

Hierarchical Dirichlet Processes for Tracking Maneuvering Targets

Emily B. Fox

Massachusetts Institute of Technology
77 Massachusetts Ave.
Cambridge, MA 02139
Email: ebfox@mit.edu

Erik B. Sudderth

UC Berkeley
527 Soda Hall
Berkeley, CA 94720-1776
Email: sudderth@eecs.berkeley.edu

Alan S. Willsky

Massachusetts Institute of Technology
77 Massachusetts Ave.
Cambridge, MA 02139
Email: willsky@mit.edu

Abstract— We consider the problem of state estimation for a dynamic system driven by unobserved, correlated inputs. We model these inputs via an uncertain set of temporally correlated dynamic models, where this uncertainty includes the number of modes, their associated statistics, and the rate of mode transitions. The dynamic system is formulated via two interacting graphs: a hidden Markov model (HMM) and a linear-Gaussian state space model. The HMM’s state space indexes system modes, while its outputs are the unobserved inputs to the linear dynamical system. This Markovian structure accounts for temporal persistence of input regimes, but avoids rigid assumptions about their detailed dynamics. Via a hierarchical Dirichlet process (HDP) prior, the complexity of our infinite state space robustly adapts to new observations. We present a learning algorithm and computational results that demonstrate the utility of the HDP for tracking, and show that it efficiently learns typical dynamics from noisy data.

Keywords: Tracking, hierarchical Dirichlet processes, hidden Markov model, Kalman filtering, estimation.

I. INTRODUCTION

We consider the problem of state estimation for a dynamic system which is driven by a set of unobserved, correlated inputs. This problem arises in many situations of surveillance and tracking. One such example which falls within this framework is monitoring a dynamic system with sensors whose regimes of operation change due to unobserved phenomena; this results in time-varying, correlated measurement noise statistics. Another example, which we will use throughout this paper, is that of tracking a maneuvering target. We specifically consider the scenario where the number and types of regimes or modes of operation are challenging to model accurately due to a lack of prior knowledge about the system.

The standard in state-of-the-art maneuvering target tracking algorithms is the interacting multiple model (IMM) method [1]–[3]. The IMM method relies on approximating the dynamics of the target by a finite set of distinct maneuver modes, which are known and well modeled *a priori*. The transitions between modes are assumed to follow a Markov chain with known transition probabilities. The IMM exhibits severe model sensitivity to the values of these parameters, which is problematic when these parameters cannot be fine-tuned due to model uncertainty. In addition, optimal tracking for multiple models methods requires examining an exponentially growing number of model sequences, where the rate

depends on the number of dynamic models [3], [4]. The IMM provides a method for merging hypotheses, which is a reasonable approximation in certain scenarios. A method for model set design by minimizing distribution, modal, or moment mismatch is presented in [5] for the case where the true mode set is known. We instead develop methods for learning models from data.

We examine the problem of tracking when the library of maneuver modes is unknown *a priori*, including the number of modes, their associated statistics, and the rate at which transitions occur. We formulate this dynamic system as a graphical model comprised of two interacting graphs: an infinite hidden Markov model (HMM) and a linear Gaussian state space model. The HMM’s state space indexes maneuver modes over time, while its outputs are the unobserved inputs to the linear dynamical system. As with the IMM, this Markovian structure accounts for the temporal persistence of target maneuvers and avoids rigid assumptions about their detailed dynamics; however, our formulation does not rely on prior knowledge of the number or types of maneuver modes. The phrase *infinite* HMM refers to the fact that the states of the HMM (i.e. maneuver modes) can take values in $\{1, \dots, \infty\}$. We assume that the system matrices and noise statistics of the state space dynamic model are known either through understood physical models or system identification techniques.

Estimation for this system is complicated by the fact that there is uncertainty in the maneuver mode, and that the number of these discrete hidden states is unknown. The unknown cardinality of the set of maneuver modes motivates our non-parametric Bayesian approach of placing a flexible, data-driven hierarchical Dirichlet process (HDP) prior on these hidden states. The Dirichlet process encourages simple models of target dynamics, but allows additional states to be created as new behaviors are observed, while the hierarchical structure accounts for temporal correlation of input modes. Dirichlet processes have recently been applied to the problem of state estimation in a dynamic system with an unknown noise process [6], where the noise is assumed to be uncorrelated. Our model captures the temporal correlation of the inputs by hierarchically coupling the input distributions at subsequent times. The HDP has been previously used to develop an infinite, discrete state HMM [7], [8]. For our model, however,

we do not have direct observations of the HMM output. In addition, the maneuver mode uncertainty implies that the target dynamics no longer follow a simple, linear-Gaussian model. Our goal, then, is to estimate this HMM output sequence from correlated observations of the control inputs' effect on the dynamic system, while simultaneously learning the set of HMM maneuver modes. We approach this joint learning-estimation problem by coupling a Markov chain Monte Carlo (MCMC) method with an embedded Kalman smoother.

In this paper we explore the utility of the HDP as a prior on the unknown number of unobserved input modes. When combined with an observation likelihood distribution, we obtain a HDP mixture model, which we will show provides a simple framework for efficiently learning maneuver modes from noisy data. In Sec. II we present background on nonparametric Bayesian methods relevant to this paper. In Sec. III we describe how the HDP is useful in this context. After describing the model, in Sec. IV we present an approach to learning the model parameters necessary for state estimation and in Sec. V provide results from an example problem. In Sec. VI, we briefly describe some methods for converting the batch processing of the MCMC sampler to an online tracker.

II. BACKGROUND

Nonparametric Bayesian methods avoid the often restrictive assumptions of parametric models by performing inference on infinite-dimensional spaces of functions or probability distributions. If suitably designed, these methods allow for efficient, data-driven posterior inference. In the following sections, we briefly describe two related nonparametric Bayesian methods: the Dirichlet and hierarchical Dirichlet processes.

A. Dirichlet Process Mixture Models

A Dirichlet process defines a distribution over probability measures on potentially infinite parameter spaces Θ . This stochastic process is uniquely defined by a concentration parameter, α , and base measure, H , on the parameter space Θ ; we denote it by $DP(\alpha, H)$. A tutorial on Dirichlet processes, including references to seminal work, can be found in [8], [9].

It can be shown that the Dirichlet process actually defines a distribution over discrete probability measures. With probability one, a random draw $G \sim DP(\alpha, H)$ is equivalent to,

$$G(\theta) = \sum_{k=1}^{\infty} \pi_k \delta_{\theta_k}(\theta) \quad \theta_k \sim H, \quad k = 1, 2, \dots \quad (1)$$

where the notation $\delta_{\theta_k}(\theta)$ indicates a Dirac delta at $\theta = \theta_k$. The mixture weights π_k define a random probability measure distributed according to the following stick breaking construction, denoted by GEM(α),

$$\begin{aligned} \beta_k &\sim \text{Beta}(1, \alpha) & k = 1, 2, \dots \\ \pi_k &= \beta_k \prod_{\ell=1}^{k-1} (1 - \beta_\ell) & k = 1, 2, \dots \end{aligned} \quad (2)$$

In effect, we have divided a unit-length stick by the mixture weights π_k defined over an infinite set of random parameters

θ_k . The k^{th} mixture weight is a random proportion β_k of the remaining stick after the previous $(k - 1)$ weights have been defined. From this construction we see that the concentration parameter α controls the relative proportion of the weights π , and thus controls model complexity in terms of the expected number of components.¹ When the Dirichlet process prior is combined with a likelihood distribution for the observations, as depicted by the graphs of Fig. 1(a)-(b), we have a Dirichlet process mixture model. We use either $\bar{\theta}_t$ or θ_{z_t} , with indicator variable z_t , to denote the parameter associated with the observation y_t .

Because random probability measures drawn from a Dirichlet process are discrete, there is a strictly positive probability that multiple observations will share a common parameter. In addition, there is a reinforcement property that makes it more likely to associate an observation with a parameter to which other observations have already been associated. This is described by the predictive distribution of a new assignment conditioned on all other previous assignments,

$$p(z_{T+1} = z | z_{1:T}, \alpha) = \frac{1}{\alpha + T} (\alpha \delta(z, \bar{K}) + \sum_{k=1}^K T_k \delta(z, k)), \quad (3)$$

where T is the total number of observations and T_k the number assigned to the k^{th} parameter. Here, we use the notation $\delta(z, k)$ to indicate the Kronecker delta. We see that with prior probability proportional to α the observation was generated from a new, previously unseen mode \bar{K} and with prior probability proportional to T_k , the number of assignments to mode k , the observation was generated by an existing mode k . Therefore, Dirichlet processes favor simpler models. It can be shown under mild conditions that if the data is generated by a finite mixture, then the Dirichlet process posterior is guaranteed to converge (in distribution) to that finite set of mixture parameters [11].

Dirichlet processes have been applied to the problem of state estimation in the presence of unknown noise statistics [6]. The noise is modelled as an infinite Gaussian mixture model, where the parameters θ_k represent the mean and covariance of a component of the mixture model. A graphical model of such a system is depicted in Fig. 2(a)-(b). However, this formulation does not capture temporal noise correlation. Therefore, this framework is not well suited for dynamic systems with unknown inputs that persist over time. In addition, the inference procedure described in [6] relies on sampling the explicit parameter value sequence $\{\bar{\theta}_t\}$ from the Dirichlet process prior, which may be impractical due to slow mixing rates [12].

B. Hierarchical Dirichlet Process Mixture Models

There are many scenarios in which groups of data are thought to be produced by related, yet unique, generative processes. For example, take a sensor network monitoring an environment where time-varying conditions may influence the quality of the data. Data collected under certain conditions

¹If the value of α is unknown, the model may be augmented with a gamma prior distribution on α , so that the parameter is learned from the data [10].

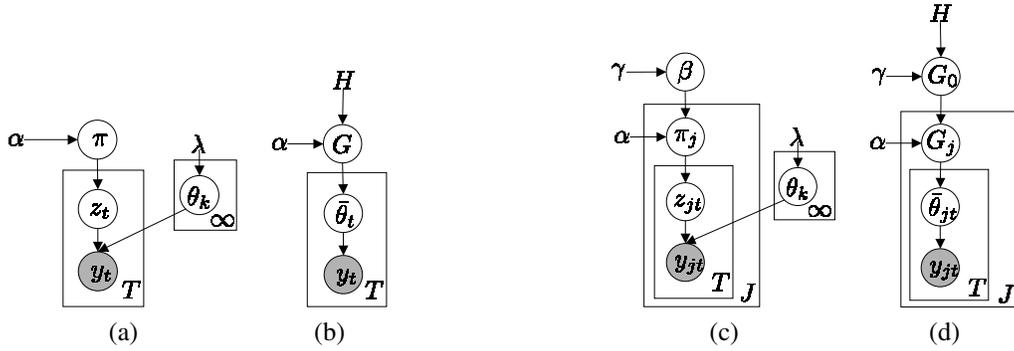


Fig. 1. Dirichlet process (left) and hierarchical Dirichlet process (right) mixture models represented by two graphs. (a) Indicator variable representation in which $\pi \sim \text{GEM}(\alpha)$, $\theta_k \sim H(\lambda)$, $z_t \sim \pi$, and $y_t \sim f(y|\theta_{z_t})$. (b) Alternative representation with $G \sim \text{DP}(\alpha, H)$, $\theta_t \sim G$, and $y_t \sim f(y|\theta_t)$. Here, the parameters θ_t are not necessarily unique. The mapping from (a) to (b) is $\bar{\theta}_t = \theta_{z_t}$. (c) Indicator variable representation in which $\beta \sim \text{GEM}(\gamma)$, $\pi_k \sim \text{DP}(\alpha, \beta)$, $\theta_k \sim H(\lambda)$, $z_{jt} \sim \pi_j$, and $y_{jt} \sim f(y|\theta_{z_{jt}})$. (d) Alternative representation with $G_0 \sim \text{DP}(\gamma, H)$, $G_j \sim \text{DP}(\alpha, G_0)$, $\theta_{jt} \sim G_j$, and $y_{jt} \sim f(y|\theta_{jt})$. Plate notation is used to compactly represent replicated variables of the graph [9].

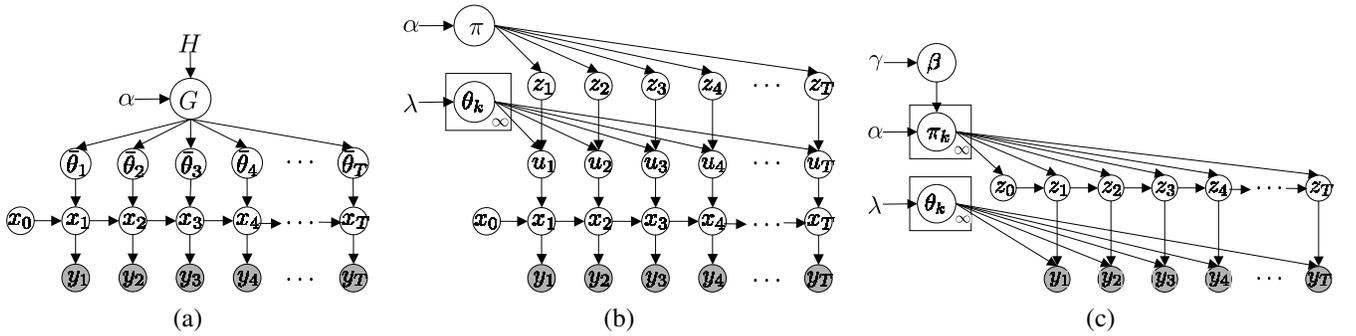


Fig. 2. Graphs of Dirichlet and hierarchical Dirichlet process priors being applied to dynamic systems. (a)-(b) Standard and indicator variable representation of a dynamic system with unknown noise distribution with a Dirichlet process prior [6]. (c) Indicator variable representation of a HDP prior on an HMM with unknown discrete state space cardinality [8]. The equivalent graphs in (a)-(b) model unobserved inputs driving a dynamic system, but do not capture temporal correlation in these inputs. The graph in (c) captures temporal correlation in modes, but does not capture continuous state space dynamics.

should be grouped and described by a similar, but disparate model from that of other data. In such scenarios we can take a hierarchical Bayesian approach and place a global Dirichlet process prior $\text{DP}(\alpha, G_0)$ on the parameter space Θ . We then draw group specific distributions $G_j \sim \text{DP}(\alpha, G_0)$, which will be discrete so that parameters are shared within the group. However, if the base measure G_0 is absolutely continuous with respect to Lebesgue measure, parameters will not be shared between groups. Only in the case where the base measure G_0 is discrete will there be a strictly positive probability of the group specific distributions having overlapping support (i.e. sharing parameters between groups.) To overcome this difficulty, the base measure G_0 should itself be a random measure distributed according to a Dirichlet process $\text{DP}(\gamma, H)$. This results in what is termed a hierarchical Dirichlet process [8].

Hierarchical Dirichlet processes can be applied as a prior on the state values of a HMM with unknown state space cardinality, as described in [7], [8]. Assume there are potentially countably infinitely many HMM state values. For each of these HMM states, there is a countably infinite transition density over the next HMM state. Let π_k be the transition density for HMM state k . Then, the model defines that $z_t \sim \pi_{z_{t-1}}$. That is to say, z_{t-1} indexes the group-specific transition density over the next HMM state value z_t . In terms of Fig. 1(c), the

observations assigned to group j are those with $z_{t-1} = j$ such that $z_t \sim \pi_j$. All groups share a common set of HMM states, but based on the previous state (i.e. group) there is a different probability density over current HMM states. The HMM state z_t determines which of the global parameters θ_k are used to generate the observation y_t . This model is termed a HDP-HMM, and is depicted by the graph in Fig. 2(c).

The HDP-HMM models correlation in time between the HMM states generating the observations. We can take the HMM states z_t to be the maneuver mode for our dynamic system. The outputs of this HMM are the control inputs u_t which drive the maneuvering target. However, we do not have direct observations of u_t , but rather correlated observations y_t of these values through the dynamic system. Therefore, our model has an extra hidden layer consisting of the state x_t of the dynamic system, as described in the following section.

III. FORMULATION

The graph of Fig. 3 represents a dynamic system with an unknown set of correlated maneuvers, or more generally, unobserved correlated inputs. Let $z_t \in \{1, 2, \dots\}$ be the unknown mode of the maneuver at time $t \in \mathbb{Z}_+$ and $u_t \in \mathbb{R}^d$ be an unobserved maneuver or control input at time t distributed according to $\mathcal{N}(\mu_{z_t}, \Sigma_{z_t})$, where $\theta_{z_t} = \{\mu_{z_t}, \Sigma_{z_t}\}$.

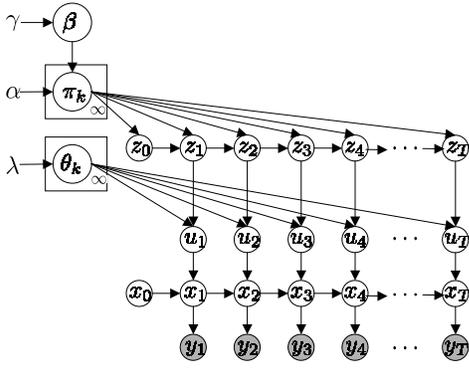


Fig. 3. Graphical model of the HDP-HMM-KF. The variables π_k define the transition densities for each of the HMM state values of the potentially countably infinite state space of z_t . We place a HDP prior on these states with base measure β and concentration parameters α and γ . The unobserved control inputs u_t that drive the linear-Gaussian dynamic system are drawn from Gaussian mixture component z_t , defined by parameters θ_{z_t} .

Therefore, θ_k defines the parameters for the distribution over the k^{th} maneuver mode, which we take to be Gaussian. We place a conjugate normal-inverse-Wishart ($\mathcal{NIW}(\kappa, \vartheta, \nu, \Delta)$ [9]) prior on θ_k . The system then evolves according to the following state space model with process noise $w_t \sim \mathcal{N}(0, Q)$ and measurement noise $v_t \sim \mathcal{N}(0, R)$.

$$\begin{aligned} x_t &= Ax_{t-1} + Bu_t(z_t) + w_t \\ y_t &= Cx_t + Du_t(z_t) + v_t. \end{aligned} \quad (4)$$

The latent maneuver states evolve according to,

$$\begin{aligned} \beta &\sim \text{GEM}(\gamma) \\ \pi_k &\sim \text{DP}(\alpha, \beta) \\ z_t &\sim \pi_{z_{t-1}}. \end{aligned} \quad (5)$$

In the case of modeling unobserved sensor regimes, we would take $D \neq 0$, though we might have $B = 0$. For the maneuvering target scenario we consider in this paper, we assume without loss of generality that $D = 0$, i.e. the unobserved input to the system solely drives the state through the system matrix B , but does not affect the sensors directly.

We place a HDP prior on the correlated maneuver modes z_t to account for the unknown dimension of the HMM. As in the HDP-HMM application, the transition densities π_k are each distributed as a Dirichlet process with concentration parameter α and base measure β , where β is distributed according to the stick breaking process with concentration parameter γ . By defining β to be a discrete probability measure, we ensure with high probability that a common set of future states are reachable from each preceding state. Because this model incorporates a linear dynamical model, which we later show leads to a Kalman filtering component in the learning algorithm, we refer to it as an *HDP-HMM-KF*.

IV. LEARNING

In order to learn the set of maneuver modes, we use a Markov chain Monte Carlo (MCMC) method, specifically Rao-Blackwellized Gibbs sampling. We briefly describe

MCMC theory and then present the Gibbs sampler for our model.

A. Markov Chain Monte Carlo

Markov chain Monte Carlo (MCMC) methods [13] are a class of algorithms used to sample from probability distributions that are challenging to sample from directly. A Markov chain is constructed whose stationary distribution is the desired density. At each iteration n , the state of the Markov chain is sampled conditioned on the previous sample at iteration $n-1$,

$$x^{(n)} \sim q(x|x^{(n-1)}) \quad n = 1, 2, \dots \quad (6)$$

After a certain “burn-in” period \bar{N} , the state evolution of this chain provides samples approximately drawn from the target distribution.

Gibbs sampling is a type of MCMC method that is well suited to state spaces with internal structure. Consider a state space with M states where we wish to draw samples from the joint distribution. With a Gibbs sampler, a sample $x_i^{(n)}$ is drawn from the conditional distribution given the previous set of samples for the other states. We iterate through every state using a specific or random node ordering τ ,

$$\begin{aligned} x &= (x_1, x_2, \dots, x_M) \\ \text{for } n &= 1 : N_{\text{iter}} \\ x_i^{(n)} &\sim p(x_i|x_{\setminus i}^{(n-1)}) \quad i = \tau(n) \\ x_j^{(n)} &= x_j^{(n-1)} \quad j \neq \tau(n) \\ \text{end} \end{aligned} \quad (7)$$

Here, we use the notation $x_{\setminus i} = \{x_j | j \neq i\}$.

In a directed graphical model, a node is conditionally independent of all other nodes given its Markov blanket, $p(x_i|x_{\setminus i}) = p(x_i|x_{MB(i)})$, where the Markov blanket consists of the node’s parents, co-parents, and children. Therefore, in the case of sparse graphs, the conditional density from which we are sampling is of much lower dimension than the joint.

B. Learning the HDP-HMM-KF

To perform state estimation, we marginalize over the infinite set of parameters θ and π and continuous hidden states $x_{1:T} \in \mathcal{X}$, and estimate the joint distribution $p(u, z, \beta | y, \gamma, \alpha, \lambda)$ using the auxiliary variable blocked Rao-Blackwellized Gibbs sampler of Algorithm 1.

Algorithm 1 u-z Blocked Gibbs Sampler

```

for n=1:N
  for t=1:T
     $(u_t^{(n)}, z_t^{(n)}) \sim p(u_t, z_t | u_{1:t-1}^{(n)}, u_{t+1:T}^{(n-1)}, z_{1:t-1}^{(n)}, z_{t+1:T}^{(n-1)}, \beta^{(n-1)})$ 
  end
   $\beta^{(n)} \sim p(\beta | z^{(n)})$ 
end

```

In Algorithm 1, we specifically examine a sequential node ordering for the Gibbs sampler to allow for simple updates, as will become clear in the following derivations. Although not explicitly written, the hyperparameters $\{\alpha, \gamma, \lambda\}$ are implied in all equations throughout this section.

We jointly sample (u_t, z_t) because a change in assignment of the mode z_t may induce a significant change in the distribution over the input u_t ; sampling these variables independently could result in different local modes which are very challenging to move between. In addition, u_t is often much lower-dimensional than x_t so by sampling the control input rather than the state, our estimates from the sampler have lower variance according to the Rao-Blackwell theorem [14]. Sampling β allows us to write the distributions in closed form.

Using the chain rule we can equivalently write,

$$p(u_t, z_t | z_{\setminus t}, u_{\setminus t}, y, \beta) = p(z_t | z_{\setminus t}, u_{\setminus t}, y, \beta) p(u_t | z_t, u_{\setminus t}, y, \beta). \quad (8)$$

The distribution in Eq.8 is a hybrid distribution: each discrete value of the indicator variable z_t corresponds to a different continuous distribution on the input u_t . We determine the weights of the multinomial distribution on z_t by integrating over the random variables of the joint distribution which are not present in the conditional distribution,

$$p(z_t = k | z_{\setminus t}, u_{\setminus t}, y, \beta) \propto \int_{\pi} \prod_j p(\pi_j | \beta) \prod_{\tau} p(z_{\tau} | \pi_{z_{\tau-1}}) d\pi \int_{u_t} \int_{\Theta_k} p(\theta_k) \prod_{\tau | z_{\tau} = k} p(u_{\tau} | \theta_k) d\theta_k \int_{\mathcal{X}} \prod_{\tau} p(x_{\tau} | x_{\tau-1}, u_{\tau}) p(y_{\tau} | x_{\tau}) dx du_t. \quad (9)$$

Similarly, we can write the conditional density of u_t for each candidate z_t as,

$$p(u_t | z_t = k, z_{\setminus t}, u_{\setminus t}, y, \beta) \propto \int_{\Theta_k} p(\theta_k) \prod_{\tau | z_{\tau} = k} p(u_{\tau} | \theta_k) d\theta_k \int_{\mathcal{X}} \prod_{\tau} p(x_{\tau} | x_{\tau-1}, u_{\tau}) p(y_{\tau} | x_{\tau}) dx. \quad (10)$$

The HDP prior provides a closed form for the marginalization of the unknown set of transition densities π . The density β determines the global probability of parameter θ_k and the expected weights in the group-dependent distributions π . If we take $z_{t-1} = j$, $z_t = k$, and $z_{t+1} = \ell$ and assume that $j \neq k$, i.e. a mode transition occurs at time t , then,

$$p(z_t = k | z_{\setminus t}, \beta) \propto \int_{\pi} \prod_j p(\pi_j | \beta, \alpha) \prod_{\tau} p(z_{\tau} | \pi_{z_{\tau-1}}) d\pi \propto (\alpha \beta_k + n_{jk}^{-t}) \left(\frac{\alpha \beta_{\ell} + n_{k\ell}^{-t}}{\alpha + n_k^{-t}} \right), \quad (11)$$

where we use the notation $n_{jk} = |\{\tau | z_{\tau} = k, z_{\tau-1} = j\}|$ to represent the number of observations assigned to maneuver mode k whose previous maneuver mode was j . The notation n_j represents the number of maneuvers from j to any other state (i.e. $n_j = \sum_k n_{jk}$) and n_{jk}^{-t} is the number of transitions

from maneuver j to maneuver k not counting the transition from z_{t-1} to z_t or from z_t to z_{t+1} . We consider each of the currently instantiated modes k in turn. We also consider the probability that z_t is assigned to a new, previously unseen mode \bar{K} , where the transition counts $n_{,k}$ and $n_{k,}$ are zero in this case. When $z_{t-1} = z_t$, a slightly more complicated expression is used (see Appendix). By placing a conjugate prior on the parameters θ_k , there is also a closed form for the marginal likelihood obtained by integrating over θ_k .

To relate the random control input taking values $u_t \in \mathcal{U}_t$ at time t to the observation sequence $y_{1:T}$ given the other inputs $u_{\setminus t}$, we integrate over the random state sequence taking values $x_{1:T} \in \mathcal{X}$. This is equivalent to a modified forward-backward filter combining the updated forward state estimate at time $t-1$ with the updated backward state estimate at time t using the dynamics in terms of u_t . We use the notation $(P_{t|t}^{-b}, P_{t|t}^{-b} x_{t|t}^b)$ and $(P_{t|t}^{-f}, P_{t|t}^{-f} x_{t|t}^f)$ as the inverse error covariance and information state for the backward and forward information filters, respectively. Because the sampler conditions on control inputs, the filter for this time-invariant system can be efficiently implemented by pre-computing the error covariances and then solely computing local Kalman updates at every time step.

In our formulation, it suffices to map the infinite dimensional distribution β of global parameter weights to the finite vector of weights associated with the currently instantiated maneuver modes and a new mode \bar{K} . Therefore, to sample $\beta \sim p(\beta | z)$ we use the auxiliary variable method of [8].

The details of the above integrations can be found in the Appendix and the resulting algorithm is summarized in Algorithm 2. Of note is that computational complexity is linear in the training sequence length, as well as the number of currently instantiated maneuver modes. There are many methods for using the samples provided by this algorithm for state estimation, as described in Sec.V.

Algorithm 2 HDP-HMM-KF learning

Initialization: **for** $t=1:T$ compute covariances $P_{t|t}^{-b}, P_{t|t}^{-f}$
for $n=1:N$
Backwards recursion: **for** $t=T-1:1$
compute information state $P_{t|t}^{-b} x_{t|t}^b$ given $\{u_{1:T}^{(n-1)}, y_{1:T}\}$
end
Forwards recursion: **for** $t=1:T$
 $z_t^{(n)} \sim p(z_t | u_{1:t-1}^{(n)}, u_{t+1:T}^{(n-1)}, z_{1:t-1}^{(n)}, z_{t+1:T}^{(n-1)}, \beta^{(n-1)}, P_{t|t}^{-b} x_{t|t}^b, P_{t-1|t-1}^{-f} x_{t-1|t-1}^f, P_{t|t}^{-b}, P_{t|t}^{-f})$
 $u_t^{(n)} \sim p(u_t | u_{1:t-1}^{(n)}, u_{t+1:T}^{(n-1)}, z_{1:t}^{(n)}, z_{t+1:T}^{(n-1)}, \beta^{(n-1)}, P_{t|t}^{-b} x_{t|t}^b, P_{t-1|t-1}^{-f} x_{t-1|t-1}^f, P_{t|t}^{-b}, P_{t|t}^{-f})$
compute $P_{t|t}^{-f} x_{t|t}^f$ given $\{u_t^{(n)}, y_t, P_{t-1|t-1}^{-f} x_{t-1|t-1}^f\}$
end
 $\beta^{(n)} \sim p(\beta | z^{(n)})$
end

V. RESULTS

The performance of the proposed HDP-HMM-KF tracking algorithm was compared to that of the IMM on a set of simulation data of a maneuvering target. We use the standard constant velocity (CV) and constant acceleration (CA) coordinate uncoupled maneuver models for the IMM with the state being x-direction position and velocity in the case of the CV model and x-direction position, velocity, and acceleration in the case of the CA model. The IMM state space equations are,

$$\begin{aligned} x_{t+1} &= A(z_t)x_t + w_t(z_t) \\ y_t &= C(z_t)x_t + v_t, \end{aligned} \quad (12)$$

where z_t indicates the mode at time t and the noise processes w_t and v_t are mutually independent zero-mean Gaussian noise processes with covariance $Q(z_t)$ and R , respectively. The system matrices for these two models are given by,

$$\begin{aligned} A_{CV} &= \begin{bmatrix} 1 & \Delta T \\ 0 & 1 \end{bmatrix} & C_{CV} &= [1 \quad 0] \\ Q_{CV} &= q_{CV} \begin{bmatrix} \frac{1}{3}\Delta T^3 & \frac{1}{2}\Delta T^2 \\ \frac{1}{2}\Delta T^2 & \Delta T \end{bmatrix} \end{aligned} \quad (13)$$

$$\begin{aligned} A_{CA} &= \begin{bmatrix} 1 & \Delta T & \frac{1}{2}\Delta T^2 \\ 0 & 1 & \Delta T \\ 0 & 0 & 1 \end{bmatrix} & C_{CA} &= [1 \quad 0 \quad 0] \\ Q_{CA} &= q_{CA} \begin{bmatrix} \frac{1}{20}\Delta T^5 & \frac{1}{8}\Delta T^4 & \frac{1}{6}\Delta T^3 \\ \frac{1}{8}\Delta T^4 & \frac{1}{3}\Delta T^3 & \frac{1}{2}\Delta T^2 \\ \frac{1}{6}\Delta T^3 & \frac{1}{2}\Delta T^2 & \Delta T \end{bmatrix} \end{aligned} \quad (14)$$

We assume that initially both IMM models are equally likely. The IMM also requires the definition of a transition matrix P defining the probability P_{ij} of transitioning to model j given current model i . Details of the CV-CA IMM implementation can be found in [3].

For the results presented in this section, we take,

$$P = \begin{bmatrix} p_{ii} & 1 - p_{ii} \\ 1 - p_{ii} & p_{ii} \end{bmatrix}, \quad (15)$$

since we have no prior bias towards the CA model versus the CV model. In Fig. 4(d) we plot the performance of the IMM as a function of p_{ii} . We see that the IMM exhibits strong model sensitivity to p_{ii} , while the HDP-HMM-KF does not depend on presetting this parameter. In the experiments of Fig. 4(a)-(b), we use $p_{ii} = 0.95$ in order to consider a ‘‘good’’ IMM.

The state space equations and associated system matrices for the HDP-HMM-KF model are as follows,²

$$\begin{aligned} x_t &= Ax_{t-1} + Bu_t(z_t) + w_t \\ y_t &= Cx_t + v_t \end{aligned} \quad (16)$$

$$A = \begin{bmatrix} 1 & \Delta T & \frac{1}{2}\Delta T^2 \\ 0 & 1 & \Delta T \\ 0 & 0 & 1 \end{bmatrix} \quad B = \begin{bmatrix} \frac{1}{2}\Delta T^2 \\ \Delta T \\ 0 \end{bmatrix} \quad C^T = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad (17)$$

²For this scenario, we take u_t to be the control input integrated into the system over the time window $t - 1$ to t . This parallels the IMM dynamics.

where u_t , v_t , and w_t are distributed as described in Sec.III.

With this formulation for the HDP-HMM-KF, we model the control inputs as realizations of a Markov jump-mean acceleration process with colored noise on a learned partition of the acceleration space. The mode parameter μ_{z_t} represents the mean of the process at time t and Σ_{z_t} allows for mode-specific variation of the control input realization u_t . If we learn a mode with a mean of zero, our model reduces to that of the constant acceleration model (i.e. zero-mean random walk on acceleration.) When, in addition, the learned covariance Σ_{z_t} and fixed process noise covariance Q are small, this model adequately describes a non-maneuvering target. By having the flexibility of learning modes with non-zero means, our model can account for fast changes in acceleration characteristic of highly maneuverable targets. While our formulation directly models the impact of the unobserved control inputs on the target dynamics, the IMM accounts for this phenomenon through a collection of different linear dynamic models. Note that although we consider 1D control inputs in this section, the formulation is general enough to learn coordinate coupled control inputs in multiple dimensions.

We use initial error covariance $P_0 = 100 * I_x$ and step size $\Delta T = 1$. For the CV-CA IMM, we take $q_{CA} = q_{CV} = 10$, while for the HDP-HMM-KF, we use $Q = 0.01 * I_x$ in order to encourage $u_{1:T}$ to capture the statistical properties of the input process. We place a $\mathcal{N}\mathcal{I}\mathcal{W}(0.001, 0, 50, I_u)$ prior on the parameters θ_k and a Gamma(1, 0.1) prior on the concentration parameters α and γ .

To compare the performance of the HDP-HMM-KF to that of the CV-CA IMM, we generated two types of simulated observations of position versus time. The first sequence is a noisy version of a modulated sinusoid starting at a random phase point with measurement noise covariance $R = 5 * 10^5 * I_y$. The underlying position sequence has continuous derivatives so that velocity and acceleration vary smoothly. The second sequence was a noisy step function generated from the Markov jump-mean model with $R = 5 * 10^9 * I_y$. The modes of the model were with means $\{-50, 0, 50\}$ and covariances $\{5, 1, 5\}$. The probability of self-transition was set to 0.99 while transitions to any other mode were equally likely. By considering both smooth and abrupt changes in acceleration, we show the flexibility of the proposed HDP-HMM-KF model.

In the following set of results we present two methods of using the HDP-HMM-KF for tracking. One method involves learning the control input sequence $u_{1:T}$ from the observation sequence $y_{1:T}$ using the MCMC samples from Algorithm 2 and then calculating Kalman smoothed estimates given the learned input sequence. For the results of this section, we simply learned $u_{1:T}$ by averaging 100 samples. The batch processing of data used by this method is impractical in many applications. Therefore, we also present an offline-training online-tracking HDP-HMM-KF approach to learning a set of dynamic models that can be used within the IMM framework. Specifically, we run the HDP-HMM-KF MCMC sampler on training data until it is well-mixed and then examine a set of 10 samples of $(u_{1:T}, z_{1:T})$. From each of these samples, we infer

a set of parameters θ_k and transition densities π_k . The resulting HDP-HMM-KF learned IMM models consist of CA dynamic models with different noise processes, both in terms of mean and covariance as determined by θ_k , and transition probabilities given by π_k . The results show the state estimates averaged over the 10 parallel HDP-HMM-KF learned IMM models, where the models were trained on random observation sequences generated from the two scenarios we aim to learn.

In Fig. 4(a), we show the track estimates of position versus time for the CV-CA IMM, HDP-HMM-KF learned IMM models, and HDP-HMM-KF smoother as well as the noisy observations. The associated average L_2 position errors versus time, averaged over 10 measurement realizations of the true target trajectory, are plotted in Fig. 4(b). These plots show the performance gain of HDP-HMM-KF methods over the CV-CA IMM. The HDP-HMM-KF learned IMM models have a 42% average decrease in total L_2 error in the modulated sinusoid case and 52% decrease in the step function case while the HDP-HMM-KF smoother has decreases of 78% and 75%.

One can analyze the complexity of the inferred HDP-HMM-KF model by looking at the number of maneuver modes to which a significant number of observations are assigned. We histogram those modes with more than 5% of the assignments over 1,000 Gibbs iterations in Fig. 4(c). When the true control inputs are drawn from a small finite set, as in the step function scenario, the HDP-HMM-KF describes the data with fewer model components than the more complicated modulated sinusoid scenario. These results emphasize the flexibility of the HDP-HMM-KF approach. A caveat, however, is that there may be several control inputs which produce similar target trajectories, and our HDP prior does not explicitly discourage fast transitions between modes. As a result, the sampler visits input sample paths that describe the observations but not the dynamic behavior we wish to capture. An area of future research is to consider priors that prefer slow switching or to learn a semi-Markov process formulation, which may better approximate the true target dynamics.

VI. ONLINE LEARNING

The batch processing of the MCMC sampler may be impractical and offline-training online-tracking infeasible for certain tracking applications. However, the developed algorithm could be converted into a recursive online implementation of joint learning-estimation using sequential Monte Carlo (SMC) methods, namely a Rao-Blackwellized particle filter [15]. Each of the particles represents a different sequence of latent control input and maneuver mode values from which a Kalman filter may be run to estimate the target state. Another approach for an online implementation is that of decayed MCMC filtering [16]. The decayed MCMC algorithm is similar to standard MCMC methods except instead of uniformly sampling the state variables the algorithm concentrates sampling activity to the recent past, since these states are the most relevant to the current state. Decayed MCMC is guaranteed to converge to the true marginal distribution given an appropriate decay function, and has provable rates of convergence.

VII. CONCLUSIONS

We have developed methods for learning models of unobserved, correlated inputs to dynamical systems. Our nonparametric approach adapts the hierarchical Dirichlet process to discover an appropriate set of input modes in a flexible, data-driven fashion. Using a Rao-Blackwellized Gibbs sampler, we may efficiently compute smoothed state estimates from noisy observation sequences. The parameters inferred by this sampler also lead to online IMM filters adapted to the structure of specific dynamical systems. Our maneuvering target tracking results demonstrate the effectiveness of this approach, and show significant gains over a fixed model set commonly used in tracking applications.

ACKNOWLEDGMENTS

This research was supported in part by MURIs funded through ARO Grant W911NF-06-1-0076 and through AFOSR Grant FA9550-06-1-0324, and by the MIT Lincoln Laboratory DMRI. E.B.F. was partially funded by an NDSEG fellowship.

APPENDIX

We analyze each of the integrals from Sec.IV-B in turn, beginning with the integral over π . Let β_k be the probability of global parameter θ_k and $\beta_{\bar{K}} = 1 - \sum_{k=1}^{K_t} \beta_k$ be the probability associated with a new, previously unseen maneuver mode \bar{K} , where K_t is the number of unique states instantiated by $z_{\setminus t}$. Then, $p(\pi_j | \beta, \alpha)$ is distributed as a Dirichlet distribution $\text{Dir}(\alpha\beta_1, \dots, \alpha\beta_{K_t}, \alpha\beta_{\bar{K}})$. In addition, $p(z | \pi)$ is distributed as a multinomial which is conjugate to the Dirichlet distribution. Using properties of these distributions and taking $z_{t-1} = j$ and $z_{t+1} = \ell$, one can show that,

$$\begin{aligned}
 p(z_t = k | z_{\setminus t}, \beta) &\propto \int_{\pi} \prod_j p(\pi_j | \beta, \alpha) \prod_{\tau} p(z_{\tau} | \pi_{z_{\tau-1}}) d\pi \\
 &\propto \begin{cases} (\alpha\beta_k + n_{jk}^{-t}) \left(\frac{\alpha\beta_{\ell} + n_{k\ell}^{-t} + \delta(j,k)\delta(k,\ell)}{\alpha + n_{k\cdot}^{-t} + \delta(j,k)} \right) & k = 1, \dots, K_t \\ \alpha\beta_{\bar{K}}\beta_{\ell} & k = \bar{K} \end{cases} \\
 &\triangleq N_k
 \end{aligned} \tag{18}$$

where n_{jk}^{-t} is defined as in Sec.IV-B.

Let us analyze the integral over θ_k . For $\theta_k = \{\mu_k, \Sigma_k\}$ defining the parameters of a Gaussian, we can place a conjugate normal-inverse-Wishart prior $\mathcal{N}\mathcal{IW}(\kappa, \vartheta, \nu, \Delta)$ so that the posterior distribution of θ_k given the control inputs currently assigned to the k^{th} mode is distributed as an updated normal-inverse-Wishart $\mathcal{N}\mathcal{IW}(\bar{\kappa}, \bar{\vartheta}, \bar{\nu}, \bar{\Delta})$. For these posterior parameter update formulas, see [9]. Marginalizing θ_k induces a multivariate Student-t predictive distribution for u_t , which can be approximated by a moment-matched Gaussian,

$$\begin{aligned}
 p(u_t | z_t = k, z_{\setminus t}, u_{\setminus t}) &\propto \int_{\Theta_k} p(\theta_k) \prod_{\tau | z_{\tau} = k} p(u_{\tau} | \theta_k) d\theta_k \\
 &\simeq \mathcal{N}\left(u_t; \bar{\vartheta}, \frac{(\bar{\kappa} + 1)\bar{\vartheta}}{\bar{\kappa}(\bar{\nu} - d - 1)} \bar{\Delta}\right) \\
 &\triangleq \mathcal{N}(u_t; \hat{\mu}_k, \hat{\Sigma}_k).
 \end{aligned} \tag{19}$$

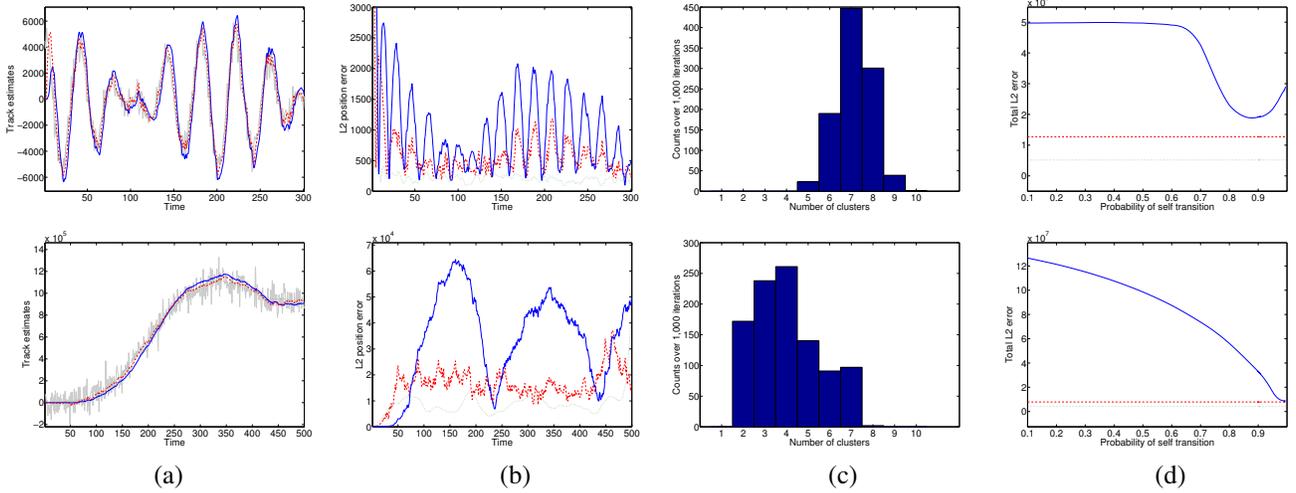


Fig. 4. Plots of (a) observation sequence (gray) with track estimates and (b) associated L2 error for the CV-CA IMM (blue,—), HDP-HMM-KF learned IMMs (red,—), and HDP-HMM-KF smoother (green,—) on a modulated sinusoid (top) and step function (bottom) control input. (c) Histogram of number of HDP-HMM-KF maneuver modes over 1,000 Gibbs iterations on the training sequence for the learned IMMs. (d) Total L2 error versus self transition probability p_{ii} depicting the model sensitivity of the IMM as compared to the HDP-HMM-KF learned IMMs or HDP-HMM-KF smoother, which do not depend on presetting this parameter.

We then marginalize the state $x_{1:T} \in \mathcal{X}$ for fixed $u_{\setminus t}$ to find the likelihood of our observations y as a function of u_t ,

$$\begin{aligned}
 p(y|u_t) &\propto \int_{\mathcal{X}} \prod_t p(x_t|x_{t-1}, u_{t-1}) \prod_t p(y_t|x_t) dx \\
 &\propto \int_{\mathcal{X}_{t-1}} \mathcal{N}(x_{t-1}; \hat{x}_{t-1|t-1}^f, P_{t-1|t-1}^f) \int_{\mathcal{X}_t} \mathcal{N}(x_t; Ax_{t-1} + Bu_t, Q) \\
 &\quad \mathcal{N}(x_t; \hat{x}_{t|t}^b, P_{t|t}^b) dx_t dx_{t-1} \\
 &= \mathcal{N}^{-1}(u_t; \phi_{u_t}, \Lambda_{u_t})
 \end{aligned} \tag{20}$$

where $\mathcal{N}^{-1}(\phi, \Lambda)$ represents the information form of a Gaussian $\mathcal{N}(\mu, P)$ with information parameters $\phi = P^{-1}\mu$ and $\Lambda = P^{-1}$. Using manipulations of Gaussian identities, we determine,

$$\begin{aligned}
 \Lambda_{u_t} &= B^T \Sigma_t^{-1} B - B^T \Sigma_t^{-1} A (A^T \Sigma_t^{-1} A + \Lambda_{t-1|t-1}^f)^{-1} A^T \Sigma_t^{-1} B \\
 \phi_{u_t} &= B^T Q^{-1} K_t^{-1} \theta_{t|t}^b - B^T \Sigma_t^{-1} A (A^T \Sigma_t^{-1} A + \Lambda_{t-1|t-1}^f)^{-1} \\
 &\quad (\theta_{t-1|t-1}^f + A^T Q^{-1} K_t^{-1} \theta_{t|t}^b)
 \end{aligned} \tag{21}$$

where $\Sigma_t = Q^{-1} + Q^{-1} K_t^{-1} Q^{-1}$ and $K_t = Q^{-1} + \Lambda_{t|t}^b$.

Joining the distributions in terms of u_t , we obtain,

$$\begin{aligned}
 p(u_t|z_t = k, z_{\setminus t}, u_{\setminus t}) p(y|u_t) \\
 &\propto C_k \mathcal{N}^{-1}(u_t; \hat{\Sigma}_k^{-1} \hat{\mu}_k + \phi_{u_t}, \hat{\Sigma}_k^{-1} + \Lambda_{u_t}) \\
 &\triangleq C_k \mathcal{N}^{-1}(u_t; \phi_{k, u_t}, \Lambda_{k, u_t}),
 \end{aligned} \tag{22}$$

where,

$$C_k = \frac{|\hat{\Sigma}_k^{-1}|^{\frac{1}{2}}}{|\Lambda_{k, u_t}|^{\frac{1}{2}}} \exp\left\{-\frac{1}{2}[\hat{\mu}_k^T \hat{\Sigma}_k^{-1} \hat{\mu}_k - \phi_{k, u_t}^T \Lambda_{k, u_t}^{-1} \phi_{k, u_t}]\right\} \tag{23}$$

We now have all of the components necessary to write the distributions of Eq.8 in closed form,

$$p(z_t = k|z_{\setminus t}, u_{\setminus t}, y, \beta) \propto N_k C_k \tag{24}$$

$$p(u_t|z_t = k, z_{\setminus t}, u_{\setminus t}, y, \beta) \propto \mathcal{N}^{-1}(u_t; \phi_{k, u_t}, \Lambda_{k, u_t}) \tag{25}$$

REFERENCES

- [1] H. Blom and Y. Bar-Shalom, "The interacting multiple model algorithm for systems with Markovian switching coefficients," *IEEE Transactions on Automatic Control*, vol. 33, no. 8, pp. 780–783, 1988.
- [2] E. Mazor, A. Averbuch, Y. Bar-Shalom, and J. Dayan, "Interacting multiple model methods in target tracking: A survey," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 34, no. 1, pp. 103–123, 1998.
- [3] X. Rong Li and V. Jilkov, "Survey of maneuvering target tracking, part v: Multiple-model methods," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1255–1321, 2005.
- [4] S. Oh and S. Sastry, "An efficient algorithm for tracking multiple maneuvering targets," in *Proc. IEEE Conference on Decision and Control, and the European Control Conference*, December 2005.
- [5] X. Rong Li, Z. Zhao, and X. Li, "General model-set design methods for multiple-model approach," *IEEE Transactions on Automatic Control*, vol. 50, no. 9, pp. 1260–1276, 2005.
- [6] F. Caron, M. Davy, A. Doucet, E. Duflos, and P. Vanheeghe, "Bayesian inference for dynamic models with Dirichlet process mixtures," in *Proc. International Conference on Information Fusion*, July 2006.
- [7] M. J. Beal, Z. Ghahramani, and C. E. Rasmussen, "The infinite hidden Markov model," in *NIPS 14*. MIT Press, 2002, pp. 577–584.
- [8] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei, "Hierarchical Dirichlet processes," *J. Amer. Stat. Assoc.*, vol. 101, no. 476, pp. 1566–1581, 2006.
- [9] E. Sudderth, "Graphical models for visual object recognition and tracking," PhD Thesis, MIT, Cambridge, MA, 2006.
- [10] M. Escobar and M. West, "Bayesian density estimation and inference using mixtures," *J Amer Stat Assoc*, vol. 90, no. 430, pp. 577–588, 1995.
- [11] H. Ishwaran and M. Zarepour, "Dirichlet prior sieves in finite normal mixtures," *Statistica Sinica*, vol. 12, pp. 941–963, 2002.
- [12] R. Neal, "Markov chain sampling methods for Dirichlet process mixture models," *Jour Comp Graph Stat*, vol. 9, no. 2, pp. 249–265, 2000.
- [13] W. Gilks, S. Richardson, and D. Spiegelhalter, Eds., *Markov Chain Monte Carlo in Practice*. Chapman & Hall, 1996.
- [14] G. Casella and C. Robert, "Rao-Blackwellisation of sampling schemes," *Biometrika*, vol. 83, no. 1, pp. 81–94, 1996.
- [15] A. Doucet, N. de Freitas, K. Murphy, and S. Russell, "Rao-Blackwellised particle filtering for dynamic Bayesian networks," in *Proc. 16th UAI*, August 2000.
- [16] B. Marthi, H. Pasula, S. Russell, and Y. Peres, "Decayed MCMC filtering," in *Proc. 18th UAI*, August 2002.