

## 1 Appendix A

2 In the following, we will provide LDA's complete conditional distributions for  $\pi_d$ ,  $\mathbf{k}_{dn}$  and  $\beta$ , but  
3 firstly, let us remind ourselves of LDA's generative process:

- 4 1. Draw topics proportions  $\pi_d \sim \text{Dir}(\alpha)$
- 5 2. For each word  $x_{dn}$ :
  - 6 (a) Draw a topic assignment  $\mathbf{k}_{dn} \sim \text{Cat}(\pi_d)$
  - 7 (b) Draw a word  $x_{dn} \sim \text{Cat}(\beta^\top \mathbf{k}_{dn})$

8 Here,  $\alpha$  is a  $K$ -dimensional vector of Dirichlet concentration parameters, where  $K$  is the number of  
9 topics,  $\pi_d$  is  $K \times 1$  vector of topic proportions, and  $\beta$  is a  $K \times V$  matrix where  $V$  is the number of  
10 words in a vocabulary. Every  $\beta_{ij}$  specifies the conditional probability of a word  $j$  given a topic  $i$ ,  
11 s.t.  $\sum_j \beta_{ij} = 1$ . We will use one-hot encoding for words and topics such that  $\mathbf{x}_{dn}$  is a  $V \times 1$  vector,  
12 and  $\mathbf{k}_{dn}$  is a  $K \times 1$  vector. We will drop the indices of documents and words,  $d$  and  $n$ , to denote the  
13 collection of variables. For example,  $\mathbf{x}$  is a collection of words across the dataset.

14 **1.** The conditional distribution  $p(\pi_d | \mathbf{k}_d)$  is the posterior of a Dirichlet-categorical model, and it can  
15 be shown to take the form:

$$\begin{aligned} p(\pi_d | \mathbf{k}_d) &= \text{Dir} \left( \pi_d; \alpha + \sum_n \mathbf{k}_{dn} \right) \\ &\propto \exp \left( \left( \alpha + \sum_n \mathbf{k}_{dn} - 1 \right)^\top \ln \pi_d \right). \end{aligned} \quad (1)$$

16 **2.** The conditional distribution over topic assignments is given by:

$$\begin{aligned} p(\mathbf{k}_{dn} | x_{dn}, \pi_d, \beta) &\propto p(x_{dn} | \pi_d, \mathbf{k}_{dn}, \beta) p(\mathbf{k}_{dn} | \pi_d, \beta) \\ &= p(x_{dn} | \mathbf{k}_{dn}, \beta) p(\mathbf{k}_{dn} | \pi_d) \\ &= (\mathbf{k}_{dn}^\top \beta \mathbf{x}_{dn}) (\mathbf{k}_{dn}^\top \pi_d) \\ &= \exp(\mathbf{k}_{dn}^\top \ln \beta \mathbf{x}_{dn} + \mathbf{k}_{dn}^\top \ln \pi_d). \end{aligned} \quad (2)$$

17 **3.** Finally, we derive the posterior distribution of  $\beta$  assuming a Dirichlet prior over the rows of  $\beta$ , i.e.,  
18  $p(\beta_k) = \text{Dir}(\eta)$  for every  $k = 1, \dots, K$ :

$$\begin{aligned} p(\beta | \mathbf{x}, \mathbf{k}) &\propto p(\mathbf{x} | \beta, \mathbf{k}) p(\beta | \mathbf{k}) \\ &= \left( \prod_{dn} \mathbf{k}_{dn}^\top \beta \mathbf{x}_{dn} \right) \left( \prod_k p(\beta_k) \right) \\ &= \left( \prod_{dn} \mathbf{k}_{dn}^\top \beta \mathbf{x}_{dn} \right) \left( \prod_k \frac{1}{\mathcal{B}(\eta)} \prod_i \beta_{ki}^{\eta_i - 1} \right) \\ &= \exp \left( \sum_{dn} \mathbf{k}_{dn}^\top \ln \beta \mathbf{x}_{dn} + (\eta - 1)^\top \sum_k \ln \beta_k - K \ln \mathcal{B}(\eta) \right). \end{aligned} \quad (3)$$

19 Here,  $\mathcal{B}$  is a multivariate beta function.

## 20 Appendix B

21 Given a set of assumptions about the conditional distributions of logistic LDA, we would like to  
22 derive the joint conditional distribution over  $\pi$ ,  $\mathbf{k}$  and  $\theta$ . For that, we can use Brook's lemma [1]:

23 **Brook's Lemma.** Let  $p(\mathbf{y}) > 0$  for all  $\mathbf{y}$ . Then for any  $\mathbf{y}$  and  $\mathbf{y}'$  we have

$$\frac{p(\mathbf{y})}{p(\mathbf{y}')} = \prod_i \frac{p(\mathbf{y}_i | \mathbf{y}_1, \dots, \mathbf{y}_{i-1}, \mathbf{y}'_{i+1}, \dots, \mathbf{y}'_M)}{p(\mathbf{y}'_i | \mathbf{y}_1, \dots, \mathbf{y}_{i-1}, \mathbf{y}'_{i+1}, \dots, \mathbf{y}'_M)}.$$

24 Logistic LDA specifies the following conditionals:

$$p(\boldsymbol{\pi}_d | \mathbf{k}_d) \propto \exp \left( \left( \boldsymbol{\alpha} + \sum_n \mathbf{k}_{dn} - 1 \right)^\top \ln \boldsymbol{\pi}_d \right), \quad (4)$$

$$p(\mathbf{k}_{dn} | \mathbf{x}_{dn}, \boldsymbol{\pi}_d, \boldsymbol{\theta}) \propto \exp \left( \mathbf{k}_{dn}^\top g(\mathbf{x}_{dn}, \boldsymbol{\theta}) + \mathbf{k}_{dn}^\top \ln \boldsymbol{\pi}_d \right), \quad (5)$$

$$p(\boldsymbol{\theta} | \mathbf{x}, \mathbf{k}) \propto \exp \left( r(\boldsymbol{\theta}) + \sum_{dn} \mathbf{k}_{dn}^\top g(\mathbf{x}_{dn}, \boldsymbol{\theta}) \right). \quad (6)$$

25 Using Brook's lemma:

$$\begin{aligned} p(\boldsymbol{\pi}, \mathbf{k}, \boldsymbol{\theta} | \mathbf{x}) &\propto \frac{p(\boldsymbol{\pi} | \mathbf{k}') p(\mathbf{k} | \boldsymbol{\pi}, \boldsymbol{\theta}', \mathbf{x}) p(\boldsymbol{\theta} | \mathbf{k}, \mathbf{x})}{p(\boldsymbol{\pi}' | \mathbf{k}') p(\mathbf{k}' | \boldsymbol{\pi}, \boldsymbol{\theta}', \mathbf{x}) p(\boldsymbol{\theta}' | \mathbf{k}, \mathbf{x})} \\ &\propto \exp \left( (\boldsymbol{\alpha} - 1)^\top \sum_d \ln \boldsymbol{\pi}_d + \sum_{dn} \mathbf{k}_{dn}^\top \ln \boldsymbol{\pi}_d + \sum_{dn} \mathbf{k}_{dn}^\top g(\mathbf{x}_{dn}, \boldsymbol{\theta}) + r(\boldsymbol{\theta}) \right) \end{aligned} \quad (7)$$

26 This is easiest to see if we chose  $\boldsymbol{\theta}'$  such that  $g(\mathbf{x}_{dn}, \boldsymbol{\theta}') = 0$ , which Brook's lemma allows us to do.

27 In case of logistic LDA with labels, we have a slightly different set of conditionals:

$$p(\mathbf{c}_d | \boldsymbol{\pi}_d) \propto \exp \left( \lambda \mathbf{c}_d^\top \ln \boldsymbol{\pi}_d \right), \quad (8)$$

$$p(\boldsymbol{\pi}_d | \mathbf{k}_d, \mathbf{c}_d) \propto \exp \left( \left( \boldsymbol{\alpha} + \sum_n \mathbf{k}_{dn} - 1 + \lambda \mathbf{c}_d \right)^\top \ln \boldsymbol{\pi}_d \right), \quad (9)$$

$$p(\mathbf{k}_{dn} | \mathbf{x}_{dn}, \boldsymbol{\pi}_d, \boldsymbol{\theta}) \propto \exp \left( \mathbf{k}_{dn}^\top g(\mathbf{x}_{dn}, \boldsymbol{\theta}) + \mathbf{k}_{dn}^\top \ln \boldsymbol{\pi}_d \right), \quad (10)$$

$$p(\boldsymbol{\theta} | \mathbf{x}, \mathbf{k}) \propto \exp \left( r(\boldsymbol{\theta}) + \sum_{dn} \mathbf{k}_{dn}^\top g(\mathbf{x}_{dn}, \boldsymbol{\theta}) \right). \quad (11)$$

28 Similarly to Eq. 7, we can derive a joint distribution as:

$$\begin{aligned} p(\boldsymbol{\pi}, \mathbf{k}, \mathbf{c}, \boldsymbol{\theta} | \mathbf{x}) &\propto \frac{p(\boldsymbol{\pi} | \mathbf{k}') p(\mathbf{k} | \boldsymbol{\pi}, \boldsymbol{\theta}', \mathbf{x}) p(\mathbf{c} | \boldsymbol{\pi}) p(\boldsymbol{\theta} | \mathbf{k}, \mathbf{x})}{p(\boldsymbol{\pi}' | \mathbf{k}') p(\mathbf{k}' | \boldsymbol{\pi}, \boldsymbol{\theta}', \mathbf{x}) p(\mathbf{c}' | \boldsymbol{\pi}) p(\boldsymbol{\theta}' | \mathbf{k}, \mathbf{x})} \\ &\propto \exp \left( (\boldsymbol{\alpha} - 1)^\top \sum_d \ln \boldsymbol{\pi}_d + \sum_{dn} \mathbf{k}_{dn}^\top \ln \boldsymbol{\pi}_d + \lambda \sum_d \mathbf{c}_d^\top \ln \boldsymbol{\pi}_d + \sum_{dn} \mathbf{k}_{dn}^\top g(\mathbf{x}_{dn}, \boldsymbol{\theta}) + r(\boldsymbol{\theta}) \right) \end{aligned}$$

## 29 Appendix C

30 We assume the following parameterizations of the variational distributions:

$$\begin{aligned} q(\mathbf{c}_d) &= \mathbf{c}_d^\top \hat{\boldsymbol{\rho}}_d \\ q(\boldsymbol{\pi}_d) &= \text{Dir}(\boldsymbol{\pi}_d; \hat{\boldsymbol{\alpha}}_d) \\ q(\mathbf{k}_{dn}) &= \mathbf{k}_{dn}^\top \hat{\boldsymbol{\rho}}_{dn} \end{aligned} \quad (12)$$

31 For now, we will fix  $\boldsymbol{\theta}$  to its point estimate  $\hat{\boldsymbol{\theta}}$  and derive the optimal variational densities for other  
32 latent variables. Our derivation is based on the coordinate ascent variational inference algorithm  
33 (CAVI) [2].

34 **1. Optimal variational density for  $\boldsymbol{\pi}_d$ :**

$$\begin{aligned} q^*(\boldsymbol{\pi}_d) &\propto \exp \left( \mathbb{E}_{q(\mathbf{k}_{dn}), q(\mathbf{c}_d)} \left[ \sum_n \mathbf{k}_{dn}^\top \ln \boldsymbol{\pi}_d + (\boldsymbol{\alpha} - 1)^\top \ln \boldsymbol{\pi}_d + \lambda \mathbf{c}_d^\top \ln \boldsymbol{\pi}_d \right] \right) \\ &= \exp \left( \left( \boldsymbol{\alpha} + \sum_n \mathbb{E}_{q(\mathbf{k}_{dn})} \mathbf{k}_{dn} + \lambda \mathbb{E}_{q(\mathbf{c}_d)} \mathbf{c}_d - 1 \right)^\top \ln \boldsymbol{\pi}_d \right) \\ &= \exp \left( \left( \boldsymbol{\alpha} + \sum_n \hat{\boldsymbol{\rho}}_{dn} + \lambda \hat{\boldsymbol{\rho}}_d - 1 \right)^\top \ln \boldsymbol{\pi}_d \right) \end{aligned} \quad (13)$$

35 Thus, the CAVI update for the variational parameter  $\hat{\alpha}$  is:

$$\hat{\alpha} = \alpha + \sum_n \hat{p}_{dn} + \lambda \hat{p}_d. \quad (14)$$

36 **2. Optimal variational density for  $\mathbf{k}_{dn}$ :**

$$\begin{aligned} q(\mathbf{k}_{dn}) &\propto \exp \left( \mathbb{E}_{\boldsymbol{\theta} \sim q(\boldsymbol{\theta}), \boldsymbol{\pi}_d \sim q(\boldsymbol{\pi}_d)} [\mathbf{k}_{dn}^\top \ln \boldsymbol{\pi}_d + \mathbf{k}_{dn}^\top g(\mathbf{x}_{dn}, \boldsymbol{\theta})] \right) \\ &= \exp \left( \mathbf{k}_{dn}^\top \mathbb{E}_{q(\boldsymbol{\pi}_d)} [\ln \boldsymbol{\pi}_d] + \mathbf{k}_{dn}^\top g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}) \right) \\ &= \exp \left( \mathbf{k}_{dn}^\top (\psi(\hat{\boldsymbol{\alpha}}_d) - \psi \left( \sum_k \hat{\alpha}_{dk} \right)) + g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}) \right) \end{aligned} \quad (15)$$

37 We now see that  $\ln \hat{p}_{dn} \propto \exp(\psi(\hat{\boldsymbol{\alpha}}_d) - \psi(\sum_k \hat{\alpha}_{dk}) + g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}))$ . After renormalizing:

$$\hat{p}_{dn} = \frac{\exp(\psi(\hat{\boldsymbol{\alpha}}_d) + g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}))}{\sum_k \exp(\psi(\hat{\alpha}_{dk}) + g_k(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}))}. \quad (16)$$

38 **3. Optimal variational density for  $\mathbf{c}_d$ :**

$$q(\mathbf{c}_d) \propto \exp \mathbb{E}_{\boldsymbol{\pi}_d \sim q(\boldsymbol{\pi}_d)} [\lambda \mathbf{c}_d^\top \ln \boldsymbol{\pi}_d] = \exp \left( \lambda \mathbf{c}_d^\top \left( \psi(\hat{\boldsymbol{\alpha}}_d) - \psi \left( \sum_k \hat{\alpha}_{dk} \right) \right) \right) \quad (17)$$

39 Thus, the update for  $\hat{p}_d$  becomes:

$$\hat{p}_d = \frac{\exp(\lambda \psi(\hat{\boldsymbol{\alpha}}_d))}{\sum_k \exp(\lambda \psi(\hat{\alpha}_{dk}))}. \quad (18)$$

40 **4.** To find the updates for the neural network parameters, we now fix  $\hat{p}_{dn}$  and search for  $\hat{\boldsymbol{\theta}}$  that  
41 maximises the following:

$$\begin{aligned} q(\boldsymbol{\theta}) &\propto \exp \mathbb{E}_{\mathbf{k}_{dn} \sim q(\mathbf{k}_{dn})} \left[ r(\boldsymbol{\theta}) + \sum_{dn} \mathbf{k}_{dn}^\top g(\mathbf{x}_{dn}, \boldsymbol{\theta}) \right] \\ &= \exp \left( r(\boldsymbol{\theta}) + \sum_{dn} \mathbb{E}_{q(\mathbf{k}_{dn})} [\mathbf{k}_{dn}^\top] g(\mathbf{x}_{dn}, \boldsymbol{\theta}) \right) \\ &= \exp \left( r(\boldsymbol{\theta}) + \sum_{dn} \hat{p}_{dn}^\top g(\mathbf{x}_{dn}, \boldsymbol{\theta}) \right) \end{aligned} \quad (19)$$

## 42 Appendix D

43 In the following, we describe further approximations which will make it easier to implement our  
44 approach in practice. In particular, we would like to be able to update  $\hat{\alpha}_d$  and  $\hat{\boldsymbol{\theta}}$  after observing just  
45 a single item (or a mini batch of data points) and not have to process large collections of items or  
46 entire datasets in each iteration.

47 For every document, we keep an estimate of the expected topic counts from which we can calculate  
48  $\hat{\alpha}_d$  in Eq. 14:

$$\bar{p}_d = \sum_n \hat{p}_{dn}. \quad (20)$$

49 Assume that at training step  $t$  we update  $\hat{p}_{dn}$ , then we would also want to perform the following  
50 update:

$$\bar{p}_d^t = \bar{p}_d^{t-1} + \hat{p}_{dn}^t - \hat{p}_{dn}^{t-1} \quad (21)$$

51 Unfortunately, this requires keeping in memory the previous value of  $\hat{\mathbf{p}}_{dn}$ , which is too costly for  
 52 datasets containing a lot of items. Instead, in every update of  $\bar{\mathbf{p}}_d$  we subtract an equal fraction of all  
 53  $\hat{\mathbf{p}}_{dn}$ :

$$\bar{\mathbf{p}}_d^t = \bar{\mathbf{p}}_d^{t-1} + \hat{\mathbf{p}}_{dn}^t - \bar{\mathbf{p}}_d^{t-1}/N_d. \quad (22)$$

54 Using Jensen’s inequality, let us bring the regularizer into a form which is more amenable to stochastic  
 55 approximation:

$$\begin{aligned} r(\boldsymbol{\theta}, \mathbf{x}) &= \gamma \cdot \mathbf{1}^\top \ln \sum_{dn} \exp g(\mathbf{x}_{dn}, \boldsymbol{\theta}) \\ &= \gamma \sum_k \ln \sum_{dn} \frac{r_{dnk}}{r_{dnk}} \exp g_k(\mathbf{x}_{dn}, \boldsymbol{\theta}) \\ &\geq \gamma \sum_k \sum_{dn} r_{dnk} \ln \frac{\exp g_k(\mathbf{x}_{dn}, \boldsymbol{\theta})}{r_{dnk}} \end{aligned} \quad (23)$$

56 Eq. 23 provides us with a lower bound on the regularizer for arbitrary  $r_{dnk}$  with  $\sum_{dn} r_{dnk} = 1$  for  
 57 every  $k$ . The bound is tight when

$$r_{dnk} = \frac{\exp g_k(\mathbf{x}_{dn}, \boldsymbol{\theta})}{\sum_{dn} \exp g_k(\mathbf{x}_{dn}, \boldsymbol{\theta})}. \quad (24)$$

58 Intuitively,  $r_{dnk}$  is the responsibility of a word  $\mathbf{x}_{dn}$  for topic  $k$ . If it is large, a word is associated  
 59 more strongly with a topic than other words in the dataset.

60 With the approximation to the regularizer, we obtain a simpler form for the loss function of  $\hat{\boldsymbol{\theta}}$ :

$$\begin{aligned} -\ell(\hat{\boldsymbol{\theta}}) &= \sum_{dn} \hat{\mathbf{p}}_{dn}^\top g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}) + \gamma \cdot \mathbf{1}^\top \ln \sum_{dn} \exp g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}) \\ &\geq \sum_{dn} \hat{\mathbf{p}}_{dn}^\top g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}) + \gamma \sum_{dn} \mathbf{r}_{dn}^\top g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}) - \gamma \sum_{dn} \mathbf{r}_{dn}^\top \ln \mathbf{r}_{dn} \\ &= \sum_{dn} \hat{\mathbf{p}}_{dn}^\top g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}) + \gamma \sum_{dn} \mathbf{r}_{dn}^\top g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}) + \text{const} \end{aligned} \quad (25)$$

61 Here, we assumed that  $r_{dnk}$  is computed based on a fixed value of neural network parameters and  
 62 thus is a constant with respect to  $\boldsymbol{\theta}$ . If  $\mathbf{r}_{dn}$  is such that the bound in Eq. 23 is tight, then the gradient  
 63 at  $\hat{\boldsymbol{\theta}}$  will be the same as for the original loss.

64 Unfortunately, the denominator in Eq. 24 depends on the entire dataset. We therefore propose to  
 65 estimate the denominator with a running average. If at training step  $t$  we observe  $\mathbf{x}_{dn}$ , our estimate  $\bar{\mathbf{r}}$   
 66 is updated as follows:

$$\hat{\mathbf{r}}_{dn}^* = \exp g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}), \quad \bar{\mathbf{r}}_{t+1} = \bar{\mathbf{r}}_t + \varepsilon (M \hat{\mathbf{r}}_{dn}^* - \bar{\mathbf{r}}_t), \quad (26)$$

67 where  $M = \sum_d N_d$  is the number of words in the document. With this we can estimate responsibili-  
 68 ties as  $\hat{\mathbf{r}}_{dn} = \hat{\mathbf{r}}_{dn}^*/\bar{\mathbf{r}}$ , and stochastically approximate the loss with

$$-\ell(\hat{\boldsymbol{\theta}})/M \approx \hat{\mathbf{p}}_{dn}^\top g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}) + \gamma \cdot \hat{\mathbf{r}}_{dn}^\top g(\mathbf{x}_{dn}, \hat{\boldsymbol{\theta}}). \quad (27)$$

69 This corresponds to a cross-entropy loss, where the targets of the neural network are given by  
 70  $\hat{\mathbf{p}}_{dn} + \gamma \cdot \hat{\mathbf{r}}_{dn}$ . The full algorithm for stochastic training of logistic LDA (supervised version) is given  
 71 in Algorithm 1. An inference procedure which can be used at test time is given in Algorithm 2.

---

**Algorithm 1** Training step given a single word  $x_{dn}$  and a document label  $c_d$ . Before training,  $\bar{p}_d = 0$  and  $\bar{r} = \frac{M}{K}\mathbf{1}$ .  $N_d$  is the number of words in each document and  $M$  is the total number of words in the dataset.

---

**Require:**  $x_{dn}, c_d, \hat{\theta}, \bar{p}_d, \bar{r}, N_d, M$

```

% Compute logits and fix their values
 $f_{dn} \leftarrow f(x_{dn}, \hat{\theta})$ 

% Compute beliefs over latent variables
 $\hat{\alpha}_d \leftarrow \alpha + \bar{p}_d + \lambda c_d$ 
 $\hat{p}_{dn} \leftarrow \text{softmax}(f_{dn} + \psi(\hat{\alpha}_d))$ 

% Compute responsibilities
 $\hat{r}_{dn}^* \leftarrow \text{softmax}(f_{dn})$ 
 $\hat{r}_{dn} \leftarrow \hat{r}_{dn}^* / \bar{r}$ 

% Update neural network's parameters
 $\hat{\theta} \leftarrow \hat{\theta} - \varepsilon \nabla_{\theta} \text{crossent}(\hat{p}_{dn} + \gamma \cdot \hat{r}_{dn}, f(x_{dn}, \hat{\theta}))$ 

% Update statistics
 $\bar{p}_d \leftarrow \bar{p}_d + \hat{p}_{dn} - \bar{p}_d / N_d$ 
 $\bar{r} \leftarrow \bar{r} + \hat{r}_{dn} - \bar{r} / M$ 

return  $\hat{\theta}, \bar{p}_d, \bar{r}$ 

```

---



---

**Algorithm 2** Stochastic inference, estimating topic proportions by looking at words one by one. This algorithm is expected to be called many times with words sampled from the corpus.

---

**Require:**  $x_{dn}, \bar{p}_d$

```

% Beliefs over class labels
 $\hat{p}_d \leftarrow \mathbf{1} / K$ 

% Beliefs over topic proportions
for  $i$  from 1 to  $N_{\text{iter}}$  do
     $\hat{\alpha}_d \leftarrow \alpha + \bar{p}_d + \lambda \hat{p}_d$ 
     $\hat{p}_d \leftarrow \text{softmax}(\lambda \psi(\hat{\alpha}_d))$ 
end for

% Make prediction for word
 $\hat{p}_{dn} \leftarrow \text{softmax}(f(x_{dn}, \hat{\theta}) + \psi(\hat{\alpha}_d))$ 

% Update document's topic counts
 $\bar{p}_d \leftarrow \bar{p}_d + \hat{p}_{dn} - \bar{p}_d / N_d$ 

return  $\bar{p}_d, \hat{p}_d, \hat{p}_{dn}$ 

```

---

## 73 Appendix E

74 For supervised models, we used random search across the following range of hyperparameters:

- 75 – number of hidden units in up to three hidden layers:
- 76 128, 256, 512, 1024, [1024 512], [512 256], [256 128], [512 256 128]
- 77 – initial learning rate: 0.001, 0.003
- 78 – exponential learning rate decay every 2000 iterations: 0.5 to 0.9
- 79 – batch size: 8 to 128
- 80 – number of training steps: 50000 to 200000
- 81 – Dirichlet prior parameter ( $\alpha$ ): 0.5 to 10

82 For models trained with a mean-field variational inference (Section 5.1 in the main text):

- 83 –  $\lambda$ : 10.0 to 8000.0
- 84 –  $\gamma$ : 0.1 to 1000.0
- 85 – number of variational iterations ( $N_{\text{iter}}$ ): 1 to 5

86 For models trained with the empirical loss (Section 5.2 in the main text):

- 87 –  $\lambda$ : 1.0
- 88 – number of variational iterations ( $N_{\text{iter}}$ ): 1 to 20

## 89 Twitter

90 The Twitter model was trained with a stochastic approximation of mean-field variational inference as  
91 given in Algorithm 1 for approximately 20 minutes on a CPU. For all models evaluated in Table 1 of  
92 the main text, we randomly sampled 100 sets of hyperparameters. A single model was chosen based on  
93 a validation set and evaluated on the test set. The model which achieved 38.7% on the author-level  
94 community prediction used the following hyperparameters:

- 95 – number of hidden units: [256, 128]
- 96 – initial learning rate: 0.0005
- 97 – exponential learning rate decay every 2000 iterations: 0.8
- 98 – batch size (number of items): 64
- 99 – number of training iterations: 60000
- 100 –  $\lambda$ : 2000.0
- 101 –  $\alpha$ : 0.5
- 102 –  $\gamma$ : 0.25

### 103 **Pinterest**

104 The Pinterest model was trained with a stochastic approximation of mean-field variational inference  
105 as given in Algorithm 1 for 5 hours on a CPU with the following hyperparameters:

- 106 – number of hidden units: [256, 128]
- 107 – initial learning rate: 0.001
- 108 – exponential learning rate decay every 2000 iterations: 0.8
- 109 – batch size (number of items): 8
- 110 – number of training iterations: 50000
- 111 –  $\lambda$ : 8000.0
- 112 –  $\alpha$ : 5.0
- 113 –  $\gamma$ : 20.0

114 The unsupervised logistic LDA model for Pinterest had quite a different range of suitable hyperpa-  
115 rameters. We used the model with the following hyperparameters:

- 116 – number of topics: 14
- 117 – number of hidden units: [128]
- 118 – initial learning rate: 0.0001
- 119 – exponential learning rate decay every 2000 iterations: 0.6
- 120 – batch size (number of items): 4
- 121 – number of training iterations: 60000
- 122 –  $\lambda$ : 1.
- 123 –  $\alpha$ : 1.0
- 124 –  $\gamma$ : 1000.0

### 125 **20-Newsgroups**

126 The logistic LDA model that achieves 15.6% error rate on 20-Newsgroups was trained for 11h on a  
127 CPU with the following hyperparameters:

- 128 – number of hidden units: [1024 512]
- 129 – initial learning rate: 0.001
- 130 – exponential learning rate decay every 2000 iterations: 0.7
- 131 – number of documents in a batch: 16
- 132 – number of training iterations: 100000
- 133 – number of variational iterations ( $N_{\text{iter}}$ ): 20
- 134 –  $\alpha, \lambda$ : 1.

135 For the unsupervised logistic LDA trained with a general mean-field VI objective, the range of suitable  
136 hyperparameters is quite different. For example, we noticed that it requires lower learning rates and  
137 stronger regularization in the form of  $\gamma$  from Eq. 27. Thanks to the stochastic approximations, the  
138 unsupervised model takes minutes to train on a CPU. To produce the results in Table 3 from the main  
139 text and tables in Appendix H, we used logistic LDA with the following hyperparameters:

- 140 – number of hidden units: 128
- 141 – initial learning rate: 0.0005
- 142 – exponential learning rate decay every 2000 iterations: 0.8
- 143 – number of documents in a batch: 16
- 144 – number of training iterations: 50000
- 145 –  $\alpha$ : 5.
- 146 –  $\lambda$ : 1.
- 147 –  $\gamma$ : 100.

## 148 **Appendix F**

### 149 **Twitter**

150 For the training data, we first collected English tweets from 200 000 Twitter users with more than  
151 5,000 followers. The creation dates of the tweets ranged from October 8 to October 10, 2018. Tweets  
152 with less than 3 tokens were omitted. Authors were clustered and assigned to communities based  
153 on their follower graph. The communities were manually annotated based on a sample of tweets  
154 from the community. We retained 300 communities containing 66,455 authors and discarded the rest,  
155 leaving a dataset of 1.45 million tweets.

156 The test data was generated by sampling 30,000 tweets created between July 1 and August 1, 2018.  
157 As for the training data, tweets with fewer than 3 tokens were filtered out. Each tweet is manually  
158 annotated with a topic from a taxonomy of over 400 topics. After only retaining tweets corresponding  
159 to one of the 300 communities, we were left with 18,864 tweets. For each of the 18,765 authors of  
160 these tweets, we sampled a number of additional unlabeled tweets, yielding a total of 3.14 million  
161 tweets.

162 Tweet text was preprocessed by removing accents, replacing certain characters (e.g., "æ" becomes  
163 "ae"), and lowercasing words, but no stemming was applied.

### 164 **Pinterest**

165 The Pinterest dataset [5] contains around 46,000 boards divided almost equally among 468 categories.  
166 We noticed that board labels are often noisy and sometimes incorrect. Therefore, we manually selected  
167 a subset of 90 categories, which we merged into 14 broader classes: food, animals, women\_fashion,  
168 men\_fashion, garden, architecture, wedding, cars, hair, home, travel, DIY, fitness, and fashion. This  
169 resulted in a set of 8,989 boards, which we splitted into 7,191 and 1,798 boards for training and  
170 testing. On average, every board contains 58 pinned images. We only kept images with a height to  
171 width ratio between 0.4 and 2.0, and rescaled every image to  $224 \times 224$  pixels.

### 172 **20-Newsgroups**

173 We used the preprocessed ‘all-terms’ version of 20 Newsgroups [3], where attachments, PGP keys,  
174 duplicates and empty messages were removed. This version was chosen for our results to be  
175 comparable with the results of Dai and Le [4]. Note that scores can differ greatly when using other  
176 variants of 20 Newsgroups, e.g., versions with headers, footers or quotes removed. Unfortunately,  
177 the preprocessing done by Cardoso-Cachopo [3] also removed special characters, punctuation,  
178 capitalization and numbers. This prevented us to restore sentences and, for instance, reduce the noise  
179 coming from mishandling contractions. In order for us to use GloVe word embeddings, we further  
180 processed the documents by applying tokenization and stemming.




Figure 1: Sample of pins from test Pinterest boards covering most of the topics. Board labels and predictions are on the left. Pins predictions and probabilities from logistic LDA and MLP are given below every image. There are some examples where MLP outputs seemingly improper predictions, e.g. ‘food’ in the top row, due to its lack of context. The same predictions we would get from our model if we let  $\hat{p}_{dn} \propto f(\mathbf{x}_{dn}, \theta)$  instead of  $\hat{p}_{dn} \propto f(\mathbf{x}_{dn}, \theta) + \psi(\hat{\alpha}_d)$ . Thus, having a context-dependent bias in the form of  $\hat{\alpha}_d$  leads to more meaningful predictions.

182 **Appendix H**

183 An unsupervised logistic LDA trained on the Pinterest data discovers topics as illustrated in Figure 2.  
184 Here, we classified every pin as belonging to one of the 14 topics, sorted the pins according to  
185  $p(\mathbf{k} | \mathbf{x})$ , and selected the top-9 pins per topic. The occasional duplicates come from the fact that the  
186 same image can be pinned multiple times.




Figure 2: Examples of nine topics discovered by unsupervised logistic LDA in the Pinterest dataset.

187 **Appendix I**

188 In the unsupervised 20-Newsgroups experiments, we used the vocabulary provided by Srivastava and  
189 Sutton [7]. Miao et al. [6] used identical settings. For direct comparison, we also chose the same  
190 number of topics, namely 50. We cannot use perplexity or topic coherence to compare logistic LDA  
191 with ProdLDA [7] and GSM Miao et al. [6] since our model does not model the distribution of inputs  
192 needed for those metrics. Therefore, we can only qualitatively evaluate the models by inspecting the  
193 topics. For ProdLDA [7] and GSM Miao et al. [6], a topic is a distribution over words, so the top-10  
194 words in Table 3 and Table 4 are sorted based on  $p(\mathbf{x} | \mathbf{k})$ . On the other hand, logistic LDA models  
195 the probability of topics given a word  $p(\mathbf{k} | \mathbf{x})$ . Thus, we can classify every word in the vocabulary

196 as belonging to some topic and sort the words within each topic according to the topic probability. In  
 197 Table 1 and Table2 we list top-10 of those words.

Table 1: Selected topics discovered by unsupervised logistic LDA represented by top-10 words

1	bmw, motor, car, honda, motorcycle, auto, mg, engine, ford, bike
2	christianity, prophet, atheist, religion, holy, scripture, biblical, catholic, homosexual, religious, atheist
3	spacecraft, orbit, probe, ship, satellite, rocket, surface, shipping, moon, launch
4	user, computer, microsoft, monitor, programmer, electronic, processing, data, app, systems
5	congress, administration, economic, accord, trade, criminal, seriously, fight, responsible, future

Table 2: Five *randomly* selected topics discovered by logistic LDA represented by top-10 words

1	homicide, city, south, area, national, north, department, since, highest, disease
2	marry, mother, father, love, lord, mary, soul, hell, author, life
3	angeles, philadelphia, chicago, vancouver, toronto, pittsburgh, colorado, detroit, university, nhl
4	science, editor, english, letter, journal, art, together, first, award, study
5	gif, jpeg, pixel, font, postscript, unix, email, server, binary, directory

Table 3: Five *randomly* selected topics discovered by ProdLDA [7] represented by top-10 words

1	motherboard, meg, printer, quadra, hd, windows, processor, vga, mhz, connector
2	armenian, genocide, turks, turkish, muslim, massacre, turkey, armenians, armenia, greek
3	voltage, nec, outlet, circuit, cable, wiring, wire, panel, motor, install
4	season, nhl, team, hockey, playoff, puck, league, flyers, defensive, player
5	israel, israeli, lebanese, arab, lebanon, arabs, civilian, territory, palestinian, militia

Table 4: Selected topics discovered by GSM [6] represented by top-10 words

1	space, satellite, april, sequence, launch, president, station, radar, training, committee
2	god, atheism, exist, atheist, moral, existence, marriage, system, parent, murder
3	encryption, device, technology, protect, americans, chip, use, privacy, industry, enforcement
4	player, hall, defensive, team, average, career, league, play, bob, year
5	science, theory, scientific, universe, experiment, observation, evidence, exist, god, mistake

## 198 References

- 199 [1] Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal*  
 200 *Statistical Society. Series B (Methodological)*, 36(2):pp. 192–236.
- 201 [2] Bishop, C. M. (2006). *Pattern Recognition and Machine Learning (Information Science and Statistics)*.  
 202 Springer-Verlag, Berlin, Heidelberg.
- 203 [3] Cardoso-Cachopo, A. (2007). Improving Methods for Single-label Text Categorization. PhD Thesis,  
 204 Instituto Superior Tecnico, Universidade Tecnica de Lisboa.
- 205 [4] Dai, A. M. and Le, Q. V. (2015). Semi-supervised sequence learning. In *Advances in Neural Information*  
 206 *Processing Systems 28*, pages 3079–3087.
- 207 [5] Geng, X., Zhang, H., Bian, J., and Chua, T. (2015). Learning image and user features for recommendation  
 208 in social networks. In *2015 IEEE International Conference on Computer Vision*, pages 4274–4282.
- 209 [6] Miao, Y., Grefenstette, E., and Blunsom, P. (2017). Discovering discrete latent topics with neural variational  
 210 inference. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2410–2419.
- 211 [7] Srivastava, A. and Sutton, C. A. (2016). Neural variational inference for topic models. In *Bayesian deep*  
 212 *learning workshop*.