UNSUPERVISED TRACK CLASSIFICATION BASED ON HIERARCHICAL DIRICHLET PROCESSES

Jüri Sildam[1], Paolo Braca[1], Kevin LePage[1], Peter Willett[2]

[1]NATO STO CMRE, La Spezia, Italy, email: {sildam,braca,lepage}@cmre.nato.int
[2]University of Connecticut, Storrs CT, USA, email: willett@engr.uconn.edu

ABSTRACT

An unsupervised track classification approach is applied to sonar multistatic multitarget tracking. Appropriate discriminative and aggregative features are derived from beamformed and normalized matched-filtered data as recorded from a linear array towed behind an AUV. A clustering algorithm based on Hierarchical Dirichlet Processes is proposed for unsupervised classification of tracks. Overall improvement of target tracking is demonstrated via the Optimal Subpattern Assignment metric.

Index Terms— clustering, hierarchical Dirichlet processes, tracking, infinite HMM

1. INTRODUCTION

In the context of anti-submarine warfare (ASW) one of the central goals of autonomous underwater vehicles (AUVs) is detection and tracking of underwater targets. Traditionally, in an active sonar framework, formation of tracks is based on contacts obtained from a detection procedure. In littoral environments signal reflection and scattering occurs due to objects like rocks, soft bottom, underwater and surface vehicles. A tracker typically does not discriminate among different detection types: all contacts, even those that are clutter, can form a track.

A question we are trying to answer is whether we can improve on existing acoustical data partitioning aimed at echo-repeater detection and tracking, using extra data taken at beam-directions next to the beam-directions associated with the detections.

Existing signal detectors are constrained by relatively short duration of the source waveform (e.g. one second) that in turn constrains the number of independent data points available for target detection at independent beam directions. On the other hand, the possible non-stationarity of reference time-series constrains the size of time-window available for background clutter estimation. Such data scarcity causes significant errors in required data partitioning. While many works try to improve the solutions leading to data partitioning based on statistical analysis of time snippets obtained at the beam of detection, in this work we extract information from time series available at neighboring beams of detection and postpone the final partitioning until the track termination.

There are a number of reasons why information about clutter and target can be available at bearings next to the bearing of detection. For example while clutter could contribute significantly to signal energy at neighboring beams, so does imperfect spatial filtering of recorded signal or change of array shape not accounted for in beam-forming. Resolving the respective causes via direct modeling may be problematic under constraints of real-time processing. Instead, we try to estimate relative variability of sets of normalized data amplitudes around detections by introducing the Maximum Mean Discrepancy (MMD) test in the next section, and incorporate the results of this test into a data clustering model.

Another reason for using the MMD test is to perform a many-to-one mapping invariant to probability distribution function (pdf) of signal envelope amplitudes. Such a mapping is motivated by requirements of field sampling. For example, the probability function incorporating information about target aspect dependency is not properly sampled when detections of a target represent samples from just one rather than all target aspect angles.

Grouping and analyzing the distributions of MMD aggregative mapping along tracks is still relatively complex for real-time implementation. Discretizing the results of MMD mapping using a small dictionary, and estimating the entropy of the resulting, we end up associating each detection with a discrete scalar that has only a limited number of possible values. Such compression decreases the computational complexity of the consequent calculations and meets the most stringent constraints of underwater acoustic communications.

Assuming a discrete aggregative discriminative feature – construction of which is described in the next two sections – one can define for it a probabilistic generative model. And armed with such a generative model, distributions of underlying latent model parameters can be inferred, these then used to classify sets of detections, grouped via tracking or otherwise.

The paper is organized as follows. In sections 3-5 we give the description of two Bayesian generative models based on Dirichlet Processes (DP). To evaluate the impact of

unsupervised track classification on tracking performance, we use the Optimal Subpattern Assignment (OSPA) tracking metric [2], which is described in section 6. Finally we show the results of track classification on data collected in the framework of active bi-static measurements during *Generic Littoral Interoperable Network Technology* sea trials in 2011 (GLINT11).

2. CONTACT DETECTION IN BEAMNUMBER BI-STATIC TRAVEL TIME SPACE

In this work, detections are formed using matched-filtered beam-formed normalized data. A typical display of normalized data is shown in Figure 1. White diamonds represent contacts with the five highest signal-to-noise (SNR) values. For purposes of illustration, we surround one of them by five "reference" rectangles, one at the center, two in along- and two in across-beam directions. These rectangles correspond to the data windows, divided into lower and upper cells, used to form data sets tested for similarity using the statistical test described below. Note the anisotropy of the energy distribution around the contacts, and its variation from contact to contact.
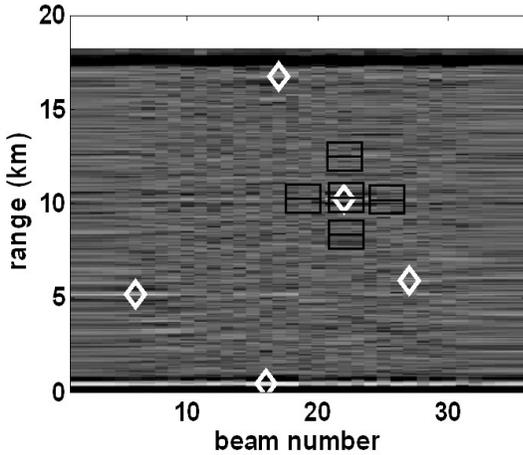

Figure 1: A typical display of normalized data scan with 4 detections (white diamonds), and 5 MMD test areas (rectangles where horizontal middle line separates two sets, here enlarged for better visibility) centered at one of the detections.

The Maximum Mean Discrepancy (MMD) test [3] is a non-parametric test defining a distance between a pair of probability measures embedded into reproducing kernel Hilbert spaces (RKHS). Such embedding allows comparison of the respective probability measures based on the distance between the respective embeddings without the explicit estimation of the probability measures [2, 4].

The MMD test is defined by a class of smooth functions F, defined in a RKHS

$$MMD[F, p, q] := \sup_{f \in F}(E_{x \sim p}[f(x)] - E_{y \sim q}[f(y)]) \quad (2.1)$$

where $E_{x \sim p}[f(x)]$ and $E_{y \sim q}[f(y)]$ denote the expectations under the distributions p and q, respectively. This class must be "rich enough" so that the outcome of this test is positive if and only if probabilities p and q underlying two data sets are equal [3]. The unit balls (i.e. the sets of all vectors with norm less than or equal to one) in characteristic reproducing kernel Hilbert spaces satisfy this property and can be used to produce the empirical estimate of the MMD test that converges quickly to its expectation with the increase of sample size. The computational cost of the MMD test is $O(m+n)^2$ time, where n and m correspond to the number of samples of the first and the second datasets, respectively.

We apply the Maximum Mean Discrepancy test (henceforth *dissimilarity*) on a pair of interleaved bearing-time cells in a time-bearing (TB) window (TBW) with a predefined non-dimensional range $\hat{N} = \lfloor f_s R/(2c) \rfloor$, (where f_s is the normalised data sampling frequency, R is the expected length of target, and c is sound speed) and number of beams ($\hat{M} = 3$) support. We apply this test to quantify dissimilarity of the interleaving *TB* cells.

An empirical biased estimate of MMD defined for the pair of TB cells $\hat{Z}$ and $\tilde{Z}$ in the TBW can be written as

$$d[\hat{Z}, \tilde{Z}] = [\frac{1}{M^2} \sum_{s,o}^{M} k(\hat{z}_s, \hat{z}_o) - \frac{2}{MN} \sum_{s,o}^{M,N} k(\hat{z}_s, \tilde{z}_o) + \frac{1}{N^2} \sum_{s,o}^{N} k(\tilde{z}_s, \tilde{z}_o)], \quad (2.2)$$

where $k(\hat{z}_s, \tilde{z}_o)$ is a kernel function, $\hat{z}_s$ and $\tilde{z}_o$ are vectors of the TBW cells, and N and M correspond to the numbers of vectors in the respective two adjacent cells of the TB window. We used the Gaussian radial basis function $k(\hat{z}_s, \tilde{z}_o) = \exp(-\frac{\|\hat{z}_s - \tilde{z}_o\|^2}{\sigma^2})$, where σ^2 is a scaling parameter that after some testing was set to 0.01.

Normalization of $\hat{z}_s$ data cells is given by:
$$\hat{z}_s = \frac{z_s}{\sum_i^M \sum_j^N \|z_{i,j}\|}$$

where M and N are the number of data points in the TBW in the range and bearing direction respectively i.e. $N = |\hat{z}_i| = |\tilde{z}_j|$ and $M = |\hat{Z}| = |\tilde{Z}|$. Note that $\hat{Z} = \{\hat{z}_1, \hat{z}_2, \hat{z}_3\}$, $\tilde{Z} = \{\tilde{z}_1, \tilde{z}_2, \tilde{z}_3\}$, and $\hat{z}_i$ and $\tilde{z}_j$ are time snippets.

Thus the difference in signal spread required for the classification is estimated at a low computational cost using only three grid points in each bearing and range direction. That is, by moving the TB window relatively to $\hat{z}_{i,j}$ in range space at a constant bearing b_j and in bearing space at a constant range r_i, we obtain two sets of dissimilarity indexes $d_\tau = \{d_{i-1,j} d_{i,j} d_{i+1,j}\}$ and $d_b = \{d_{i,j-1} d_{i,j} d_{i,j+1}\}$.

Each of the sets of the dissimilarity indexes of a single contact can be used to estimate a three-bin histogram:

$$p(d_r) = \frac{W_r}{M} \qquad (2.3)$$

where W_r is the number of d_r values (counted either in $\{d_{i-1,j}, d_{i,j}, d_{i+1,j}\}$ or in $\{d_{i,j-1}, d_{i,j}, d_{i,+1j}\}$) falling within the r-th bin, and M=3 is the overall number of values used in the histogram estimation. A probability mass function in the range and the bearing directions can be estimated from the respective normalised histograms such that

$$\sum_{r=1}^{M} p_\tau(d_r) = 1, \qquad (2.4)$$
$$\sum_{r=1}^{M} p_b(d_r) = 1. \qquad (2.5)$$

The entropy at constant bearing can be then estimated as:

$$h_\tau = -\sum_{r=1}^{M} p_\tau(d_r) \log(p_\tau(d_r)). \qquad (2.6)$$

Similarly the entropy at constant range can be then estimated as:

$$h_b = -\sum_{r=1}^{M} p_b(d_r) \log p_b(d_r). \qquad (2.7)$$

Finally, the entropy difference, which is the final feature used below for clustering and classification, is given as

$$\Delta h = h_\tau - h_b. \qquad (2.8)$$

3. THE SAMPLING STRATEGY

The coupled sample sets of high dimensional vectors of normalized data snippets with the support shown in Fig.1 (the upper and lower rectangles respectively) can be seen as random samples from a Dirichlet process. The statistical test performed on these samples, having multinomial outcome eq. (2.3)-(2.5), maps and aggregates high-dimensional samples respectively onto- and in 3D space.

The entropy of the probability mass function eq. (2.6)-(2.7) estimated on three bins using three samples is again a three-dimensional multinomial. Therefore the entropy difference estimated between across- and along-beam directions, is a seven dimensional multinomial. One can expect that entropy difference can be used to discriminate the reflecting objects that have different signal spread in across- and along-beam directions.

Having performed data discriminative aggregation, we may view a track as a set $\Delta h_{1:N} = \{\Delta h_1, \Delta h_2, \ldots, \Delta h_N\}$. Note that given the exchangeability assumption the order of detections can be ignored. We model the distribution of Δh as a mixture, where each component specifies a multinomial over Δh, which is shared among different tracks. We wish to find a probabilistic model that places significant probability not only over the observed but also over future unobserved tracks if they are "similar" to the tracks already observed.

4. CHOICE OF A TRACK CLASSIFICATION MODEL

4.1. Latent Dirichlet Allocation

A parametric approach to the track classification problem can be provided by Latent Dirichlet Allocation [5]. Initially this approach has been proposed for the probabilistic description of documents. We identify the multinomial features as words drawn from a vocabulary of seven words. "Documents" are the tracks that consist of N estimated multinomial features $\Delta h_{1:N}$. Finally "topics" are virtual reflecting objects (VRO). LDA can be adopted to describe track feature generation.

According to LDA, the proportions of the mixture model are drawn on track-specific basis from a Dirichlet distribution. Each detection (i.e. feature) is an independent draw from a mixture model conditioned on the mixing proportions.

Uncertainty in the number of mixture components can be addressed using the framework of hierarchical collections of Dirichlet Processes (HDP) [6], which can be seen as the nonparametric version of LDA where the well-known clustering property of DP is applied via placing nonparametric prior on the number of mixture components.

4.2. Bypassing the track clustering problem

If the LDA is used for track classification then track clustering problem needs to be addressed. Due to sharing of mixture components in the HDP framework [6], a different feature grouping approach, not just kinematic tracking, can be used for training of a track classifier.

We sort detections according to their SNR and form a single set of detections that includes L contacts from all experiments collected during GLINT11: $\Delta h_{1:L} = \{\Delta h_j\}, j = 1, \ldots, L$. Here j is a detection count such that $j = 1$ corresponds to the contact with the highest SNR of the scan obtained at time t_1, $j = 2$ corresponds to the highest SNR of the scan at time t_2, $j = n_1$ corresponds to the contact with the highest SNR at the last scan of experiment one, $j = n_1 + 1$ corresponds to the contact with the highest SNR of the first scan of experiment two and so on, until all contacts of all experiments with the highest SNR have been included. This cycle is repeated again until the contacts with the second, the third, up to the twentieth SNR have been included in the set, so that $L = \sum_{i=1}^{O} n_i$, where O is the number of field experiments, and n_i is the number of scans of the i[th] experiment.

Now the HDP Hidden Markov Model can be used.

5. THE INFINITE HIDDEN MARKOV MODEL

The HMM can be seen as a set of mixing models, one for each latent state, and in our case, one for each target class.

A DP mixture model can be used to learn a mixture model with a countably infinite number of mixture components. However, to accommodate a countably infinite number of mixture models one needs a mechanism to couple the respective DP models [6]. Such a mechanism is the hierarchical DP and the resulting HMM model is called HDP-HMM or infinite HMM.

Coupling across transitions can be obtained using hierarchical Bayesian formalism by introducing the Dirichlet priors with the shared parameters β_k, and a higher level prior γ, and a base measure H [6]: $\pi_k \sim Dirichlet(\alpha, \beta)$,

$\beta \sim GEM(\gamma)$, where π_k corresponds to the probability of state k, β is generated via stick-breaking construction (GEM denotes stick-breaking process [7]).

The hierarchy of DPs is given by the following equations [6]: $G_0 \sim DP(\gamma, H)$, $G_j \sim DP(\alpha, G_0)$. Stick breaking presentation gives $G_0 = \sum_{k'=1}^{\infty} \beta_{k'} \delta_{\theta_{k'}}$, $G_j = \sum_{k'=1}^{\infty} \pi_{kk'} \delta_{\theta_{k'}}$, where $\pi_{kk'}$ is the transmission (or mixing) parameter, and $\theta_{k'}$ is the distribution emission parameter. $\pi_k|\beta_k \sim DP(\alpha, \beta)$, $\phi_k|\theta_k \sim DP(H)$, $s_t|s_{t-1} \sim multinomial(\pi_{s_{t-1}})$, and $h_t|\phi_{s_t} \sim F(\phi_{s_t})$, $s = \{s_1, ..., s_n\}$ is the state sequence, $F(\phi_{s_t})$ denotes the distribution of Δh_t given the factor ϕ_{s_t}.

Inference of the HDP-HMM can be carried out by a Gibbs sampler, which converges to the true posterior. The implementation in its current form suffers from "slow mixing behavior" when applied to strongly correlated time series. An approach, coined by the authors the *beam sampler* [8], overcomes the problem of slow mixing.

The beam sampler introduces an auxiliary variable u_t such that conditioned on u the number of trajectories in the HMM is finite. As a result, such an approach adaptively truncates (i.e. only the paths that have large than u_t transition matrix values are used) the infinitely large transition matrix, and makes possible to use dynamic programming in the forward calculation. In the backward calculation the whole sequence is re-sampled.

The basic steps of beam-sampler calculation are given by the following block-scheme [8]

- Initialize hidden states and parameters
- While (enough samples)

  a) Sample $p(u|s)$: $u_t \sim uniform\left(0, \pi_{s_{t-1}, s_t}\right)$

  b) Sample the whole trajectory of s in two steps: first, forward filtering, second, backward sampling.

    a. Initialize DP, $p(s_o = 1) = 1$
    b. For each t=1,...,T

  $p(s_t|\Delta h_{1:t}, u_{1:t}) \propto$
  $p(\Delta h_t|s_t) \sum_{s_{t-1}: u \le \pi_{s_{t-1}, s_t}} p(s_{t-1}|\Delta h_{1:t-1}, u_{1:t-1})$

- Sample T $p(s_T|\Delta h_{1:T})$

    c. Sample t=T-1,...,1
  $p(s_t|\Delta h_{1:T}) \propto$
  $p(s_{t+1}|s_t)(s_t|\Delta h_{1:t})$

- Resample $\pi, \gamma, \alpha|s$ using dynamic programming.

Finally, tracks can be classified using:

$$p(s_t = k|\Delta h_{1:t}) = p(\Delta h_t|s_t) \sum_{s_{t-1}} p(s_t|s_{t-1}) p(s_{t-1}|\Delta h_{1:t-1})$$

$$\underset{k}{\arg\max} \, p(s_{1:T} = k|\Delta h_{1:T}) = \prod_{t=1}^{T} p(s_t = k|\Delta h_{1:t})$$

6. TRACKING PERFORMANCE METRIC

OSPA [2] compares two sets of tracks - the set of tracks of targets (usually given by the target's own navigation system, here called ground truth) X, and the set of tracks Y estimated by the tracking algorithm. OSPA does not require explicit labelling of the tracks.

OSPA combines minimum track-to-target distance with the scaled difference of cardinality of sets of tracks |Y| and targets |X|. Removal of a subset of tracks $\hat{Y}$ such that $|\tilde{Y}| < |Y|$, where $\tilde{Y} = Y \backslash \hat{Y}$, and $\hat{Y} \subset Y$ based on some track labelling approach, we are guaranteed to lower the cardinality of Y, but we are not guaranteed to reduce the OSPA metric. Therefore one can expect that the overall improvement of the OSPA metric evaluated based only on labelled tracks as opposed to the OSPA results using all tracks will indicate usefulness of the applied labelling.
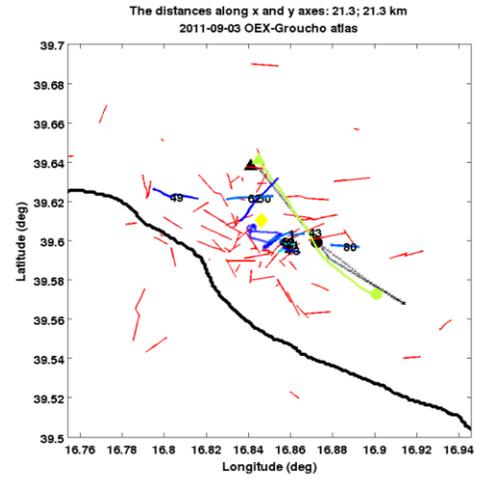

Figure 2: Sea trial 2011-09-03. Output of tracker and classifier. Legend: Coastline (black thick), AUV trajectory (blue), ER trajectory (black), rejected tracks (red), classified tracks (thick colored lines).

7. RESULTS

The tracking classification results are shown in figures 2-4. The tracks not classified (or rejected by the classifier) are shown in red. The tracks with a probability of classification exceeding 0.95 are numbered and indicated by thick colored lines. The colors span linearly on red-green-blue light scale, scaled by a number of classes i.e. the IHMM states. Black thick, and blue lines, and black line with crosses correspond to the coastline, the AUV track, and the echo-repeater (ER) track respectively. The yellow-filled diamond corresponds to the position of the static source. The only difference between figures 2 and 3 is that while in the first case all tracks are shown, in the second case the tracks that remained unclassified (the red tracks in fig. 2) have been removed. In

these figures, one can see that obviously the track colored in green corresponds to the ER.

Fig. 4 quantifies tracking performance using the OSPA metric. The top and middle graphs correspond to the two components of OSPA, minimum target-track distance and the number of tracks (offset by one) estimated at any given scan time respectively. The lines correspond to the OSPA metric components (the top and middle panels) and OSPA metric (bottom panel), estimated for all tracks (A, blue line), for all classified tracks (AC, black line), and for the class believed to correspond to the ER (ACS, red line). The respective lines are shown only when the respective tracks were present and their distance was below the maximum distance threshold on 6 km.
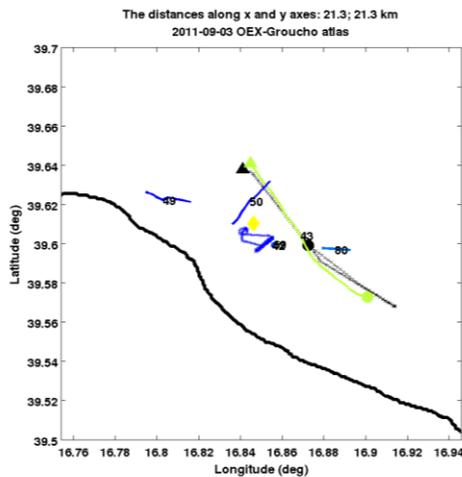

Figure 3: Sea trial 2011-09-03. Equal to Figure 2 but without the unclassified tracks. Legend: Coastline (black thick), AUV trajectory (blue), ER trajectory (black), classified tracks (thick colored lines).

8. DISCUSSION AND CONCLUSIONS

As opposed to the problem of kinematic tracking, track classification has increased complexity, which is required when one needs not just to improve tracking performance in terms of track duration and accuracy but also needs to know what kind of object has been tracked. In this work we have tried to mimic the feature generation process, used in the clustering models, via a feature construction approach. Namely, we have performed two steps of feature aggregation while still discriminating between targets of interest and other targets, which in our view provides better grounds for using the assumptions underpinning the unsupervised HDP clustering. The presented Bayesian approaches provide a consistent way to incorporate uncertainty into a hierarchy of parameters governing track classification. A degree of feature aggregation and discrimination should be a balance between the amount of non-redundant data and the requirements of target classification. Although the results presented in this work have already been extensively tested on field data

collected in a number of cruises, improvement may be achieved with online classification, which is an interesting future topic.
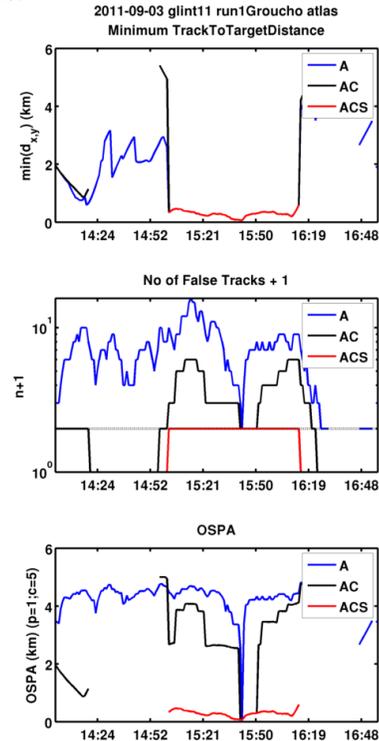

Figure 4: Sea trial 2011-09-03. Top plot: distributions of track to target minimum distance. Middle plot: number +1 of active tracks. Bottom plot: OSPA metric.

9. REFERENCES

[1] Sildam J, and F. Ehlers, "Supervised Track Classification In Support of AUV Decisions", UAM Conference Proceedings, 2011.
[2] Schuhmacher, D., Vo, B.-T. & Vo, B.-N., "A Consistent Metric for Performance Evaluation of Multi-Object Filters.", *IEEE Transactions on Signal Processing* **56** (8-1) , 3447-3457, 2008.
[3] Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B. and Smola, A. J. "A Kernel Two-Sample Test," *Journal of Machine Learning Research*, 723-773, 13 (2012).
[4] Sriperumbudur, B. K., Gretton, A., Fukumizu, K., Schölkopf, B., and Lanckriet, G. R. G., "Hilbert Space Embeddings and Metrics on Probability Measures," *Journal of Machine Learning Research* **11**, 1517-1561, 2010.
[5] Blei, D.M., Ng, A.Y. and Jordan, M.I., "Latent Dirichlet allocation," *The Journal of Machine Learning Research*, 993—1022, 3 (2003).
[6] Teh, Y. W., Jordan, M. I., Beal, M. J. and Blei, D. M., "Hierarchical Dirichlet Processes," *Journal of the American Statistical Association* **101** (476), 1566—1581, 2006.
[7] J. Sethuraman, "A constructive defnition of Dirichlet priors," *Statistica, Sinica,* 4, 639—650, 1994.
[8] Gael, J. V.; Saatci, Y.; Teh, Y. W. and Ghahramani, Z., "Beam sampling for the infinite hidden Markov model," *in* William W. Cohen; Andrew McCallum, and Sam T. Roweis, ed., 'ICML' , ACM, pp. 1088-1095, 2008.