



Slow mixing for Latent Dirichlet Allocation

Johan Jonasson*

Chalmers University of Technology, Sweden
University of Gothenburg, Sweden



ARTICLE INFO

Article history:

Received 13 January 2017

Accepted 15 May 2017

Available online 24 May 2017

MSC 2010:

60J10

Keywords:

Mixing time

MCMC

Gibbs sampler

Topic model

ABSTRACT

Markov chain Monte Carlo (MCMC) algorithms are ubiquitous in probability theory in general and in machine learning in particular. A Markov chain is devised so that its stationary distribution is some probability distribution of interest. Then one samples from the given distribution by running the Markov chain for a “long time” until it appears to be stationary and then collects the sample. However these chains are often very complex and there are no theoretical guarantees that stationarity is actually reached. In this paper we study the Gibbs sampler of the posterior distribution of a very simple case of Latent Dirichlet Allocation, an attractive Bayesian unsupervised learning model for text generation and text classification. It turns out that in some situations, the mixing time of the Gibbs sampler is exponential in the length of documents and so it is practically impossible to properly sample from the posterior when documents are sufficiently long.

© 2017 Published by Elsevier B.V.

1. Introduction

Markov chain Monte Carlo (MCMC) is a powerful tool for sampling from a given probability distribution on a very large state space, where direct sampling is difficult, in part because of the size of the state space and in part because of normalizing constants that are difficult to compute.

In machine learning in particular, MCMC algorithms are extremely common for sampling from posterior distributions of Bayesian probabilistic models. The posterior distribution given observed data is then difficult to sample from for the reasons above. One then designs an (irreducible aperiodic) Markov chain whose stationary distribution is precisely the targeted posterior. This is usually fairly easy since the posterior is usually easy to compute up to the normalizing constant (the denominator in Bayes formula). One usually uses Gibbs sampling or the related Metropolis–Hastings algorithm.

Gibbs sampling in these situations can be described as follows. Our state space is a finite set of random variables $X = \{X_a\}_{a \in A}$, where $X_a \in T$ for some measurable space T , so that $X \in T^A$ and the targeted distribution is a given probability measure \mathbb{P} on T^A . In order to sample from \mathbb{P} one starts a Markov chain on T^A whose updates are given by first choosing an index $a \in A$ at random and then choosing a new value of X_a according to the conditional distribution of X_a given all X_b , $b \in A \setminus \{a\}$. Under mild conditions, this chain converges in distribution to \mathbb{P} . The Markov chain is then run for a “long time” and then a sample, hopefully approximately from \mathbb{P} , is collected. A key question here is for how long the chain actually has to be run, in order for the distribution after that time to be a good approximation of \mathbb{P} . Since A is usually large, the number of steps needed should at least be no more than polynomial in the size of A for Gibbs sampling to be feasible. In almost all practical cases, the structure of the sample space and the probability measure \mathbb{P} is so complex that is virtually impossible to make a rigorous analysis of the mixing rate. However it may be possible to consider some very simplified special cases. In

* Correspondence to: Chalmers University of Technology, Sweden.

E-mail address: jonasson@chalmers.se.

this paper, we will analyze a special case of Latent Dirichlet allocation, henceforth LDA for short, and demonstrate for such a simple special case that mixing can indeed be a problem.

LDA is model used to classify documents according to their topics introduced by Blei et al. (2003). We have a large corpus of documents and we want to determine for each word in each document which topic it comes from. Knowing this we can also classify the documents according to the proportion of words of the different topics it contains. The setup in LDA is that one has a fixed number D of documents of lengths N_d a fixed set of topics t_1, t_2, \dots, t_s and a fixed set of words w_1, w_2, \dots, w_v . These are specified in advance. The number of topics is usually not large, whereas the number of words is. Next, for each document $d = 1, \dots, D$, a multinomial distribution $\theta_d = (\theta_d(1), \dots, \theta_d(s))$ over topics is chosen according to a Dirichlet prior with a known parameter $\alpha = (\alpha_1, \dots, \alpha_s)$. For each topic t_i a multinomial distribution $\phi_i = (\phi_i(1), \dots, \phi_i(v))$ according to a Dirichlet prior with parameter $\beta = (\beta_1, \dots, \beta_v)$ independently of each other and of the θ_d 's. Given these, the corpus is then generated by for each position $p = 1, \dots, N_d$ in each document d picking a topic $z_{d,p}$ according to θ_d and then picking the word at that position according to $\phi_{z_{d,p}}$, doing this independently for all positions. Note that the LDA model is a so called “bag of words” model, i.e. it is invariant under permutations within each document.

In this paper, we will consider the very simplified case with $D = s = v = 2, N_1 = N_2 = m$ and $\alpha = \beta = (1, 1)$ and study the mixing time asymptotics as $m \rightarrow \infty$. To simplify the notation, denote the two topics by A and B and the two words by 1 and 2. Define n_{ij} as the number of occurrences of the word j in document $i, i, j = 1, 2$ and write $n_i = n_{i1} + n_{i2}$ (which by assumption equals m) and $n_j = n_{1j} + n_{2j}$ and $n_{..} = \sum_{i,j} n_{ij} = 2m$. We consider the mixing time for Gibbs sampling of the posterior in a seemingly typical case, namely that the number of 1:s in the first document is $3m/10$ and in the second document $6m/10$. Our result is the following.

Theorem 1.1. Consider the case $n_{11} = 3m/10$ and $n_{21} = 6m/10$. Then there exists a $\lambda > 0$ such that for each $0 < \kappa < 3/4$,

$$\tau_{\text{mix}}(\kappa) > e^{\lambda n}.$$

Remark. The point of this paper is to demonstrate that the mixing time issue can be a real problem for Bayesian inference in machine learning in general and not for LDA in particular. There are also other methods for estimation that seem to work well for LDA, in particular variational inference, see Hoffman et al. (2013). Furthermore, experimental results seem to be fairly well in line with what one would expect from a topic classification algorithm. There are various extensions of LDA. For a model that seems to perform well, see Andrews and Vigliocco (2010), where jumps between topics follow a Markov chain. For an overview of probabilistic topic models, see Blei (2011).

Before moving on to the proof of Theorem 1.1, we formally introduce the concept of mixing time. Let $\{X_t\}_{t=0}^\infty$ be a discrete time Markov chain on the finite state space S and for $s \in S$, let \mathbb{P}_s be the underlying probability measure under $X_0 = s$.

Definition 1.2. Let μ and ν be two probability measures on S . The total variation distance between μ and ν is given by

$$\|\mu - \nu\|_{TV} = \max_{A \subset S} (\mu(A) - \nu(A)) = \frac{1}{2} \sum_{s \in S} |\mu(s) - \nu(s)|.$$

Definition 1.3. For each $\kappa \in (0, 1)$, the κ -mixing time of $\{X_t\}$ is given by

$$\tau_{\text{mix}}(\kappa) = \max_{s \in S} \min\{t : \|\mathbb{P}_s(X_t \in \cdot) - \pi\|_{TV} \leq \kappa\}.$$

The essence of Theorem 1.1 is that even after an exponentially long time, the distribution of the state of the Gibbs sampler is concentrated on a set whose probability mass according to the targeted posterior is at most $1/4$.

2. Proof of Theorem 1.1

The following two combinatorial lemmas will be needed.

Lemma 2.1. Let X be a standard uniform random variable and $0 \leq k \leq n$. Then

$$\mathbb{E}[X^k(1 - X)^{n-k}] = \frac{1}{(n + 1) \binom{n}{k}}.$$

Proof. First recall that for any $m = 1, 2, \dots$,

$$\mathbb{E}[X^m] = \frac{1}{m + 1}.$$

Next observe that the result holds true for any n with $k = n$. We want to prove that the claimed result holds for $(n, k) = (p, r)$ for some arbitrary $0 \leq r \leq p$. We do this by induction. We may then assume that the result is true with $(n, k) = (m, l)$ for

Download English Version:

<https://daneshyari.com/en/article/5129716>

Download Persian Version:

<https://daneshyari.com/article/5129716>

[Daneshyari.com](https://daneshyari.com)