mode sequence. Switched affine and piecewise affine models, which we do not consider in this article, instead allow mode transitions to depend on the continuous state of the dynamical system [17].

STATE-SPACE MODELS, VAR PROCESSES, AND FINITE MARKOV JUMP LINEAR SYSTEMS

A state-space model provides a general framework for analyzing many dynamical phenomena. The model consists of an underlying state, $x_t \in \mathbb{R}^n$, with linear dynamics observed via $y_t \in \mathbb{R}^d$. A linear time-invariant state-space model, in which the dynamics do not depend on time, is given by

$$x_t = Ax_{t-1} + e_t \qquad y_t = Cx_t + w_t, \tag{12}$$

where e_t and w_t are independent, zero-mean Gaussian noise processes with covariances Σ and R, respectively. The graphical model for this process is equivalent to that of the HMM depicted in Figure 1(a), replacing z_t with x_t.

An order r VAR process, denoted by VAR(r), with observations $y_t \in \mathbb{R}^d$, can be defined as

$$y_t = \sum_{i=1}^{r} A_i y_{t-i} + e_t \qquad e_t \sim \mathcal{N}(0, \Sigma). \tag{13}$$

Here, the observations depend linearly on the previous r observation vectors. Every VAR(r) process can be described in state-space form by, for example, the following transformation:

$$x_t = \begin{bmatrix} A_1 & A_2 & \dots & A_r \\ I & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \dots & I & 0 \end{bmatrix} x_{t-1} + \begin{bmatrix} I \\ 0 \\ \vdots \\ 0 \end{bmatrix} e_t$$
$$y_t = \begin{bmatrix} I & 0 & \dots & 0 \end{bmatrix} x_t. \tag{14}$$

On the other hand, not every state-space model may be expressed as a VAR(r) process for finite r [18].

Building on the HMM of the section "Finite HMM," we define an SLDS by

$$z_t \sim \pi_{z_{t-1}}$$
$$x_t = A^{(z_t)} x_{t-1} + e_t(z_t) \quad y_t = Cx_t + w_t. \tag{15}$$

Here, we assume the process noise $e_t(z_t) \sim \mathcal{N}(0, \Sigma^{(z_t)})$ is mode-specific, while the measurement mechanism is not. This assumption could be modified to allow for both a mode-specific measurement matrix $C^{(z_t)}$ and noise $w_t(z_t) \sim \mathcal{N}(0, R^{(z_t)})$. However, such a choice is not always necessary nor appropriate for certain applications, and can have implications on the identifiability of the model. We similarly define a switching VAR(r) process by

$$z_t \sim \pi_{z_{t-1}}$$
$$y_t = \sum_{i=1}^{r} A_i^{(z_t)} y_{t-i} + e_t(z_t). \tag{16}$$

Both the SLDS and the switching VAR process are contained within the class of MJLS, with graphical model representations shown in Figure 1(b) and (c). Compare to that of the HMM in Figure 1(a).

HDP-AR-HMM AND HDP-SLDS

In the formulations of the MJLS mentioned in the section "Markov Jump Linear Systems," it was assumed that the number of dynamical modes was known. However, it is often desirable to relax this assumption to provide more modeling flexibility. It has been shown that in such cases, the sticky HDP-HMM can be extended as a Bayesian nonparametric approach to learning both SLDS and switching VAR processes [19], [20]. Specifically, the transition distributions are defined just as in the sticky HDP-HMM. However, instead of independent observations, each mode now has conditionally linear dynamics. The generative processes for the resulting HDP-AR-HMM and HDP-SLDS are summarized in (17), shown at the bottom of the page, with an example HDP-AR-HMM observation sequence depicted in Figure 4(d).

Here, π_j is as defined in (10). The issue, then, is in determining an appropriate prior on the dynamic parameters. In [19], a conjugate matrix-normal inverse-Wishart (MNIW) prior [21] was proposed for the dynamic parameters $\{A^{(k)}, \Sigma^{(k)}\}$ in the case of the HDP-SLDS, and $\{A_1^{(k)}, \dots, A_r^{(k)}, \Sigma^{(k)}\}$ for the HDP-AR-HMM. The HDP-SLDS additionally assumes an inverse-Wishart prior on the measurement noise R; however, the measurement matrix, C, is fixed for reasons of identifiability. The MNIW prior assumes knowledge of either the autoregressive order r, or the underlying state dimension n, of the switching VAR process or SLDS, respectively. Alternatively, Fox et al. [20] explore an automatic relevance determination (ARD) sparsity-inducing prior [22]–[24] as a means of learning MJLS with variable order structure. The ARD prior penalizes nonzero components of the model through a zero-mean Gaussian prior with gamma-distributed precision. In the context of the HDP-AR-HMM and HDP-SLDS, a maximal autoregressive order or SLDS state dimension is assumed. Then, a structured version of the ARD prior is employed so as to drive entire VAR lag blocks $A_i^{(k)}$ or columns of the SLDS matrix $A^{(k)}$ to zero, allowing one to infer nondynamical components of a given dynamical mode.

	HDP-AR-HMM	HDP-SLDS	
Mode dynamics	$z_t \sim \pi_{z_{t-1}}$	$z_t \sim \pi_{z_{t-1}}$	
Observation dynamics	$y_t = \sum_{i=1}^{r} A_i^{(z_t)} y_{t-i} + e_t(z_t)$	$x_t = A^{(z_t)} x_{t-1} + e_t(z_t)$.	(17)
		$y_t = Cx_t + w_t$	