Sharing Features among Dynamical Systems with Beta Processes

Emily B. Fox
Electrical Engineering & Computer Science, Massachusetts Institute of Technology
ebfox@mit.edu

Erik B. Sudderth
Computer Science, Brown University
sudderth@cs.brown.edu

Michael I. Jordan
Electrical Engineering & Computer Science and Statistics, University of California, Berkeley
jordan@cs.berkeley.edu

Alan S. Willsky
Electrical Engineering & Computer Science, Massachusetts Institute of Technology
willsky@mit.edu

Abstract

We propose a Bayesian nonparametric approach to the problem of modeling related time series. Using a beta process prior, our approach is based on the discovery of a set of latent dynamical behaviors that are shared among multiple time series. The size of the set and the sharing pattern are both inferred from data. We develop an efficient Markov chain Monte Carlo inference method that is based on the Indian buffet process representation of the predictive distribution of the beta process. In particular, our approach uses the sum-product algorithm to efficiently compute Metropolis-Hastings acceptance probabilities, and explores new dynamical behaviors via birth/death proposals. We validate our sampling algorithm using several synthetic datasets, and also demonstrate promising results on unsupervised segmentation of visual motion capture data.

1 Introduction

In many applications, one would like to discover and model dynamical behaviors which are shared among several related time series. For example, consider video or motion capture data depicting multiple people performing a number of related tasks. By jointly modeling such sequences, we may more robustly estimate representative dynamic models, and also uncover interesting relationships among activities. We specifically focus on time series where behaviors can be individually modeled via temporally independent or linear dynamical systems, and where transitions between behaviors are approximately Markovian. Examples of such *Markov jump processes* 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, remote target tracking, and human motion capture. Our approach envisions a large *library* of behaviors, and each time series or *object* exhibits a subset of these behaviors. We then seek a framework for discovering the set of dynamic behaviors that each object exhibits. We particularly aim to allow flexibility in the number of total and sequence-specific behaviors, and encourage objects to share similar subsets of the large set of possible behaviors.

One can represent the set of behaviors an object exhibits via an associated list of *features*. A standard featural representation for N objects, with a library of K features, employs an $N \times K$ binary matrix $F = \{f_{ik}\}$. Setting $f_{ik} = 1$ implies that object i exhibits feature k. Our desiderata motivate a Bayesian nonparametric approach based on the *beta process* [10, 22], allowing for infinitely many