

---

# Minimum Volume Topic Modeling

---

**Byoungwook Jang**  
Department of Statistics  
University of Michigan  
Ann Arbor, MI, 48109  
bwjang@umich.edu

**Alfred Hero**  
Department of EECS  
University of Michigan  
Ann Arbor, MI, 48109  
hero@umich.edu

## Abstract

We propose a new topic modeling procedure that takes advantage of the fact that the Latent Dirichlet Allocation (LDA) log likelihood function is asymptotically equivalent to the logarithm of the volume of the topic simplex. This allows topic modeling to be reformulated as finding the probability simplex that minimizes its volume and encloses the documents that are represented as distributions over words. A convex relaxation of the minimum volume topic model optimization is proposed, and it is shown that the relaxed problem has the same global minimum as the original problem under the separability assumption and the sufficiently scattered assumption introduced by Arora et al. (2013) and Huang et al. (2016). A locally convergent alternating direction method of multipliers (ADMM) approach is introduced for solving the relaxed minimum volume problem. Numerical experiments illustrate the benefits of our approach in terms of computation time and topic recovery performance.

## 1 Introduction

Since the introduction by Blei et al. (2003) and Pritchard et al. (2000), the Latent Dirichlet Allocation (LDA) model has remained an important tool to explore and organize large corpora of texts and images. The goal of topic modeling can be summa-

rized as finding a set of topics that summarizes the observed corpora, where each document is a combination of topics lying on the topic simplex. There are many extensions of LDA, including a nonparametric extension based on the Dirichlet process called Hierarchical Dirichlet Process (Teh et al., 2005), a correlated topic extension based on the logistic normal prior on the topic proportions (Lafferty and Blei, 2006), and a time-varying topic modeling extension (Blei and Lafferty, 2006). There are two main approaches for estimation of the parameters of probabilistic topic models: the variational approximation popularized by Blei et al. (2003) and the sampling based approach studied by Pritchard et al. (2000). These inference algorithms either approximate or sample from the posterior distributions of the latent variable representing the topic labels. Therefore, the estimates do not necessarily have a meaningful geometric interpretation in terms of the topic simplex - complicating assessment of goodness of fit to the model. In order to address this problem, Yurochkin and Nguyen (2016) introduced Geometric Dirichlet Mean (GDM), a novel geometric approach to topic modeling. It is based on a geometric loss function that is surrogate to the LDA's likelihood and builds upon a weighted k-means clustering algorithm, introducing a bias correction. It avoids excessive redundancy of the latent topic label variables and thus improves computation speed and learning accuracy. This geometric viewpoint was extended to a nonparametric setting (Yurochkin et al., 2017).

LDA-type models also arise in the hyperspectral unmixing problem. Similar to the documents in topic modeling, hyperspectral image pixels are assumed to be mixtures of a few spectral signatures, called endmembers (equivalent to topics). Unmixing procedures aim to identify the number of endmembers, their spectral signatures, and their abundances at each pixel (equivalent to topic proportions). One dif-

ference between topic modeling and unmixing is that hyperspectral spectra are not normalized. Nonetheless, algorithms for hyperspectral unmixing are similar to topic model algorithms, and similar models have been applied to both problems. Geometric approaches in the hyperspectral unmixing literature take advantage of the fact that linearly mixed vectors also lie in a simplex set or in a positive cone. One of the early geometric approaches to unmixing was introduced in Nascimento and Dias (2005) and Bioucas-Dias (2009), which aim to first identify the  $K$ -dimensional subspace of the data and then estimate the endmembers that minimize the volume of the simplex spanned by these endmembers. Bioucas-Dias (2009) estimates the endmembers through minimizing the log determinant of the endmember matrix, as the log-determinant is proportional to the volume of the simplex defined by the endmembers. This idea of minimizing the simplex volume motivated the algorithm proposed in this paper for topic modeling. In Bioucas-Dias (2009), however, the authors experience an optimization issue as their formulation is highly non-convex. It was found that the local minima of the objective in Bioucas-Dias (2009) may be unstable.

The topic modeling problem also has similarities to matrix factorization. In particular, nonnegative matrix factorization, while it does not enforce a sum-to-one constraint, is directly applicable to topic modeling (Deerwester et al., 1990; Xu et al., 2003; Anandkumar et al., 2012; Arora et al., 2013; Fu et al., 2018). Recover KL, recently introduced by Arora et al. (2013), provides a fast algorithm that identifies the model under a separability assumption, which is the assumption that the sample set includes the vertices of the true topic model (pure endmembers). As the separability assumption is often not satisfied in practice, Fu et al. (2018) introduced a weaker assumption called the sufficiently scattered assumption. We provide a theoretical justification of our geometric minimum value method under this weaker assumption.

### 1.1 Contribution

We propose a new geometric inference method for LDA that is formulated as minimizing the volume of the topic simplex. The estimator is shown to be identifiable under the separability assumption and the sufficiently scattered assumption. Compared to Bioucas-Dias (2009), our geometric objective involves  $\log \det \beta \beta^T$  instead of  $\log |\det \beta|$ , making our objective function convex. At the same time, the  $\log \det \beta \beta^T$  term remains proportional to the volume

enclosed by the topic matrix  $\beta$  and simplifies the optimization. In particular, we propose a convex relaxation of the minimization problem whose global minimization is equivalent to the original problem. This relaxed objective function is minimized using an iterative augmented Lagrangian approach, implemented using the alternating direction method of multipliers (ADMM), that is shown to be locally convergent.

### 1.2 Notation

We use the following notations. We are given a corpus  $W \in \mathbb{D}^{M \times V}$  with  $M$  documents,  $K$  topics, vocabulary size  $V$  and  $N_m$  words in document  $m$  for  $m = 1, \dots, M$ . Let  $\mathbb{D}^{n \times p}$  be the space of  $n \times p$  row-stochastic matrices. Then, our goal is to decompose  $W$  as  $W = \theta \beta$ , where  $\theta \in \mathbb{D}^{M \times K}$  is the matrix of topic proportions, and  $\beta \in \mathbb{D}^{K \times V}$  is the topic-term matrix. Finally,  $\Delta^d$  represents the  $d$ -dimensional simplex. It is assumed that the documents in the corpus obey the following generative LDA model.

1. For each topic  $\beta_i$  for  $i = 1, \dots, K$ 
  - (a) Draw a topic distribution  $\beta_i$
2. For  $j$ -th document  $w^{(j)} \in \mathbb{R}^V$  in the corpus  $W$  for  $j = 1, \dots, M$ 
  - (a) Choose the topic proportion  $\theta_w \sim \text{Dir}(\alpha)$
  - (b) For each word  $\delta_n$  in the document  $w^{(j)}$ 
    - i. Choose a topic  $z_n \sim \text{Mult}(\theta)$
    - ii. Choose a word  $\delta_n \sim \beta_{z_n}$

## 2 Proposed Approach

We assume that the number  $K$  of topics is known in advance and is much smaller than size of the vocabulary, i.e.  $K \ll V$ . Furthermore, since LDA models the document as being inside the topic simplex, it is advantageous to represent the documents in a  $K$ -dimensional subspace basis.

Let  $E_K = [e_1, \dots, e_K]$  be a matrix of dimension  $V \times K$  with  $K$  orthogonal directions spanning the document subspace. Specifically, we define  $E_K$  as the set of  $K$  eigenvectors of the sample covariance matrix of the documents  $w^{(i)}$ ,  $i = 1, \dots, N$ .

Most of the paper focuses on working with  $\tilde{w}^{(i)} = w^{(i)} E_K \in \mathbb{R}^K$ , which corresponds to the coordinates of  $w^{(i)}$  in  $\text{colspan}(E_K)$ . Note that we can recover the projected documents in the original  $V$ -dimensional space by

$$\begin{aligned} \hat{w}^{(i)} &= \bar{w} + (w^{(i)} - \bar{w}) E_K E_K^T \\ &= \bar{w} + (\tilde{w}^{(i)} - \bar{w} E_K) E_K^T \in \mathbb{R}^V \end{aligned}$$

where  $\bar{w}$  is the sample average of the observed documents. Therefore,

$$W = \theta\beta \Rightarrow (W - \theta\beta)E_K = 0$$

where  $\Theta$  belongs to the simplex  $\Delta^K$ . This  $K$ -dimensional probability simplex is defined by the topic distributions, which are the rows of  $\beta E_K \in \mathbb{R}^{K \times K}$ . For the rest of the paper, given  $\omega \in \mathbb{R}^V$ , we denote  $\tilde{\omega}$  as the corresponding coordinates in the projected subspace and  $\hat{\omega}$  as the projected vector in the original  $V$ -dimensional space.

## 2.1 Topic Estimation

Let  $\gamma = (\beta E_K)^{-1}$ . Then, it follows that  $\theta = (W E_K)\gamma$ . We know that  $\beta E_K$  is invertible as we assume that there are  $K$  distinct topics, and the rank of the topic matrix  $\beta$  is  $K$ . Then, as noted in Nascimento and Bioucas-Dias (2012), the likelihood w.r.t.  $\Theta$  can be written as

$$\begin{aligned} l(\theta, \beta|W) &= \sum_{i=1}^M p(w^{(i)}|\beta, \alpha) \\ &= \sum_{i=1}^M \log \left( p(\theta^{(i)} = (w^{(i)} E_K)\gamma|\beta, \alpha) \cdot |\det(\gamma)| \right) \\ &= \sum_{i=1}^M \log \left( p(\theta^{(i)} = (w^{(i)} E_K)\gamma|\beta, \alpha) \right) \\ &\quad + N \log |\det(\gamma)| \end{aligned} \quad (1)$$

This formulation gives a nice geometric interpretation.

**Geometric Interpretation of log likelihood:** As we increase the number of documents  $M \rightarrow \infty$ , the dominant term is  $\log |\det(\beta)|$ . That is,

$$\begin{aligned} &\lim_{N \rightarrow \infty} \arg \max_{\beta} l(\theta, \beta|W) \\ &= \lim_{M \rightarrow \infty} \arg \max_{\tilde{\beta}} \sum_{i=1}^M \log p(W_i|\theta, \beta) \\ &\approx \arg \min_{\beta} \log |\det(\beta E_K)| \\ &= \arg \min_{\gamma} -\log |\det \gamma| \end{aligned} \quad (2)$$

Note that  $\log |\det(\beta E_K)|$  is proportional to the volume enclosed by the row vectors of  $\beta E_K$ , i.e. the topic simplex in the projected subspace. In other words, the estimated topic matrix  $\beta$  that minimizes its intrinsic volume is asymptotically equivalent to the asymptotic form of the log-likelihood (1). This is the main motivation for our proposal to minimize the volume of topic simplex.

## 3 Minimum Volume Topic Modeling

In the remote sensing literature, Nascimento and Bioucas-Dias (2012) proposed to work with the likelihood (1) by modeling  $\theta$  as a Dirichlet mixture. However, their endmembers are spectra and do not necessarily satisfy the sum-to-one constraints on the endmember matrix; constraints which are fundamental to topic modeling. These additional constraints on the endmember complicate the minimization of (1). The first difficulty arises from the  $\log |\det \beta|$  term, as  $\beta$  is not a symmetric matrix, which makes the log likelihood (1) non-convex. Due to this non-convexity issue, Nascimento and Bioucas-Dias (2012) propose using a second-order approximation to the  $\log \det \beta$  term. Yet, no rigorous justification has been provided to their approach. In contrast, we propose using  $\log \det \beta \beta^T$  instead of  $\log |\det \beta|$ , prove identifiability under the sufficiently scattered assumption, and derive an ADMM update.

As we are optimizing  $(\beta E_K)^{-1}$  directly, we use the notation  $\gamma = (\beta E_K)^{-1}$  in the sequel. We can then rewrite the objective (2) as follows

$$\begin{aligned} \hat{\gamma} &= \arg \min_{\gamma \in \mathbb{R}^{K \times K}} -\log |\det(\gamma \gamma^T)| \\ \text{s.t. } &\theta > 0 \quad \theta \mathbf{1} = \mathbf{1} \quad \theta = (W E_K)\gamma \\ &\beta > 0 \quad \beta \mathbf{1} = \mathbf{1} \end{aligned} \quad (3)$$

where  $\beta = \bar{w} + (\gamma^{-1} - \bar{w} E_K) E_K^T$ . The first set of constraints corresponds to the sum-to-one and nonnegative constraint on the topic proportions  $\theta = (W E_K)\gamma$ , and the second constraint imposes the same conditions on  $\beta$ . Thus, the problem (3) provides an exact solution to the asymptotic estimation of (1). However, this is not a convenient formulation of the optimization problem, as it involves the constraint on the inverse of  $\gamma$ . Note that as we assume  $\beta \in \mathbb{R}^V$  intrinsically lives in a  $K$ -dimensional subspace, there is a one-to-one mapping between  $\beta$  and  $\gamma = (\beta E_K)^{-1}$ . Throughout this paper, we will make use of this relationship between  $\beta$  and  $\gamma$ . Here, working with a geometric interpretation of the second set of constraints we propose a relaxed version of (3).

**Sum-to-one constraint on  $\beta$ :** Combined with the non-negativity constraint, the sum-to-one constraint  $\beta \mathbf{1} = \mathbf{1}$  forces the rows of  $\beta$  to lie in the  $K$ -dimensional topic simplex within the word simplex. To be specific,  $\beta \mathbf{1} = \mathbf{1}$  narrows our search space to be in an affine subspace, which is accomplished with a projection of the documents onto this  $K$ -dimensional affine subspace. This projection takes care of the sum-to-one constraint in the objective (3).

**Non-negativity constraint on  $\beta$ :** We propose relaxing the non-negativity constraint to the following

$$\sigma_{\min}(\gamma) \geq R^{-1}$$

where  $\sigma_{\min}(\gamma)$  is the minimum singular value of  $\gamma$ . As illustrated in Figure 1, this is interpreted as replacing the non-negativity constraint on the elements of the matrix  $\beta$  with a radius  $R$  ball constraint on the rows of the matrix  $\beta$ . As noted before, there is a mapping between  $\gamma$  and  $\beta$  through  $\beta = \bar{w} + (\gamma^{-1} - \bar{w}E_K)E_K^T$ . Thus, if  $\gamma^t$  is the current iterate of an iterative optimization algorithm, to be specified below, then we can represent the corresponding  $i$ -th topic vector in the projected space as  $b_i^t = (\gamma^t)^{-1}[i, :] = (\beta^t E_K)[i, :]$ . It follows that

$$\begin{aligned} \|b_i^t\|^2 &= \frac{\text{tr}((b_i^t)^T b_i^t) \lambda_{\min}(\gamma \gamma^T)}{\lambda_{\min}(\gamma \gamma^T)} \leq \frac{\text{tr}((b_i^t)^T b_i^t \gamma \gamma^T)}{\lambda_{\min}(\gamma \gamma^T)} \\ &= \frac{\text{tr}(b_i^t \gamma (b_i^t \gamma)^T)}{\lambda_{\min}(\gamma \gamma^T)} = \frac{\text{tr}(e_i e_i^T)}{\sigma_{\min}(\gamma)^2} \leq R^2 \end{aligned} \quad (4)$$

Then, imposing  $\sigma_{\min}(\gamma) \geq R^{-1}$  results in  $\|b_i\|^2 \leq R$ . The first inequality in (4) comes from the fact that  $\lambda_{\min}(A) \text{tr}(B) \leq \text{tr}(AB) \leq \lambda_{\max}(A) \text{tr}(B)$  for positive semidefinite matrices  $A$  and  $B$ . The second equality in (4) comes from the definition that  $b_i$  is the  $i$ -th row of  $(\gamma^t)^{-1}$

With this spectral relaxation of the non-negativity constraint, the relaxed version of the problem (3) becomes

$$\begin{aligned} \hat{\gamma} &= \arg \min_{\gamma} -\log |\det(\gamma \gamma^T)| \\ \text{s.t. } &\theta > 0 \quad \theta \mathbf{1} = \mathbf{1} \quad \theta = (WE_K)\gamma \\ &\sigma_{\min}(\gamma) \geq R^{-1} \end{aligned} \quad (5)$$

Intuitively, as shown in Figure 1, the optimization problem (3) and (5) are equivalent to each other except that the ball relaxation has expanded the solution space beyond the feasible space.




Figure 1: Visualization of the difference between the feasible space of (3) (blue triangle) and that of (5) (pink circle).

The blue triangle represents the set of feasible points

for problem (3), and the red circle corresponds to the solution space in (5).

### 3.1 Identifiability

Here we establish identifiability of the model obtained by solving problem (5). Identifiability gained interests in the topic modeling literature (Arora et al. (2013) and Fu et al. (2018)). We show the identifiability under the sufficiently scattered condition. We first state the following lemma.

**Lemma 3.1.** *Let  $\hat{\gamma}$  be a solution to the problem (5). If  $\text{rank}(W) = K$ , we have that  $\hat{\gamma} \in \Gamma$ , where*

$$\Gamma = \{\gamma \in \mathbb{R}^{K \times K} : \beta = \bar{w} + (\gamma^{-1} - \bar{w}E_K)E_K^T \in \mathbb{D}^{K \times V} \text{ and } \exists \theta \in \mathbb{D}^{M \times K} \text{ s.t. } \theta = (WE_K)\gamma\}$$

Intuitively, Lemma 3.1 tells us that we cannot have the solution outside of the blue triangle in Figure 2. If there was a solution outside of the triangle (Figure 2a), we could find the projection (Figure 2b) onto the word simplex (blue triangle) that still satisfies the constraint yet has a smaller volume, which is a contradiction.




Figure 2: Visualization of the proof for Lemma 3.1. The set  $\Gamma$  corresponds to the blue triangle, which represents the feasible set of the problem (3). The red circle represents the feasible set of the relaxed problem (5). Given a potential solution  $\hat{\beta}$  for (5), we can always argue that the projection of  $\hat{\beta}$ , namely  $\text{Proj}_{\Delta^V}(\hat{\beta})$ , is a better solution to (5) as illustrated in Panel 2b.

*Proof.* We prove this statement by contradiction. Suppose  $\hat{\gamma} \notin \Gamma$ . Then, as  $\hat{\gamma}$  is an optimal solution to the problem (5), we have that  $\theta = (WE_K)\hat{\gamma} \in \mathbb{D}^{M \times K}$ . Furthermore, since  $W \in \Delta^V$  and  $E_K$  is obtained from PCA, we have that  $E_K^T \mathbf{1}_V = 0$ . Thus, it follows that

$$\begin{aligned} \hat{\beta} \mathbf{1}_V &= \bar{w} \mathbf{1}_V + (\hat{\gamma}^{-1} - \bar{w}E_K)E_K^T \mathbf{1}_V \\ &= \mathbf{1}_K - (\hat{\gamma}^{-1} - \bar{w}E_K) \mathbf{0}_V = \mathbf{1}_K \end{aligned}$$

Therefore, the only constraint that  $\hat{\gamma}$  could possibly violate is non-negativity of  $\hat{\beta}$ . Let  $\text{Proj}_{\Delta^V}(\hat{\beta})$  be

the projection of  $\widehat{\beta}$  onto the simplex  $\Delta^V$  and let  $\widehat{\gamma}_{\text{proj}} = \text{Proj}_{\Delta^V}(\widehat{\beta})E_K$ . Then,  $\widehat{\gamma}_{\text{proj}} \in \Gamma$  satisfies all the constraints in the optimization problem (5), but we also have that

$$-\log(\det \widehat{\gamma}_{\text{proj}} \widehat{\gamma}_{\text{proj}}^T) < -\log(\det \widehat{\gamma} \widehat{\gamma}^T)$$

since the volume of  $\text{Proj}_{\Delta^V}(\widehat{\beta})$  is smaller than that of  $\widehat{\beta}$ . This is a contradiction as  $\widehat{\gamma}$  is the optimal solution to the problem (5). Thus, it follows that  $\widehat{\gamma} \in \Gamma$ .  $\square$

We now state the sufficiently scattered assumption from Huang et al. (2016).

**Assumption 1:** (*sufficiently scattered condition* (Huang et al. (2016))) Let  $\text{cone}(\beta)^* = \{x : \beta x \geq 0\}$  be the polyhedral cone of  $\beta$  and  $S = \{x : \|x\|_2 \leq 1^T x\}$  be the second order cone. Matrix  $\beta$  is called sufficiently scattered if it satisfies:

- 1)  $\text{cone}(\beta)^* \subset S$
- 2)  $\text{cone}(\beta)^* \cap \text{bd}(S) = \{ae_k : a \geq 0, k = 1, \dots, K\}$ , where  $\text{bd}S$  denotes the boundary of  $S$ .

The sufficiently scattered assumption can be interpreted as an assumption that we observe a sufficient number of documents on the faces of the topic simplex. In real-world topic model applications, such an assumption is not unreasonable since there are usually documents in the corpora having sparse representations.

**Proposition 3.1.** *Let  $\gamma_*$  be the optimal solution to the problem (5) and  $\beta_* = \bar{w} + (\gamma_*^{-1} - \bar{w}E_K)E_K^T$  be the corresponding topic matrix. If the true topic matrix  $\beta$  is sufficiently scattered and  $\text{rank}(\widetilde{W}) = K$ , then  $\beta_* = \beta\Pi$ , where  $\Pi$  is a permutation matrix.*

The proof structure is similar to the one in Huang et al. (2016), and we include here for completeness.

*Proof.* Given a corpus  $W \in \mathbb{D}^{M \times K}$ , let  $\beta \in \mathbb{D}^{K \times V}$  be the true topic-word matrix. Suppose  $\text{rank}(W) = K$  and  $\beta$  is sufficiently scattered. Let  $\gamma_*$  be the solution to the problem (5). Then, by Lemma 3.1, we have that  $\gamma_* \in \Gamma$ . Furthermore, since  $\text{rank}(W) = K$ , we have that  $\text{rank}(\beta_*) = K$  as  $W = \theta\beta_*$  where  $\theta = (WE_K)\gamma_*$ . It also follows that  $\text{rank}(\beta) = K$  due to the constraint  $W = \theta\beta$ . Therefore,  $|\det \beta|$  and  $|\det \beta_*|$  are strictly positive. In other words, we cannot have a trivial solution to (5) as the objective is bounded. As  $\beta$  and  $\beta_*$  are full row rank, there exists an invertible matrix  $Z \in \mathbb{R}^{K \times K}$  such that  $\beta_* = Z\beta$ .

Also, as  $\gamma_* \in \Gamma$ , it follows that  $\beta_* = Z\beta \geq 0$  and

$$\begin{aligned} \beta_* \mathbf{1}_V &= Z\beta \mathbf{1}_V = \mathbf{1}_K \\ &\Rightarrow Z \mathbf{1}_K = \mathbf{1}_K \end{aligned}$$

The inequality constraint  $Z\beta \geq 0$  tells us that rows of  $Z$  are contained in  $\text{cone}(\beta)^*$ . As  $\beta$  is sufficiently scattered, it follows that

$$Z[k, :] \in \text{cone}(\beta)^* \subset S \quad (6)$$

by the first condition of (A1). Then, by the definition of the second order cone  $S$ , it follows that

$$\begin{aligned} |\det Z| &= |\det Z^T| \leq \prod_{k=1}^K \|Z[k, :]\|_2 \\ &\leq \prod_{k=1}^K Z[k, :]\mathbf{1}_K = 1 \end{aligned} \quad (7)$$

The first inequality comes from the Hadamard inequality, which states that the equality holds if and only if the vectors  $Z[k, :]$ 's are orthogonal to each other. The second inequality holds when  $\|Z[k, :]\|_2 = Z[k, :]\mathbf{1}_K \forall k = 1 \dots K$ . In other words, when  $Z[k, :] \in \text{bd}(S) \forall k$ . Then, together with (6), it follows that

$$\begin{aligned} Z[k, :] &\in \text{cone}(\beta)^* \cap \text{bd}S \\ &= \{\lambda e_k | \lambda \geq 0, k = 1, \dots, K\} \end{aligned} \quad (8)$$

Thus, it follows that the  $|\det Z|$  achieves its maximum at 1, when  $Z \in \text{cone}(\beta)^* \cap \text{bd}S$  sums to one and is an orthogonal matrix, i.e. when  $Z$  is a permutation matrix.

Furthermore, since  $\gamma_* = (\beta_* E_K)^{-1} = Z^{-1}(\beta E_K)^{-1} = Z^{-1}\gamma$ , we have that

$$\begin{aligned} -\det(\gamma_* \gamma_*^T) &= -\det(Z^{-1}\gamma\gamma^T(Z^{-1})^T) \\ &= -|\det Z^{-1}| \det(\gamma\gamma^T) |\det Z^{-1}| \\ &= -|\det Z|^{-2} \det(\gamma\gamma^T) \\ &\geq -|\det(\gamma\gamma^T)| \end{aligned}$$

where the equality holds when  $|\det Z| = 1$ . In other words, the minimum is achieved when  $Z$  is a permutation matrix. Therefore, our solution  $\gamma_*$  to the problem (5) and the corresponding  $\beta_*$  are equal to the true topic-word matrix up to permutation.  $\square$

**Assumption 2** (*Separability assumption* from Arora et al. (2013)) There exists a set of indices  $\Lambda = \{i_1, \dots, i_K\}$  such that  $\beta(\Lambda, :) = \text{Diag}(c)$ , where  $c \in \mathbb{R}_+^K$ .

The separability assumption, also known as the anchor-word assumption, states that every topic  $k$  has a unique word  $w_k$  that only shows up in topic  $k$ . These words are also referred to as the anchor words as introduced in Arora et al. (2013).

**Remark:** The identifiability statement in Proposition 3.1 holds true under the separability assumption as well, as the sufficiently scattered assumption is a weaker version of the separability assumption.

### 3.2 Augmented Lagrangian Formulation

With  $\mu \geq 0$ , we work with the following augmented Lagrangian version of the constrained optimization problem (5)

$$\begin{aligned} \hat{\gamma} &= \arg \min_{\gamma} -\log |\det(\gamma\gamma^T)| + \mu \|\widetilde{W}\gamma\|_h \\ \text{s.t. } \gamma \mathbf{1}_K &= (\widetilde{W}^T \widetilde{W})^{-1} \widetilde{W}^T \mathbf{1}_V \quad \sigma_{\min}(\gamma) \geq \zeta \end{aligned} \quad (9)$$

where  $\widetilde{W} = WE_K$ ,  $\zeta = R^{-1}$ , and  $\|X\|_h = \sum_{i,j} \max(-X_{i,j}, 0)$  is a hinge loss that captures the non-negativity constraint on  $\theta$ . Furthermore, the linear constraint is converted to  $\gamma \mathbf{1}_K = (\widetilde{W}^T \widetilde{W})^{-1} \widetilde{W}^T \mathbf{1}_V$ , which is the same constraint as  $\widetilde{W}\gamma \mathbf{1}_K = \mathbf{1}_V$ . For simplicity, we define  $\mathbf{a} = (\widetilde{W}^T \widetilde{W})^{-1} \widetilde{W}^T \mathbf{1}_V$ .

The Lagrangian objective function in (9) can be written as

$$f(\gamma) = -\log |\det(\gamma\gamma^T)| + \mu \|\widetilde{W}\gamma\|_h + \mathbf{1}(\sigma_{\min}(\gamma) > \zeta) \quad \text{s.t. } \gamma \mathbf{1}_K = \mathbf{a} \quad (10)$$

Introducing the auxiliary optimization variables  $V_1 \in \mathbb{R}^{n \times k}$  and  $V_2 \in \mathbb{R}^{k \times k}$ , we reformulate (5)

$$\begin{aligned} \hat{\gamma} &= \arg \min_{\gamma, V_1, V_2} \left\{ -\log |\det \gamma\gamma^T| + \mu \|V_1\|_h + \right. \\ &\quad \left. + \mathbf{1}(\sigma_{\min}(\gamma) > \zeta) \right\} \\ \text{s.t. } V_1 &= \widetilde{W}\gamma \quad \gamma = V_2 \quad V_2 \mathbf{1}_K = \mathbf{a} \end{aligned} \quad (11)$$

For a penalty parameter  $\rho > 0$  and Lagrange multiplier matrix  $\Lambda \in \mathbb{R}^{n \times k}$ , we consider the augmented Lagrangian of this problem

$$\begin{aligned} \mathcal{L}(\gamma, V_1, V_2, \Lambda_1, \Lambda_2) &= -\log |\det \gamma\gamma^T| + \mu \|V_1\|_h + \mathbf{1}(\sigma_{\min}(\gamma) > \zeta) \\ &\quad + \frac{\rho}{2} \|\widetilde{W}\gamma - V_1\|_F^2 + \langle \Lambda_1, \widetilde{W}\gamma - V_1 \rangle \\ &\quad + \frac{\rho}{2} \|\gamma - V_2\|_F^2 + \langle \Lambda_2, \gamma - V_2 \rangle \quad \text{s.t. } V_2 \mathbf{1}_K = \mathbf{a} \end{aligned} \quad (12)$$

This function can be minimized using an iterative ADMM update scheme on the arguments  $\gamma$ ,  $V_1$ ,  $V_2$ ,

$\Lambda_1$ , and  $\Lambda_2$ . The update for  $V_1$  and  $V_2$  can be accomplished by standard proximal operators that implement soft-thresholding and a projection. Furthermore, the  $\gamma$ -update can be derived in a closed form by solving a quadratic equation in its singular values. The details of the ADMM updates are included in the supplement. First, consider the  $\gamma$ -subproblem without the linear constraint  $\gamma \mathbf{1} = \mathbf{a}$ . Then, as derived in the supplement, the resulting update equation for  $\gamma$  is

$$\begin{aligned} \gamma^\dagger &= \arg \min_{\gamma \in \mathbb{R}^{k \times k}} \left\{ -\log |\det \gamma^T \gamma| + \frac{\rho}{2} \|C^{1/2}(\gamma - A)\|_F^2 \right\} \\ &= U \widehat{D} W^T \end{aligned} \quad (13)$$

where  $\widehat{D}$  is defined in the supplement. Using  $\gamma_+$ , we obtain a closed-form solution to the  $\gamma$  sub-problem in (12) as follows

$$\gamma^{t+1} = \gamma_+ - (\gamma_+ \mathbf{1} - \mathbf{a})(\mathbf{1}^T C^{-1} \mathbf{1})^{-1} \mathbf{1}^T C^{-1}$$

This solution to the linear constrained problem can be easily derived as a stationary point of the convex function that is minimized in (13). Note that, by construction,  $\gamma^{t+1} \mathbf{1} = \mathbf{a}$ .

---

#### Algorithm 1: Minimum volume topic modeling

---

**Input:**  $W, E_K, \gamma_0, \rho > 0, \mu > 0$

**Output:**  $\hat{\beta}$

Initialize  $V_2^0 = \gamma^0$ ,  $V_1 = \widetilde{W}\gamma^0$ ,  $\Lambda_1^0 = \mathbf{0}$ ,  $\Lambda_2^0 = \mathbf{0}$  ;

Calculate  $C = I + \widetilde{W}^T \widetilde{W}$  ;

Calculate the projected documents  $\widetilde{W}$  ;

**while not converged do**

$$V_1^{t+1} = \text{Prox}_{\|\cdot\|_h, \mu/\rho} \left( \frac{\rho \widetilde{W}\gamma^t + \Lambda_1^t}{\rho} \right)$$

$$V_2^{t+1} = \text{Proj}_{G_R} \left( \frac{\rho \gamma^t + \Lambda_2^t}{\rho} \right)$$

$$\gamma^{t+1} = \gamma_+ - (\gamma_+ \mathbf{1} - \mathbf{a})(\mathbf{1}^T C^{-1} \mathbf{1})^{-1} \mathbf{1}^T C^{-1}$$

where  $\gamma_+$  is defined in (13)

$$\Lambda_1^{t+1} = \Lambda_1^t + \rho(\widetilde{W}\gamma^{t+1} - V_1^{t+1})$$

$$\Lambda_2^{t+1} = \Lambda_2^t + \rho(\gamma^{t+1} - V_2^{t+1})$$

**end**

---

In the nonnegative matrix factorization literature, Liu et al. (2017) used a large-cone penalty that constrains either the volume or the pairwise angles of the simplex vertices. However, this does not impose a sum-to-one constraint on the topics, and the optimization is performed over  $\beta$ . Furthermore, our formulation has an advantage over the problem in Liu et al. (2017) as we directly work with the latent topic proportions  $\theta$ . This is possible in our formulation as we decoupled  $\beta$  from  $\theta$  using the ADMM mechanism.

### 3.3 Convergence

The following proposition shows that Algorithm 1 converges to a stationary point of (10).

**Proposition 3.2.** *For any limit point  $(\gamma^*, V_1^*, V_2^*, \Lambda_1^*, \Lambda_2^*)$  of Algorithm 1,  $\gamma^*$  is also a stationary point of (10).*

This follows by applying a standard convergence proof of the ADMM algorithm (Algorithm 1) based on the KKT condition. The proposition states that our ADMM formulation converges to a stationary point. However, while the unconstrained objective function in (10) is convex, the constraint on the minimum singular value makes the constrained optimization function non-convex. Thus, our algorithm is only guaranteed to converge to a stationary point of (10).

Figure 3 demonstrates the convergence of our algorithm with synthetic data generated from an LDA model with parameters  $\alpha = 0.1, \eta = 0.1, V = 1200, K = 3, M = 1000$ , and  $N_m = 1000$ .




Figure 3: Experimental runs using Algorithm 1. The data was simulated from an LDA model with  $\alpha = 0.1, \eta = 0.1, V = 1200, K = 3, M = 1000, N_m = 1000$ . The algorithm was initialized with  $\gamma$  equal to the identity matrix. The left panel shows the relative Frobenius error between the iterates  $\gamma^t$  and the true  $\gamma$ . The right panel shows the convergence in terms of the objective values.

## 4 Performance Comparison

To demonstrate the performance of the proposed minimum volume topic model (MVTM) estimation algorithm (Algorithm 1), we generate the LDA data with the parameters  $\eta = 0.1, V = 1200, K = 3, M = 1000, N_m = 1000$  with varying  $\alpha$ , which is Dirichlet hyperparameter for the topic proportion  $\theta$ . For ease of visualization, the first two dimensions of the projected documents and the estimated topics are used. The first scenario ( $\alpha = 0.1$ ) in Figure 4 shows the performance of our algorithm is comparable to the vertex based method GDM (Yurochkin and Nguyen




Figure 4: Visualization of Minimum Volume Topic Modeling (MVTM) with the observed documents in black, optimization path of MVTM in the gradient of red (dark red = beginning, light red = end), and the final estimate in yellow. The ground-truth topic vertices are plotted in cyan. The Dirichlet parameter for the topic proportion was set at  $\alpha = 0.1$ , and MVTM was initialized at the identity matrix.

(2016)), when there are plenty of observed documents around the vertices. While there is no anchor word in the generated dataset, we observe enough documents around the vertices. In other words, the separability assumption is slightly violated. With higher values of  $\alpha$ , however, Figure 5 shows the advantages of our method, denoted as MVTM. Note that the higher




Figure 5: Visualization of the proposed MVTM algorithm with the observed documents in black, optimization path of MVTM in the gradient of red (dark red = beginning, light red = end), and the final estimate in yellow under different values of  $\alpha$ . The ground-truth topic vertices are plotted in cyan, and the final estimate of GDM is plotted in green for comparison. MVTM was initialized at the identity matrix.

values of  $\alpha$  correspond to the situation where the sufficiently scattered condition is satisfied, but the separability condition is violated. Thus, we can see the vertex based method (GDM) starts to suffer in the oracle performance. In contrast, with an appropriate choice of  $\mu$  for the hinge loss, our method recovers the correct topics even for the well mixed scenario where  $\alpha = 5$ . Figure 5b shows that there is a kink in the optimization path, where MVTM is finding the right orientation of the true simplex. Furthermore, there is a lack of loops in the optimization

path, illustrating the identifiability of MVTM.

Lastly, we explore the asymptotic behavior by varying document lengths  $N_m$  with  $M = 1000$ ,  $K = 5$ ,  $V = 1200$ ,  $\eta = 0.1$ ,  $\alpha = 0.1$  and 100 held-out documents. MVTM is all initialized at the identity matrix, and VEM had 10 restarts as the objective for the variational method is nonconvex.




Figure 6: Perplexity of the held-out data and the corresponding time complexity of each method at varying values of the number words per document  $N_m$  with  $M = 1000$ ,  $K = 5$ ,  $V = 1200$ ,  $\eta = 0.1$  and  $\alpha = 0.1$

Figure 6 tells us that 1) Gibbs sampling and MVTM have comparable performance in terms of perplexity, 2) MVTM and VEM both show the computational advantages over the Gibbs sampling method, and 3) VEM suffers from the statistical performance due to the nature of the non-convex objective function of VEM. Additional simulation results can be found in the supplement.

#### 4.1 NIPS dataset

To illustrate the performance of MVTM on a real-world data, we apply our algorithm to NIPS dataset. We preprocess the raw data using a standard stop word list and filter the resulting data through a stemmer. After preprocessing, words that appeared more than 25 times across the whole corpus are retained. Then, we further remove the documents that have less than 10 words. The final dataset contained 4492 unique words and 1491 documents with mean document length of 1187. We compare our algorithm’s performance to GDM and Gibbs sampling at  $K=5$ ,

10, 15, and 20. The perplexity score is used to perform the comparison in Table 1. The additional time

|      | MVTM | GDM  | RecoverKL | Gibbs |
|------|------|------|-----------|-------|
| K=5  | 1483 | 1602 | 1569      | 1336  |
| K=10 | 1387 | 1441 | 1507      | 1192  |
| K=15 | 1293 | 1344 | 1438      | 1109  |
| K=20 | 1273 | 1294 | 1574      | 1068  |

Table 1: Perplexity score of the geometric algorithms and the Gibbs sampling for analyzing the NIPS dataset. The proposed algorithm MVTM is performing better than the vertex methods (GDM and RecoverKL) in terms of perplexity as it only requires the documents lie on the face of the topic simplex. GDM provides a similar performance to MVTM.

comparison and top 10 words of top 10 learned topics for MVTM, GDM, and Gibbs sampling are provided in the supplement.

## 5 Discussion

This paper presents a new estimation procedure for LDA topic modeling based on the minimization of the volume of the topic simplex  $\beta$ . Such formulation can be thought of as an asymptotic estimation to the LDA model. The proposed minimum volume topic model (MVTM) algorithm differs from moment-based methods including RecoverKL and the vertex based method such as the GDM. We proved the identifiability of MVTM under the sufficiently scattered assumption introduced in Huang et al. (2016). When the sufficiently scattered assumption is satisfied and the separability assumption is violated, MVTM continues to perform well with an appropriate choice of the hinge loss parameter.

There are open questions on the statistical convergence of our estimator in terms of the document length and the number of documents. Such relationships have been explored in the work of Tang et al. (2014), and it would be interesting to see if these could be applied to the proposed MVTM. The understanding of the statistical behavior of MVTM will provide us with the theoretical guidance on the choice the hinge loss parameter. Besides the theoretical questions, MVTM also has some potential modeling extensions. The immediate extension includes the nonparametric setting, where one would also estimate the number of topics  $K$ .

## Acknowledgments

This research was partially supported by grant ARO W911NF-15-1-0479.



## References

- Anandkumar, A., Foster, D. P., Hsu, D. J., Kakade, S. M., and Liu, Y.-K. (2012). A spectral algorithm for latent dirichlet allocation. In *Advances in Neural Information Processing Systems*, pages 917–925.
- Arora, S., Ge, R., Halpern, Y., Mimno, D., Moitra, A., Sontag, D., Wu, Y., and Zhu, M. (2013). A practical algorithm for topic modeling with provable guarantees. In *International Conference on Machine Learning*, pages 280–288.
- Bioucas-Dias, J. M. (2009). A variable splitting augmented lagrangian approach to linear spectral unmixing. In *Hyperspectral Image and Signal Processing: Evolution in Remote Sensing, 2009. WHISPERS'09. First Workshop on*, pages 1–4. IEEE.
- Blei, D. M. and Lafferty, J. D. (2006). Dynamic topic models. In *Proceedings of the 23rd international conference on Machine learning*, pages 113–120. ACM.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American society for information science*, 41(6):391–407.
- Fu, X., Huang, K., Sidiropoulos, N. D., Shi, Q., and Hong, M. (2018). Anchor-free correlated topic modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Huang, K., Fu, X., and Sidiropoulos, N. D. (2016). Anchor-free correlated topic modeling: Identifiability and algorithm. In *Advances in Neural Information Processing Systems*, pages 1786–1794.
- Lafferty, J. D. and Blei, D. M. (2006). Correlated topic models. In *Advances in neural information processing systems*, pages 147–154.
- Liu, T., Gong, M., and Tao, D. (2017). Large-cone nonnegative matrix factorization. *IEEE transactions on neural networks and learning systems*, 28(9):2129–2142.
- Nascimento, J. M. and Bioucas-Dias, J. M. (2012). Hyperspectral unmixing based on mixtures of dirichlet components. *IEEE Transactions on Geoscience and Remote Sensing*, 50(3):863–878.
- Nascimento, J. M. and Dias, J. M. (2005). Vertex component analysis: A fast algorithm to unmix hyperspectral data. *IEEE transactions on Geoscience and Remote Sensing*, 43(4):898–910.
- Pritchard, J. K., Stephens, M., and Donnelly, P. (2000). Inference of population structure using multilocus genotype data. *Genetics*, 155(2):945–959.
- Tang, J., Meng, Z., Nguyen, X., Mei, Q., and Zhang, M. (2014). Understanding the limiting factors of topic modeling via posterior contraction analysis. In *International Conference on Machine Learning*, pages 190–198.
- Teh, Y. W., Jordan, M. I., Beal, M. J., and Blei, D. M. (2005). Sharing clusters among related groups: Hierarchical dirichlet processes. In *Advances in neural information processing systems*, pages 1385–1392.
- Xu, W., Liu, X., and Gong, Y. (2003). Document clustering based on non-negative matrix factorization. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pages 267–273. ACM.
- Yurochkin, M., Guha, A., and Nguyen, X. (2017). Conic scan-and-cover algorithms for nonparametric topic modeling. In *Advances in Neural Information Processing Systems*, pages 3881–3890.
- Yurochkin, M. and Nguyen, X. (2016). Geometric dirichlet means algorithm for topic inference. In *Advances in Neural Information Processing Systems*, pages 2505–2513.