

Topic Analysis for Text with Side Data

Biyi Fang

Department of Engineering Science and Applied Mathematics

Northwestern University

Evanston, IL 60208

biyifang2021@u.northwestern.edu

Kripa Rajshekhar

Metonymize

kripa@metonymize.com

Diego Klabjan

Department of Industrial Engineering and Management Sciences

Northwestern University

Evanston, IL 60208

d-klabjan@northwestern.edu

August 16, 2021

Abstract

Although latent factor models (e.g., matrix factorization) obtain good performance in predictions, they suffer from several problems including cold-start, non-transparency, and suboptimal recommendations. In this paper, we employ text with side data to tackle these limitations. We introduce a hybrid generative probabilistic model that combines a neural network with a latent topic model, which is a four-level hierarchical Bayesian model. In the model, each document is modeled as a finite mixture over an underlying set of topics and each topic is modeled as an infinite mixture over an underlying set of topic probabilities. Furthermore, each topic probability is modeled as a finite mixture over side data. In the context of text, the neural network provides an overview distribution about side data for the corresponding text, which is the prior distribution in LDA to help perform topic grouping. The approach is evaluated on several different datasets, where the model is shown to outperform standard LDA and Dirichlet-multinomial regression (DMR) in terms of topic grouping, model perplexity, classification and comment generation.

Keywords. LDA, NN

1 Introduction

As the conjoint and massive knowledge in the forms of news, blogs, web pages, etc., continues to be digitized and stored, discovery becomes more and more challenging and together with it the main underlying topics. A Bayesian multinomial mixture model, latent Dirichlet allocation (LDA) [5], has recently gained much popularity due to its simplicity, and usefulness in stratifying a large collection of documents by projecting every document to a low dimensional space which is spanned by a set of bases capturing the semantic aspects of the collection. However, as the acquisition of information becomes more convenient, text data can be accompanied by extra side data. For example, when customers post their comments for products or restaurants, they usually associate a comment with a rating score or a thumb-up or thumb-down opinion, and retailers usually provide categorical labels for the products in question. In addition, customer loyalty data can be pulled in. Taking such side data into account improves the ability of LDA to discover patterns and topics among the documents. Currently, there are two types of existing models that combine side data: (1) downstream topic models and (2) upstream topic models. The downstream models assume that the text content and side data are generated simultaneously given latent topics, while upstream models assume that the text content is generated conditioned on the side data as well. Our model belongs to the family of upstream topic models where we accommodate much more complex interactions between side data and text by means of deep neural networks, when compared to other upstream topic models, i.e. DMR [13].

In this paper, we propose a new LDA-style topic model, namely hybrid neural network LDA (nnLDA), based on LDA and a neural network. Our model captures not only text content of the dataset, but also useful high-level content and secondary, non-dominant, and more salient statistical patterns from side data. Formally, the model represents the document-topic distribution as mixtures of feature-specific distributions. The prior distribution over topics is the output of a neural network whose input is side data, therefore, it is specific to each distinct combination of side data. Moreover, the neural network is optimized together with the rest of the model in a stochastic EM sampling scheme to better interpret the collection of the documents. The expectation step corresponds to finding the optimal word group and topic group while the maximization step aims to find the optimal neural network parameters and the topic-word distribution. In standard LDA, the prior is fixed while in nnLDA, it depends on a sample.

In this paper, we not only propose a more general model, nnLDA, but also present a complete technical proof confirming that nnLDA performs at least as well as plain LDA in terms of log likelihood. Furthermore, we provide an efficient variational EM algorithm for nnLDA. Lastly, we demonstrate our approach on a few real-world datasets. In summary, we make the following contributions.

- We provide a new topic model for text datasets with side data.
- We prove that the lower bound of log likelihood of nnLDA is greater than or equal to the lower bound of log likelihood of LDA for any dataset.
- We provide an efficient variational EM algorithm for nnLDA.

- We present numerical results showing that nnLDA outperforms LDA and DMR in terms of topic grouping, model perplexity, classification and text generation.

The paper is organized as follow. In the next section, we review several related works about incorporating side data in generative topic models. In Section 3, we present the hybrid neural network model, followed by our analyses of log likelihood of nnLDA and an efficient variational EM algorithm. In Section 4, we present experimental results comparing nnLDA with plain LDA and DMR.

2 Related Work

There are a large amount of extensions of the plain LDA model, however, a full retrospection of this immense literature exceeds the scope of this work. In this section, we state several kinds of variations of LDA which are most related to our new model and interpret the relationships among them.

LDA: Plain LDA, a probabilistic latent aspect model, has been widely applied on text documents [5] and [17], images [11], and network entities [1] on account of its convenience and functionality in reducing the dimensionality of the data and generating interpretable and semantically coherent topics. It is an unsupervised model which is typically constructed on a distinct bag of words from input contents. Nevertheless, in many practical applications, besides the document contents, useful side information can be easily obtained. Furthermore, such side information often provides useful high-level or direct summarization of the content, while it is not directly involved in the plain LDA model to affect topic inference. In contrast, nnLDA incorporates such information into latent aspect modeling by applying a neural network to discover secondary, non-dominant and more salient statistical patterns that may be more interesting and related to the user’s goal.

Downstream Topic Models: One approach of incorporating side data in generative topic models is to generate both the text content and side data simultaneously given latent topics. More precisely, for this type of models, each hidden topic not only has a distribution over words but also has another distribution over side data. As a consequence, in training, the loss function in optimization-based learning is the joint likelihood of the content and side data. Examples of such “downstream” models are correspondence LDA (Corr-LDA) [3], mixed-membership model for authorship [8], Group-Topic model [19], Topics over Time model (TOT) [18], and maximum entropy discrimination LDA (MedLDA) [21].

One of the most flexible downstream models is the supervised LDA (sLDA) model [4], which has variants including multi-class sLDA [16] and TOT [18]. sLDA generates side data such as customers’ ratings by maximizing the joint likelihood of the content data and the responses, where the likelihood-based objective is a generalized linear model (GLM) incorporating a proper link function with an exponential family dispersion function specified by the modeler for different types of side data. Compared to our model, in order to explicitly estimate probability distributions over all different side data, sLDA has to fully specify the link function and dispersion functions for the GLM, which increases the modeling complexity significantly. In essence, only a relatively small number of distinct side data vectors are allowed. Our model has no such restriction by applying a completely different approach.

Upstream Topic Model: In a “downstream” model, the side data is predicted based on the latent topics of the dataset, whereas in an “upstream” topic model, the side data is being conditioned on to generate the latent topics of the dataset. Another distinct difference from “downstream” topic models is the choice of the likelihood-based loss in optimization-based learning. More precisely, instead of maximizing the joint likelihood of the content and side data, an “upstream” topic model maximizes the conditional likelihood. Examples of such “upstream” topic models are Discriminative LDA (DiscLDA) [10], the scene understanding models [15] and the author-topic model [14]. In the author-topic model, words are generated by first selecting one author uniformly from an observed author list and then selecting a topic from the topic distribution with respect to that specific author. Then, given a topic, words are selected from the topic-word distribution of that topic. This model assumes that each word is generated only by one author. There are a few extensions of the author-topic model which allow a mixture of latent topics for one document and one author, i.e. [14], [12], [7]. However, these aforementioned models cannot accommodate combinations of modalities of side data, for example, the aforementioned models cannot handle categorical data and continuous data at the same time. In addition, the side data used in these models are either ratings or labels which essentially is the final intention of learning. Different from the previous models, DMR can handle combinations of modalities of side data, which uses the dot product to project the impact of side data onto the prior. An advanced version introduced in [2] is collective supervision of topic models, where aggregate-level labels are provided for groups of documents instead of individual documents, followed by a deterministic relationship between the labels and the priors. Similar to DMR, although the advanced version in [2] uses group-level labels and can handle missing side data, it also employs dot product to directly project the impact of the side data onto the prior. Compared to DMR and collective supervision of topic models, nnLDA provides a more comprehensive and flexible learning of side data by applying a neural network when compared to the dot product employed in DMR and the normal distribution assumption in collective supervision of topic models.

3 Model and Algorithm

We first present notation and the setting. We use the language of text collections throughout the paper, referring to terminologies such as “words,” “documents” and “corpus” since it makes the concepts more intuitive to understand. In general, similar to plain LDA, nnLDA is not restricted to text datasets, and can also be applied on other kinds of datasets, i.e. image datasets.

- A *word*, defined as an item from a vocabulary indexed by $\{1, \dots, V\}$, is applied one-hot encoding. More precisely, using superscripts to denote components, the v 'th word in the vocabulary is represented by a V -vector w such that $w^v = 1$ and $w^u = 0$ for all $u \neq v$.
- A *document* is a set of N words denoted by $d = \mathbf{w} = \{w_1, w_2, \dots, w_N\}$ if it only contains textual data. Similarly, if a document contains q different kinds of side data together with the aforementioned textual data, we denote it by $d = (\mathbf{w}, \mathbf{s}) = (\{w_1, w_2, \dots, w_N\}, (s_1, s_2, \dots, s_q))$ where $\mathbf{s} \in \mathbb{R}^q$.

- A *corpus* is a collection of M documents denoted by $D = \{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_M\}$ for textual only documents and $(D, S) = \{(\mathbf{w}_1, \mathbf{s}_1), (\mathbf{w}_2, \mathbf{s}_2), \dots, (\mathbf{w}_M, \mathbf{s}_M)\}$ for documents containing both side and textual data.

The main goal of nnLDA is to find a probabilistic model of a corpus that, by involving high-level summarization from side data, not only assigns high probability to documents in this corpus but also assigns high probability to other similar documents based on side data.

3.1 Generative Model

We propose the nnLDA model to explain the generative process of a document d with textual data \mathbf{w} (containing N words) and side data (structural data) \mathbf{s} , the steps of which can be summarized as follows.

1. Choose $N \sim \text{Poisson}(\xi)$.
2. Choose $\mathbf{s} \sim \mathcal{N}(\mu, \sigma^2 I)$.
3. Choose $\alpha_d = g(\gamma; \mathbf{s})$
4. Choose $\theta \sim \text{Dir}(\alpha_d)$.
5. For each of the N words w_n :
 - (a) Choose a topic $z_n \sim \text{Multinomial}(\theta)$.
 - (b) Choose a word w_n from $p(w_n | z_n, \beta)$, a multinomial probability conditioned on the topic z_n .

Notation ‘‘Poisson,’’ ‘‘Dir’’ represents the Poisson and Dirichlet distribution, respectively. In step 3, g refers to a parametric model to generate α . In summary, the model has two trainable parameters: γ , the parameters of g for side info \mathbf{s} ; β , the topic-word distribution. In the meanwhile, there are three hyper parameters: μ and σ^2 , the mean and the variance of the probability distribution for side data \mathbf{s} ; and K , which does not explicitly appear in the generative process, the number of topics.

Step 1 is independent of the remaining steps, which determines the number of words in the document. Then, for each document, step 2 provides a representation of side data \mathbf{s} by using a normal distribution with mean μ and variance σ^2 . Then, applying a model with input \mathbf{s} in step 3 provides the prior α_d for the Dirichlet distribution. Next, the random parameter of a multinomial distribution over topics, θ , is generated by the Dirichlet distribution. Finally, for the n 'th word in the document, step 5(a) first selects a topic z_n among the K different topics by the multinomial distribution with parameter θ , and then step 5(b) generates a word w_n based on the topic-word distribution β specific to topic z_n . Step 5 follows standard LDA.

3.2 Analysis

Note that a Dirichlet random vector $\theta = (\theta_1, \theta_2, \dots, \theta_K)$ has the following probability density:

$$p(\theta | \alpha) = \frac{\Gamma\left(\sum_{i=1}^K \alpha_i\right)}{\prod_{i=1}^K \Gamma(\alpha_i)} \theta_1^{\alpha_1-1} \dots \theta_K^{\alpha_K-1},$$

where K is the number of topic groups, α is the prior of the Dirichlet distribution and θ takes values in the $(K-1)$ -simplex. Then, the generative process implies that the conditional distribution of the nnLDA model of a document $d = (\mathbf{w}, \mathbf{s})$ is

$$\begin{aligned} P_1(\mathbf{w} | \mu, \sigma, \gamma, \beta) &= \tilde{P}_1(\mathbf{w} | \mathbf{s}, \gamma, \beta) \\ &= \int \tilde{p}(\theta | \mathbf{s}, \gamma) \left(\prod_{n=1}^N \sum_{z_k} \tilde{p}(z_k | \theta) \tilde{p}(w_n | z_k, \beta) \right) d\theta \\ &= \int \tilde{p}(\theta | \mu, \sigma, \gamma) \left(\prod_{n=1}^N \sum_{z_k} \tilde{p}(z_k | \theta) \tilde{p}(w_n | z_k, \beta) \right) d\theta \\ &= \int \tilde{p}(\theta | \mu, \sigma, \gamma) \left(\prod_{n=1}^N \sum_{i=1}^K \prod_{j=1}^V (\theta_i \beta_{ij})^{w_n^j} \right) d\theta, \end{aligned}$$

which in turn yields

$$\begin{aligned} P_1(D | \mu, \sigma, \gamma, \beta) &= \mathbb{E} \left[\int \tilde{p}(\theta_d | \mu, \sigma, \gamma) \left(\prod_{n=1}^N \sum_{z_{d_k}} \tilde{p}(z_{d_k} | \theta_d) \tilde{p}(w_{d_n} | z_{d_k}, \beta) \right) d\theta_d \right] \\ &= \mathbb{E} \left[\int \tilde{p}(\theta_d | \mu, \sigma, \gamma) \left(\prod_{n=1}^N \sum_{i=1}^K \prod_{j=1}^V (\theta_i \beta_{ij})^{w_n^j} \right) d\theta_d \right], \end{aligned}$$

where $\tilde{p}(\theta_d | \mu, \sigma, \gamma) = \tilde{p}(\theta_d | \mathbf{s}, \gamma) = p(\theta_d | g(\gamma; \mathbf{s})) = p(\theta_d | \alpha_d)$ for a corpus D .

The nnLDA model represented above is a probabilistic graphical model with three levels. Parameters μ , σ , γ and β are corpus-level parameters, which are assumed to be sampled once in the generative process of a corpus. Variables α_d and θ_d are document-level variables, which are sampled once per document. Finally, w_{d_n} and z_{d_k} are word-level variables, sampled once for each word in each document.

In the rest of this section, we provide an analytical comparison of standard LDA and nnLDA.

Compared to standard LDA, nnLDA employs an extra neural network g to generate document-level variable α_d . Since nnLDA is “richer” than LDA, we expect that it should produce a higher likelihood. Without assumptions on $g(\gamma; \cdot)$ this does not hold since, for example, $g(\gamma; \cdot)$ can map everything to a constant vector different from the prior used by LDA. As a result, in order for the statement to hold the network must be expressive. The question to consider is whether a neural network is capable of memorizing arbitrary side data of a given size. We tackle this question by introducing the concept of finite sample expressivity which is an extension of a similar definition in [20]. Given the definition, if $g(\gamma; \cdot)$ has finite sample expressivity, nnLDA at least can find the optimal α^* used in standard LDA.

Definition 1. Function $g(\gamma; \cdot)$ has finite sample expressivity if for all inputs $x_i \in \mathbb{R}^{d_x}, 1 \leq i \leq N$ and for all $y_i \in [-M, +M]^{d_y}, 1 \leq i \leq N$ for some constant $M > 0$, there exists a parameter γ such that $g(\gamma; x_i) = y_i$ for every $1 \leq i \leq N$.

Based on Definition 1, Theorem 3.1 shown in [20] provides a specific set of constraints, i.e. any 3-layer (i.e., 2-hidden-layer) ReLU FCNN with hidden layer widths d_1 and d_2 can fit any arbitrary dataset if $d_1 d_2 \geq 4Nd_y$, where d_y and N are the dimension of the label and the number of samples, respectively. By extending the aforementioned theorem, Proposition 3.4 and Theorem 4.1 in [20] argue that any FCNN given constraints on the number of neurons in each layer is able to have finite sample expressivity.

In the following, we assume that $g(\gamma; \cdot)$ has finite sample expressivity. Therefore, given K and any α^* representing the number of topic groups and optimal parameters in LDA, since $\alpha^* \in [-M, +M]^K$ for some constant M , there exists a γ_1 such that, for all inputs \mathbf{s}_i and α^* , $g(\gamma_1; \mathbf{s}_i) = \alpha^*$ for all $1 \leq i \leq N$.

We next prove that the optimized probability of nnLDA is at least as good as that of plain LDA. Let α^* and β^* be optimal solutions to $P_2 = \max_{\alpha, \beta} P(D | \alpha, \beta)$ of LDA, meanwhile, let μ^*, σ^* and γ^* be optimal solutions to $P_1 = \max_{\mu, \sigma, \gamma} P_1(D | \mu, \sigma, \gamma, \beta^*)$ of nnLDA (see Appendix A for formal definitions).

Theorem 1. If α^*, β^* are optimal solutions to LDA, then there exists optimal solutions μ^*, σ^* and γ^* to nnLDA such that

$$P_1(D | \mu^*, \sigma^*, \gamma^*, \beta^*) \geq P_2(D | \alpha^*, \beta^*).$$

Proof. See Appendix B. □

While Theorem 1 asserts that when it comes to model fit nnLDA fits the data better than LDA, it does not provide a gap statement. If the side data provides positive influence during the learning process by a constant C , then, due to the independence of words, topics and documents, we are able to argue that the optimized probability is at least improved by $C - 1$.

Theorem 2. For any document $(\mathbf{w}, \mathbf{s}) \in (D, S)$, if $\hat{p}(w_i | \alpha^*, \beta^*) \neq 0$ for all i , and there exists a positive constant $C > 1$ such that $\prod_{i=1}^N \tilde{p}(w_i | \gamma^*, \beta^*, \mu^*, \sigma^*) \geq C \prod_{i=1}^N \hat{p}(w_i | \alpha^*, \beta^*)$ for every $w_i \in \mathbf{w}$, and if D in P_1 and D in P_2 follow the same distribution, then

$$\frac{P_1(D | \mu^*, \sigma^*, \gamma^*, \beta^*) - P_2(D | \alpha^*, \beta^*)}{P_2(D | \alpha^*, \beta^*)} \geq C - 1.$$

Proof. See Appendix C for a formal proof. □

The assumption on $\hat{p}(w_i | \alpha^*, \beta^*)$ in Theorem 2 is reasonable since it indicates that all documents are not randomly generated. The positive constant C in the assumption captures the improvement given by the side data. In other words, as long as the side data has positive impact on the text data, this assumption holds. Next, we link the existence of C to lift from data mining. Let us define lift as

$$l(d) = \frac{P(\mathbf{w})P(\mathbf{s})}{P(\mathbf{w}, \mathbf{s})}$$

with $d = (\mathbf{w}, \mathbf{s})$. Lift measures the dependency level of words \mathbf{w} and side data \mathbf{s} . If $l(d) < 1$ for d with N words and $P(\mathbf{s}) > 0$, we have

$$P(\mathbf{s}) \prod_{n=1}^N P(w_n) = P(\mathbf{w})P(\mathbf{s}) < P(\mathbf{w}, \mathbf{s}) = P(\mathbf{s}) \prod_{n=1}^N P(w_n|\mathbf{s}),$$

and in turn

$$\prod_{n=1}^N P(w_n) < \prod_{n=1}^N P(w_n|\mathbf{s}),$$

and

$$\prod_{n=1}^N \hat{p}(w_n | \alpha^*, \beta^*) < \prod_{n=1}^N \tilde{p}(w_n | \gamma^*, \beta^*, \mu^*, \sigma^*).$$

This implies that there exists $C > 1$. In summary, when $l(d) < 1$ and $P(\mathbf{s}) > 0$ for each d in the corpus, Theorem 2 holds. Lift essentially measures the dependency of \mathbf{w} and \mathbf{s} , which is widely used in data mining. The condition indicates that the side data helps to link the words to the documents they are more likely to be in.

Informally, in the proof, due to the independence assumption of words, topics and documents in nnLDA, the generative probability of nnLDA for a corpus can be reformulated as a product of $\tilde{p}(\theta_d | \mu^*, \sigma^*, \gamma^*)$ and conditional probability of words $\tilde{p}(w_n | \theta_d, \beta^*)$. Likewise, the same property holds for plain LDA. Lastly, given a relationship between documents $d = \mathbf{w}$ and $d = (\mathbf{w}, \mathbf{s})$ as an expression of the conditional probability of words, we are able to build a connection of the optimized probabilities between nnLDA and LDA.

3.3 Variational Inference with EM Algorithm

We train the nnLDA model using a stochastic EM sampling scheme, in which we alternate between sampling topic assignments from the current prior distribution conditioned on the observed words and side data, and optimizing the parameters given the topic assignments.

Details are similar to those in [5]. In this section, instead of showing all the details, we only point out the differences from the derivation of plain LDA. By applying the Jensen's inequality and KL divergence between the variational posterior probability and the true posterior probability, which is a formally stated technique in [5], a lower bound of log likelihood reads

$$\begin{aligned} L(\xi, \phi; \gamma, \beta) &= \mathbb{E}_q [\log p(\theta | g(\gamma; \mathbf{s}))] + \mathbb{E}_q [\log p(z | \theta)] + \mathbb{E}_q [\log p(\mathbf{w} | z, \beta)] \\ &\quad - \mathbb{E}_q [\log q(\theta)] - \mathbb{E}_q [\log q(z)], \end{aligned} \tag{1}$$

where ξ, ϕ are variational parameters of θ and z , respectively, and $q(\cdot)$ represents the variational distribution. Then, the iterative algorithm is

1. (E-step) For each document, find the optimizing values of the variational parameters ξ and ϕ of z and θ , respectively.

2. (M-step) Maximize the resulting lower bound of log likelihood with respect to the model parameters γ and β .

The E-step is similar to the E-step in [5] except replacing prior α by $g(\gamma; \mathbf{s})$. We run the E-step until it converges for each document. The M-step is finding a maximum likelihood estimation with expected sufficient statistics for each document under the approximate posterior parameters ξ and ϕ , which are computed in the E-step. Likewise, since the log likelihood objective related to β does not involve $g(\gamma; \mathbf{s})$, we are allowed to directly borrow the update rule of β from [5], which is

$$\beta_{ij} \propto \sum_{d=1}^M \sum_{n=1}^N \phi_{dni}^* w_{dn}^j.$$

In contrast, for the neural network parameter γ , we resort to log likelihood objective related to γ as follows,

$$L_{[\gamma]} = \sum_{d=1}^M \left(\log \Gamma \left(\sum_{j=1}^K [g(\gamma; \mathbf{s}_d)]_j \right) - \sum_{i=1}^K \log \Gamma([g(\gamma; \mathbf{s}_d)]_i) + \sum_{i=1}^K \left(([g(\gamma; \mathbf{s}_d)]_i - 1) \left(\Psi(\xi_{di}) - \Psi \left(\sum_{j=1}^K \xi_{dj} \right) \right) \right) \right),$$

where M is the number of documents in the corpus, and Ψ is the digamma function, the first derivative of the log Gamma function. Then, applying the backpropagation approach provides the derivative and the update rule for parameter γ .

4 Experimental Study

In this section, we compare the nnLDA model with standard LDA and the DMR model introduced in [5] and [13], respectively. We conduct experiments on five different-size datasets among which one is a synthetic dataset and the remaining four are real-world datasets. For these datasets, we study the performance of topic grouping, perplexity, classification and comment generation for nnLDA, plain LDA and DMR models. For each of the tasks, some datasets are not eligible to be examined due to lack of information. The synthetic dataset is publically available at https://github.com/biyifang/nnLDA/blob/main/syn_file.csv while the real-world datasets are proprietary.

4.1 Datasets and Training Details

The first dataset we use is a synthetic dataset of 2,000 samples. Each sample contains a customer’s feedback with respect to his or her purchase along with the characteristics of the product. More precisely, there are two different categories, which are product and description. In the product category, it can either be TV or burger; similarly, in the description category, the word can either be price or quality. In order to generate comments, we assign a bag of words to each combination of product and description as shown in Table 1.

Category combination	Bag of words
(burger, price)	value, pricey, ouch, steep, cheap, value, reason, accept, unreason, unacceptable
(burger, quality)	nasty, fantastic, delicious, tasty, juicy, unreason, unacceptable, reason, accept, fresh
(TV, price)	promotion, affordable, value, increase, expensive, tasty, economical, fancy, okay
(TV, quality)	fabulous, fantastic, promising, sharp, large, clear, eco friendly, fresh, pixilated

Table 1: Synthetic Dataset

After randomly selecting one category combination from the four combinations, a comment is generated containing at least one word and at most five words with an average 2.97 words, by selecting a certain number of words at random from the corresponding bag.

The second dataset is a real-world dataset, PTS for short, which has 795 samples. Each sample contains a customer’s short feedback and rating with respect to his or her purchase along with the characteristics of the product. Additionally, the category (side data) selected for nnLDA corresponds to sectors, which are generalizations of products. In this dataset, there is only only 1 word in the shortest comment, while the longest comment in the dataset contains 49 words. Overall, the average length of the comments is 10.6 words. For example, a customer, who bought a product belonging to sector Baby, leaves a comment “Cheap& Soft” with a rating of 3.

The third dataset WIP is a medium-size dataset with 3,451 samples. Each sample contains a customer’s short feedback and rating with respect to his or her purchase along with the characteristics of the product. The sector attribution is again side data when training models with one feature. The other attribution counted for models with two features is channel. The most concrete comment in the dataset has 138 words, while the briefest comment has only 1 word. In the meanwhile, the average length of the comments in the dataset is 8.9 words.

DCL is another medium-size dataset of 5,427 samples. Different from the PTS and WIP datasets, each sample in DCL contains a customer’s long feedback and rating with respect to his or her purchase along with the characteristics of the product. Additionally, the side data selected for nnLDA corresponds to groups of products. The smallest number of words for a comment in this dataset is 1, while the largest is 988. Overall, the average length of the comments is 61.7 words. A short sample comment is “quick points that will be all that matters to a buyer wanting accurate metrics to buy by tinny sound but plenty of audio hookups.”

The last dataset is RR, which has 100,000 samples, from which we randomly select 10,000 samples. Each sample contains a customer’s feedback with respect to his or her purchase along with the characteristics of the product. Additionally, the side data for nnLDA corresponds to the category, which can be grocery, health and personal care, furniture, kitchen, etc. The longest comment has 418 words while the shortest comment has only 1 word as the previous datasets, and the average length of the comments is 69.5 words. A short sample comment reads “great tasting oil and made the most excellent gluten free chocolate

cake.”

Due to the lack of some information from the certain datasets, we are unable to study all tasks of interest for all of these datasets. For the topic grouping task, we examine the ability of nnLDA, plain LDA and DMR to assign the comments from the same topic group into the same correct topic group. For this task, we only conduct experiments on the synthetic dataset since only the topic groups of the synthetic dataset are clear. For the perplexity task, we compute the logarithm of the perplexity of all the words in the corresponding dataset. We do not study the performance of perplexity for the synthetic dataset since we know the true number of topic groups. For the classification task, we use the probability vector generated by the topic models to predict the rating for that comment. Since the RR dataset does not have ratings, we are unable to examine the classification ability of the topic models on the RR dataset. The last task tests the performance of the topic models on generating new comments. For this task, we only conduct experiments on the two smallest real world datasets since it is of interest how topic models perform given a small number of samples. Table 2 presents the tasks of interest for each dataset.

Dataset	Topic grouping	Perplexity	Classification	Comment generation
Synthetic dataset	Yes	No	No	No
PTS	No	Yes	Yes	Yes
WIP	No	Yes	Yes	Yes
DCL	No	Yes	Yes	No
RR	No	Yes	No	No

Table 2: Tasks of Interest

For all of these datasets, we employ a two-layer fully connected neural network as $g(\gamma; \cdot)$ in nnLDA. Furthermore, we set the number of neurons to be 20 in the first layer, the number of neurons of the second layer to be the number of topic groups assigned in the beginning and the batch size to be 64. All features of the side data are categorical and are one-hot encoded. Additionally, all weights in $g(\gamma; \cdot)$ are initialized by Kaiming Initialization [9]. We apply the ADAM algorithm with the learning rate of 0.001 and weight decay being 0.1. Meanwhile, we train all the models using EM with exactly the same stopping criteria of stopping E-step and M-step when the average change over the whole training dataset in the expected log likelihood becomes less than 0.01%. We vary the number of topic groups from 4 to 30. For DMR, we use the same values for the parameters as those in [13]. All the algorithms are implemented in Python with Pytorch and trained on a single GPU card.

4.2 Experimental Results

In this section, we present all the results based on the tasks of interest.

Overall, nnLDA outperforms plain LDA and DMR in all datasets in terms of topic grouping, classification, perplexity and comment generation. Meanwhile, based on the fact that the last two datasets have many more words and more intrinsic concepts in their comments

when compared to the first three datasets, nnLDA exceeds the performance of plain LDA and DMR dramatically when a document contains several topics or it is more comprehensive.

4.2.1 Topic Grouping

Table 3 shows the most frequent 5 words in each topic group generated by plain LDA, DMR and nnLDA when setting the number of topic groups to be 4 in the synthetic dataset. The topic groups generated by plain LDA and DMR are very vague and it is very hard to distinguish which topic group is describing what combination of product and description, while the topic groups given by nnLDA are very distinguishable, i.e. topic group 1 is about (burger, quality), topic group 2 is about (TV, price), topic group 3 is about (TV, quality) and topic group 4 is about (burger, price). It identifies correctly the seed topics. Therefore, nnLDA outperforms plain LDA in grouping.

	plain LDA	DMR	nnLDA
Topic group 1	promising, rebate, sharp, increase, outstanding	pricey, unacceptable, juicy, pixilated	unreason, unacceptable, juicy delicious, nasty
Topic group 2	unreason, value, okay, steep, ecofriendly	ouch, steep, tasty, unreason, promotion	promotion, increase, tasty, economical, okay
Topic group 3	reason, accept, promotion, large, unacceptable	accept, fantastic, value reason, affordable	fresh, promising, fantastic, large, eco friendly
Topic group 4	fresh, reason, outstanding, ecofriendly, fantastic	sharp, delicious, accept, fresh, clear	reason, accept, value, steep, cheap

Table 3: Top words of groups generated by LDA, DMR and nnLDA

	macro-precision	macro-recall	macro-F1	micro-F1
LDA	0.7238	0.7272	0.7211	0.7240
DMR	0.7238	0.7460	0.7313	0.7392
nnLDA	0.7401	0.7919	0.7536	0.7905
relative improvement from LDA	2.25%	8.90%	4.51%	9.19%
relative improvement from DMR	2.25%	6.15%	3.05%	6.94%

Table 4: Precision, recall and relative improvement of the synthetic dataset generated by LDA, DMR and nnLDA

Additionally, based on the top words of topics generated by LDA, DMR and nnLDA, we are able to assign the most related category combination to a comment with respect to a model. Since we have the category combination of each comment, Table 4 shows the macro-recall, macro-precision and macro-F1 scores and micro-F1 of LDA, DMR and nnLDA, respectively, when training on the synthetic dataset, and the overall relative improvement of nnLDA. As the table shows, nnLDA outperforms plain LDA and DMR, which implies that nnLDA assigns more samples correctly to the right topic group. Therefore, in general, nnLDA improves the recall, precision and F1 scores.

In conclusion, nnLDA outperforms standard LDA and DMR in terms of the ability of topic grouping.

4.2.2 Perplexity

Figures 1 and 2 represent the $\log(\text{perplexity})$ of plain LDA, DMR and nnLDA on the PTS and WIP datasets, respectively. Additionally, in Figure 2, for DMR and nnLDA, we not only conduct experiments on the dataset with the single feature (sector) as the side data, denoted as “DMR with single feature” and “nnLDA with single feature,” but also on the dataset with two features (sector and channel) as side data, denoted as “DMR with two features” and “nnLDA with two features,” respectively. The smallest $\log(\text{perplexity})$ values generated by plain LDA and DMR are competitive to those of nnLDA for these two datasets. In Figure 1, the $\log(\text{perplexity})$ value generated by plain LDA increases as the number of topic groups grows, while the $\log(\text{perplexity})$ values generated by DMR and nnLDA decrease first and then increase as the number of topic groups increases on the PTS dataset. As it is shown in Figure 2, the $\log(\text{perplexity})$ values generated by plain LDA and DMR increase as the number of topic groups grows on the WIP dataset. However, the $\log(\text{perplexity})$ values generated by nnLDA decrease first and then increase as the number of topic groups increases on both of the aforementioned datasets. Moreover, we examine DMR and nnLDA models with two features on the WIP dataset, which take both sector and channel attributions as side data into account, in Figure 2. As we can observe, the minimum $\log(\text{perplexity})$ generated by nnLDA with two features (sector and channel attributions) is better than that of nnLDA with the single feature (sector attribution), although the optimal number of topic groups occurs at a different point since more side data is provided. Consequently, plain LDA does not learn the datasets, and DMR is able to learn the small datasets. In contrast, nnLDA starts learning the datasets as the $\log(\text{perplexity})$ value decreases in the beginning and finds an optimal number of topic groups, then it gets confused since the number of topic groups are more than needed. Furthermore, nnLDA with two features provides better $\log(\text{perplexity})$ than nnLDA with the single feature. Therefore, nnLDA is more capable of understanding the datasets; both small and medium-size datasets with short comments.

When learning more complex datasets, the advantage of the nnLDA model becomes more pronounced. Figure 3 represents the $\log(\text{perplexity})$ of plain LDA, DMR and nnLDA on the DCL dataset, while Figure 4 shows the same on the RR dataset. In these figures we observe that the $\log(\text{perplexity})$ generated by plain LDA and DMR blows up as the number of topic groups increases, while the $\log(\text{perplexity})$ generated by nnLDA decreases first and then increases as the number of topic groups grows. Furthermore, the $\log(\text{perplexity})$ values of nnLDA are much smaller than those of plain LDA and DMR. Consequently, nnLDA performs as well as plain LDA and DMR in small and medium-size datasets with short comments, and at the same time, nnLDA explains the datasets better than plain LDA and DMR in medium and large size datasets with long comments. There is also a trade-off between the accuracy and running time as shown in Table 5. In the table, we compare the running time of plain LDA, DMR with one feature and nnLDA with one feature on three different datasets. We observe that nnLDA spends more time than both DMR and plain LDA on training. In conclusion, nnLDA performs better on learning while it requires a slightly longer training time. It is less than 10% slower than DMR.

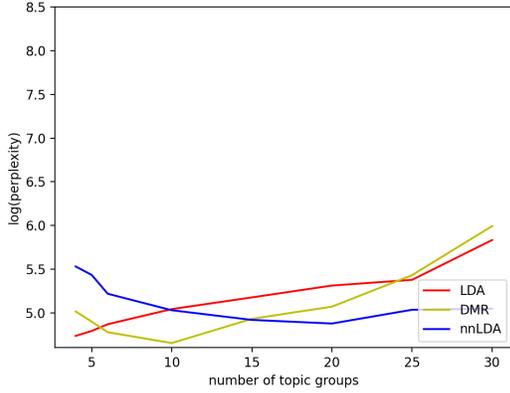


Figure 1: PTS dataset

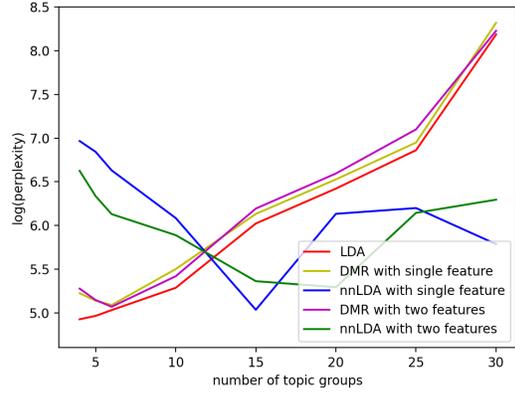


Figure 2: WIP dataset

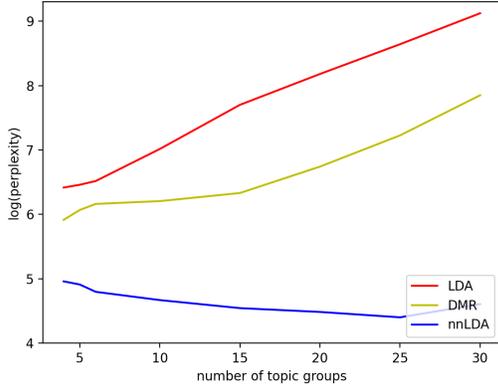


Figure 3: DCL dataset

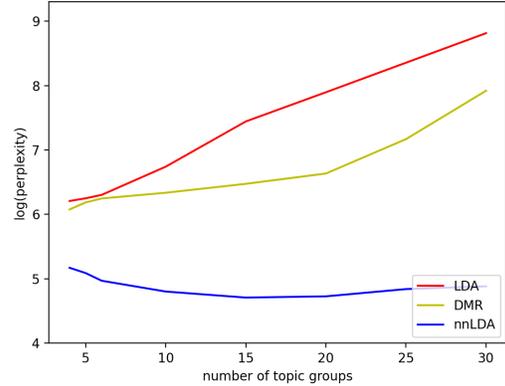


Figure 4: RR dataset

running time(s)	plain LDA	DMR	nnLDA
PTS	3	4	4
WIF	19	24	26
DCL	138	179	191

Table 5: Running time of different models on different datasets

In the following section, we study the classification problem of predicting the rating of each sample. In all the cases, we use 10-fold cross validation, which holds out 10% of the data for test purposes and trains the models on the remaining 90%. We apply nnLDA, plain LDA and DMR to find the probability of each sample to be assigned to each topic group and treat it as the feature matrix. Lastly, we train a classification model (xgboost [6]) on the feature matrix with the rating labels as the ground truth.

4.2.3 Classification

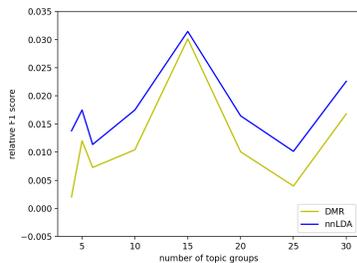


Figure 5: PTS

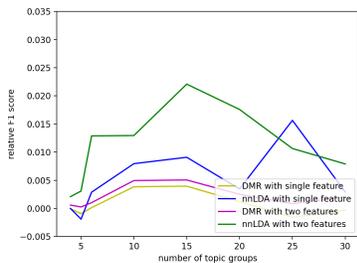


Figure 6: WIP

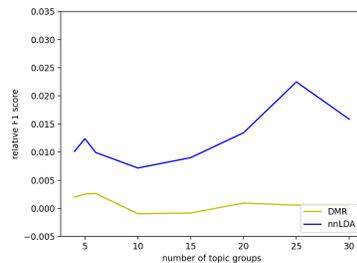


Figure 7: DCL

Figures 5, 6 and 7 depict the relative F1 scores of DMR and nnLDA with respect to plain LDA on the PTS, WIP, and DCL datasets, respectively. In Figure 5, the most distinguishable difference of F1 scores occurs when the number of topic groups is 15, where nnLDA has a gap of 0.032. In the meanwhile, DMR achieves its best performance at the same point with a gap of 0.030. Moreover, this chart shows that nnLDA outperforms plain LDA and DMR no matter what the number of topic groups is. In Figure 6, when using the single feature (sector attribution), the biggest gaps of F1 scores happen when the number of topic groups is 15 for DMR and 25 for nnLDA. The biggest gap between nnLDA and plain LDA is 0.016, while the largest gap between DMR and plain LDA is 0.003. Considering models using two features (sector and channel attributions) as the side data, the highest relative F1 score given by nnLDA with two features is 0.022 with 15 topic groups, compared with 0.004 produced by DMR with 10 topic groups. Although plain LDA provides a slightly higher F1 score than nnLDA when applying 5 topic groups, nnLDA outperforms plain LDA and DMR significantly given any other number of topic groups. In Figure 7, the highest relative F1 score given by nnLDA is 0.022 with 25 topic groups, compared with 0.003 given by DMR for 6 topic groups. Moreover, this figure shows that nnLDA outperforms plain LDA dramatically whatever the number of topic groups is.

Therefore, nnLDA performs better than plain LDA and DMR when predicting the rating given customer’s comments and product information in all datasets.

4.2.4 Comment Generation

In this section, we compare the comments generated by nnLDA with plain LDA and DMR. We set the number of topic groups to be 5 since all of plain LDA, DMR and nnLDA have relatively low perplexity scores based on Figures 1 and 2, and comparable F1 scores based on Figures 5 and 6 on the PTS and WIP datasets. A comment is generated based on the topic-document probability of the sample and the topic-word distribution. More precisely, for DMR and LDA, the prior α is generated based on the side data (sector) first while α is fixed in plain LDA. Next, a comment is created by selecting the top words which have the highest score computed by adding the products of the topic-document probability and topic-word for each word. Then, we randomly pick 50 comments that contain a certain level of information, for example, we rule out comments like “N/A.” Meanwhile, in order

to evaluate the quality of comment generation, we employed 50 PhD students. Each one of them assessed a pair of comments (one based on plain LDA or DMR, and the other one based on nnLDA) for the same side data and they provided an assessment as to which one is better.

	Number of generated comments	
	PTS	WIP
plain LDA < nnLDA	15	22
plain LDA > nnLDA	11	9
plain LDA \sim nnLDA	24	19
DMR < nnLDA	16	20
DMR > nnLDA	11	10
DMR \sim nnLDA	23	20

Table 6: Comparison of the generated comments on different datasets

The upper left three values in Table 6 show the comparison of the generated comments given by plain LDA and nnLDA on the PTS dataset. Based on the table, among all these 50 samples, nnLDA generates more accurate comments in 15 samples, while plain LDA does better in 11 samples, and the two are tied for the remaining 24 samples. The lower left three values in Table 6 show the comparison of the generated comments given by DMR and nnLDA on the PTS dataset. Based on the table, among all these 50 samples, nnLDA generates more accurate comments in 16 samples, while DMR does better in 11 samples, and the two are tied for the remaining 23 samples. On the PTS dataset, nnLDA generates in $\frac{15-11}{50} = 8\%$ more reasonable comments compared to plain LDA, and in $\frac{16-11}{50} = 10\%$ more comparing to DMR.

The right column in Tables 6 shows the comparison of the generated comments given by plain LDA and nnLDA, and DMR and nnLDA on the WIP dataset, respectively. The observations and conclusions are similar. Furthermore, the advantage in number is more obvious on the WIP dataset, i.e. the improvement of nnLDA compared to plain LDA is as large as $\frac{22-9}{50} = 26\%$ and the improvement from DMR to nnLDA is $\frac{20-10}{50} = 20\%$. Therefore, taking generated comments into consideration, nnLDA generates more reasonable comments than plain LDA and DMR for both small and medium-sized datasets.

References

- [1] Edoardo M. Airoldi, David M. Blei, Stephen E. Fienberg, and Eric P. Xing. Mixed membership stochastic blockmodels. *Journal of machine learning research*, 9:1981–2014, 2008.
- [2] Adrian Benton, Michael J. Paul, Braden Hancock, and Mark Dredze. Collective supervision of topic models for predicting surveys with social media. In *AAAI*, 2016.
- [3] David M. Blei and Michael I. Jordan. Modeling annotated data. In *SIGIR*, 2003.
- [4] David M. Blei and Jon D. McAuliffe. Supervised topic models. In *NIPS*, 2007.
- [5] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent Dirichlet allocation. *Journal of machine learning research*, 3:993–1022, 2003.
- [6] Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD*, 2016.
- [7] Laura Dietz, Steffen Bickel, and Tobias Scheffer. Unsupervised prediction of citation influences. In *ICML*, 2007.
- [8] Elena A. Erosheva, Stephen E. Fienberg, and John Lafferty. Mixed-membership models of scientific publications. *Proceedings of the National Academy of Sciences of the United States of America*, 101:5220 – 5227, 2004.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *Proceedings of the IEEE international conference on computer vision*, 2015.
- [10] Simon Lacoste-Julien, Fei Sha, and Michael I. Jordan. DiscLDA: Discriminative learning for dimensionality reduction and classification. In *NIPS*, 2008.
- [11] Fei-Fei Li and Pietro Perona. A Bayesian hierarchical model for learning natural scene categories. In *ICCV*, 2005.
- [12] Andrew McCallum, Xuerui Wang, and Andrés Corrada-Emmanuel. Topic and role discovery in social networks with experiments on Enron and academic email. *Journal of artificial intelligence research*, 30:249–272, 2007.
- [13] David Mimno and Andrew McCallum. Topic models conditioned on arbitrary features with Dirichlet-multinomial regression. In *UAI*, 2008.
- [14] Michal Rosen-Zvi, Thomas L. Griffiths, Mark Steyvers, and Padhraic Smyth. The author-topic model for authors and documents. *ArXiv*, abs/1207.4169, 2004.
- [15] Erik B. Sudderth, Antonio Torralba, William T. Freeman, and Alan S. Willsky. Learning hierarchical models of scenes, objects, and parts. In *ICCV*, 2005.

- [16] Chong Wang, David M. Blei, and Fei-Fei Li. Simultaneous image classification and annotation. In *ICCV*, 2009.
- [17] Xingyu Wang, Lida Zhang, and Diego Klabjan. Keyword-based topic modeling and keyword selection. *ArXiv*, abs/2001.07866, 2020.
- [18] Xuerui Wang and Andrew McCallum. Topics over time: a non-Markov continuous-time model of topical trends. *Proceedings of the 12th ACM SIGKDD*, 2006.
- [19] Xuerui Wang, Natasha Mohanty, and Andrew McCallum. Group and topic discovery from relations and their attributes. In *NIPS*, 2005.
- [20] Chulhee Yun, Suvrit Sra, and Ali Jadbabaie. Small ReLU networks are powerful memorizers: a tight analysis of memorization capacity. In *NIPS*, 2019.
- [21] Jun Zhu, Amr Ahmed, and Eric P. Xing. MedLDA: maximum margin supervised topic models. *Journal of machine learning research*, 13:2237–2278, 2012.

5 Appendix

A Probability Distribution of LDA

Given the generative process of LDA, which is formally presented in [5], we obtain the marginal distribution of a document $d = \mathbf{w}$ with text only as

$$\begin{aligned} P_2(\mathbf{w} \mid \alpha, \beta) &= \int \hat{p}(\theta \mid \alpha) \left(\prod_{n=1}^N \sum_{z_k} \hat{p}(z_k \mid \theta) \hat{p}(w_n \mid z_k, \beta) \right) d\theta \\ &= \int \hat{p}(\theta \mid \alpha) \left(\prod_{n=1}^N \sum_{i=1}^K \prod_{j=1}^V (\theta_i \beta_{ij})^{w_n^j} \right) d\theta, \end{aligned}$$

which in turn yields

$$\begin{aligned} P_2(D \mid \alpha, \beta) &= \mathbb{E} \left[\int \hat{p}(\theta_d \mid \alpha) \left(\prod_{n=1}^N \sum_{z_{d_k}} \hat{p}(z_{d_k} \mid \theta_d) \hat{p}(w_{d_n} \mid z_{d_k}, \beta) \right) d\theta_d \right] \\ &= \mathbb{E} \left[\int \hat{p}(\theta_d \mid \alpha) \left(\prod_{n=1}^N \sum_{i=1}^K \prod_{j=1}^V (\theta_i \beta_{ij})^{w_n^j} \right) d\theta_d \right], \end{aligned}$$

where $\hat{p}(\theta_d \mid \alpha) = p(\theta_d \mid \alpha)$.

B Proof of Theorem 1

Proof. By finite sample expressivity of $g(\gamma; \cdot)$, there exists a model with parameters γ_1 such that

$$g(\gamma_1; \mathbf{s}) = \alpha^*,$$

which in turn yields

$$\tilde{\hat{p}}(\theta \mid \mathbf{s}, \gamma_1) = \hat{p}(\theta \mid g(\gamma_1; \mathbf{s})) = \hat{p}(\theta \mid \alpha^*).$$

Therefore,

$$P_2(D \mid \alpha^*, \beta^*) = \tilde{\tilde{P}}_1(D \mid S, \gamma_1, \beta^*) = P_1(\mu^*, \sigma^*, \gamma_1, \beta^*).$$

Since nnLDA also optimizes over the network parameter γ , we have

$$P_1(D \mid \mu^*, \sigma^*, \gamma^*, \beta^*) \geq P_1(D \mid \mu^*, \sigma^*, \gamma_1, \beta^*),$$

and thus,

$$P_1(D \mid \mu^*, \sigma^*, \gamma^*, \beta^*) \geq P_2(D \mid \alpha^*, \beta^*).$$

□

C Proof of Theorem 2

Proof. Note that

$$\begin{aligned}
& \frac{P_1(D \mid \mu^*, \sigma^*, \gamma^*, \beta^*) - P_2(D \mid \alpha^*, \beta^*)}{P_2(D \mid \alpha^*, \beta^*)} \\
&= \frac{\mathbb{E} \left[\int \tilde{p}(\theta_d \mid \mu^*, \sigma^*, \gamma^*) \left(\prod_{n=1}^N \sum_{z_{d_k}} \tilde{p}(z_{d_k} \mid \theta_d) \tilde{p}(w_{d_n} \mid z_{d_k}, \beta^*) \right) d\theta_d \right]}{\mathbb{E} \left[\int \hat{p}(\theta_d \mid \alpha^*) \left(\prod_{n=1}^N \sum_{z_{d_k}} \hat{p}(z_{d_k} \mid \theta_d) \hat{p}(w_{d_n} \mid z_{d_k}, \beta^*) \right) d\theta_d \right]} \\
&= \frac{\mathbb{E} \left[\int \tilde{p}(\theta_d \mid \mu^*, \sigma^*, \gamma^*) \left(\prod_{n=1}^N \sum_{z_{d_k}} \tilde{p}(z_{d_k} \mid \theta_d) \tilde{p}(w_{d_n} \mid z_{d_k}, \beta^*) \right) d\theta_d \right]}{\mathbb{E} \left[\int \hat{p}(\theta_d \mid \alpha^*) \left(\prod_{n=1}^N \sum_{z_{d_k}} \hat{p}(z_{d_k} \mid \theta_d) \hat{p}(w_{d_n} \mid z_{d_k}, \beta^*) \right) d\theta_d \right]} \tag{2}
\end{aligned}$$

Since

$$\begin{aligned}
& \tilde{p}(\theta_d \mid \mu^*, \sigma^*, \gamma^*) \left(\prod_{n=1}^N \sum_{z_{d_k}} \tilde{p}(z_{d_k} \mid \theta_d) \tilde{p}(w_{d_n} \mid z_{d_k}, \beta^*) \right) \\
&= \tilde{p}(\theta_d \mid \mu^*, \sigma^*, \gamma^*) \left(\prod_{n=1}^N \tilde{p}(w_{d_n} \mid \theta_d, \beta^*) \right) = \prod_{n=1}^N \tilde{p}(w_{d_n} \mid \gamma^*, \beta^*, \mu^*, \sigma^*)
\end{aligned}$$

and

$$\begin{aligned}
& \hat{p}(\theta_d \mid \alpha^*) \left(\prod_{n=1}^N \sum_{z_{d_k}} \hat{p}(z_{d_k} \mid \theta_d) \hat{p}(w_{d_n} \mid z_{d_k}, \beta^*) \right) \\
&= \hat{p}(\theta_d \mid \alpha^*) \left(\prod_{n=1}^N \hat{p}(w_{d_n} \mid \theta_d, \beta^*) \right) = \prod_{n=1}^N \hat{p}(w_{d_n} \mid \alpha^*, \beta^*),
\end{aligned}$$

equation (2) could be further simplified as

$$\begin{aligned}
& \frac{P_1(D \mid \mu^*, \sigma^*, \gamma^*, \beta^*) - P_2(D \mid \alpha^*, \beta^*)}{P_2(D \mid \alpha^*, \beta^*)} \\
&= \frac{\mathbb{E} \left[\int \prod_{n=1}^N \tilde{p}(w_{d_n} \mid \mu^*, \sigma^*, \gamma^*, \beta^*) d\theta_d \right] - \mathbb{E} \left[\int \prod_{n=1}^N \hat{p}(w_{d_n} \mid \alpha^*, \beta^*) d\theta_d \right]}{\mathbb{E} \left[\int \prod_{n=1}^N \hat{p}(w_{d_n} \mid \alpha^*, \beta^*) d\theta_d \right]} \\
&\geq \frac{\mathbb{E} \left[\int C \prod_{n=1}^N \hat{p}(w_{d_n} \mid \alpha^*, \beta^*) d\theta_d \right] - \mathbb{E} \left[\int \prod_{n=1}^N \hat{p}(w_{d_n} \mid \alpha^*, \beta^*) d\theta_d \right]}{\mathbb{E} \left[\int \prod_{n=1}^N \hat{p}(w_{d_n} \mid \alpha^*, \beta^*) d\theta_d \right]} \\
&= \frac{C \cdot \mathbb{E} \left[\int \prod_{n=1}^N \hat{p}(w_{d_n} \mid \alpha^*, \beta^*) d\theta_d \right] - \mathbb{E} \left[\int \prod_{n=1}^N \hat{p}(w_{d_n} \mid \alpha^*, \beta^*) d\theta_d \right]}{\mathbb{E} \left[\int \prod_{n=1}^N \hat{p}(w_{d_n} \mid \alpha^*, \beta^*) d\theta_d \right]} = C - 1.
\end{aligned}$$

□