# Slow mixing for Latent Dirichlet Allocation

Johan Jonasson\*†

January 11, 2017

#### 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.

AMS Subject classification : 60J10 Key words and phrases: mixing time, MCMC, Gibbs sampler, topic model Short title: Slow mixing for LDA

### **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.

<sup>\*</sup>Chalmers University of Technology and University of Gothenburg

<sup>&</sup>lt;sup>†</sup>Research supported by the Knut and Alice Wallenberg Foundation

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 this paper, we will analyse 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. 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  $t_1, t_2, \ldots, t_s$  and a fixed set of words  $w_1, w_2, \ldots, w_v$ . These are specified in advance. The number of topics is usually not large, whereas the number of words is. Next, for each document  $d = 1, \ldots, D$ , a multinomial distribution  $\theta_d = (\theta_d(1), \ldots, \theta_d(s))$  over topics is chosen according to a Dirichlet prior with a known parameter  $\alpha = (\alpha_1, \ldots, \alpha_s)$ . For each topic  $t_i$  a multinomial distribution  $\phi_i = (\phi_i(1), \ldots, \phi_i(v))$  according to a Dirichlet prior with parameter  $\beta = (\beta_1, \ldots, \beta_v)$  independently of each other and of the  $\theta_d$ :s. Given these, the corpus is then generated by for each position  $p = 1, ..., N_d$  in each document d picking a topic  $z_{d,p}$  according to  $\theta_d$  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  $m \to \infty$ . 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  $0 < \kappa < 3/4$ ,

$$au_{\min}(\kappa) > e^{\lambda n}$$

**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. Furthermore, experimental results seem to be fairly well in line with what one would expect from a topic classification algorithm.

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  $\mu$  and  $\nu$  be two probability measures on S. The the total variation distance between  $\mu$  and  $\nu$  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  $\kappa$ -mixing time of  $\{X_t\}$  is given by

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

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  $0 \le k \le n$ . Then

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

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

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

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  $0 \le r \le p$ . We do this by induction. We may then assume that the result is true with (n, k) =(m, l) for any m and l with either m < p or m = p and l > r. Then, taking  $Y = X^r (1 - X)^{p-r-1}$ ,

$$\begin{split} \mathbb{E}[X^{r}(1-X)^{p-r}] &= \mathbb{E}[Y(1-X)] \\ &= \mathbb{E}[Y] - \mathbb{E}[YX] \\ &= \mathbb{E}[X^{r}(1-X)^{p-r-1}] - \mathbb{E}[X^{r+1}(1-X)^{p-r}] \\ &= \frac{1}{p\binom{p-1}{r}} - \frac{1}{(p+1)\binom{p}{r+1}} \\ &= \frac{1}{(p+1)\binom{p}{r}}, \end{split}$$

where the fourth equality follows form the induction hypothesis and the last equality holds because

$$\frac{1}{(p+1)\binom{p}{r}} + \frac{1}{(p+1)\binom{p}{r+1}} = \frac{\binom{p+1}{r+1}}{(p+1)\binom{p}{r}\binom{p}{r+1}} = \frac{\frac{1}{r+1}\binom{p}{r}}{\binom{p}{r+1}} = \frac{1}{p\binom{p-1}{r}}.$$

Lemma 2.2 For any nonnegative integers a, b, c and d,

$$\binom{a+b+c+d}{a+c}\binom{a+c}{a}\binom{b+d}{d} = \binom{a+b+c+d}{a+b}\binom{a+b}{a}\binom{c+d}{d}.$$

*Proof.* Both sides of the equality are equal to the multinomial coefficient

$$\binom{a+b+c+d}{a,b,c,d}.$$

Let  $\mathbf{W} = (w_{11}, w_{12}, \dots, w_{1m}, w_{21}, \dots, w_{2m})$  be the words in our corpus, let  $\mathbf{Z} = (z_{11}, z_{12}, \dots, z_{2m})$  be the latent topics and  $\mathbf{Z}_d$  be the latent topics in document d. Let also  $\theta_i$  be the probability that  $z_{d1} = A$ , d = 1, 2 and let  $\phi_t$  be the conditional probability that  $w_{dj} = 1$  given that  $z_{dj} = t$ , t = A, B. In the case under study, these four quantities are all independent standard uniform random variables. We begin by determining the posterior distribution.

$$\mathbb{P}(\mathbf{Z} = \mathbf{z} | \mathbf{W} = \mathbf{w}) \propto \mathbb{P}(\mathbf{W} = \mathbf{w} | \mathbf{Z} = \mathbf{z}) \mathbb{P}(\mathbf{Z} = \mathbf{z}).$$

Now

$$\mathbb{P}(\mathbf{Z} = \mathbf{z}) = \mathbb{E}[\mathbb{P}(\mathbf{Z} = \mathbf{z}|\theta_1, \theta_2)] \\ = \mathbb{E}[\theta_1^{k_{1.}}(1 - \theta_1)^{n_{1.} - k_{1.}}]\mathbb{E}[\theta_2^{k_{2.}}(1 - \theta_2)^{n_{2.} - k_{2.}}] \\ = \frac{1}{(n_{1.} + 1)(n_{2.} + 1)\binom{n_{1.}}{k_{1.}}\binom{n_{2.}}{k_{2.}}}.$$

Here we have used the notation  $k_{dj}$  for the number words in document d with the topic being A and the word being j and the same dot notation for the k:s as for the n:s. The last equality is Lemma 2.1. For the second factor we have analogously, again using Lemma 2.1,

$$\mathbb{P}(\mathbf{W} = \mathbf{w} | \mathbf{Z} = \mathbf{z}) = \mathbb{E}[\phi_A^{k_1} (1 - \phi_A)^{k_{..} - k_{.1}}] \mathbb{E}[\phi_B^{n_1 - k_{.1}} (1 - \phi_B)^{n_{..} - (n_{.1} - k_{.1})}]$$
$$= \frac{1}{(k_{..} + 1)(n_{..} - k_{..} + 1)\binom{k_{..}}{k_{.1}}\binom{n_{..} - k_{..}}{n_{.1} - k_{.1}}}.$$

Hence, ignoring factors that do not depend on the k:s and using Lemma 2.2 for the second equality

$$\mathbb{P}(\mathbf{Z} = \mathbf{z} | \mathbf{W} = \mathbf{w}) \propto \left( (k_{..} + 1)(n_{..} - k_{..} + 1) \binom{n_{1.}}{k_{1.}} \binom{n_{2.}}{k_{2.}} \binom{k_{..}}{k_{.1}} \binom{n_{..} - k_{..}}{n_{.1} - k_{.1}} \right)^{-1} \\ = \frac{\binom{n_{..}}{k_{..}}}{(k_{..} + 1)(n_{..} - k_{..} + 1)\binom{n_{1.}}{k_{1.}} \binom{n_{2.}}{k_{2.}} \binom{n_{..}}{k_{..}}}.$$

This expression only depends on  $\mathbf{z}$  via  $\mathbf{k} = \mathbf{k}(\mathbf{z}) := (k_{11}, k_{12}, k_{21}, k_{22})$ . Identifying all  $\mathbf{z}$  having the same  $\mathbf{k}(\mathbf{z})$ , we have (regarding ,with some abuse of notation, a  $\mathbf{k}$  also as the equivalence class consisting of all  $\mathbf{z}$ :s having that particular  $k_{dj}$ :s) that for any  $\mathbf{k}_1$  and  $\mathbf{k}_2$ , all  $\mathbf{z} \in \mathbf{k}_1$  have the same probability of transitioning into  $\mathbf{k}_2$ . Hence the process where we only record the  $\mathbf{k}$ :s is a lumped Markov chain and it suffices to regard this chain, whose state space is  $[n_{11}] \times [n_{12}] \times [n_{21}] \times [n_{22}]$ . In particular the Gibbs sampler does not mix any faster than the lumped Markov chain; this follows from the definition of the total variation norm and the triangle inequality. The lumped Markov chain has the stationary distribution

$$f(\mathbf{k}) = f_{\mathbf{w}}(\mathbf{k}) := \mathbb{P}(\mathbf{K} = \mathbf{k} | \mathbf{W} = \mathbf{w}) = C_0 \frac{\binom{n_{..}}{k_{..}}\binom{n_{11}}{k_{11}}\binom{n_{12}}{k_{21}}\binom{n_{21}}{k_{21}}\binom{n_{22}}{k_{22}}}{(k_{..} + 1)(n_{..} - k_{..} + 1)\binom{n_{1}}{k_{1.}}\binom{n_{2}}{k_{2.}}\binom{n_{.1}}{k_{2.}}\binom{n_{.2}}{k_{2.}}}$$

where of course  $\mathbf{K} = \mathbf{k}(\mathbf{Z})$  and  $C_0$  is the normalizing constant.

To prove Theorem 1.1, we will show that there are small neighborhoods of the two states  $\mathbf{k}_0 = (3m/10, 0, 6m/10, 0)$  and  $\mathbf{k}_1 = (3m/10, 7m/10, 0, 0)$  that are extremely difficult to leave.

**Remark.** By switching the names of the two topics, it follows that the states (0, 7m/10, 0, 4m/10) and (0, 0, 6m/10, 4m/10) are equally difficult to leave. One can on good grounds argue that these are really the same as  $\mathbf{k}_0$  and  $\mathbf{k}_1$  respectively. Doing so, the result of Theorem 1.1 is still valid, at least for  $\kappa < 1/2$ , which is "bad enough".

Define

$$h(x) = (x^{x}(1-x)^{1-x})^{-1}, x \in (0,1).$$

Then, for  $a \in (0, 1)$ , by Stirling's formula,

$$\binom{m}{am} = C_1 h(a)^m,$$

where  $C_1$  is of order  $m^{-1/2}$ . Hence

$$f(\mathbf{k}_0) = C_0 C_2 \frac{h(9/20)^{2m}}{h(3/10)^m h(3/5)^m} = C_0 C_2 \left(\frac{h(9/20)^2}{h(3/10)h(3/5)}\right)^m,$$

where  $C_2$  is of order  $m^{1/2}$ . Now consider f in a small neighborhood of  $\mathbf{k}_0$ : let a, b, c, d be small positive numbers, say at most 1/10. Then

$$\begin{aligned} f(3m/10 - am, bm, 3m/5 - cm, dm) \\ &= C_0 C_3 \left( \frac{h(\frac{9}{20} + \frac{(b+d-a-c)}{2})^2 h(\frac{10a}{3})^{3/10} h(\frac{10b}{7})^{7/10} h(\frac{5c}{3})^{3/5} h(\frac{5d}{2})^{2/5}}{h(\frac{10(a+c)}{9})^{9/10} h(\frac{10(b+d)}{11})^{11/10} h(\frac{3}{10} + b - a) h(\frac{3}{5} + d - c)} \right)^m \\ &=: C_0 C_3 g(a, b, c, d)^m, \end{aligned}$$

where  $C_3$  is of order between  $m^{-1/2}$  and  $m^{1/2}$ . We claim that all partial derivatives of g are strictly negative at the origin. We do this with respect to a, leaving the analogous calculations to the reader. Now

$$g(a,0,0,0) = \frac{h(\frac{9}{20} - \frac{a}{2})^2 h(\frac{10a}{3})^{3/10}}{h(\frac{10a}{9})^{9/10} h(\frac{3}{10} - a)}$$

so that

$$\log g(a, 0, 0, 0) = 2 \log h\left(\frac{9}{20} - \frac{a}{2}\right) + \frac{3}{10} \log h\left(\frac{10a}{3}\right) - \frac{9}{10} \log h\left(\frac{10a}{9}\right) - \log h\left(\frac{3}{10} - a\right).$$

It is readily checked that

$$\frac{d}{dx}\log h(x) = \log \frac{1-x}{x}.$$

Hence

$$\frac{\partial}{\partial a}\log g(a,0,0,0) = \log \frac{(\frac{9}{20} - \frac{a}{2})(1 - \frac{10a}{3})\frac{10a}{9}(\frac{7}{10} + a)}{(\frac{11}{20} + \frac{a}{2})\frac{10a}{3}(1 - \frac{10a}{9})(\frac{3}{10} - a)}.$$

For a = 0, the right hand side equals  $-\log(11/7)$  which is smaller than -9/20. Analogously, the partial derivatives of  $\log g$  with respect to b, c and d are less than -1/4 (we do not get exactly the same numbers here). It follows that for a sufficiently small  $\epsilon > 0$ ,

$$f(\mathbf{k}) < e^{-\frac{\epsilon}{4}m} f(\mathbf{k}_0)$$

whenever  $\|\mathbf{k} - \mathbf{k}_0\|_{\infty} = \epsilon m$ . For such a k, let  $V_k$  be the number of visits to k of the lumped Markov chain starting at  $\mathbf{k}_0$  until it returns to  $\mathbf{k}_0$ . Then by an intuitive and well known fact in Markov theory (see e.g. [1], Lemma 6 of Chapter 2)

$$\mathbb{E}[V_{\mathbf{k}}] = \frac{f(\mathbf{k})}{f(\mathbf{k}_0)} < e^{-\frac{\epsilon}{4}m}.$$

Hence

$$\mathbb{P}(V_{\mathbf{k}} > 0) = \frac{\mathbb{E}[V_{\mathbf{k}}]}{\mathbb{E}[V_{\mathbf{k}}|V_{\mathbf{k}} > 0]} < e^{-\frac{\epsilon}{4}m}.$$

By a Bonferroni bound we get

$$\mathbb{P}\left(\bigcup_{\mathbf{k}:\|\mathbf{k}-\mathbf{k}_0\|_{\infty}} \{V_{\mathbf{k}}>0\}\right) < 4(\epsilon m)^3 e^{-\frac{\epsilon}{4}m} < e^{-\frac{\epsilon}{8}m}$$

for sufficiently large m ( $\epsilon m > 170$  is sufficient). Hence, starting at  $\mathbf{k}_0$ , the probability of getting more than distance  $\epsilon m$  from  $\mathbf{k}_0$  in  $e^{\frac{\epsilon}{9}m}$  steps is at most  $e^{-\frac{\epsilon}{72}m}$  which goes to 0 as  $m \to \infty$ .

Next, we make an analogous analysis of the neighborhood of  $k_1$ . We get

$$f(\mathbf{k}_1) = C_0 C_1 \left( \frac{h(\frac{1}{2})^2}{h(\frac{1}{3})^{9/10} h(\frac{4}{11})^{11/10}} \right)^m,$$

where  $C_1$  is of order  $n^{1/2}$ , and

$$\begin{split} f(3m/10-am,bm,3/5-cn,dm) \\ = C_0 C_3 \left( \frac{h(\frac{1}{2} + \frac{(c+d-a-b)}{2})^2 h(\frac{10a}{3})^{3/10} h(\frac{10b}{7})^{7/10} h(\frac{5c}{3})^{3/5} h(\frac{5d}{2})^{2/5}}{h(\frac{1}{3} + \frac{10(c-a)}{9})^{9/10} h(\frac{7}{10} + \frac{10(d-b)}{11})^{11/10} h(a+b) h(c+d)} \right)^m \\ =: C_0 C_3 g(a,b,c,d)^m \end{split}$$

This entails that

$$g(a,0,0,0) = \frac{h(\frac{1}{2} - \frac{a}{2})^2 h(\frac{10a}{3})^{3/10}}{x}$$

Taking the logarithm and differentiating with respect to a gives

$$\frac{\partial}{\partial a}\log g(a,0,0,0) = \log\left(\frac{\frac{1}{2} - \frac{a}{2}\left(1 - \frac{10a}{3}\right)\left(\frac{2}{3} + \frac{10a}{9}\right)a}{\left(\frac{1}{2} + \frac{a}{2}\right)\frac{10a}{3}\left(\frac{1}{3} - \frac{10a}{9}\right)(1 - a)}\right)$$

and taking a = 0, this becomes  $-\log(5/3) < -1/2$ . Repeating the above argument, it follows for some  $\nu > 0$  that starting from  $\mathbf{k}_1$ , the probability of getting more than distance  $\nu m$  from  $\mathbf{k}_1$  in  $e^{\frac{\nu}{9}m}$  steps is at most  $e^{-\frac{\nu}{72}m}$ . Combining this with the analogous result for  $\mathbf{k}_0$ , this proves Theorem 1.1.

#### References

- [1] David Aldous and James A. Fill, Reversible Markov Chains and Random Walks on Graphs, Unfinished monograph at http://www.stat.berkeley.edu/ aldous/RWG/book.html
- [2] Mark Andrews and Gabriella Vigliocco, The Hidden Markov Topic Model (2010), A Probabilistic Model of Semantic Representation, *Topics in Cognitive Science* 2, 101-113.
- [3] Davi M. Blei, Introduction to Probabilistic Topic Models (2011), Princeton University.
- [4] David M. Blei, Andrew Y. Ng and Michael I. Jordan, Latent Dirichlet Allocation (2003), *Journal of Machine Learning Research* **93**, 993-1022.
- [5] Matt Hoffman, David M. Blei, Chong Wang and John Paisley (2013), *Journal of Machine Learning Research* **14**, 1303-1347.