In recent independent work, Saria, Koller and Penn (2010) developed an alternative model for multiple time series via the HDP-HMM. Their *time series topic model* (TSTM) describes coarse-scale temporal behavior using a finite set of "topics," which are themselves distributions on a common set of autoregressive dynamical models. Each time series is assumed to exhibit all topics to some extent, but with unique frequencies and temporal patterns. Alternatively, the mixed HMM [Altman (2007)] uses generalized linear models to allow the state transition and emission distributions of a finite HMM to depend on arbitrary external covariates. In experiments, this is used to model the differing temporal dynamics of a small set of known time series classes.

More broadly, the problem we address here has received little previous attention, perhaps due to the difficulty of treating combinatorial relationships with parametric models. There are a wide variety of models which capture correlations among multiple aligned, interacting univariate time series, for example, using Gaussian state space models [Aoki and Havenner (1991)]. Other approaches cluster time series using a parametric mixture model [Alon et al. (2003)], or a Dirichlet process mixture [Qi, Paisley and Carin (2007)], and model the dynamics within each cluster via independent finite HMMs.

Dynamic Bayesian networks [Murphy (2002)], such as the factorial HMM [Ghahramani and Jordan (1997)], define a structured representation for the latent states underlying a single time series. Factorial models are widely used in applied time series analysis [Duh (2005), Lehrach and Husmeier (2009)]. The infinite factorial HMM [Van Gael, Teh and Ghahramani (2009)] uses the IBP to model a single time series via an infinite set of latent features, each evolving according to independent Markovian dynamics. Our work instead focuses on discovering behaviors shared across *multiple* time series.

Other approaches do not explicitly model latent temporal dynamics and instead aim to align time series with consistent global structure [Aach and Church (2001)]. Motivated by the problem of detecting temporal anomalies, Listgarten et al. (2006) describe a hierarchical Bayesian approach to modeling shared structure among a known set of time series classes. Independent HMMs are used to encode nonlinear alignments of observed signal traces to latent reference time series, but their states do not represent dynamic behaviors and are not shared among time series.

9. Motion capture experiments. The linear dynamical system is a common model for describing simple human motion [Hsu, Pulli and Popović (2005)], and the switching linear dynamical system (SLDS) has been successfully applied to the problem of human motion synthesis, classification, and visual tracking [Pavlović, Rehg and MacCormick (2000), Pavlović et al. (1999)]. Other approaches develop nonlinear dynamical models using Gaussian processes [Wang, Fleet and Hertzmann (2008)] or are based on a collection of binary latent features [Taylor, Hinton and Roweis (2006)]. However, there has been little effort in jointly segmenting