using only discrete assignment variables (marginalizing away continuous HMM parameters), and demonstrate that annealing the Hastings term in the acceptance ratio can dramatically improve performance.

Our presentation is organized as follows. Section 2 introduces motion capture data. In Section 3 we present our proposed beta-process-based model for multiple time series. Section 4 provides a formal specification of all prior distributions, while Section 5 summarizes the model. Efficient posterior computations based on an MCMC algorithm are developed in Section 6. The algorithm does not rely on model truncation; instead, we exploit the finite dynamical system induced by a fixed set of features to sample efficiently, while using data-driven reversible jump proposals to explore new features. Section 7 introduces our novel split-merge proposals, which allow the sampler to make large-scale improvements across many variables simultaneously. In Section 8 we describe related work. Finally, in Section 9 we present results on unsupervised segmentation of data from the CMU motion capture database [CMU (2009)]. Further details on our algorithms and experiments are available in the supplemental article [Fox et al. (2014)].

2. Motion capture data. Our data consists of motion capture recordings taken from the CMU MoCap database (http://mocap.cs.cmu.edu). From the available set of 62 positions and joint angles, we examine 12 measurements deemed most informative for the gross motor behaviors we wish to capture: one body torso position, one neck angle, two waist angles, and a symmetric pair of right and left angles at each subject's shoulders, wrists, knees, and feet. As such, each recording provides us with a 12-dimensional time series. A collection of several recordings serves as the observed data which our model analyzes.

An example data set of six sequences is shown in Figure 1. This data set contains three sequences from Subject 13 and three from Subject 14. These sequences were chosen because they had many exercises in common, such as "squat" and "jog," while also containing several unique behaviors appearing in only one sequence, such as "side bend." Additionally, we have human annotations of these sequences, identifying which of 12 exercise behaviors was present at each time step, as shown in Figure 1. These human segmentations serve as ground-truth for assessing the accuracy of our model's estimated segmentations (see Section 9). In addition to analyzing this small data set, we also consider a much larger 124 sequence data set in Section 9.

3. A featural model for relating multiple time series. In our applications of interest, we are faced with a *collection* of N time series representing realizations of related dynamical phenomena. Our goal is to discover dynamic behaviors shared between the time series. Through this process, we can infer how the data streams relate to one another as well as harness the shared structure to pool observations from the same behavior, thereby improving our estimates of the dynamic parameters.