of latent dynamic behaviors which are individually modeled via temporally independent linear dynamical systems. Examples include the hidden Markov model (HMM), switching vector autoregressive (VAR) process, and switching linear dynamical system (SLDS). These models have proven useful in such diverse fields as speech recognition, econometrics, neuroscience, remote target tracking, and human motion capture. In this paper, we focus our attention on the descriptive yet computationally tractable class of switching VAR processes. Here, the state of the underlying Markov process encodes the behavior exhibited at a given time step, and each dynamic behavior defines a VAR process. That is, conditioned on the Markov-evolving state, the likelihood is simply a VAR model with time-varying parameters.

To discover the dynamic behaviors shared between multiple time series, we propose a feature-based model. The entire collection of time series can be described by a globally shared set of possible behaviors. Individually, however, each time series will only exhibit a subset of these behaviors. The goal of joint analysis is to discover which behaviors are shared among the time series and which are unique. We represent the behaviors possessed by time series i with a binary *feature vector* $\mathbf{f}_i$, with $f_{ik} = 1$ indicating that time series i uses global behavior k (see Figure 1). We seek a prior for these feature vectors which allows flexibility in the number of behaviors and encourages the sharing of behaviors. Our desiderata motivate a feature-based Bayesian nonparametric approach based on the *beta process* [Hjort (1990), Thibaux and Jordan (2007)]. Such an approach allows for *infinitely* many potential behaviors, but encourages a sparse representation. Given a fixed feature set, our model reduces to a collection of finite Bayesian VAR processes with partially shared parameters.

We refer to our model as the *beta-process autoregressive hidden Markov model*, or BP-AR-HMM. We also consider a simplified version of this model, referred to as the BP-HMM, in which the AR emission models are replaced with a set of conditionally independent emissions. Preliminary versions of these models were partially described in Fox et al. (2009) and in Hughes, Fox and Sudderth (2012), who developed improved Markov chain Monte Carlo (MCMC) inference procedures for the BP-AR-HMM. In the current article we provide a unified and comprehensive description of the model and we also take further steps toward the development of an efficient inference algorithm for the BP-AR-HMM. In particular, the *unbounded* nature of the set of possible behaviors available to our approach presents critical challenges during posterior inference. To efficiently explore the space, we introduce two novel MCMC proposal moves: (1) split-merge moves to efficiently change the feature structure for many sequences at once, and (2) data-driven reversible jump moves to add or delete features unique to one sequence. We expect the foundational ideas underlying both contributions (split-merge and data-driven birth–death) to generalize to other nonparametric models beyond the time-series domain. Building on an earlier version of these ideas in Hughes, Fox and Sudderth (2012), we show how to perform data-driven birth–death proposals