We begin by describing a model for the dynamics of each individual time series. We then describe a mechanism for representing dynamics which are shared between multiple data streams. Our Bayesian nonparametric prior specification plays a key role in this model, by addressing the challenge of allowing for uncertainty in the number of dynamic behaviors exhibited within and shared across data streams.

3.1. *Per-series dynamics.* We model the dynamics of each time series as a *Markov switching process* (MSP). Most simply, one could consider a hidden Markov model (HMM) [Rabiner (1989)]. For observations $\mathbf{y}_t \in \mathbb{R}^d$ and hidden state z_t, the HMM assumes

$$\begin{aligned} z_t \mid z_{t-1} &\sim \pi_{z_{t-1}}, \\ \mathbf{y}_t \mid z_t &\sim F(\theta_{z_t}), \end{aligned} \tag{1}$$

for an indexed family of distributions $F(\cdot)$. Here, π_k is the state-specific *transition distribution* and θ_k the *emission parameters* for state k.

The modeling assumption of the HMM that observations are conditionally independent given the latent state sequence is insufficient to capture the temporal dependencies present in human motion data streams. Instead, one can assume that the observations have *conditionally linear dynamics*. Each latent HMM state then models a single linear dynamical system, and over time the model can switch between dynamical modes by switching among the states. We restrict our attention in this paper to switching vector autoregressive (VAR) processes, or *autoregressive HMMs* (AR-HMMs), which are both broadly applicable and computationally practical.

We consider an AR-HMM where, conditioned on the latent state z_t, the observations evolve according to a state-specific order-r VAR process:[3]

$$\mathbf{y}_t = \sum_{\ell=1}^{r} A_{\ell,z_t} \mathbf{y}_{t-\ell} + \mathbf{e}_t(z_t) = \mathbf{A}_k \tilde{\mathbf{y}}_t + \mathbf{e}_t(z_t), \tag{2}$$

where $\mathbf{e}_t(z_t) \sim \mathcal{N}(0, \Sigma_{z_t})$ and $\tilde{\mathbf{y}}_t = [\,\mathbf{y}_{t-1}^T \quad \cdots \quad \mathbf{y}_{t-1}^T\,]^T$ are the aggregated past observations. We refer to $\mathbf{A}_k = [\,A_{1,k} \quad \cdots \quad A_{r,k}\,]$ as the set of *lag matrices*. Note that an HMM with zero-mean Gaussian emissions arises as a special case of this model when $\mathbf{A}_k = \mathbf{0}$ for all k. Throughout, we denote the VAR parameters for the kth state as $\theta_k = \{\mathbf{A}_k, \Sigma_k\}$ and refer to each VAR process as a *dynamic behavior*. For example, these parameters might each define a linear motion model for the behaviors *walking*, *running*, *jumping*, and so on; our time series are then each modeled as Markov switches between these behaviors. We will sometimes refer to k itself as a "behavior," where the intended meaning is the VAR model parameterized by θ_k.

[3]We denote an order-r VAR process by VAR(r).