed in two separate simulation studies. Two empirical illustrations of the methods are provided, including one where a Dirichlet Process Mixture model with a novel posterior dependence structure is repeatedly updated in the context of predicting the future behaviour of vehicles on a stretch of the US Highway 101.

Keywords: Importance Sampling, Forecasting, Clustering, Dirichlet Process Mixture, Variational Inference

JEL Classifications: C11, G18, G39.

1 Introduction

Time series data often arrives in high frequency streams in applications that may require a response within a very short period of time. For example, self-driving vehicles may need to constantly monitor the position of each surrounding vehicle, predict or infer the behaviour of their human drivers, and react accordingly. In this context, the most recently received data can be highly informative for very short-term predictions, if the inferred models can be processed very quickly in an online fashion. In order to account for uncertainty in the models or predictions, Bayesian updating methods may be employed by targeting a sequence of posterior distributions, each conditioned on an

*Department of Econometrics and Business Statistics, Monash University, VIC, 3800, Australia. Email: nathaniel.tomasetti@monash.edu. Code available at https://github.com/NTomasetti/UVB_Code

†Corresponding author. Department of Econometrics and Business Statistics, Monash University, VIC, 3800, Australia. Email: catherine.forbes@monash.edu. This author acknowledges financial support under the Australian Research Council Discovery Grant No. DP150101728 and the National Science Foundation Grant SES-1921523.

‡Department of Econometrics and Business Statistics, Monash University, VIC, 3800, Australia. Email: anastasios.panagiotelis@monash.edu.

expanding dataset. The computational demands of such an algorithm may be improved if the incorporation of additional data does not require the re-use of any observations that have previously been conditioned upon.

In many empirical settings, the desired Bayesian posterior distributions are not analytically tractable. In such cases posterior inference may be obtained using Markov chain Monte Carlo (MCMC) methods, which result in a (dependent) sample from the posterior. Unfortunately, this approach typically involves relatively slow algorithms that are incompatible with the time frames demanded by streaming data, as the computation of each posterior update involves the entire currently observed dataset. Further, while particle filtering methods for sequential posterior updating have been developed both for static parameter models (Chopin, 2002) and dynamic latent variable models (e.g. Doucet et al., 2001), these available methods appear to be too slow for practical online use. This is particularly the case when they require use of the entire dataset to avoid particle degeneracy and/or when the number of inferred parameters is large. For a recent review of particle filtering methods, see Doucet and Lee (2018). An alternative approach appears in Jasra et al. (2010) and Del Moral et al. (2015), who apply Approximate Bayesian Computation (ABC) for sequential posterior updating, however this involves an embedded particle filter that similarly scales poorly to higher dimensional models. Bhattacharya and Wilson (2018) provide a sequential method to update parameter inference, however their grid-based posterior evaluation is suitable only for low dimensions. Taking a different approach, Chen et al. (2019) learn the parameters of a so-called flow operator, a neural network that approximates a function which maps a set of particles from a posterior distribution at one time period, and additional data, to a set of particles belonging to an updated posterior distribution.

An alternative approach that has grown in popularity in the recent literature for high dimensional models is the so-called Variational Bayes (VB) method (see Zhang et al., 2017, for a review). VB approximates the posterior with a tractable family of distributions, and chooses a member of this family by minimising a particular loss function with respect to auxiliary parameters. Early work in VB found an optimal approximation with coordinate descent algorithms in exponential family models, an approach widely known as Mean Field Variational Bayes (MFVB, see Jordan et al., 1999, Attias, 1999 Ghahramani and Beal, 2000, Wainwright and Jordan, 2008). Recent developments in VB consider gradient based algorithms (Ranganath et al., 2014, Kingma and Welling, 2014), which allow for a much richer class of models and approximating distributions to be utilised. These gradient based approaches are stochastic, and target the true gradient of a given loss function with an unbiased estimator. We refer to this approach as Stochastic Variational Bayes (SVB).

There is a rich tradition of using only a subset of the complete dataset for certain aspects of VB inference, such as for gradient estimation. Hoffman et al. (2010) and Wang et al. (2011) propose MFVB algorithms for Dirichlet Process Mixture (DPM) models where the optimisation of a subset of the auxiliary parameter vector occurs through gradient based approaches, using a subsample of the complete data at each iteration. Hoffman et al. (2013) and Titsias and Lázaro-Gredilla (2014) implement this data subsampling into the fully gradient based SVB approaches. Alternatively, Sato (2001) considers an alternative loss function defined as the expected value of the Kullback-Leibler (KL) divergence, with respect to the data generating process. Any realisation from the data generating process may be used within the MFVB coordinate descent algorithm, which is applied online with newly observed data substituted in as it becomes available. However, each of these approaches results in only a single posterior distribution conditioned on data up to some pre-specified time period T_n , and do not provide a mechanism for the approximation to be updated at a later time period T_{n+1} following the availability of additional observations.

Smidl (2004) and Broderick et al. (2013) each consider VB approximations for Bayesian updating, resulting in a progressive sequence of approximate posterior distributions that each condition on data up to any given time period T_n . Their approaches update to the time T_{n+1} by substitution of the time T_n posterior with MFVB approximations, which are feasibly obtained due to assuming the model and approximation each adhere to a suitably defined exponential family form. In these special settings, MFVB is able to linearly combine the available optimally converged auxiliary parameters. While Smidl (2004) is concerned with state space models, Broderick et al. (2013) considers application to a latent Dirichlet allocation problem, and shows it performs favourably compared to the approach of Hoffman et al. (2010) in terms of log predictive score and computational time.

In this paper we formalise and extend the prior approximation approach, developing an algorithm that we call Updating Variational Bayes (UVB). This algorithm can be seen as an application of the framework of Broderick et al. (2013) to the time series setting. UVB can be applied to sequentially update posterior distributions, and in a manner suitable for applications of streaming data. UVB treats data as arriving in a sequence, with the production of recursive, but approximate, posterior distributions obtained from conditioning on past information at nominated time points according to a Bayesian updating scheme. The approach delivers the approximate posterior distributions to the user at any desired point in time, with each new update using only the data observed since the previous update time. UVB requires an optimisation step for each update, which may be too slow for practical use.

To reduce the computational load of repeated updates, we extend UVB to a second, and completely novel algorithm, called Updating Variational Bayes with Importance Sampling (UVB-IS). Significant gains in computation speed per update can be achieved, albeit with some potential cost in gradient estimator variance and subsequently accuracy. Our proposed UVB-IS shares some similarities with the gradient estimator of Sakaya and Klami (2017), however the important distinction is that our proposed UVB-IS is developed for the sequential updating setting. To the best of our knowledge our approach is the first to exploit Importance Sampling in an online Variational Bayes framework.

We provide two simulation studies: a small scale time series forecasting application, and a larger application clustering time series, to compare the approximation error of each of SVB, UVB, and UVB-IS relative to exact inference obtained using MCMC. We also compare the computational time required by each of the variational approximations, and show that UVB-IS is substantially faster than either UVB or SVB, while incurring only a minor cost in performance, dependent on the application. We also demonstrate the application of UVB and UVB-IS to a simple hierarchical model to re-analyse the eight schools problem of Gelman et al. (1997), and measure the increased approximation error from the updating approaches relative to SVB in this setting.

Finally we demonstrate the application of UVB to the problem of updating posterior inference in the context of a DPM model. Here the aim is to provide Bayesian inference and prediction regarding the heterogeneous behaviour of 500 drivers from the New Generation Simulation dataset (FHWA, 2017), according to the distribution of their lateral lane position. In this context, data arrives rapidly. We introduce a new class of dependent approximating distributions, and show that the DPM model with UVB based inference is able to provide more accurate forecasts than those achieved using a standard MFVB based approach. UVB in this case has accuracy comparable to

repeated use of (full data) SVB, but benefits from an ability to process updates sequentially as additional data arrives.

The paper is arranged as follows: in Section 2 we review standard VB methods and the available gradient algorithms commonly employed. In Section 3 we propose our main UVB approach, with the UVB-IS extension detailed in Section 4. Next, Section 5 contains simulation studies for time series data and a mixture distribution, while Section 6 details applications of the newly proposed methods to the Eight Schools hierarchical model of Gelman et al. (2014). UVB is applied to a vehicle DPM model in Section 7, and Section 8 concludes the paper.

2 Background on Variational Bayes

Before introducing our novel algorithms for recursively updating approximations to the posterior, the main ideas associated with the implementation of an SVB approach are introduced. A more detailed description of SVB can be found in Blei et al. (2017), with further references provided therein.

The usual target of Bayesian inference is the posterior distribution for a potentially vector-valued static parameter $\boldsymbol{\theta}$, as characterised by its probability density function (pdf) denoted by $p(\boldsymbol{\theta}|\mathbf{y}_{1:T})$. Here $\mathbf{y}_{1:T}$ denotes data observed from time 1 to T and the posterior pdf is obtained using Bayes' theorem, given by

$$p(\boldsymbol{\theta}|\mathbf{y}_{1:T}) = \frac{p(\mathbf{y}_{1:T}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{\int_{\boldsymbol{\theta}} p(\mathbf{y}_{1:T}|\boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta}}, \quad (1)$$

where $p(\boldsymbol{\theta})$ denotes the pdf for the prior distribution that characterises belief about $\boldsymbol{\theta}$ prior to the observation of $\mathbf{y}_{1:T}$. Although MCMC algorithms are commonly used to produce a (typically dependent) sample from this posterior distribution, these can be computationally intensive.

As an alternative to MCMC, VB aims to approximate the pdf in (1) with another density of given parametric form, denoted by $q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$. Here $\boldsymbol{\lambda}$ is a vector of auxiliary parameters associated with the approximation, to be selected via optimisation. We note that the distribution q is explicitly parameterised by $\boldsymbol{\lambda}$, and any evaluation of this distribution does not require $\mathbf{y}_{1:T}$. The approximating distribution depends on $\mathbf{y}_{1:T}$ only through the choice of $\boldsymbol{\lambda}$. We include the conditioning notation to reinforce that the selected $q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$, which is an approximation to the posterior distribution in (1). This point is particularly relevant to the process of updating VB as is shown in Section 3.

In the SVB context, the family of the approximating distribution $q_{\boldsymbol{\lambda}}$ is held fixed, with the member of that family indexed by the parameter vector $\boldsymbol{\lambda}$ selected to minimise a given loss function. Typically the KL divergence (Kullback and Leibler, 1951) from $q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$ to $p(\boldsymbol{\theta}|\mathbf{y}_{1:T})$, denoted as $KL[q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T}) || p(\boldsymbol{\theta}|\mathbf{y}_{1:T})]$, is used, with

$$KL[q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T}) || p(\boldsymbol{\theta}|\mathbf{y}_{1:T})] = E_q [\log(q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})) - \log(p(\boldsymbol{\theta}|\mathbf{y}_{1:T}))]. \quad (2)$$

We note that the KL divergence is not symmetric, and reversing $q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$ and $p(\boldsymbol{\theta}|\mathbf{y}_{1:T})$, leads to similarities with assumed density filtering, independently proposed in statistics (Lauritzen, 1992) and artificial intelligence (Boyan and Koller, 2013; Opper and Winther, 1998) and control (Maybeck, 1982). Often in practice, the KL divergence in (2) is intractable, with $p(\boldsymbol{\theta}|\mathbf{y}_{1:T})$ only known up to a proportionality constant due to the difficulties involved in the evaluation of the integral in the denominator of (1). Nevertheless, it has been shown that an equivalent problem to minimising the

KL divergence is to maximise the so-called *evidence lower bound* (ELBO Attias, 1999), given by

$$\mathcal{L}(q, \boldsymbol{\lambda}) = E_q [\log(p(\boldsymbol{\theta}, \mathbf{y}_{1:T})) - \log(q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T}))]. \quad (3)$$

A further complication that typically arises when attempting to implement SVB is that an analytical expression for the expectation in (3) may not be available. In this case, maximisation of the ELBO may be achieved via stochastic gradient ascent (SGA, Bottou, 2010). To apply SGA to the problem of maximising the ELBO, an initial value $\boldsymbol{\lambda}^{(1)}$ is selected and is recursively modified to $\boldsymbol{\lambda}^{(m)}$, for $m = 1, 2, \dots$, according to

$$\boldsymbol{\lambda}^{(m+1)} = \boldsymbol{\lambda}^{(m)} + \rho^{(m)} \left. \frac{\partial \widehat{\mathcal{L}}(q, \boldsymbol{\lambda})}{\partial \boldsymbol{\lambda}} \right|_{\boldsymbol{\lambda}=\boldsymbol{\lambda}^{(m)}} \quad (4)$$

with the converged value of $\boldsymbol{\lambda}^{(*)}$ obtained when the change from $\mathcal{L}(q, \boldsymbol{\lambda}^{(m)})$ to $\mathcal{L}(q, \boldsymbol{\lambda}^{(m+1)})$ falls below some pre-specified threshold (Hoffman et al., 2013).

The adjustment term in (4) is made of two factors, the so-called *learning rate*, $\rho^{(m)}$, and an estimate of the gradient of the ELBO, $\frac{\partial \widehat{\mathcal{L}}(q, \boldsymbol{\lambda})}{\partial \boldsymbol{\lambda}}$. A popular estimator of this gradient is the score-based estimator (Ranganath et al., 2014), given by

$$\frac{\partial \widehat{\mathcal{L}}(q, \boldsymbol{\lambda})}{\partial \boldsymbol{\lambda}}_{SC} = \frac{1}{S} \sum_{j=1}^S \frac{\partial \log(q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}^{(j)}|\mathbf{y}_{1:T}))}{\partial \boldsymbol{\lambda}} \left(\log(p(\mathbf{y}_{1:T}, \boldsymbol{\theta}^{(j)})) - \log(q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}^{(j)}|\mathbf{y}_{1:T})) - \widehat{\mathbf{a}} \right), \quad (5)$$

where the simulated values $\{\boldsymbol{\theta}^{(j)}, \text{ for } j = 1, 2, \dots, S\}$ are drawn from the presiding approximating density $q_{\boldsymbol{\lambda}^{(m)}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$, and $\widehat{\mathbf{a}}$ is a vector of control variates with

$$\widehat{a}_k = \frac{\widehat{\text{Cov}} \left(\frac{\partial \log(q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T}))}{\partial \lambda_k} (\log(p(\mathbf{y}_{1:T}, \boldsymbol{\theta})) - \log(q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T}))), \frac{\partial \log(q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T}))}{\partial \lambda_k} \right)}{\widehat{\text{Var}} \left(\frac{\partial \log(q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}|\mathbf{y}_{1:T}))}{\partial \lambda_k} \right)}. \quad (6)$$

As (5) results in an unbiased estimator of the gradient of the ELBO, it is known that the SGA procedure will converge in probability to a local maximum (Robbins and Monro, 1951), provided that the learning rate sequence¹ satisfies

$$\sum_{m=1}^{\infty} \rho^{(m)} = \infty \quad (7)$$

and

$$\sum_{m=1}^{\infty} (\rho^{(m)})^2 < \infty. \quad (8)$$

We note that although SGA is itself a recursive procedure, the result in the VB context is the one-time posterior pdf approximation $q_{\boldsymbol{\lambda}^*} \approx p(\boldsymbol{\theta}|\mathbf{y}_{1:T})$, where $\boldsymbol{\lambda}^* = \boldsymbol{\lambda}^{(M)}$ is the optimal parameter.

¹The learning rate used for all implementations of SGA in this paper is provided by the Adaptive Moment (Adam) algorithm of Kingma and Ba (2014).

2.1 Dependence in the Approximation

Considering the vector $\boldsymbol{\theta} = (\theta_1, \theta_2)'$, the application of SVB often employs the so-called Mean Field approximation (Bishop, 2006) where the approximating distribution is factorised as

$$q_{\boldsymbol{\lambda}}(\theta_1, \theta_2 | \mathbf{y}_{1:T}) = q_{\boldsymbol{\lambda}}(\theta_1 | \mathbf{y}_{1:T}) q_{\boldsymbol{\lambda}}(\theta_2 | \mathbf{y}_{1:T}). \quad (9)$$

However SVB allows more general forms of the approximating distribution that may include dependence. In this paper we also consider approximation families that can exploit cases where the posterior of a subset of parameters conditional on remaining parameters is known. In this case the full posterior can be approximated by

$$q_{\boldsymbol{\lambda}}(\theta_1, \theta_2 | \mathbf{y}_{1:T}) = q_{\boldsymbol{\lambda}}(\theta_1 | \mathbf{y}_{1:T}) p(\theta_2 | \theta_1, \mathbf{y}_{1:T}). \quad (10)$$

where the second term on the right hand side is known (hence the use of p rather than q) and only θ_1 requires approximation. For many models the posterior can be decomposed in this manner. An example is the model we consider in Section 7 where we explicitly include the exact conditional distribution. To our knowledge, the potential to exploit this conditional approximation structure - and in particular to include an exact component within that structure - appears to be a novel contribution to the literature.

3 Updating Variational Bayes

We now introduce the proposed algorithm for updating VB when data is observed in an online setting. Let T_1, T_2, \dots be a sequence of time points, from which a sequence of posterior distributions $p(\boldsymbol{\theta} | \mathbf{y}_{1:T_1}), p(\boldsymbol{\theta} | \mathbf{y}_{1:T_2}), \dots$, is desired. Now suppose that the (exact) posterior distribution for the governing (static) parameter vector $\boldsymbol{\theta}$ is available, as given by its pdf $p(\boldsymbol{\theta} | \mathbf{y}_{1:T_n})$. Our objective is to update this posterior distribution, after observing data up to, and including, time T_{n+1} , when the additional $T_{n+1} - T_n$ data points have become available. The pdf of the resulting updated posterior distribution is denoted as $p(\boldsymbol{\theta} | \mathbf{y}_{1:T_{n+1}})$. In an online setting, where new data continues to appear, we will want to repeat this updating procedure sequentially, each time updating the past posterior to reflect all of the data, including the latest available.

The usual application of Bayes' rule at a given time T_{n+1} involves a likelihood made up of T_{n+1} factors. However with the availability of the posterior at time T_n , given by its density $p(\boldsymbol{\theta} | \mathbf{y}_{1:T_n})$, the updated time T_{n+1} posterior is given by

$$p(\boldsymbol{\theta} | \mathbf{y}_{1:T_{n+1}}) \propto p(\mathbf{y}_{T_n+1:T_{n+1}} | \mathbf{y}_{1:T_n}, \boldsymbol{\theta}) p(\boldsymbol{\theta} | \mathbf{y}_{1:T_n}), \quad (11)$$

where $p(\mathbf{y}_{T_n+1:T_{n+1}} | \boldsymbol{\theta}, \mathbf{y}_{1:T_n})$ on the right hand side of (11) is comprised of only $T_{n+1} - T_n$ factors.

We propose the Updating Variational Bayes (UVB) algorithm for use when the evaluation of the online posterior updating is computationally demanding. Our UVB algorithm, detailed in Algorithm 1, is initialised by forming the variational approximation at a given time T_1 as $q_{\boldsymbol{\lambda}_1^*}(\boldsymbol{\theta} | \mathbf{y}_{1:T_1})$ where

$$\boldsymbol{\lambda}_1^* = \arg \min_{\boldsymbol{\lambda}_1} KL[q_{\boldsymbol{\lambda}_1}(\boldsymbol{\theta} | \mathbf{y}_{1:T_1}) || p(\boldsymbol{\theta} | \mathbf{y}_{1:T_1})]. \quad (12)$$

At this first stage, we simply approximate the first posterior $p(\boldsymbol{\theta} | \mathbf{y}_{1:T_1})$ with the optimised distribution as in SVB, namely $q_{\boldsymbol{\lambda}_1^*}(\boldsymbol{\theta} | \mathbf{y}_{1:T_1})$. Importantly, this first update depends only on the first set

of observations, $y_{1:T_1}$, through the choice of the optimal parameter, λ^* .

In general then, after approximating the posterior at time T_n with $q_{\lambda_n^*}(\theta|\mathbf{y}_{1:T_n})$ and observing additional data up to time T_{n+1} , UVB replaces the posterior construction described by (11) with the available approximation,

$$\tilde{p}(\theta|\mathbf{y}_{1:T_{n+1}}) \propto p(\mathbf{y}_{T_n+1:T_{n+1}}|\theta)q_{\lambda_n^*}(\theta|\mathbf{y}_{1:T_n}). \quad (13)$$

This defines an alternate target distribution, $\tilde{p}(\theta|\mathbf{y}_{1:T_{n+1}})$, referred to as the ‘pseudo-posterior’ at time T_{n+1} .

The objective for each update is to find λ_{n+1}^* (and hence $q_{\lambda_{n+1}^*}(\theta|\mathbf{y}_{1:T_{n+1}})$) through the minimisation of the KL divergence to the corresponding pseudo-posterior, resulting in

$$\lambda_{n+1}^* = \arg \min_{\lambda_{n+1}} KL[q_{\lambda_{n+1}}(\theta|\mathbf{y}_{1:T_{n+1}}) || \tilde{p}(\theta|\mathbf{y}_{1:T_{n+1}})], \quad (14)$$

for each $n = 1, 2, \dots$. The sequence of distributional families $q_{\lambda_1}, q_{\lambda_2}, \dots$, may differ at each time period, though we note it is convenient to hold the family fixed.

Algorithm 1: Updating Variational Bayes (UVB)

Input: Prior, Likelihood.

Result: Posterior approximation at T_τ .

Observe $\mathbf{y}_{1:T_1}$.

Minimises $KL[q_{\lambda_1}(\theta|\mathbf{y}_{1:T_1}) || p(\theta|\mathbf{y}_{1:T_1})]$ using SGA via (5).

for n in $1, \dots, \tau - 1$ **do**

 Observe next data $\mathbf{y}_{T_n+1:T_{n+1}}$.

 Use $q_{\lambda_n}(\theta|\mathbf{y}_{1:T_n})$ and (13) to construct the UVB pseudo-posterior up to proportionality.

 Minimise $KL[q_{\lambda_{n+1}}(\theta|\mathbf{y}_{1:T_{n+1}}) || \tilde{p}(\theta|\mathbf{y}_{1:T_{n+1}})]$ using SGA via (5).

end

We note some important features of the proposed UVB algorithm compared with an SVB implementation. First, at time T_{n+1} an SVB implementation would target the exact posterior $p(\theta|\mathbf{y}_{1:T_{n+1}}) \propto p(\mathbf{y}_{1:T_{n+1}}|\theta)p(\theta)$ whereas UVB instead targets an alternate pseudo-posterior distribution in (13). Second, at time T_{n+1} , the evaluation for UVB corresponding to (5), which conditions on all available data $\mathbf{y}_{T_{n+1}}$, is composed of only $T_{n+1} - T_n$ factors, since the earlier data $\mathbf{y}_{1:T_n}$ has already been incorporated in the previous update which forms the new prior, as shown in (13). Hence, the computational complexity of UVB has rate $O(T_{n+1} - T_n)$ rather than rate $O(T_{n+1})$, i.e. computing UVB is not increasing in the number of observations for equally spaced intervals, as is the case for SVB. Third, unlike SVB, the UVB algorithm can begin even when only part of the data has been observed, making it well-suited to online applications. Further, the prevailing optimal value of λ_n , denoted λ_n^* , could be used as the UVB starting value for the optimisation at time T_{n+1} as long as the class of approximating distributions q is the same for each update. This may reduce the number of SGA iterations required for the UVB algorithm to converge. However, even with frequent updating via UVB, there will be some loss of accuracy in the posterior approximation at time T_{n+1} , $q_{\lambda_{n+1}}(\theta|\mathbf{y}_{1:T_{n+1}})$, relative to the corresponding approximation obtained by SVB. The extent of this loss will be context specific.

Clearly there will be a trade-off between computational speed and accuracy. These issues are investigated in a simulation setting in Section 5. Before exploring these aspects, we introduce a modified approach whereby the computational speed may be further improved, albeit potentially with some additional loss in accuracy. This modified approach, referred to as UVB with Importance Sampling (UVB-IS), is described in the next section.

4 UVB with Importance Sampling

An application of UVB up to some time T_n involves SVB inference at time T_1 followed by $n - 1$ updates, for a total of n applications of SGA optimisation. Repeated updates may incur a significant computational overhead relative to SVB, which applies only a single SGA algorithm using all data up to time T_n . In this section we address this problem and explore the possibility of achieving large computational gains per update through the incorporation of ideas from importance sampling. (For a general overview of importance sampling, see Gelman et al. (2014).) Before introducing our UVB with Importance Sampling (UVB-IS) algorithm, we briefly review the incorporation of importance sampling into SGA, as introduced by Sakaya and Klami (2017).

Temporarily suppressing the subscript n on the given time period T , the m^{th} iteration in the SGA algorithm for a given target VB posterior changes $\boldsymbol{\lambda}^{(m)}$ to $\boldsymbol{\lambda}^{(m+1)}$ via S simulations of $\boldsymbol{\theta}^{(m)}$ from $q_{\boldsymbol{\lambda}^{(m)}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$ as per (5). For each of these simulations, the log-likelihood, log-prior, and additional terms involving $q_{\boldsymbol{\lambda}^{(m)}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$ must be evaluated. Note that, for large scale applications this computation is dominated by the T terms in the log-likelihood.

In the subsequent SGA iteration from $\boldsymbol{\lambda}^{(m+1)}$ to $\boldsymbol{\lambda}^{(m+2)}$, the evaluation of the log-likelihood requires a new set of S simulations $\boldsymbol{\theta}^{(m+1)}$ from $q_{\boldsymbol{\lambda}^{(m+1)}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$. Sakaya and Klami (2017) note that as the change from $\boldsymbol{\lambda}^{(m)}$ to $\boldsymbol{\lambda}^{(m+1)}$ is likely to be small, the distributions $q_{\boldsymbol{\lambda}^{(m)}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$ and $q_{\boldsymbol{\lambda}^{(m+1)}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$ will likely be similar. Using this motivation, an alternative gradient estimator is suggested for each iteration $k = m + 1, m + 2, \dots, m + r$ via an importance sampler that uses $q_{\boldsymbol{\lambda}^{(m)}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$ as a proposal distribution, rather than generating new draws of $\boldsymbol{\theta}$ from each $q_{\boldsymbol{\lambda}^{(k)}}(\boldsymbol{\theta}|\mathbf{y}_{1:T})$. This approach retains the set of samples $\boldsymbol{\theta}^{(m)}$ and their associated log-likelihood values, only resampling $\boldsymbol{\theta}$ and re-evaluating the corresponding log-likelihood at iteration $m + r + 1$. In the SVB context the value of r should not be taken to be too large, as substantial differences between $\boldsymbol{\lambda}^{(m)}$ and $\boldsymbol{\lambda}^{(m+r)}$ may lead to a corresponding increase in the variance of the resulting gradient estimator.

In the context of UVB, we sequentially update the posterior approximation at each time T_n via repeated applications of SGA. As before UVB-IS holds the family of the approximating distribution $q_{\boldsymbol{\lambda}}$ fixed between each update, and sets the initial value of the parameter vector at time T_{n+1} equal to the optimal value from the previous update, i.e. we set $\boldsymbol{\lambda}_{n+1}^{(1)} = \boldsymbol{\lambda}_n^*$. During the subsequent application of SGA, the sequence of parameter vectors $\boldsymbol{\lambda}_{n+1}^{(1)}, \boldsymbol{\lambda}_{n+1}^{(2)}, \dots, \boldsymbol{\lambda}_{n+1}^*$ corresponds to a sequence of distributions moving from $q_{\boldsymbol{\lambda}_n^*}(\boldsymbol{\theta}|\mathbf{y}_{1:T_n})$ to $q_{\boldsymbol{\lambda}_{n+1}^*}(\boldsymbol{\theta}|\mathbf{y}_{1:T_{n+1}})$. For repeated updates with small values of $T_{n+1} - T_n$, the new information about $\boldsymbol{\theta}$ in $\mathbf{y}_{T_n+1:T_{n+1}}$ will typically be relatively small, and unless there is a structural change in the data process, we expect the approximating distributions will become similar.

The above observation motivates the addition of an importance sampling gradient estimator to be applied for each update. In each update using the SGA algorithm at time T_{n+1} all of the requisite gradients are estimated via importance sampling, using the previous UVB posterior $q_{\boldsymbol{\lambda}_n^*}(\boldsymbol{\theta}|\mathbf{y}_{1:T_n})$

as the (identical) proposal distribution. The consequence of this approach is that only S $\boldsymbol{\theta}$ samples are required for the entire SGA algorithm, and the likelihood is evaluated S times in total, rather than S times per iteration (or S times per r iterations in the case of Sakaya and Klami (2017)).

Suppressing the SGA iteration superscript index (m), the UVB-IS gradient estimator is derived from the score-based estimator implied by (5). In this case, the updated joint distribution, given by $p(\mathbf{y}_{T_n+1:T_{n+1}}, \boldsymbol{\theta} | \mathbf{y}_{1:T_n})$, is replaced by an expression proportional to (13), with

$$\frac{\partial \mathcal{L}(q, \boldsymbol{\lambda}_{n+1})}{\partial \boldsymbol{\lambda}_{n+1}} = \int_{\boldsymbol{\theta}} q_{\boldsymbol{\lambda}_{n+1}}(\boldsymbol{\theta} | \mathbf{y}_{1:T_{n+1}}) \underbrace{\frac{\partial \log(q_{\boldsymbol{\lambda}_{n+1}} | \mathbf{y}_{1:T_{n+1}})}{\partial \boldsymbol{\lambda}_{n+1}} \left(\log \left(\frac{\tilde{p}(\mathbf{y}_{T_n+1:T_{n+1}}, \boldsymbol{\theta} | \mathbf{y}_{1:T_n})}{q_{\boldsymbol{\lambda}_{n+1}}(\boldsymbol{\theta} | \mathbf{y}_{1:T_{n+1}})} \right) - \hat{\mathbf{a}} \right)}_{f(\boldsymbol{\theta})} d\boldsymbol{\theta}. \quad (15)$$

Multiplication and division of the integrand in (15) by $q_{\boldsymbol{\lambda}_n^*}(\boldsymbol{\theta} | \mathbf{y}_{1:T_n})$ allows it to be written as an expectation with respect to $q_{\boldsymbol{\lambda}_n^*}(\boldsymbol{\theta} | \mathbf{y}_{1:T_n})$,

$$\frac{\partial \mathcal{L}(q, \boldsymbol{\lambda}_{n+1})}{\partial \boldsymbol{\lambda}_{n+1}} = \int_{\boldsymbol{\theta}} q_{\boldsymbol{\lambda}_n^*}(\boldsymbol{\theta} | \mathbf{y}_{1:T_n}) \frac{q_{\boldsymbol{\lambda}_{n+1}}(\boldsymbol{\theta} | \mathbf{y}_{1:T_{n+1}})}{q_{\boldsymbol{\lambda}_n^*}(\boldsymbol{\theta} | \mathbf{y}_{1:T_n})} f(\boldsymbol{\theta}) d\boldsymbol{\theta}, \quad (16)$$

where we note the factor $f(\boldsymbol{\theta})$ is defined using the underbrace in (15). Hence, (16) may be estimated via a Monte Carlo average,

$$\frac{\partial \widehat{\mathcal{L}}(q, \boldsymbol{\lambda}_{n+1})}{\partial \boldsymbol{\lambda}_{n+1}} \underset{IS}{=} \frac{1}{S} \sum_{j=1}^S w(\boldsymbol{\theta}^{(j)}) f(\boldsymbol{\theta}^{(j)}) \quad (17)$$

since $\boldsymbol{\theta}^{(j)} \sim q_{\boldsymbol{\lambda}_n^*}(\boldsymbol{\theta} | \mathbf{y}_{1:T_n})$ and

$$w(\boldsymbol{\theta}^{(j)}) = \frac{q_{\boldsymbol{\lambda}_{n+1}}(\boldsymbol{\theta}^{(j)} | \mathbf{y}_{1:T_{n+1}})}{q_{\boldsymbol{\lambda}_n^*}(\boldsymbol{\theta}^{(j)} | \mathbf{y}_{1:T_n})}, \quad (18)$$

with $\hat{\mathbf{a}}$ estimated as per Equation (6).

Since only the value of $\boldsymbol{\lambda}_{n+1}$ changes in each iteration of SGA, and the S samples $\boldsymbol{\theta}^{(j)}$ are held fixed, only the terms involving $\boldsymbol{\lambda}_{n+1}$, namely $\frac{\partial}{\partial \boldsymbol{\lambda}_{n+1}} \log(q_{\boldsymbol{\lambda}_{n+1}}(\boldsymbol{\theta}^{(j)} | \mathbf{y}_{1:T_{n+1}}))$ and $q_{\boldsymbol{\lambda}_{n+1}}(\boldsymbol{\theta}^{(j)} | \mathbf{y}_{1:T_{n+1}})$, are required to be calculated.

The variance of the UVB-IS gradient estimator is increased relative to the score-based gradient estimator in (5) due to the presence of the importance sampling weights. This increased variance may result in a reduction in the accuracy of $q_{\boldsymbol{\lambda}_{n+1}}(\boldsymbol{\theta} | \mathbf{y}_{1:T_{n+1}})$. This is due to the fact that the algorithm stopping criterion, which is a sufficiently small value of $|\mathcal{L}(q, \boldsymbol{\lambda}^{(m+1)}) - \mathcal{L}(q, \boldsymbol{\lambda}^{(m)})|$ can only be evaluated approximately by a noisy estimator, also produced via an importance sampler. As the computation per iteration is extremely small, S may be set to a larger value to reduce the variance, thereby allowing the user the capacity to balance the inevitable trade-off between computational time and approximation accuracy to suit their requirements. Provided there is no major structural change in the data generating process, it is expected that the distributions $q_{\boldsymbol{\lambda}_n^*}(\boldsymbol{\theta} | \mathbf{y}_{1:T_n})$ and $q_{\boldsymbol{\lambda}_{n+1}^*}(\boldsymbol{\theta} | \mathbf{y}_{1:T_{n+1}})$ become more similar as n increases, subsequently reducing the UVB-IS gradient estimator variance.

The proposed UVB-IS algorithm is summarised in Algorithm 2. Figure 1 provides a diagram to help illustrate the differences between the approach of Sakaya and Klami (2017) to UVB-IS. In

panel (a) of Figure 1, each block indicates r separate iterations of SGA, each undertaken over an entire sample of length T , with arrows indicating that the final iteration of each block is used as an importance sampling proposal distribution for the entire next block. That is, there is one SGA algorithm applied for all data, but the importance sampling distribution changes every r^{th} iteration until convergence is reached. In panel (b) of Figure 1, three distributional updates using UVB-IS are depicted. In this case, the posterior itself is updated periodically, as indicated by arrows and corresponding to times T_1, T_2 , and T_3 , with the same importance sampling distribution used for *all* SGA iterations needed to complete a single distributional update.

Algorithm 2: UVB with Importance Sampling (UVB-IS)

Input: Prior, Likelihood.

Result: Approximating distribution at T_τ .

Observe $\mathbf{y}_{1:T_1}$.

Minimises $KL[q_{\lambda_1}(\boldsymbol{\theta}|\mathbf{y}_{1:T_1}) || p(\boldsymbol{\theta}|\mathbf{y}_{1:T_1})]$ using SGA via (5).

for n in $1, \dots, \tau - 1$ **do**

 Observe next data $\mathbf{y}_{T_n+1:T_{n+1}}$.

 Sample $\boldsymbol{\theta}^{(j)} \sim q_{\lambda_n^*}(\boldsymbol{\theta}|\mathbf{y}_{1:T_n})$ for $j = 1, 2, \dots, S$.

 Evaluate $p(\mathbf{y}_{T_n+1:T_{n+1}}|\boldsymbol{\theta}^{(j)})$ and $q_{\lambda_n^*}(\boldsymbol{\theta}^{(j)}|\mathbf{y}_{1:T_n})$ for each j .

 Set $\lambda_{n+1}^{(1)}$ to λ_n^* .

 Minimise $KL[q_{\lambda_{n+1}}(\boldsymbol{\theta}|\mathbf{y}_{1:T_{n+1}}) || \tilde{p}(\boldsymbol{\theta}|\mathbf{y}_{1:T_{n+1}})]$ using SGA via (17).

end

5 Simulation Study

To investigate the trade-off between the computational efficiency and accuracy of different methods we consider two simulated examples. The first is a time series forecasting application, while the second is a clustering example based on a mixture model. As well as considering both of the proposed algorithms (i.e. UVB and UVB-IS) we also consider a standard SVB approach and an exact MCMC algorithm, based on a Random Walk Metropolis Hastings strategy (see Gilks et al., 1995, Gilks et al., 1995, and Garthwaite et al., 2016), employed using all data observed up to each relevant time point.

5.1 Time Series Forecasting

In this first simulation study, we consider $R = 500$ replications of time series data, with each comprised of $T = 500$ observations simulated from the following auto-regressive order 3 (AR3) model,

$$y_t = \mu + \phi_1(y_{t-1} - \mu) + \phi_2(y_{t-2} - \mu) + \phi_3(y_{t-3} - \mu) + e_t \quad (19)$$

where $e_t \sim N(0, \sigma^2)$. For each replication, the true values of the parameters are obtained by drawing μ and each auto-regressive coefficient, ϕ_1, ϕ_2 , and ϕ_3 from an independent $N(0, 1)$ distribution, accepting only draws where each ϕ lies in the AR3 stationary region. The precision parameter, σ^{-2} , is drawn from a Gamma distribution with both shape and rate equal to five.

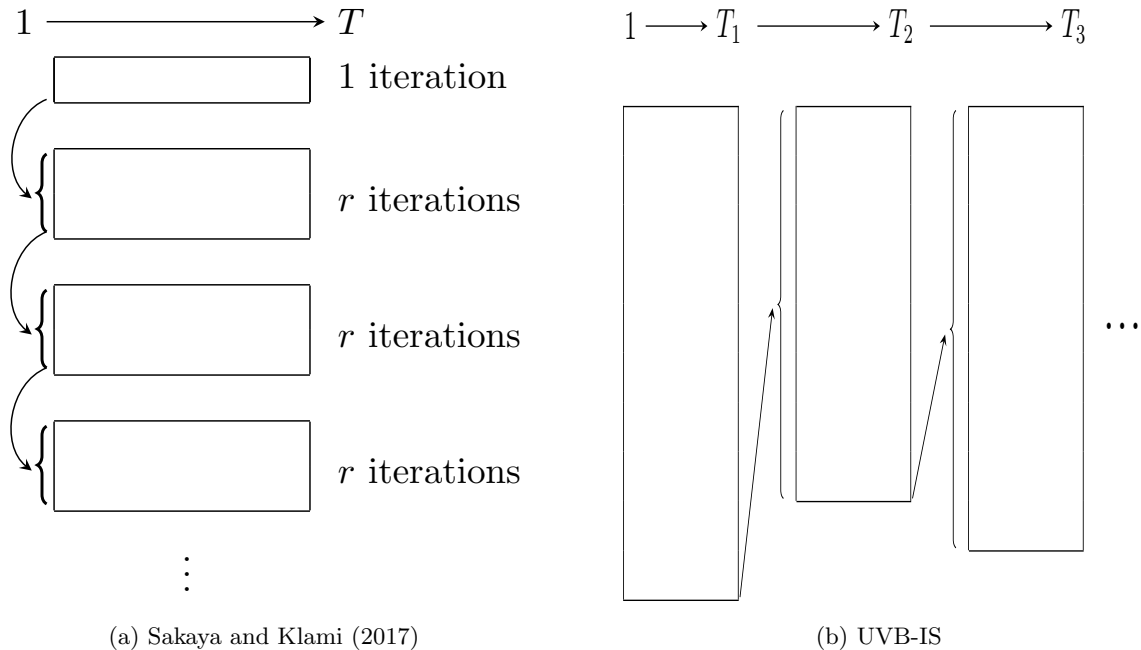


Figure 1: Graphical illustrations for importance sampling in SVB algorithms. (a): The approach of Sakaya and Klami (2017). Each block indicates r iterations of a single implementation of the SGA algorithm, with arrows indicating that the final iteration of each block is used as an importance sampling proposal distribution for the next r iterations contained in the subsequent block. (b): The UVB-IS algorithm, where each block indicates that SGA is applied three times, once for each of three distributional updates corresponding to an increase in data. For update, indicated by an arrow, a sample from the pseudo-posterior distribution corresponding to the previous update is used as proposal draws in every iteration of the SGA algorithm.

The inferential objective is to progressively produce the one-step ahead predictive densities, each based on a UVB approximation to the target posterior distribution that results from assuming data arises from the AR3 model above, with a prior distribution specified for $\boldsymbol{\theta} = \{\mu, \phi_1, \phi_2, \phi_3, \log(\sigma^2)\}$. The prior distribution for the parameter vector is taken as $\boldsymbol{\theta} \sim N(\mathbf{0}_5, 10\mathbb{I}_5)$, where $\mathbf{0}_d$ and \mathbb{I}_d denote, respectively, the d -dimensional zero vector and identity matrix. In particular, we aim to produce UVB-based approximate one-step ahead predictive distributions progressively, using at time T_n all (and only) data up to and including time period T_n , recursively for each of the 21 time periods given by $T_n = 100, 125, 150, \dots, 500$. That is, the first target predictive distribution is given by $p(y_{101}|\mathbf{y}_{1:100})$, followed by $p(y_{126}|\mathbf{y}_{1:125})$, and continuing on to the final predictive $p(y_{501}|\mathbf{y}_{1:500})$. For each update, predictive distributions are approximated with $q_{\boldsymbol{\lambda}}$ taken as a K -component mixture of multivariate normal distributions. Diagonal covariance matrices for each normal are assumed; alternatively, a sparse structure (Tan and Nott, 2018) or a factor structure (Ong et al., 2018) could be employed. The results are compared using three different choices of K , with $K = 1, 2$ and 3 . This strategy allows us to compare the approximation accuracy of the simple $K = 1$ distribution that may not adequately capture the entire posterior distribution as well as more complex approximations. In all cases the convergence criterion is to compare the mean of objective function from the last five iterations to the five iteration before that, and to stop if the difference is less than 10^{-4} times the number of parameters.

For the cases involving SVB and UVB, the score-based gradient estimator (5) uses $S = 25$ draws of $\boldsymbol{\theta}$, however we use a larger number of draws for UVB-IS to offset the increased variance, setting

$S = 100$. Finally the MCMC benchmark comparison is based on 15000 posterior draws, with the first 10000 discarded for ‘burn in’. In each approach we allow $\{\phi_1, \phi_2, \phi_3\}$ to take any value in \mathbb{R}^3 , so the posterior distribution for these parameters is not restricted to the AR3 stationary region.

Under the posterior given by $p(\boldsymbol{\theta}|\mathbf{y}_{1:T_n})$ together with the conditional predictive densities implied by (19), the one-step ahead predictive density is given by

$$p(y_{T_n+1}|\mathbf{y}_{1:T_n}) = \int_{\boldsymbol{\theta}} p(y_{T_n+1}|\mathbf{y}_{1:T_n}, \boldsymbol{\theta})p(\boldsymbol{\theta}|\mathbf{y}_{1:T_n})d\boldsymbol{\theta}. \quad (20)$$

Given our UVB approximation to the posterior at time T_n , we approximate the integral in (20) using M draws $\boldsymbol{\theta}^{(1)} \dots \boldsymbol{\theta}^{(M)} \sim q_{\lambda_n^*}(\boldsymbol{\theta}|\mathbf{y}_{1:T_n})$, with the resulting marginal predictive density estimate given by

$$\hat{p}(y_{T_n+1}|\mathbf{y}_{1:T_n}) \approx \frac{1}{M} \sum_{i=1}^M p(y_{T_n+1}|\mathbf{y}_{1:T_n}, \boldsymbol{\theta}^{(i)}). \quad (21)$$

The forecast accuracy associated with the resulting approximate predictive densities is measured using the cumulative predictive log score (*CLS*), given by


$$CLS_n = \sum_{j=1}^n \log(\hat{p}(y_{T_j+1}^{(obs)}|\mathbf{y}_{1:T_j})), \quad (22)$$

for $n = 1, 2, \dots, 21$, where $y_{T_n+1}^{(obs)}$ denotes the realised (observed) value of y_{T_n+1} . In particular, we compare the mean CLS (MCLS) over the R replications, for each approximation method and each given value of K , at consecutive update times T_n with $n = 1, 2, \dots, 21$. The results are displayed in Figure 2, where each row indicates a different (known) value of K . Panel (a), on the left-hand side, the MCLS value is displayed relative to the MCLS value for MCMC inference, against the corresponding incremental values of $T_n + 1 = 101, 126, \dots, 501$. As greater values of MLS indicate better forecast accuracy, it is not surprising to find that each of the approximate VB method produces a lower MLS relative to exact (MCMC) inference. Amongst the approximate methods, UVB performs the best, in terms of MLS, followed by SVB and UVB-IS at $K = 1$, though all of the variational methods improve as K increases.

Panel (b) of Figure 2 displays the relative cumulative mean runtime (RCMR) for each VB method using data up to T_n , for $T_n = 100, 125, 150, \dots, 500$, each calculated by dividing the CMR by the mean run time of the SVB algorithm fitting a single mixture (i.e. $K = 1$) at $T_n = 100$. Note that the SVB approximation at each update time T_n requires an application of the SGA algorithm using all data from time $T = 1$ up to time T_n , while each of the updating methods are comprised of an SVB approximation at T_1 , followed by $n - 1$ progressive updates each using only the new data since the last update period. As can be seen in the top row of Panel (b), the RCMRs are all identical and equal to one at the first update time $T_n = 100$ and all increase with consecutive updates. While all three methods show an increase in RCMR with each update, the UVB method appears least efficient, while significant improvements in computational time accrue from using UVB-IS.

In this setting the amount of data in each update is relatively small, and UVB increases the runtime compared to SVB. This is due to the computational overhead of n SGA applications not being offset by a reduction in the number of log-likelihood calculations. In contrast, UVB-IS achieves sizeable computational gains despite showing minimal loss in the corresponding MLS for $K > 1$.

To illustrate the reduced variability in subsequent UVB gradient estimators, Figure 3 displays the median variance of the gradient estimator for the posterior mean parameter of μ , with UVB at $S = 25$, and UVB-IS at each of $S = 25, 50, 100$ and 200 , for each of six selected update periods. The algorithms at T_1 are applications of SVB with arbitrary starting values for $\lambda_1^{(1)}$. This causes extreme, but declining, variance until convergence is reached. This pattern is typical for SVB inference at all time periods, parameters, and values of K . In subsequent time periods each updating method sets the starting value at $\lambda_n^{(1)} = \lambda_{n-1}^*$, noting that, for example, $\lambda_5^{(1)}$ is equal to the omitted λ_4^* . The estimated variance is subsequently orders of magnitude smaller than SVB. For small values of n the distributions $q_{\lambda_n^*}(\theta|\mathbf{y}_{1:T_n})$ and $q_{\lambda_{n+1}^{(m)}}(\theta|\mathbf{y}_{1:T_n})$ may differ as m increases, causing a reduction in the effective sample size associated with the gradient estimator, and an increase in the UVB-IS estimator variance. This effect is visible at times T_2, T_3 , and T_5 , though the UVB-IS estimator variance is low relative to SVB despite this inefficiency.



(a) Forecast accuracy (MCLS)

(b) Computational efficiency (RCMR)

Figure 2: AR3 Simulation (a): Forecast accuracy, indicated by one-step-ahead mean cumulative predictive log scores (MCLS), corresponding to incremental updates under competing methods (SVB, UVB and UVB-IS) relative to MCMC. Higher values of MCLS indicate better forecast accuracy. (b): Computational efficiency, indicated by relative cumulative mean runtime (RCMR), again corresponding to incremental updates under competing method (SVB, UVB and UVB-IS), each reported relative to the mean runtime for SVB when $K = 1$ and $T_n = 100$.

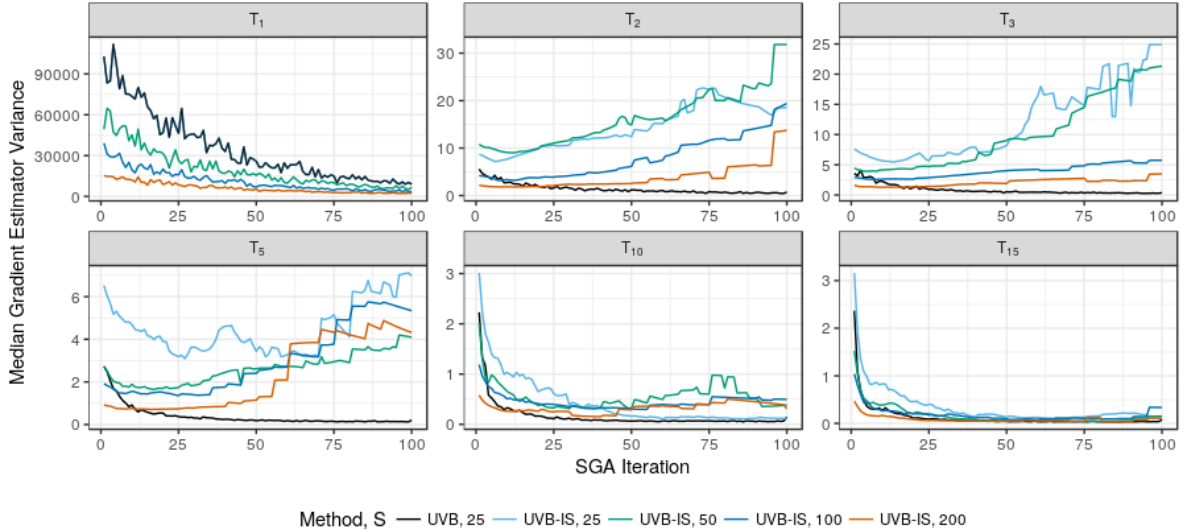


Figure 3: Median gradient estimator variance for the first 100 SGA iterations, with colour indicated by the acronym (either UVB or UVB-IS) followed by the value of S . Both UVB and UVB-IS algorithms have arbitrary starting values at T_1 , denoted as $\lambda_1^{(1)}$, where the estimated gradient exhibits high variance. Note the difference in the y-axis scales at subsequent update times, where the starting value at time T_n is set to the previous optimal value, i.e. $\lambda_n^{(1)} = \lambda_{n-1}^*$, noting several time periods have associated λ_n^* values, but are omitted in the figure. Consequently, we the variance of the UVB gradient estimator is reduced relative to SVB, though the UVB-IS variance increases slightly for small n and large iteration index m .

5.2 Mixture Model Clustering

In the second simulated example we consider the case where repeated measurements are simulated on $N = 100$ cross-sectional units at each of $T = 100$ times. The measurements for a given unit follows one of two possible DGPs, with the objective being to cluster the units into the correct groups, according to the underlying DGP, with additional observations of each cross-sectional unit accumulating in an online fashion as time increases. Each of these scenarios was then replicated $R = 500$ times.

For each independent replication, we generate data $y_{i,t}$ as the measurement of unit i at time t , for $i = 1, 2, \dots, N$ and $t = 1, 2, \dots, T$ as follows. We first define the cluster indicator for unit i as k_i , and generate these for a given probability $0 < \pi < 1$ according to

$$k_i | \pi \stackrel{i.i.d.}{\sim} \text{Bernoulli}(\pi), \quad (23)$$

where *i.i.d* abbreviates *independent and identically distributed*. Then, conditional on k_i we let

$$y_{i,t} | (k_i = j), \mu_j, \sigma_j^2 \stackrel{ind}{\sim} N(\mu_j, \sigma_j^2), \quad (24)$$

for $j = 0, 1$, with *ind* short for *independent*. For this exercise, we set $\pi = 0.5$, with the replicated values of μ_0 and μ_1 independently drawn from an $N(0, 0.25)$ distribution, while σ_0^2 and σ_1^2 are independently drawn from a uniform distribution over the interval $(1, 2)$.

Having simulated the data, the actual values k_i are retained for each replication. We then use the UVB algorithm of the described model with the simulated data, as if all N units were being

observed online at increasing times $T_n = 10, 20, 30, \dots, 100$. The aim of the exercise is to cluster the units into two groups aligning with the true, but ‘unobserved’ value of k_i .

The Bayesian updating analysis proceeds as follows. Denoting the collective parameter vector as $\boldsymbol{\theta} = \{\log(\sigma_0^2), \log(\sigma_1^2), \mu_0, \mu_1\}$, the joint prior for $\boldsymbol{\theta}$ and π used at T_1 is given by independent components

$$\boldsymbol{\theta} \sim N(\mathbf{0}_4, 10\mathbb{I}_4), \text{ and} \quad (25)$$

$$\pi \sim \text{Beta}(\alpha, \beta). \quad (26)$$

Note that the model for π in (23) and the prior in (26) imply that the k_i are independent *a priori*, with marginal probabilities given by

$$\Pr(k_i = j) = \frac{\mathcal{B}(j + \alpha, \beta - j + 1)}{\mathcal{B}(\alpha, \beta)}, \quad (27)$$

for $j = 0, 1$, where $\mathcal{B}(\cdot, \cdot)$ denotes the Beta function. Hence we have marginalised out the ‘unknown’ value of π , and can now proceed to updating the prior in (27), for each $i = 1, 2, \dots, N$, on the basis of information at times $T_n = 10, 20, \dots, 100$.

Denoting $\mathbf{y}_{i,1:T_n} = \{y_{i,t} | t = 1, \dots, T_n\}$ and $\mathbf{y}_{1:N,1:T_n} = \{\mathbf{y}_{i,1:T_n}; i = 1, \dots, N\}$, the initial *augmented* posterior distribution is given by

$$p(\boldsymbol{\theta}, \mathbf{k}_{1:N} | \mathbf{y}_{1:N,1:T_1}) \propto p(\boldsymbol{\theta}) \prod_{i=1}^N p(\mathbf{y}_{i,1:T_1} | \boldsymbol{\theta}, k_i = j) \Pr(k_i = j), \quad (28)$$

with each likelihood $p(\mathbf{y}_{i,1:T_1} | \boldsymbol{\theta}, k_i = j)$ given by the product of densities associated with (24) and the value of j .

Due to the conditional independence of the components of $\boldsymbol{\theta}$ and the cluster indicators, subsequent posteriors at times T_{n+1} are approximated by

$$\widehat{p}(\boldsymbol{\theta}, \mathbf{k}_{1:N} | \mathbf{y}_{1:N,1:T_{n+1}}) \propto \prod_{i=1}^N p(\mathbf{y}_{i,T_n+1:T_{n+1}} | \boldsymbol{\theta}, k_i = j) \widehat{\Pr}(k_i = j | \mathbf{y}_{1:N,1:T_n}) p(\boldsymbol{\theta} | \mathbf{y}_{1:N,1:T_n}), \quad (29)$$

where the latent class probabilities, $\Pr(k_i = j | \mathbf{y}_{1:N,1:T_n})$, are estimated before updating with

$$\widehat{\Pr}(k_i = j | \mathbf{y}_{1:N,1:T_n}) \propto \frac{1}{M} \sum_{l=1}^M p(\mathbf{y}_{1:N,1:T_n} | \boldsymbol{\theta}^{(l)}, k_i = j) \Pr(k_i = j), \quad (30)$$

with $\boldsymbol{\theta}^{(l)} \sim p(\boldsymbol{\theta} | \mathbf{y}_{1:N,1:T_n})$ for $l = 1, 2, \dots, M$.

As in Section 5.1, the UVB and UVB-IS algorithms are compared to standard SVB and MCMC. Each of these approaches utilises an approximation to the augmented posterior of the form

$$q_{\lambda_{n+1}}(\boldsymbol{\theta}, \mathbf{k}_{1:N} | \mathbf{y}_{1:N,1:T_{n+1}}) = q_{\lambda_{n+1}}(\boldsymbol{\theta} | \mathbf{y}_{1:N,1:T_{n+1}}) \prod_{i=1}^N \widehat{\Pr}(k_i = j | \mathbf{y}_{1:N,1:T_n}), \quad (31)$$

where $q_{\lambda_{n+1}}(\boldsymbol{\theta} | \mathbf{y}_{1:N,1:T_{n+1}})$ is a $K = 1, 2$, or 3 component mixture of multivariate normal distributions and the $\boldsymbol{\theta}^{(l)}$ samples used to estimate (30) are simulated from the previous approximation

$q_{\lambda_n}(\boldsymbol{\theta}|\mathbf{y}_{1:N,1:T_n})$.

The form of the approximation used in (31) is chosen due to the fact that the gradient of the augmented divergence, $KL[q_{\lambda_n}(\boldsymbol{\theta}, \mathbf{k}_{1:N}|\mathbf{y}_{1:N,1:T_n}) \parallel \widehat{p}(\boldsymbol{\theta}, \mathbf{k}_{1:N}|\mathbf{y}_{1:N,1:T_n})]$ is equivalent to the gradient of the marginal divergence, $KL[q_{\lambda_n}(\boldsymbol{\theta}|\mathbf{y}_{1:N,1:T_n}) \parallel \widehat{p}(\boldsymbol{\theta}|\mathbf{y}_{1:N,1:T_n})]$, and hence the same approximation can be found by instead targeting the marginal posterior distribution,

$$p(\boldsymbol{\theta}|\mathbf{y}_{1:N,1:T_1}) \propto p(\boldsymbol{\theta}) \prod_{i=1}^N \left(\sum_{j=0}^1 p(\mathbf{y}_{i,1:T_1}|\boldsymbol{\theta}, k_i = j) \Pr(k_i = j) \right), \quad (32)$$

or its updated form

$$\widehat{p}(\boldsymbol{\theta}|\mathbf{y}_{1:N,1:T_{n+1}}) \propto \prod_{i=1}^N \left(\sum_{j=0}^1 p(\mathbf{y}_{i,T_n+1:T_{n+1}}|\boldsymbol{\theta}, k_i = j) \widehat{\Pr}(k_i = j|\mathbf{y}_{1:N,1:T_n}) p(\boldsymbol{\theta}|\mathbf{y}_{1:N,1:T_n}) \right). \quad (33)$$

At each update we estimate class labels for k_i according to

$$\widehat{k}_{i,n} = \arg \max_j \widehat{\Pr}(k_i = j|\mathbf{y}_{1:N,1:T_n}), \quad (34)$$

and assign a classification accuracy (CA) score at T_n , given by


$$CA_n = \max \left(\frac{1}{N} \sum_{i=1}^N I(\widehat{k}_{i,n} = k_i), \frac{1}{N} \sum_{i=1}^N I(\widehat{k}_{i,n} \neq k_i) \right), \quad (35)$$

the proportion of successful classifications up to label switching.


SVB and UVB gradients are estimated from $S = 25$ samples of $\boldsymbol{\theta}$ per iteration, while UVB-IS sets $S = 100$.

The results for this problem are displayed in Figure 4, where each row corresponds to a different value of K . Panel (a) displays the mean classification accuracy (MCA), corresponding to updates at times $T_n = 10, 20, \dots, 100$ and across R replications. As in the previous study, each variational approximation reduces accuracy relative to exact inference. In this example both UVB and UVB-IS are more accurate than SVB, with little change apparent in any variational approach between different values of K .

As in the previous section, Panel (b) of Figure 4 displays the relative cumulative mean runtime (RCMR) for each VB method using data up to T_n , for $T_n = 10, 20, \dots, 100$, calculated relative to the mean run time of the SVB algorithm fitting a single mixture at the initial update time, when $T_n = 10$. As the problem features a large number of cross-sectional units, the computational cost of calculating the log-likelihood dominates the gradient estimation. Processing smaller amounts of data, and having a reduced gradient variance lead to reduced computational time for both UVB and UVB-IS relative to SVB, particularly in the case of UVB-IS. Despite the updating methods consisting of 10 SGA applications while SVB uses only one, UVB and UVB-IS require, on average, 14.7%, and 4.6% of the computational time of SVB, respectively, in the top right panel when $K = 1$ at time $T_{10} = 100$.



(a) Clustering Accuracy (MCA)



(b) Computational Efficiency (RCMR)

Figure 4: (a): Mean Classification Accuracy (MCA) for each inference method, higher is better. (b): Average SVB runtime for one approximation at time T_n and average cumulative UVB and UVB-IS runtimes fit to T_1 , then updated $n - 1$ times to T_n . Runtimes are relative to the time required for the T_1 fit. UVB and UVB-IS perform better than SVB and are also much faster, as computation of the data likelihood is a large part of the gradient calculation in this scenario.

6 Eight Schools Example

In this section the co-called ‘Eight Schools’ problem described in Gelman et al. (2014) is considered. This problem analyses the effectiveness of a short term coaching program, implemented independently by each of eight studied schools, for the SAT-V test². For students $i = 1, 2, \dots, N_j$ in each school $j = 1, 2, \dots, 8$, consider the linear regression

$$SAT-V_{i,j} = \beta_{0,j} + \beta_{1,j}Coach_{i,j} + \beta_{2,j}PSAT-V_{i,j} + \beta_{3,j}PSAT-M_{i,j} + \epsilon_{i,j} \quad (36)$$

where $Coach_{i,j}$ is a dummy variable indicating a student’s inclusion (or not) in a coaching program run by their school, alongside control variables $PSAT-V_{i,j}$ and $PSAT-M_{i,j}$, corresponding to each student’s scores in the verbal and mathematical preliminary SAT, respectively.

²The SAT-V is a standardised aptitude test commonly taken by high school students in the USA.

Following Gelman et al. (2014), the estimated school-level coaching coefficients that correspond to the ordinary least squares estimators are taken as the observations, $y_j = \widehat{\beta}_{1,j}$, for $j = 1, 2, \dots, 8$, and have approximate sampling distributions given by

$$y_j | \theta_j, \sigma_j^2 \sim \mathcal{N}(\theta_j, \sigma_j^2), \quad (37)$$

where θ_j is the latent ‘true’ effectiveness of school j ’s coaching program. The standard deviation of the sampling distribution, σ_j , is assumed to be known and is held fixed at the standard error estimated by the relevant regression, with each having taken account of the individual school sample size N_j .

Again following Gelman et al. (2014), we apply a hierarchical prior to the population mean values in (37), assuming that the θ_j themselves are random and *iid* from a Student- t distribution,

$$\frac{\theta_j - \mu}{\tau} \sim t(\nu) \quad (38)$$

where ν is the degrees of freedom, fixed at $\nu = 4$. The hierarchical model also employs the uninformative hyper-prior

$$p(\mu, \tau) \propto 1, \quad (39)$$

over positive values of τ , and both positive and negative values of μ .

Collecting the unknown school means together and denoting by $\boldsymbol{\theta}_{1:8} = \{\theta_1, \theta_2, \dots, \theta_8\}$, the posterior distribution of all unknowns and based on the observed values from all schools is then given by

$$p(\boldsymbol{\theta}_{1:8}, \mu, \tau | \mathbf{y}_{1:8}) \propto \prod_{j=1}^8 p(y_j | \theta_j, \sigma_j^2) p(\theta_j | \tau, \mu). \quad (40)$$

It is feasible to obtain this posterior exactly, via MCMC, for example using the algorithm provided in the statistical modelling platform *Stan* (Stan Development Team, 2018).

Our aim here is to demonstrate the application of UVB and UVB-IS to this hierarchical model, where with each update we sequentially ‘observe’ an additional school, as indicated by the inclusion of an additional observation y_j . Each variational algorithm approximates the progressive posterior by the multivariate normal distribution $q_{\boldsymbol{\lambda}_n}(\boldsymbol{\theta}_{1:n}, \mu, \tau | \mathbf{y}_{1:n})$, for $n = 1, 2, \dots, 8$. The initial distribution approximation at T_1 for UVB and UVB-IS is given by the multivariate normal distribution $q_{\boldsymbol{\lambda}_1^*}(\boldsymbol{\theta}_1, \mu, \tau | y_1)$ where

$$\boldsymbol{\lambda}_1^* = \arg \min_{\boldsymbol{\lambda}_1} KL[q_{\boldsymbol{\lambda}_1}(\boldsymbol{\theta}_1, \mu, \tau | y_1) || p(\boldsymbol{\theta}_1, \mu, \tau | y_1)]. \quad (41)$$

Updates at further ‘times’ $T_{n+1} = n + 1$, for $n = 1, 2, \dots, 7$, involves sequentially adding schools to the model targeting the pseudo-posterior distribution, given by the decomposition

$$\tilde{p}(\boldsymbol{\theta}_{1:n+1}, \mu, \tau | \mathbf{y}_{1:n+1}) \propto p(y_{n+1} | \theta_{n+1}) p(\theta_{n+1} | \mu, \tau) q_{\boldsymbol{\lambda}_n^*}(\boldsymbol{\theta}_{1:n}, \mu, \tau | \mathbf{y}_{1:n}). \quad (42)$$

Either UVB or UVB-IS then may be used to obtain the updated approximate posterior, given by $q_{\boldsymbol{\lambda}_{n+1}^*}(\boldsymbol{\theta}_{1:n+1}, \mu, \tau | \mathbf{y}_{1:n+1})$, with

$$\boldsymbol{\lambda}_{n+1}^* = \arg \min_{\boldsymbol{\lambda}_{n+1}} KL[q_{\boldsymbol{\lambda}_{n+1}}(\boldsymbol{\theta}_{1:n+1}, \mu, \tau | \mathbf{y}_{1:n+1}) || \tilde{p}(\boldsymbol{\theta}_{1:n+1}, \mu, \tau | \mathbf{y}_{1:n+1})], \quad (43)$$

for $n = 1, 2, \dots, y$. As each update adds a new variable θ_{n+1} to the model, the optimal vector λ_{n+1}^* updates the auxiliary parameters associated with the pseudo-posterior distribution for θ_{n+1} together with the previously included variables μ, τ , and $\theta_{1:n}$. We note that our implementation of UVB-IS here employs a hybrid strategy utilising importance sampled gradients (17) for simulations of μ, τ , and $\theta_{1:n}$ from the previous $q_{\lambda_n^*}(\theta_{1:n}, \mu, \tau | \mathbf{y}_{1:n+1})$, and score-based gradients for θ_{n+1} , as per (5). The score-based gradients use samples generated from $\theta_{n+1} \sim q_{\lambda_{n+1}^*}(\theta_{n+1} | \mathbf{y}_{1:n+1}, \mu, \tau, \theta_{1:n})$, which is available as this variational approximation was chosen to be a multivariate normal distribution.

We compare approximations that result from using UVB and UVB-IS, relative to the sequential implementation of SVB, as each new school is added. As the ordering of the inclusion of schools is arbitrary in this example, we report results that are averaged over a randomly selected 100 of the $8! = 40,320$ possible permutations of school sequences. For each ordering, the KL divergence $KL[q_{\lambda}(\theta_{1:n+1}, \mu, \tau | \mathbf{y}_{1:n+1}) || p(\theta_{1:n+1}, \mu, \tau | \mathbf{y}_{1:n+1})]$ is calculated and tallied, using the exact posterior $p(\theta_{1:n+1}, \mu, \tau | \mathbf{y}_{1:n+1})$ in (40), each calculated using 10,000 MCMC sample draws retained following a burn-in period of 10,000 iterations.

These results are shown in Table 1, where the row labelled UVB contains the average difference between the KL divergence calculated from the UVB approximation and the KL divergence calculated from the SVB approximation. These values are calculated independently for each margin of the distribution, and for the entire joint distribution. Each KL divergence was obtained using simulated draws from each relevant distribution. Marginal densities were estimated using the density function in R (R Core Team, 2017), with the corresponding joint density estimated using an additional vine copula (Dißmann et al., 2013). The row labelled UVB-IS contains the average difference between the corresponding UVB-IS and SVB KL divergences.

	τ	μ	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	θ_7	θ_8	Joint
UVB	1.04	0.11	0.26	0.05	0.16	0.08	0.16	0.12	0.16	0.14	2.67
UVB-IS	0.16	0.64	0.32	0.34	0.38	0.22	0.17	0.27	0.36	0.42	5.00

Table 1: Results of the Schools example. Average KL divergence measures: $KL_{UVB} - KL_{SVB}$ (row 2) and $KL_{UVB-IS} - KL_{SVB}$ (row 3), with each KL divergence calculated from the VB approximation to MCMC samples. Lower is better.

7 Lane Position Example

Vehicle drivers may exhibit a tendency to move laterally (i.e. side-to-side) within their designated lane on a highway. Figure 5 displays this notion, by plotting the trajectory of five drivers as they travel along a section of the US Route 101 Highway, as taken from the Next Generation Simulation (NGSIM, FHWA (2017)) dataset. In this figure, the vehicles - indicated in black - are travelling towards the right, with each (estimated) lane centre line given by the red dashed line. Drivers likely adapt their position in real time, in at least partial response to the perceived position of vehicles that are travelling nearby.

The aim of this section is to apply the UVB methodology to analyse a model of the lateral position of vehicles. The model incorporates driver heterogeneity, while the analysis itself produces sequential, per-vehicle distributional forecasts of a large number of future car positions. The methodology suggests that a smart vehicle (i.e. one without a human driver) may be able to repeatedly ‘observe’

neighbouring vehicle positions, predict their positions in real time as they travel along the road, and appropriately respond to those forecasts.

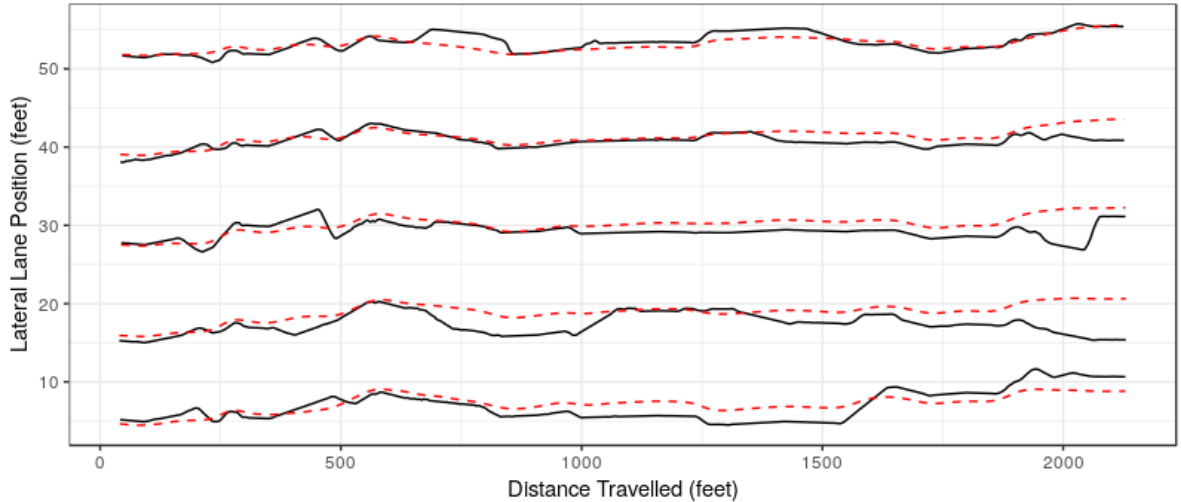


Figure 5: The path of five selected vehicles from the NGSIM dataset, travelling from left to right, with each black line representing a unique vehicle, with estimated lane centre lines in red. This section of US Route 101 is comprised of five main lanes, with a sixth entry/exit lane not shown.

To set up the scenario, we randomly select, from the NGSIM dataset, trajectories associated with $N = 500$ vehicles that do not change lane. We note that the NGSIM dataset is the result of a project conducted by the US Federal Highway Administration (FHWA), and includes data recorded from 6101 vehicles traveling along a 2235 foot long section of the US 101 freeway in Los Angeles, California from 7:50 am to 8:35 am on June 15th, 2005. Though initially collected by static cameras, the data was then processed by Cambridge Systematics Inc. to produce coordinates of the centre of the front of each vehicle at 100 millisecond intervals.

7.1 A Hierarchical Model

In developing a model for the position of cars we consider a number of issues. First, each vehicle/driver is likely to have its own idiosyncratic behaviour, captured by its own parameter values. Let $y_{i,t}$ denote the lateral deviation from the lane centre of vehicle i at time t , with details on calculating the lateral deviation provided in Appendix A. For $i = 1, \dots, N$ and $t = 1, \dots, T$, we assume

$$y_{i,t} \mid \mu_i, \sigma_i^2 \stackrel{ind}{\sim} \mathcal{N}(\mu_i, \sigma_i^2), \quad (44)$$

where μ_i and σ_i^2 are parameters specific to vehicle i . For simplicity, we collect the individual vehicle-specific parameters into a single vector, θ_i , by defining $\theta_i = (\mu_i, \log(\sigma_i^2))$, for $i = 1, 2, \dots, N$. We note that alternative parametric models could be used here, including a time series model for vehicle i , with little loss in generality.

Multiple cars may display similar behaviour, a phenomenon that can be modeled by allowing different cross sectional units to share parameters. This structure, whereby cross sectional units belong to mixture components, leads to predictions that ‘borrow strength’ from the full sample of vehicles. To make this idea explicit, let k_i denote an indicator variable such that vehicle i belongs to mixture

component j if $k_i = j$. All vehicles within the same mixture component share parameters, that is $\theta_i = \theta_j^*$, for all i such that $k_i = j$. Note that the star superscript and j subscript are generally used to index the mixture component that the parameters belong to, while the subscript i is generally used to index the cross-sectional unit, i.e. vehicle.

Since the number of components are unknown and since there is a possibility that a new vehicle will be observed with behaviour that cannot be well described by any of the prevailing parameters, we consider an infinite mixture model. In particular, we use an infinite mixture model induced by a Dirichlet Process (DP) Prior for the distribution of the parameters. The DP prior is given by

$$G \sim DP(\alpha, G_0), \quad (45)$$

where G_0 is the DP base distribution, assumed here to be $N(\mathbf{0}_2, 10\mathbb{I}_2)$, and the DP concentration parameter α is fixed here and equal to one. The prior for the collection of θ_i values represent a draw from the DP, with

$$\theta_i | G \stackrel{iid}{\sim} G, \text{ for } i = 1, 2, \dots, N. \quad (46)$$

Combining (44), (45) and (46) leads to the hierarchy

$$\begin{aligned} G &\sim DP(\alpha, G_0) \\ \theta_i | G &\stackrel{iid}{\sim} G, \text{ for } i = 1, 2, \dots, N \\ y_{i,t} | \theta_i &\stackrel{ind}{\sim} \mathcal{N}(\mu_i, \sigma_i^2), \text{ for } i = 1, 2, \dots, N \text{ and } t = 1, 2, \dots, T. \end{aligned} \quad (47)$$

We note that the DP prior induces clustering on the observation sequences, as described by the Chinese Restaurant Process (CRP, Aldous, 1985) representation. The CRP provides a mechanism for drawing from the prior of $\theta_1, \dots, \theta_n$, marginal of the random G , via the introduction of discrete variables that act as component indicators. Define s_i as the number of unique values in k_1, k_2, \dots, k_i , and let $n_{ij} = \sum_{m=1}^i I(k_m = j)$. Then, the indicator variables can be simulated from $p(k_i = j | \alpha, \mathbf{k}_{1:i-1})$ where

$$p(k_1 = 1 | \alpha, G_0) = 1, \quad (48)$$

$$p(k_i = j | \alpha, G_0, \mathbf{k}_{1:i-1}) = \begin{cases} \frac{n_{i-1,j}}{\alpha + i - 1} & \text{for } j = 1, 2, \dots, s_{i-1} \\ \frac{\alpha}{\alpha + i - 1} & \text{for } j = s_{i-1} + 1, \end{cases} \quad (49)$$

for $i = 2, \dots, N$. Note that although simulation of the indicators does not require knowledge of G_0 , we include explicit conditioning on both α and G_0 in (48) and (49) to emphasise the marginalisation over G . Under the CRP, unique values of θ_i , denoted as θ_j^* , for $j = 1, 2, \dots, s_N$ are drawn from the base distribution G_0 , and if we set $\theta_i = \theta_j^*$ for all i such that $k_i = j$, then $(\theta_1, \theta_2, \dots, \theta_N)$ is a draw from the hierarchical setup in (47). Note that although the model is an infinite component mixture model, under the CRP the maximum number of unique clusters, s_N , can be no greater than the number of vehicles in the sample, N . For simplicity we retain the full vector $\theta_{1:N}^*$, noting that some values $\theta_{1:N}^*$ may not be associated with any vehicle.

The overall model may be seen as a Dirichlet Process Mixture (DPM) model for the lane deviations. Background material regarding Bayesian analysis of DPM models, including many references and detailed discussions relating to MCMC-based techniques for sampling from the relevant posterior, is given in Müller et al. (2015). Online VB-based inference for DPMS has been established using a Mean Field approach - see, e.g., Hoffman et al. (2010), Wang et al. (2011), and Kabisa et al.

(2016)). In contrast, our approach incorporates SVB, which allows for greater flexibility regarding the form of the approximating posterior distribution. Another important distinction between our analysis and this literature is that our setting involves multiple observations over time, on each cross-sectional unit. Rather than updating as new cross-sectional units are observed, we update parameters relating to the same cross-sectional units observed periodically over a period of time.

7.2 Implementation of SVB at time T_1

Before discussing how UVB is applied to this problem it is instructive to discuss how SVB is implemented for the DPM in (47) that targets the posterior conditional on all cross-sectional units N over just the first time period from $t = 1$ to $t = T_1$. For notational convenience, conditional dependence on α and G_0 is suppressed in all notation for the remainder of this section. The objective is to minimise the KL divergence between a suitable variational approximation and a posterior that is augmented by indicator variables. To implement SVB, we must evaluate

$$p(\mathbf{y}_{1:N,1:T_1}, \boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N}) = \left[\prod_{i=1}^N \prod_{t=1}^{T_1} p(y_{i,t} | \boldsymbol{\theta}_{1:N}^*, k_i) \right] \left[\prod_{i=1}^N p(k_i | \mathbf{k}_{1:i-1}) \right] p(\boldsymbol{\theta}_{1:N}^*) \quad (50)$$

for given values of $\mathbf{y}_{1:N,1:T_1}$, $\boldsymbol{\theta}_{1:N}^*$, and $\mathbf{k}_{1:N}$. Each of the three main components on the right hand side of (50) can be computed from the hierarchical structure in (47) and the CRP, as Sethuraman (1994) shows that the unique values $\boldsymbol{\theta}_{1:N}^*$ are *a priori* independent and identically distributed according to the base distribution G_0 .

A second required input into SVB is an approximate posterior density structure, given by q_λ , and for this we propose

$$q_\lambda(\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N} | \mathbf{y}_{1:N,1:T_1}) = \left[\prod_{j=1}^N q_j(\theta_j^* | \mathbf{y}_{1:N,1:T_1}) \right] \left[\prod_{i=1}^N p(k_i | \mathbf{y}_{1:N,1:T_1}, \boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:i-1}) \right]. \quad (51)$$

Each $q_j(\cdot)$ on the right hand side is a bivariate normal distribution with unique means, variances and covariances for each $i = 1, 2, \dots, N$, leading to a total of $5N$ auxiliary parameters in the approximation. In the second product term on the right hand side of (51), the notation p is used instead of q since $p(k_i | \mathbf{y}_{1:N,1:T_1}, \mathbf{k}_{1:i-1}, \boldsymbol{\theta}_{1:N}^*)$ is known exactly and can be computed recursively using

$$p(k_i = j | \mathbf{y}_{1:N,1:T_1}, \mathbf{k}_{1:i-1}, \boldsymbol{\theta}_{1:N}^*) \propto p(k_i = j | \mathbf{k}_{1:i-1}) p(\mathbf{y}_{i,1:T_1} | \boldsymbol{\theta}_{1:N}^*, k_i), \quad (52)$$

for $i = 1, 2, \dots, N$.

The use of the so-called full conditional distribution for $\mathbf{k}_{1:N}$, given by the second product in (51), is a novel inclusion that enables our model to capture some of the dependence structure of the posterior. In contrast, a MFVB approximation would force posterior independence between each k_i and every θ_j^* , as in, for example, Wang et al. (2011).

Furthermore, in addition to minimising the KL divergence to the augmented posterior, our choice has the benefit of ensuring minimisation of the KL divergence to the corresponding marginal posterior. That is, the augmented gradients are given by

$$\frac{\partial KL[q_\lambda(\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N} | \mathbf{y}_{1:N,1:T_1}) || p((\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N} | \mathbf{y}_{1:N,1:T_1})]}{\partial \lambda_1} \quad (53)$$

and are equal to the marginal gradients

$$\frac{\partial KL[q_{\lambda_1}(\boldsymbol{\theta}_{1:N}^*|\mathbf{y}_{1:N,1:T_1}) || p(\boldsymbol{\theta}_{1:N}^*|\mathbf{y}_{1:N,1:T_1})]}{\partial \lambda_1}, \quad (54)$$

and so the optimisation procedure is equivalent to one where the indicator variables used to construct the DPM have been marginalised out.

The proof of this result is shown in Appendix B.

7.3 Iterating UVB

Using data up to time T_1 , the first UVB posterior is obtained using SVB, as described in Section 7.2. For updating at time T_{n+1} , we construct a pseudo-posterior using information from the previous variational approximation $q_{\lambda_n}(\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N}|\mathbf{y}_{1:N,1:T_n})$ in two distinct ways. First, the base distribution in the DP as the prior distribution for $\boldsymbol{\theta}_{1:N}^*$ is updated to reflect the clustering present in the previously obtained posterior, and so is replaced with $q_{\lambda_n}(\boldsymbol{\theta}_{1:N}^*|\mathbf{y}_{1:N,1:T_n})$. Second, retaining the form of the approximation in (51) for the update is complicated by the use of the full conditional distribution for k_i , given by

$$p(k_i = j|\mathbf{y}_{1:N,1:T_{n+1}}, \boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:i-1}) \propto p(\mathbf{y}_{i,T_n+1:T_{n+1}}|\boldsymbol{\theta}_{1:N}^*, k_i)p(k_i = j|\mathbf{y}_{1:N,1:T_n}, \boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:i-1}), \quad (55)$$

as all currently observed data up to time T_{n+1} is required for each new $\boldsymbol{\theta}_{1:N}^*$ value simulated within the SGA algorithm. Instead our approach is to marginalise the variational distribution using

$$q(k_i = j|\mathbf{y}_{1:N,1:T_n}, \mathbf{k}_{1:i-1}) = \int_{\boldsymbol{\theta}_{1:N}^*} q_{\lambda_n}(\boldsymbol{\theta}_{1:N}^*|\mathbf{y}_{1:N,1:T_n})p(k_i = j|\mathbf{y}_{1:N,1:T_n}, \boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:i-1})d\boldsymbol{\theta}_{1:N}^*, \quad (56)$$

before each update, estimating (56) from a sample average of $p(k_i = j|\mathbf{y}_{1:N,1:T_n}, \boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:i-1})$ using M samples $\boldsymbol{\theta}_{1:N}^*$ and $\mathbf{k}_{1:i-1}$ simulated from the available approximate distribution. This requires use of all observed data at T_n , for each of the M samples, but is independent of $\boldsymbol{\theta}_{1:N}^*$ and thus data up to T_n is not required as new $\boldsymbol{\theta}_{1:N}^*$ values are simulated in the SGA algorithm. The component of the variational approximation for k_i is then replaced by

$$\widehat{p}(k_i = j|\mathbf{y}_{1:N,1:T_{n+1}}, \boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:i-1}) \propto p(\mathbf{y}_{i,T_n+1:T_{n+1}}|\boldsymbol{\theta}_{1:N}^*, k_i = j)q(k_i = j|\mathbf{y}_{1:N,1:T_{n+1}}, \mathbf{k}_{1:i-1}), \quad (57)$$

which may be calculated using only the newly observed data $\mathbf{y}_{1:N,T_n+1:T_{n+1}}$ in the SGA algorithm. Note that the marginalisation step for all updates uses the exact full conditional distribution from the CRP representation, $p(k_i = j|\mathbf{y}_{1:N,1:T_n}, \boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:i-1})$, rather than the marginalised form $\widehat{p}(k_i = j|\mathbf{y}_{1:N,1:T_{n+1}}, \boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:i-1})$ from the previous update.

The targeted pseudo-posterior distribution for the update at T_{n+1} is given by

$$\widehat{p}(\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N}|\mathbf{y}_{1:N,1:T_{n+1}}) \propto \prod_{i=1}^N [p(\mathbf{y}_{i,T_n+1:T_{n+1}}|\boldsymbol{\theta}_{1:N}^*, k_i)q(k_i|\mathbf{y}_{1:N,1:T_n}, \mathbf{k}_{1:i-1})] q_{\lambda_n}^*(\boldsymbol{\theta}_{1:N}^*|\mathbf{y}_{1:N,1:T_n}), \quad (58)$$

where the base distribution of the DP posterior in the DPM (and its corresponding CRP) is replaced with its associated variational approximation at time T_n . The approximating distribution for the update at time T_{n+1} is given by

$$q_{\lambda_{n+1}}(\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N}|\mathbf{y}_{1:N,1:T_{n+1}}) = \prod_{j=1}^N q_{j,n+1}(\theta_j^*|\mathbf{y}_{1:N,1:T_{n+1}}) \prod_{i=1}^N \widehat{p}(k_i|\mathbf{y}_{1:N,1:T_{n+1}}, \boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:i-1}). \quad (59)$$

Given the pseudo-posterior (50), form of approximating distribution (59), and components of the time T_n approximation: $q_{\lambda_n}(\boldsymbol{\theta}_{1:N}^*|\mathbf{y}_{1:N,1:T_n})$ and $q(k_i = j|\mathbf{y}_{1:N,1:T_n}, \mathbf{k}_{1:i-1})$, the optimal parameter vector at time T_{n+1} , λ_{n+1}^* , may be obtained via Algorithm 3.

Algorithm 3: UVB for the DPM

Input: DP base distribution G_0 or updated approximating distribution at T_n .

Result: Approximating distribution at T_{n+1} .

Calculate (56) for all i .

Observe $\mathbf{y}_{1:N, T_n+1:T_{n+1}}$.

Set $\mathcal{L}(q, \lambda_{n+1}^{(0)}) = -\infty$.

Set initial values $\lambda_{n+1}^{(1)}$.

Set $m = 1$.

while $|\mathcal{L}(q, \lambda_{n+1}^{(m)}) - \mathcal{L}(q, \lambda_{n+1}^{(m-1)})| < \epsilon$ **do**

 Simulate $\boldsymbol{\theta}_{1:N}^{*(s)} \sim q_{\lambda_{n+1}^{(m)}}(\boldsymbol{\theta}_{1:N}^*|\mathbf{y}_{1:N,1:T_{n+1}})$ for $s = 1, 2, \dots, S$.

 Simulate $\mathbf{k}_{1:N}^{(s)}$ with probabilities (52) or (57).

 Evaluate $\tilde{p}(\mathbf{y}_{1:N,1:T_{n+1}}, \boldsymbol{\theta}_{1:N}^{*(s)}, \mathbf{k}_{1:i-1}^{(s)})$.

 Evaluate $q_{\lambda_{n+1}^{(m)}}(\boldsymbol{\theta}_{1:N}^{*(s)}, \mathbf{k}_{1:N}^{(s)}|\mathbf{y}_{1:N,1:T_{n+1}})$.

 Evaluate $\partial q_{\lambda_{n+1}^{(m)}}(\boldsymbol{\theta}_{1:N}^{*(s)}, \mathbf{k}_{1:N}^{(s)}|\mathbf{y}_{1:N,1:T_{n+1}})/\partial \lambda_{n+1}$.

 Update auxiliary parameter $\lambda_{n+1}^{(m+1)} = \lambda_{n+1}^{(m)} + \rho^{(m)} \frac{\partial \widehat{\mathcal{L}}(q, \lambda_{n+1})}{\partial \lambda_{n+1}} \Big|_{\lambda_{n+1} = \lambda_{n+1}^{(m)}}$.

 Calculate $\mathcal{L}(q, \lambda_{n+1}^{(m+1)})$.

 Set $m = m + 1$.

end

7.4 Predicting Lane Positions

Given a posterior approximation $q_{\lambda_n}(\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N}|\mathbf{y}_{1:N,1:T_n})$ we may obtain the approximate predictive distribution for vehicle i at some future time $T_n + h$ as

$$q(y_{i, T_n+h}|\mathbf{y}_{1:N,1:T_n}) = \int p(y_{i, T_n+h}|\boldsymbol{\theta}_{1:N}^*, k_i) q_{\lambda_n}^*(\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N}|\mathbf{y}_{1:N,1:T_n}) d\boldsymbol{\theta}_{1:N}^* d\mathbf{k}_{1:N}. \quad (60)$$

After obtaining this distribution from samples $\{\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N}\}^{(j)} \sim q_{\lambda_n}^*(\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N}|\mathbf{y}_{1:N,1:T_n})$, for $j = 1, 2, \dots, M$, we calculate the predictive log score (LS),

$$LS_{i,n,h} = \log(q(y_{i, T_n+h}^{(obs)}|\mathbf{y}_{1:N,1:T_n})), \quad (61)$$

where $y_{i, T_n+h}^{(obs)}$ is the observed value of y_{i, T_n+h} . The performance of the UVB algorithm is evaluated by comparing its cumulative predictive log scores relative to those produced by competing methods.

We also infer the DPM model via MFVB using the so-called ‘stick-breaking’ representation of the Dirichlet Process, as in Wang et al. (2011). This approach estimates the fully factorised posterior approximation, given by

$$q_{\lambda}(\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N}|\mathbf{y}_{1:N,1:T_n}) = \prod_{j=1}^N q(\boldsymbol{\theta}_j^*|\mathbf{y}_{1:N,1:T_n}) q(k_j|\mathbf{y}_{1:N,1:T_n}). \quad (62)$$

This may be used to build a predictive distribution in the same manner as (60). Details of the MFVB approximation are provided in Appendix C.

To illustrate the benefits of including posterior dependence in the approximation, we also introduce a parametric and *independent* model, which retains a normal likelihood for each vehicle, i.e.

$$y_{i,t} \sim \mathcal{N}(\mu_i, \sigma_i^2) \quad (63)$$

and assumes for each vehicle an independent uninformative prior, given by

$$p(\mu_i, \sigma_i^2) \propto \sigma_i^{-2}. \quad (64)$$

For this model the predictive distribution for vehicle i is analytically available as

$$p(y_{i,T_n+h} | \mathbf{y}_{i,1:T_n}) = \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2}) \sqrt{\frac{\pi\nu(T_n+1)s_{i,n}^2}{T_n}}} \left(1 + \frac{T_n (y_{i,T_n+h} - \bar{y}_{i,n})^2}{\nu(T_n+1)s_{i,n}^2} \right)^{-\frac{(\nu+1)}{2}} \quad (65)$$

a location-scale transform of the usual Student- t distribution with $\nu = T_n - 1$ degrees of freedom, where $\bar{y}_{i,n}$ and $s_{i,n}^2$ denote the sample mean and variance of $\mathbf{y}_{i,1:T_n}$, respectively. Note that this model ignores any information from all other vehicles, and is similarly evaluated by the corresponding cumulative predictive log score.

7.5 Analysis of the NGSIM Data

We now discuss the empirical application results from the UVB algorithm for the DPM model described above for the NGSIM data. The posterior updates for both the cluster locations, $\boldsymbol{\theta}_{1:N}^*$, and the indicator variables, $\mathbf{k}_{1:N}$, occur at a sequence of pre-determined time periods, given by $T_1 = 50, T_2 = 75, T_3 = 100, T_4 = 125, T_5 = 150$, and $T_6 = 175$.

Consider first the two graphs shown in the top panel (panel (a)) of Figure 6. In each graph, the approximate marginal posterior distributions for each unique value μ_j^* (on the left) and $\sigma_j^{2,*}$ (on the right). Noting there are $N = 500$ marginal densities for each of μ^* and $\sigma^{2,*}$, the plotted densities for each parameter are weighted according to the proportion of vehicles in a sample of $M = 100$ draws of $(\boldsymbol{\theta}_{1:N}^*, \mathbf{k}_{1:N})$ obtained from the UVB approximation. That is, the weights are calculated according to

$$w_j = \sum_{m=1}^M \sum_{i=1}^N \frac{I(k_i^{(m)} = j)}{MN}, \quad (66)$$

so that w_j represents the proportion of the MN many sampled k_i values, denoted by $k_i^{(m)}$ for $i = 1, 2, \dots, N$ and $m = 1, 2, \dots, M$, that correspond to the given value of j . The weights suggest that only six of the $\boldsymbol{\theta}_j^*$ values account for the majority of the vehicles, with the six weighted densities associated with μ_k^* and σ_k^* most prominent in the figures shown in panel (a). In contrast, the sample of $\boldsymbol{\theta}_j^*$ values that are seldom (if ever) allocated to a vehicle and hence receive little or no weight appear in these figures as flat lines indistinguishable from zero.

Now turning to panel (b) of Figure 6, a predictive distribution for new values of y is estimated for each cluster location j , using the M previously simulated values $\boldsymbol{\theta}_{1:N}^*$. The mean of each predictive distribution is plotted against the corresponding predictive standard deviation, with the size of

each point given by w_j . The fifty pairs of means and standard deviations shown correspond to 80% of all simulated k_i values, with the results showing that the majority of vehicles belong to a relatively small number of large and cohesive groups, each associated with a distinct predictive mean value coupled with low predictive standard deviation. Members of these groups appear to stay in the same region of their lane, but with these regions spread across both sides of the centre line. There are also many smaller groups, having predictive means closer to zero but with larger standard deviations, perhaps describing idiosyncratic vehicle positioning in the region of the centre lane.

The bottom panel plots, in grey, the individual predictive densities associated with fifty randomly selected vehicles, with the average predictive density over all $N = 500$ vehicles in the sample shown in dark blue. Note that the predictive distribution associated with an individual vehicle will typically itself be comprised of a mixture of components. Importantly, many of the individual predictive densities display reduced uncertainty, relative to the overall average.

Using data up to each time period T_n , we predict the future position y_{i,T_n+h} , $h = 1, 2, \dots, 50$ for each vehicle using four different predictive distributions described in Section 7.4:


1. The DPM predictive distribution (60), with approximate inference provided via UVB.
2. The DPM predictive distribution (60), with approximate inference provided via MFVB,
3. The DPM predictive distribution (60), with approximate inference provided via SVB,
4. The independent model predictive distribution (65), with exact inference.

The mean cumulative predictive log scores (MCLS), averaged across each of the $N = 500$ vehicles, and associated with each of the four types of predictive distributions for individual cars enumerated above, are plotted in Figure 7.


The results show that, while in each case both approximate implementations of the DPM model outperform the analytically exact independent model, the posterior dependency in the SVB and UVB approximations greatly improves forecasts relative to MFVB. The UVB and SVB lines coincide, and there is no evidence of accumulating approximation error through the UVB recursion relative to the single model fit of SVB. As the amount of data increases, the MFVB and independent model log scores similarly increase. In contrast, the UVB inference MCLS stays at the same level: the $N \times (T_6 - T_1) = 62,500$ additional observations included in T_6 has not provided much marginal information to improve forecasts relative to the original T_1 fit with $N \times T_1 = 25,000$ observations. By construction the DPM shares information between vehicles, so forecasts of vehicle i are accurate even with only $T_1 = 50$ observations of that particular vehicle. When MFVB inference is employed forecasts are only slightly stronger than the fully independent model that does not share information, implying that the MFVB implementation did not successfully include behaviour of other vehicles.

8 Conclusions

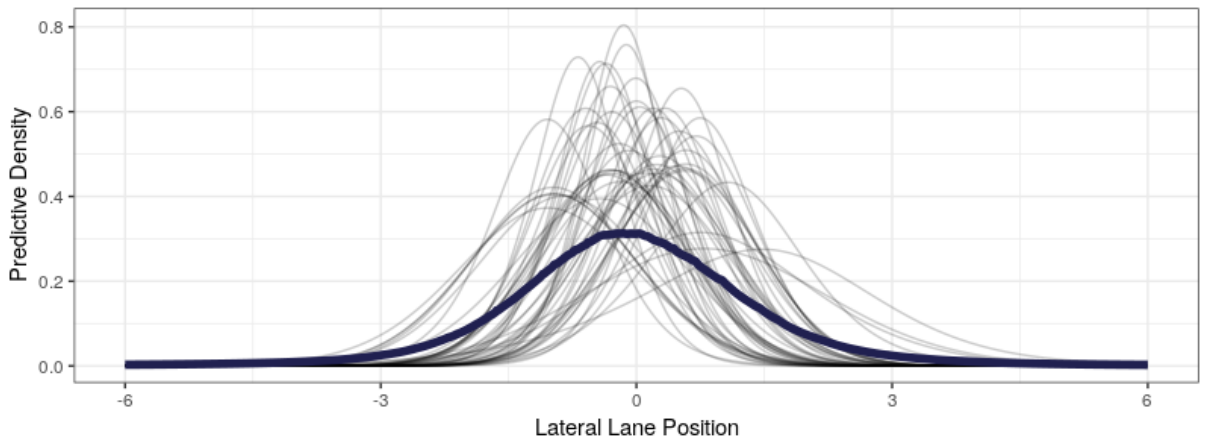
This paper proposes a framework to extend the use of SVB inference to a sequential posterior updating setting. UVB is a variational analogue to exact Bayesian updating, where the previous posterior distribution, taken as the prior for the update, is replaced with an approximation itself



(a) Weighted θ^* marginal posterior approximation, based on UVB, at time T_6 .



(b) UVB Predictive moments for high probability groups, at time T_6 .



(c) Individual vehicle and average predictive densities from UVB at time T_6 .

Figure 6: (a): Posterior approximation for each θ^* , weighted by proportion of $\mathbf{k}_{1:N}$ draws. Two groups have high posterior precision with numerous groups showing more uncertainty. (b): Posterior predictive distribution means and standard deviations, sized according to the top 80% of $\mathbf{k}_{1:N}$ draws. (c): Averaged predictive distribution for all groups in dark blue, with a random subset of fifty per vehicle distributions in grey.




Figure 7: Mean cumulative predictive log scores (MCLS) for each model averaged across $N = 500$ vehicles. Each model is fit using data up to T_n , then forecasts are made for each of the following fifty observations. The SVB and UVB implementations are nearly identical, while the MFVB implementation performs only slightly better than the fully independent model.

derived from an earlier SVB approximation. This allows a sequence of posterior distributions to be available in an online data setting, before observation of the complete dataset. Repeated updates may incur a large computational overhead from multiple stochastic optimisation algorithms, and so UVB-IS is proposed as an extension to UVB, as it significantly reduces the computational load of updates by using draws from previous posterior approximations, as well as their associated log-likelihoods, in an importance sampler.

We demonstrate the application of UVB and UVB-IS in two simulation settings, namely for time series forecasting and clustering, as well as in the empirical ‘Eight Schools’ example involving a hierarchical model. In each case the approximation error introduced by UVB and UVB-IS is assessed relative to both exact inference, and a standard implementation of SVB, where all past observations are used in a single optimisation. We find that, despite UVB introducing additional approximations, there are mixed results in terms of accuracy and computational speed between UVB and SVB, with SVB having less error and being faster in scenarios where the log-likelihood component of the gradient estimation is small. However, in the clustering model, where the log-likelihood computation dominates the gradient estimation, UVB outperforms SVB in terms of both accuracy and speed. In each case UVB-IS introduces a small amount of additional error relative to UVB, but nevertheless still achieves large computational gains. In addition, the variance of the gradient estimator in UVB and UVB-IS is shown to be orders of magnitude less than that of SVB, which may partially offset the approximation error in the prior used by the updating methods.

The proposed UVB and UVB-IS algorithms are highly suited to large time-series data streams where up-to-date inference is required at all times to inform decisions, including those based on forecasts and where data arrives so rapidly as to render MCMC infeasible. To illustrate this type of situation, an empirical illustration regarding observed lane positions of vehicles on the US-101 Highway is presented using a Dirichlet Process Mixture. In this implementation of UVB, an ap-

proximating distributional family that exploits dependence between cluster locations and indicator variables is detailed. Forecasts of future lane positions produced using UVB are comparable to an SVB approach. Posterior dependence is induced by exploiting the known full conditional distribution for the discrete indicator variables by using these as a component of the approximating distribution. Inferring the model through UVB and SVB outperform inference using MFVB, as this method requires an independent posterior approximation. Future research involves the application of UVB to build a more sophisticated heterogeneous model to provide forecasts of vehicle movement from this dataset in an online fashion – where UVB facilitates model updates and forecasts in a short time-frame after data arrives.

References

- Aldous, D. J. (1985). Exchangeability and related topics. In Ecole d’Ete de Probabilites de Saint-Flour XIII 1983.
- Attias, H. (1999). A variational Bayesian framework for graphical models. In Advances in Neural Information Processing Systems 12.
- Bhattacharya, A. and S. P. Wilson (2018). Sequential Bayesian inference for static parameters in dynamic state space models. Computational Statistics and Data Analysis 127, 187 – 203.
- Bishop, C. M. (2006). Pattern Recognition and Machine Learning. Springer.
- Blei, D. M., A. Kucukelbir, and J. D. McAuliffe (2017). Variational inference: A review for statisticians. Journal of the American Statistical Association 112(518), 859–877.
- Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. In Proceedings of COMPSTAT’2010.
- Boyen, X. and D. Koller (2013, jan). Tractable Inference for Complex Stochastic Processes.
- Broderick, T., N. Boyd, A. Wibisono, A. C. Wilson, and M. I. Jordan (2013). Streaming variational Bayes. In Advances in Neural Information Processing Systems 26.
- Chen, X., H. Dai, and L. Song (2019). Meta particle flow for sequential Bayesian inference. arXiv:1902.00640 [cs.LG].
- Chopin, N. (2002). A sequential particle filter method for static models. Biometrika 89, 539–551.
- Del Moral, P., A. Jasra, A. Lee, C. Yau, and X. Zhang (2015). The alive particle filter and its use in particle markov chain monte carlo. Stochastic Analysis and Applications 33(6), 943–974.
- Dißmann, J., E. C. Brechmann, C. Czado, and D. Kurowicka (2013). Selecting and estimating regular vine copulae and application to financial returns. Computational Statistics and Data Analysis 59, 52–69.
- Doucet, A., N. de Freitas, and N. Gordon (2001). Sequential Monte Carlo Methods in Practice. Springer-Verlag.
- Doucet, A. and A. Lee (2018). Sequential Monte Carlo methods, Chapter 7, pp. 165–189. Chapman and Hall.

- FHWA (2017). Next Generation Simulation (NGSIM) Vehicle Trajectories and Supporting Data. Available online at <https://data.transportation.gov/Automobiles/Next-Generation-Simulation-NGSIM-Vehicle-Trajectory/8ect-6jqj>.
- Garthwaite, P. H., Y. Fan, and S. A. Sisson (2016). Adaptive optimal scaling of Metropolis–Hastings algorithms using the Robbins–Monro process. Communications in Statistics - Theory and Methods 45(17), 5098–5111.
- Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin (2014). Bayesian Data Analysis (3 ed.). CRC Press.
- Gelman, A., W. R. Gilks, and G. O. Roberts (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. Annals of Applied Probability 7, 110–120.
- Ghahramani, Z. and M. J. Beal (2000). Propagation algorithms for variational Bayesian learning. In Advances in Neural Information Processing Systems 13.
- Gilks, W. R., N. G. Best, and K. K. C. Tan (1995). Adaptive rejection Metropolis sampling within Gibbs sampling. Journal of the Royal Statistical Society. Series C (Applied Statistics) 44(4), 445–472.
- Gilks, W. R., S. Richardson, and D. Spiegelhalter (1995). Markov Chain Monte Carlo in Practice. Chapman and Hall.
- Hoffman, M., F. R. Bach, and D. M. Blei (2010). Online learning for latent Dirichlet allocation. In Advances in Neural Information Processing Systems 23.
- Hoffman, M., D. M. Blei, C. Wang, and J. Paisley (2013). Stochastic variational inference. Journal of Machine Learning Research 14, 1303–1347.
- Jasra, A., S. S. Singh, and J. S. Martin (2010). Filtering via approximate Bayesian computation. Statistics and Computing 22, 1223–1237.
- Jordan, M. I., Z. Ghahramani, T. S. Jaakola, and L. K. Saul (1999). An introduction to variational methods for graphical models. Machine Learning 37, 183–233.
- Kabisa, S. T., D. B. Dunson, and J. S. Morris (2016). Online variational Bayes inference for high-dimensional correlated data. Journal of Computational and Graphical Statistics 25, 426–444.
- Kingma, D. P. and J. L. Ba (2014). [arXiv:1412.6980v9 \[cs.LG\]](https://arxiv.org/abs/1412.6980).
- Kingma, D. P. and M. Welling (2014). Auto-encoding variational Bayes. [arXiv:1312.6114 \[stat.ML\]](https://arxiv.org/abs/1312.6114).
- Kullback, S. and R. A. Leibler (1951). On information and sufficiency. The Annals of Mathematical Statistics 22(1), 79–86.
- Lauritzen, S. L. (1992). Propagation of probabilities, means, and variances in mixed graphical association models. Journal of the American Statistical Association.
- Llera, A. and C. F. Beckmann (2016). Estimating an inverse gamma distribution. [arXiv:1605.01019 \[stat.ME\]](https://arxiv.org/abs/1605.01019).
- Maybeck, P. S. (1982). Stochastic models, estimation, and control. Academic press.

- Miller, J. W. (2018). An elementary derivation of the chinese restaurant process from Sethuraman’s stick-breaking process. [arXiv:1801.00513 \[math.ST\]](https://arxiv.org/abs/1801.00513).
- Müller, P., F. A. Quintana, A. Jara, and T. Hanson (2015). Bayesian Nonparametric Data Analysis. Springer.
- Ong, V. M.-H., D. J. Nott, and M. S. Smith (2018). Gaussian variational approximation with a factor covariance structure. Journal of Computational and Graphical Statistics 27(3), 465–478.
- Opper, M. and O. Winther (1998). A bayesian approach to on-line learning. On-line learning in neural networks, 363–378.
- R Core Team (2017). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing. Available online at <https://www.R-project.org/>.
- Ranganath, R., S. Gerrish, and M. Blei, David (2014). Black box variational inference. In Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics.
- Robbins, H. and S. Monro (1951). A stochastic approximation method. The Annals of Mathematical Statistics 22(3), 400–407.
- Sakaya, J. and A. Klami (2017). Importance sampled stochastic optimization for variational inference. In Uncertainty in Artificial Intelligence.
- Sato, M. (2001). Online model selection based on variational Bayes. Neural Computation 13, 1649–1681.
- Sethuraman, J. (1994). A constructive definition of Dirichlet priors. Statistica Sinica 4, 639–650.
- Smidl, V. (2004). The Variational Bayes Approach in Signal Processing. Ph. D. thesis, Trinity College, The University of Dublin.
- Stan Development Team (2018). RStan Getting Started. Available online at <https://github.com/stan-dev/rstan/wiki/RStan-Getting-Started>.
- Tan, L. S. and D. J. Nott (2018). Gaussian variational approximation with sparse precision matrices. Statistics and Computing 28(2), 259–275.
- Titsias, M. K. and M. Lázaro-Gredilla (2014). Doubly stochastic variational Bayes for non-conjugate inference. In Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32.
- Wainwright, M. J. and M. I. Jordan (2008). Graphical models, exponential families, and variational inference. Foundations and Trends in Machine Learning 1(1-2), 1–305.
- Wang, C., J. Paisley, and D. Blei (2011). Online variational inference for the hierarchical Dirichlet process. In Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics.
- Zhang, C., J. Bütepage, H. Kjellström, and S. Mandt (2017). Advances in variational inference. [arXiv:1711.05597 \[cs.LG\]](https://arxiv.org/abs/1711.05597).

A Calculation of Lateral Lane Deviation

Let $x_{i,t}$ denote the position of vehicle i along the direction of travel at time t , and $y_{i,t}$ denote the position across the lane, as in Figure 8 for one vehicle travelling to the right.




Figure 8: Coordinate system for one vehicle. The X -axis denotes distance travelled along the lane, and the Y -axis denotes the relative vertical position in the lane.

For each vehicle i and time t since entering the road, with travel originating at $t = 1$, the total distance travelled up to time t is given by

$$d_{i,t} = \sum_{s=2}^t \sqrt{(x_{i,s} - x_{i,s-1})^2 + (y_{i,s} - y_{i,s-1})^2}. \quad (67)$$

Using this distance measure and 100 randomly sampled vehicles per lane, the two-dimensional coordinates corresponding to the centre line of each lane are estimated via independent smoothing splines, where each coordinate is a function of the distance travelled to that point. Each smoothing spline is calculated using the ‘R stats’ package (R Core Team, 2017). The estimated centre line for lane k , is denoted by the curve $\{\hat{x}_{d,k} = f_x^k(d), \hat{y}_{d,k} = f_y^k(d)\}$, for $d \geq 0$. The fitted spline models are shown in red overlaying the raw data in Figure 9




Figure 9: The fitted spline models for each lane in red overlaid on the raw data in black.

Excluding the vehicles used to estimate the spline models, each of the vehicles in the dataset uses the relevant lane centre line estimate fit from the spline model associated with its starting lane to

calculate relative coordinates $\{x_{i,t}^*, y_{i,t}^*\}$. $x_{i,t}^*$ denotes the distance travelled along the road, and $y_{i,t}^*$ denotes the deviation from the lane centre line, and are calculated by

$$\widehat{d}_{i,t} = \arg \min_d \sqrt{(x_{i,t} - f_x^k(d))^2 + (y_{i,t} - f_y^k(d))^2} \quad (68)$$

$$\widehat{x}_{i,t} = f_x^k(\widehat{d}_{i,t}) \quad (69)$$

$$\widehat{y}_{i,t} = f_y^k(\widehat{d}_{i,t}) \quad (70)$$

$$y_{i,t}^* = \text{sign} \left(\tan \left(\frac{\widehat{x}_{i,t} - x_{i,t}}{\widehat{y}_{i,t} - y_{i,t}} \right)^{-1} - \tan \left(\frac{f_x^{\prime,k}(\widehat{d}_{i,t})}{f_y^{\prime,k}(\widehat{d}_{i,t})} \right)^{-1} \right) \sqrt{(x_{i,t} - \widehat{x}_{i,t})^2 + (y_{i,t} - \widehat{y}_{i,t})^2} \quad (71)$$

The coordinate pair $(\widehat{x}_{i,t}, \widehat{y}_{i,t})$ denotes the closest position of the spline model to the actual vehicle position given by the pair $(x_{i,t}, y_{i,t})$. Lateral deviation $y_{i,t}^*$ has magnitude equal to that of the vector from $(\widehat{x}_{i,t}, \widehat{y}_{i,t})$ to $(x_{i,t}, y_{i,t})$. A negative sign on $y_{i,t}^*$ indicates that the vehicle is to the left of the lane centre, and a positive sign indicates that the vehicle is to the right of the lane centre.

B Equivalence of augmented and marginal KL Divergence gradients

Consider the augmented posterior distribution

$$p(\theta, k|y) \propto p(y|\theta, k)p(k|\theta)p(\theta) \quad (72)$$

and variational approximation given by

$$q_\lambda(\theta, k|y) = q_\lambda(\theta|y)p(k|y, \theta). \quad (73)$$

The corresponding KL divergence, $KL[q_\lambda(\theta, k|y) || p(\theta, k|y)]$, is indirectly minimised using the gradient

$$\frac{\partial KL[q_\lambda(\theta, k|y) || p(\theta, k|y)]}{\partial \lambda} = - \frac{\partial \mathcal{L}(q, \lambda)}{\partial \lambda}, \quad (74)$$

where the gradient $\partial \mathcal{L}(q, \lambda)/\partial \lambda$ is the score-based gradient of the ELBO, given by

$$\frac{\partial \mathcal{L}(q, \lambda)}{\partial \lambda} = \int_{\theta, k} q_\lambda(\theta, k|y) \frac{\partial \log(q_\lambda(\theta, k|y))}{\partial \lambda} (\log(p(y, \theta, k)) - \log(q_\lambda(\theta, k|y))) d\theta dk. \quad (75)$$

Next, consider the associated marginal posterior distribution,

$$p(\theta|y) \propto p(y|\theta)p(\theta), \quad (76)$$

and consider using as the variational approximation $\tilde{q}_\lambda(\theta|y)$ given by the first component (only) on the right hand side of (73), i.e. $\tilde{q}_\lambda(\theta|y) \equiv q_\lambda(\theta|y)$. Note that, as a consequence, $\log q_\lambda \equiv \log \tilde{q}_\lambda$ and $\frac{\partial \log q_\lambda}{\partial \lambda} \equiv \frac{\partial \log \tilde{q}_\lambda}{\partial \lambda}$. The KL divergence in this case, $KL[\tilde{q}_\lambda(\theta|y) || p(\theta|y)]$, has gradient given by

$$\frac{\partial KL[\tilde{q}_\lambda(\theta|y) || p(\theta|y)]}{\partial \lambda} = - \frac{\partial \mathcal{L}(\tilde{q}, \lambda)}{\partial \lambda}, \quad (77)$$

where $\partial \mathcal{L}(\tilde{q}, \lambda)/\partial \lambda$ is

$$\frac{\partial \mathcal{L}(\tilde{q}, \lambda)}{\partial \lambda} = \int_{\theta} \tilde{q}_\lambda(\theta|y) \frac{\partial \log(\tilde{q}_\lambda(\theta|y))}{\partial \lambda} (\log(p(y, \theta)) - \log(\tilde{q}_\lambda(\theta|y))) d\theta. \quad (78)$$

Here we show that (75) is equal to (78), and hence the gradient of both KL divergences are equal, and must share local minima.

Begin by expanding each joint density in (75) by (72) and (73),

$$\begin{aligned} \frac{\partial \mathcal{L}(q, \lambda)}{\partial \lambda} &= \int_{\theta, k} q_\lambda(\theta|y)p(k|\theta, y) \frac{\partial(\log(q_\lambda(\theta|y)) + \log(p(k|\theta, y)))}{\partial \lambda} \\ &\quad \times (\log(p(\theta)p(y|k, \theta)p(k|\theta)) - \log(q_\lambda(\theta|y)p(k|\theta, y))) d\theta dk \end{aligned} \quad (79)$$

$$\begin{aligned} &= \int_{\theta, k} q_\lambda(\theta|y)p(k|\theta, y) \frac{\partial \log(q_\lambda(\theta|y))}{\partial \lambda} \\ &\quad \times \left(\log(p(\theta)) + \log\left(\frac{p(y|k, \theta)p(k|\theta)p(y|\theta)}{p(y|\theta)}\right) - \log(q_\lambda(\theta|y)) - \log(p(k|\theta, y)) \right) d\theta dk \end{aligned} \quad (80)$$

as the term $\log(p(k|\theta, y))$ is independent of λ . Then

$$\begin{aligned} \frac{\partial \mathcal{L}(q, \lambda)}{\partial \lambda} &= \int_{\theta, k} q_\lambda(\theta|y)p(k|\theta, y) \frac{\partial \log(q_\lambda(\theta|y))}{\partial \lambda} \\ &\quad \times (\log(p(\theta)) + \log(p(y|\theta)) + \log(p(k|y, \theta)) - \log(q_\lambda(\theta|y)) - \log(p(k|\theta, y))) d\theta dk \end{aligned} \quad (81)$$

by Bayes' Rule. Cancelling $\log(p(k|y, \theta))$ results in

$$\frac{\partial \mathcal{L}(q, \lambda)}{\partial \lambda} = \int_{\theta, k} q_\lambda(\theta|y)p(k|\theta, y) \frac{\partial \log(q_\lambda(\theta|y))}{\partial \lambda} (\log(p(\theta)) + \log(p(y|\theta)) - \log(q_\lambda(\theta|y))) d\theta dk \quad (82)$$

$$= \int_{\theta} \left(\int_k p(k|\theta, y) dk \right) q_\lambda(\theta|y) \frac{\partial \log(q_\lambda(\theta|y))}{\partial \lambda} (\log(p(y, \theta)) - \log(q_\lambda(\theta|y))) d\theta. \quad (83)$$

The final expression (83) is equivalent to the marginal model gradient (78) and the proof is complete.

C Mean Field Variational Bayes implementation of the Dirichlet Process Mixture

Implementation of MFVB for this model follows the offline coordinate ascent approach of Wang et al. (2011), employing the stick-breaking construction of the Dirichlet Process as

$$\theta_j^* \stackrel{iid}{\sim} G_0, \quad (84)$$

$$\beta_j' \stackrel{iid}{\sim} \text{Beta}(1, \alpha), \quad (85)$$

$$\beta_N' = 1, \quad (86)$$

$$\beta_j = \beta_j' \prod_{l=1}^{j-1} (1 - \beta_l'), \quad (87)$$

$$G = \sum_{j=1}^N \beta_j \delta(\theta_j^*), \quad (88)$$

where δ is the Dirac Delta function. The stick-breaking construction is equivalent to the CRP representation of the DP, after marginalisation over β (Miller, 2018), and is similarly augmented with the set of indicator variables $k_{1:N}$. In this case the prior distribution is given by

$$k_i \sim \text{Multinomial}(\boldsymbol{\beta}_{1:N}). \quad (89)$$

The contribution to the likelihood from observation i is then determined by

$$\mathbf{y}_{i,1:T_1} | \boldsymbol{\theta}_{1:N}^*, k_i \sim \mathcal{N}(\mu_{k_i}^*, \sigma_{k_i}^{2*}). \quad (90)$$

To maintain the analytical tractability of the MFVB approximation, we replace the base distribution G_0 with a conjugate prior for the normal likelihood,

$$\mu^* | G_0 \sim \mathcal{N}(0, 10) \quad (91)$$

$$\sigma^{2*} | G_0 \sim \text{InverseGamma}(\text{shape} = \alpha_0, \text{scale} = \kappa_0) \quad (92)$$

where α_0 and κ_0 are chosen to be the MLE values for the inverse gamma distribution, estimated from 100,000 samples of the implied lognormal(0, 10) distribution for σ^{2*} that was used in the SVB and UVB approaches. These values are estimated by the second algorithm of Llera and Beckmann (2016) as

$$\alpha_0 = 0.15275 \quad (93)$$

$$\kappa_0 = 0.00102. \quad (94)$$

The variational approximation employed is of the form

$$q_{\lambda_n}(\mathbf{k}_{1:N}, \boldsymbol{\beta}'_{1:N}, \boldsymbol{\theta}_{1:N}^*) = \prod_{i=1}^N q(k_i) q(\beta'_i) q(\mu_i^*) q(\sigma_i^{2*}) \quad (95)$$

with

$$k_i \sim \text{Multinom}(\boldsymbol{\rho}_i) \quad (96)$$

$$\beta'_i \sim \text{Beta}(a_i, b_i) \quad (97)$$

$$\mu_i^* \sim \mathcal{N}(\gamma_i, \tau_i) \quad (98)$$

$$\sigma_i^{2*} \sim \text{Inv.Gamma}(\alpha_i, \kappa_i) \quad (99)$$

Coordinate ascent algorithms for MFVB consists of cycling through in a set of equations for each parameter until the change in each parameter is below some threshold. For this model the equations

are given by

$$a_j = 1 + \sum_{i=1}^N \rho_{ij}, \quad (100)$$

$$b_j = \alpha + \sum_{i=1}^N \sum_{l=j+1}^N \rho_{il}, \quad (101)$$

$$\rho_{ij} \propto -\frac{T_n}{2} E_q[\log(\sigma_j^{2*})] - \frac{\alpha_j}{2\kappa_j} \left(\sum_{t=1}^{T_n} y_{it}^2 - 2\gamma_j \sum_{t=1}^{T_n} y_{it} + T_n(\gamma_j^2 + \tau_j) \right) + E_q[\log(\beta_j)], \quad (102)$$

$$\gamma_j = \frac{10 \frac{\alpha_j}{\kappa_j} \sum_{i=1}^N \sum_{t=1}^{T_n} \rho_{ij} y_{it}}{10 T_n \frac{\alpha_j}{\kappa_j} \sum_{i=1}^N \rho_{ij} + 1}, \quad (103)$$

$$\tau_j = \frac{10}{10 T_n \frac{\alpha_j}{\kappa_j} \sum_{i=1}^N \rho_{ij} + 1}, \quad (104)$$

$$\alpha_j = \alpha_0 + \frac{T_n}{2} \sum_{i=1}^N \rho_{ij}, \quad (105)$$

$$\kappa_j = \kappa_0 + \frac{1}{2} \left(\sum_{i=1}^N \rho_{ij} \left(\sum_{t=1}^{T_n} y_{it}^2 + T_n(\gamma_j^2 + \tau_j) - 2\gamma_j \sum_{t=1}^{T_n} y_{it} \right) \right). \quad (106)$$

The expectations $E_q[\log(\beta_j)]$ are available in closed form as

$$E_q[\log(\beta_j)] = E_q[\log(\beta'_j)] + \sum_{l=1}^{j-1} E_q[\log(1 - \beta'_j)] \quad (107)$$

where

$$E_q[\log(\beta'_j)] = \Psi(a_j) - \Psi(a_j + b_j) \quad (108)$$

and

$$E_q[\log(1 - \beta'_j)] = \Psi(b_j) - \Psi(a_j + b_j) \quad (109)$$

where Ψ is the digamma function. The expectation $E_q[\log(\sigma_j^{2*})]$ does not have a closed form solution but is estimated from samples of $q(\sigma_j^{2*})$.