
Truncation-free Hybrid Inference for DPMM

Arnim Bleier

Department of Computational Social Science
Leibniz Institute for the Social Sciences
Cologne, 50667 - Germany
arnim.bleier@gesis.org

Abstract

Dirichlet process mixture models (DPMM) are a cornerstone of Bayesian non-parametrics. While these models free from choosing the number of components a-priori, computationally attractive variational inference often reintroduces the need to do so, via a truncation on the variational distribution. In this paper we present a truncation-free hybrid inference for DPMM, combining the advantages of sampling-based MCMC and variational methods. The proposed hybridization enables more efficient variational updates, while increasing model complexity only if needed. We evaluate the properties of the hybrid updates and their empirical performance in single- as well as mixed-membership models. Our method is easy to implement and performs favorably compared to existing schemas.

1 Background

To begin with, consider a model for data $\mathcal{X} = \{x_1, x_2, \dots, x_N\}$ that is assumed to be generated by a mixture of simpler component models $F(\phi_k)$. Following the single-membership assumption of each data point $x_i \in \mathcal{X}$ being explained by a single component ϕ_{z_i} we have

$$\begin{aligned} \boldsymbol{\theta} &\sim \text{Dir}\left(\frac{\boldsymbol{\alpha}}{K}\right) \\ \phi_k &\sim H(\beta) && \forall k \in [1, K] \\ z_i &\sim \text{Cat}(\boldsymbol{\theta}) && \forall i \in [1, N] \\ x_i &\sim F(\phi_{z_i}) && \forall i \in [1, N], \end{aligned} \quad (1)$$

and arrive at the infinite-dimensional DPMM for $K \rightarrow \infty$. While collapsed Gibbs sampling (CGS) is commonly used to explore the unbounded latent space of this model, consider as an alternative the collapsed variational distribution

$$q(\mathbf{z}) = \prod_{i=1}^N \text{Cat}(z_i \mid \gamma_{i1}, \dots, \gamma_{iK+1}), \quad (2)$$

similar to partially collapsed approximations [4], however with $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ integrated out. The updates to optimize the variational parameter γ_i with regards to the true distribution p are then for each observation $i \in \{1, \dots, N\}$

$$\gamma_{ik} \propto \exp(\mathbb{E}[\log p(z_i = k \mid \mathbf{z}_{-i})] + \mathbb{E}[\log p(x_i \mid \mathbf{x}_{z=k}^{-i}, \beta)]). \quad (3)$$

While exact computations are hard, second-order Taylor expansions [4] or the use of zero-order information have been suggested for estimation [1]. Using only zero-order information the optimal settings of γ are

$$\gamma_{ik} \propto \begin{cases} \frac{n_k^{-i}}{n^{-i} + \alpha} p(x_i \mid \mathbf{x}_{z=k}^{-i}, \beta) & \text{if } k \leq K \\ \frac{\alpha}{n^{-i} + \alpha} p(x_i \mid \beta) & \text{if } k = K + 1, \end{cases} \quad (4)$$

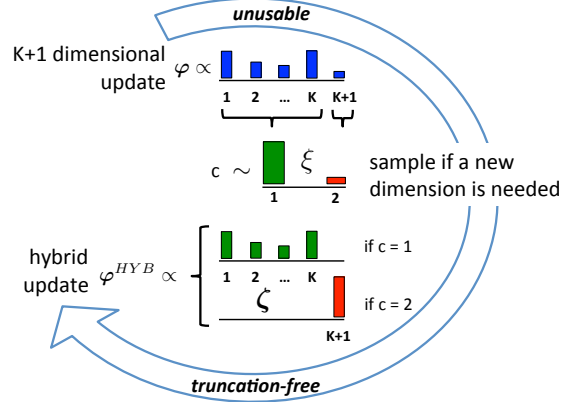


Figure 1: Illustration of hybrid updates

where we overload the notation $n_k^{-i} = \sum_{j=1}^{N-i} \gamma_{jk}$, as compared to the standard CGS, with the expected number of data points explained by the k^{th} component. Besides the simplifying assumptions made, these updates would place probability mass on each component $\gamma_{ik} \neq 0 \forall k \in \{1, \dots, K+1\}$ and introduce a new component in each step of the inference. A common way to address this problem is the use of fixed finite-dimensional variational approximations for DPMM [4]. Alternatively, Lin [5] as well as Wang et al. [7], amongst others, discuss methods for growing the truncation as part of the inference. Lin [5] introduces an additional parameter controlling the growth of the truncation. Wang et al. [7] estimate parameters for *locally collapsed variational inference* from traditional samples, losing valuable information in the updates.

We extend the idea of estimating variational parameters from samples with a method to construct the samples more efficiently. In the remainder of this paper, we start with the construction of the proposed truncation-free hybrid updates. After that we study the properties of these updates. We conclude with a short evaluation as well as a discussion of the current limitations and directions for future work.

2 Construction of Hybrid Updates

Our goal is to allow for a truncation-free variational posterior inference that keeps as much information as possible of the updates, while still being able to explore the unbounded latent space in a fashion similar to the Gibbs sampler. To reach that goal, we suggest to replace the unusable $K+1$ dimensional update of the variational parameter γ_i with a hybrid update. The suggested update has either the form of a K dimensional truncated variational parameter or a $K+1$ dimensional parameter instantiating a new component. Our hybridization depends only on local information and we will use the abbreviation $\varphi_k = \gamma_{ik}$ to refer, in this section, to the k^{th} component in the $K+1$ dimensional probability vector in Equation 4.

Let ξ be the two-dimensional parameter of the categorical distribution with the first dimension

$$\xi_1 \propto \sum_{k=1}^K \varphi_k \quad (5)$$

being proportional to the sum of the explanatory power of the first K components, and the second dimension

$$\xi_2 \propto \varphi_{K+1} \quad (6)$$

being proportional to what is explained by the yet uninstantiated components. Furthermore, let ζ_c , with $c \in \{1, 2\}$, be probability vectors representing, respectively, a truncated variational distribution and a Gibbs sample instantiating a new dimension in vector notation

$$\zeta_{1k} \propto \begin{cases} \varphi_k & \text{if } k \leq K \\ 0 & \text{if } k = K+1 \end{cases}, \quad \zeta_{2k} \propto \begin{cases} 0 & \text{if } k \leq K \\ 1 & \text{if } k = K+1. \end{cases} \quad (7)$$




Figure 2: Predictive performance in single-membership models for the Associated Press data set.

With this setup, we then sample a variable

$$c \sim \text{Cat}(\xi) \quad (8)$$

from ξ to indicate whether the truncated variational update or the probability vector instantiating a new component is selected

$$\varphi_k^{HYB} \propto \zeta_{ck} . \quad (9)$$

The probability vector φ^{HYB} is our hybrid update. The hybrid update replaces the original unusable $K + 1$ dimensional variational update $\gamma_i \leftarrow \varphi^{HYB}$, using most of the time the efficient truncated K dimensional variational distribution without introducing a new dimension in the update step, while introducing a new $K + 1^{th}$ component only if needed, similar to a Gibbs sampler. For a graphical illustration of the construction of φ^{HYB} , see Figure 1.

3 Properties of Hybrid Updates

This construction of the hybrid updates has a number of favorable properties. By definition of the parameter ξ (Equations 5, 6), the event of sampling the second category $c = 2$ has the same expectation and variance as sampling a new component in the Gibbs sampler:

$$\begin{aligned} \mathbb{E}[\mathbb{1}[c = 2]] &= \mathbb{E}[\mathbb{1}[z_i = K + 1]] \\ \text{Var}[\mathbb{1}[c = 2]] &= \text{Var}[\mathbb{1}[z_i = K + 1]] . \end{aligned} \quad (10)$$

Together with Equation 7 and 9, we see that this carries over to the hybrid update itself

$$\begin{aligned} \mathbb{E}[\mathbb{1}[\varphi_{K+1}^{HYB} = 1]] &= \mathbb{E}[\mathbb{1}[z_i = K + 1]] \\ \text{Var}[\mathbb{1}[\varphi_{K+1}^{HYB} = 1]] &= \text{Var}[\mathbb{1}[z_i = K + 1]] , \end{aligned} \quad (11)$$

making it possible to introduce new components like in a Gibbs sampler. Even more, the preservation of expectation is not limited to the $K + 1^{th}$ dimension itself, but our hybrid updates φ^{HYB} preserve the expectation, with regards to φ_k , over all $K + 1$ dimensions

$$\mathbb{E}[\varphi_k^{HYB}] = \mathbb{E}[\varphi_k] \quad \forall k \in [1, \dots, K + 1] . \quad (12)$$

Note that this is, with the exception of *locally collapsed variational inference* [7], generally not the case for variational updates in non-parametric models.

Moreover, the sum of the explanatory power of the existing K dimensions will exceed the explanatory power of introducing a new dimension $\mathbb{E}[\xi_1] > \mathbb{E}[\xi_2]$ for most data points, supporting the use of the more informative variational distribution ζ_1 in most of the updates. Finally, the computations necessary for the hybrid update φ^{HYB} are easy to implement and readily available by almost no additional computational costs from the normalization terms of φ .




Figure 3: Predictive performance in HDP-LDA models for the New York Times data set.

4 Experiments and Discussion

This section concludes our paper with an early empirical evaluation of the proposed hybrid updates. For the evaluation we used two text data sets: (1) The *Associated Press* corpus consisting of 2,250 documents, where we used a vocabulary of 10,932 distinct terms occurring over a total of 398k tokens. (2) The larger *New York Times* corpus consisting of 1,8 million articles, from which we extracted 153 million tokens using a vocabulary of 77,928 distinct terms.

The *Associated Press* corpus was used for evaluating the proposed updates in the single-membership model together with a Dirichlet-Multinomial data model for the documents. In the experiments, we held out 20% of the documents as a test set \mathcal{X}^{test} and batch-trained on the remaining \mathcal{X}^{train} documents. Next, we split each test document $x_i \in \mathcal{X}^{test}$ in two parts $x_i = (x_i^a, x_i^b)$, x_i^a consisting of 70% of the document for estimating the indicator variable and computed the perplexity of the remaining 30% x_i^b . We then compared the perplexity versus the number of iterations. Figure 2 displays the performance for hybrid updates in the zero-order collapsed variational setting (HCSVB0), CGS and truncated ($T = 40$) stick-breaking mean field variational Bayes (TSBVB) [4].

The *New York Times* corpus was used for evaluating the hybrid updates in mixed-membership HDP-LDA models, with test-train splits similar to above. However, for this larger dataset we resorted to collapsed stochastic inference using minibatches of 60 documents. We compared our method to the truncation-free *locally collapsed variational inference* (SCTFVB) [7] and the finite dimensional stochastic collapsed variational Bayesian inference for LDA (SCVB0) [3] using 40, 100 and 300 topics. We employed our hybrid updates in a setting similar to SCTFVB, however using a lower-bound approximation for the estimation of the stick-breaking weights [6, 2]. We used the same parameterizations for the update schedules in all inference schemas. Figure 3 displays the results, as a function of the number of documents processed (left) and wall-clock time in seconds (right), for the same runs. In the figure, HCSVB0 denotes the results with our hybrid updates. PCSVB0 denotes a finite dimensional variational approximation otherwise identical to HCSVB0, but truncated at 300 topics.

In this work, we desired to use the advantages of MCMC schemas and variational schemas combined in a single inference schema for Bayesian non-parametric models. We answered to this demand by presenting a novel type of hybridization that efficiently uses the full variational distribution while sampling for the introduction of new components. The proposed method is easy to implement and measurably improves the predictive performance over state of the art methods for single- as well as mixed-membership models at little additional computational cost. The current limitations of the presented work are two-fold. While we have established some favorable properties of the updates and found predictive performance improvements, we rely on approximations and have only limited theoretical arguments legitimizing our approach. The other limitation of this work is its scope. Next to a more thorough experimental evaluation and further formalization, an adaption of the hybrid updates to Wang et al. [8]’s Chinese restaurant process based variational inference for the HDP could potentially be a promising direction for future work.

References

- [1] Arthur Asuncion, Max Welling, Padhraic Smyth, and Yee Whye Teh. On smoothing and inference for topic models. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, 2009.
- [2] Arnim Bleier. Practical collapsed stochastic variational inference for the hdp. In *NIPS Workshop on Topic Models: Computation, Application, and Evaluation*, 2013.
- [3] James Foulds, Levi Boyles, Christopher DuBois, Padhraic Smyth, and Max Welling. Stochastic collapsed variational bayesian inference for latent dirichlet allocation. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2013.
- [4] Kenichi Kurihara, Max Welling, and Yee Whye Teh. Collapsed variational dirichlet process mixture models. In *IJCAI*, volume 7, 2007.
- [5] Dahua Lin. Online learning of nonparametric mixture models via sequential variational approximation. In *Advances in Neural Information Processing Systems*, 2013.
- [6] Issei Sato, Kenichi Kurihara, and Hiroshi Nakagawa. Practical collapsed variational bayes inference for hierarchical dirichlet process. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2012.
- [7] Chong Wang and David M Blei. Truncation-free online variational inference for bayesian non-parametric models. In *Advances in neural information processing systems*, 2012.
- [8] Chong Wang, John William Paisley, and David M Blei. Online variational inference for the hierarchical dirichlet process. In *AISTATS*, 2011.