

# Sparse Topic Modeling via Spectral Decomposition and Thresholding

**Huy Tran**

*Department of Statistics  
The University of Chicago  
Chicago, IL 60637, USA*

HUYDTRAN@UCHICAGO.EDU

**Yating Liu**

*Department of Statistics  
The University of Chicago  
Chicago, IL 60637, USA*

YATINGLIU@UCHICAGO.EDU

**Claire Donnat**

*Department of Statistics  
The University of Chicago  
Chicago, IL 60637, USA*

CDONNAT@UCHICAGO.EDU

**Editor:** Rajarshi Guhaniyogi

## Abstract

In probabilistic Latent Semantic Indexing (pLSI), word frequencies across document corpora are modeled through a low-rank factorization of the expected document-term matrix into topic-word and topic-document components. In this paper, we study the estimation of the topic-word matrix under a sparsity structure motivated by Zipf’s law: word frequencies within each topic exhibit a rapid empirical decay, with most probability mass concentrated on a small subset of words. Motivated by this observation, we introduce a spectral estimator that adaptively thresholds rare words prior to factorization. We show that the resulting estimator achieves an  $\ell_1$ -error rate whose dependence on the vocabulary size  $p$  is only logarithmic. Our error bounds hold across parameter regimes, including high-dimensional settings with extremely large vocabularies, a practically important scenario that has received limited theoretical attention. Unlike many existing methods, our approach does not require the *separability* (or anchor-word) assumption. Synthetic and real-data experiments demonstrate that the proposed procedure is computationally efficient, statistically reliable, and effective across domains with widely varying dimensions, sparsity levels, and document lengths.

**Keywords:** topic modeling, non-negative matrix factorization, high-dimensional statistics,  $\ell_q$ -sparsity.

## 1. Introduction

Topic modeling has proven to be a useful tool for dimensionality reduction and exploratory analysis in natural language processing. Beyond text analysis, it has also been successfully applied in areas such as population genetics (Pritchard et al., 2000; Bicego et al., 2012), social networks (Curiskis et al., 2020) and image analysis (Li et al., 2010).

### 1.1 The Statistical Model

In this paper, we focus on the probabilistic Latent Semantic Indexing (pLSI) model introduced by Hofmann (1999). This simple bag-of-words model involves three variables, namely topics (which are unobserved), words and documents.

Suppose we observe  $n$  documents written using a vocabulary of  $p$  words. For each  $1 \leq i \leq n$ , let  $N_i$  denote the length of document  $i$ . The corpus matrix  $D \in \mathbb{R}^{p \times n}$ , which is a sufficient statistic under the pLSI model and which records the empirical frequency of each word in each document, is defined by

$$D_{ji} = \frac{\text{count of word } j \text{ in document } i}{N_i} \quad \text{for all } 1 \leq i \leq n, 1 \leq j \leq p.$$

Let  $\{D_{*i} : 1 \leq i \leq n\}$  denote the columns of  $D$ , each of which contains only non-negative entries that sum up to 1. The pLSI model specifies that the raw word counts for each document  $\{N_i D_{*i} : 1 \leq i \leq n\}$  are independently generated, with

$$N_i D_{*i} \sim \text{Multinomial}(N_i, [D_0]_{*i}), \tag{1}$$

for some matrix  $D_0 \in \mathbb{R}^{p \times n}$  whose columns are  $\{[D_0]_{*i} : 1 \leq i \leq n\}$ . Here, the columns of  $D_0$  specify how words are assigned to documents, and these columns are required to be probability vectors with non-negative entries summing up to 1. Note that (1) implies  $\mathbb{E}(D) = D_0$ . If we let  $Z := D - D_0$ , we can write the observation model in a ‘‘signal plus noise’’ form:

$$D = D_0 + Z.$$

The pLSI model further assumes that, for some unobserved  $K \in \mathbb{N}$  (which denotes the number of topics), we can factorize  $D_0$  as

$$\mathbb{E}(D) = D_0 = AW, \tag{2}$$

for some matrices  $A \in \mathbb{R}^{p \times K}$  and  $W \in \mathbb{R}^{K \times n}$ . Like  $D_0$ , the columns of  $A$  and  $W$  are required to be probability vectors, so that they can only contain non-negative entries that sum up to 1.  $A$  assigns words to topics, while  $W$  assigns topics to documents. In this paper, we focus more specifically on estimating the topic-word matrix  $A$ .

One can think of (2) as equivalent to requiring that the following Bayes’ formula holds for any word  $j$  and document  $i$ :

$$\mathbb{P}(\text{word } j \mid \text{document } i) = \sum_{k=1}^K \mathbb{P}(\text{word } j \mid \text{topic } k) \cdot \mathbb{P}(\text{topic } k \mid \text{document } i).$$

In most applications,  $K \ll \min(n, p)$  and thus (2) imposes a low-rank structure on  $\mathbb{E}(D)$ . We note that the number of topics  $K$  plays a role similar to that of the number of principal components in principal component analysis. For technical reasons, we will assume throughout this paper that  $K$  is fixed as  $n, p$  and the document lengths  $N_i$ ’s vary. This is reasonable if one expects *a priori* that the number of topics covered by the corpus is small and bounded.

## 1.2 Related Work and Unaddressed Issues

Before outlining our contributions in Section 1.3, it is important to provide context by discussing previous works that are relevant to the estimation of  $A$  under the pLSI model. In particular, we want to highlight some of the unaddressed issues from prior papers that our work aims to resolve.

### 1.2.1 THE SEPARABILITY CONDITION

We first present the definition of anchor words and the separability condition.

**Definition 1 (Anchor words and separability)** *We call word  $j$  an anchor word for topic  $k$  if row  $j$  of  $A$  has exactly one nonzero entry at column  $k$ . The separability condition is said to be satisfied if there exists at least one anchor word for each topic  $k \in \{1, \dots, K\}$ .*

Observe that the decomposition  $D_0 = AW$  in general may not be unique, but under the separability condition,  $A$  is identifiable. The separability condition was first introduced by Donoho and Stodden (2003) to ensure uniqueness in the Non-negative Matrix Factorization (NMF) framework. The interpretation in our context is that, for each topic, there exist some words which act as unique signatures for that topic.

The separability condition greatly simplifies the problem of estimating  $A$ , as one can identify the anchor words for each topic as a first step. Prior works exploiting anchor words mainly differ in how anchor words are used to estimate the remaining non-anchor rows of  $A$ . Arora et al. (2012) start from the *word co-occurrence matrix*  $DD^T$  and apply a successive projection algorithm to rows of  $DD^T$  to find one anchor word per topic. The matrix  $DD^T$  is then re-arranged into four blocks where the top left  $K \times K$  block corresponds to the anchor words identified, and  $A$  is estimated by taking advantage of the special structure of this block partition. More recently, Bing et al. (2020b) consider a matrix  $B \in \mathbb{R}^{p \times K}$  obtained from  $A$  via multiplication by diagonal matrices. Unlike  $A$ , all rows of  $B$  sum up to 1, so anchor rows of  $B$  are simply canonical basis vectors in  $\mathbb{R}^K$ . The non-anchor rows of  $B$  are then obtained via regression given the anchor rows of  $B$ . The topic matrix  $A$  can subsequently be recovered through an appropriate normalization of  $B$ .

A major drawback of these methods is that they rely heavily on the separability assumption, which suffices for uniqueness of the decomposition (2) but is far from necessary. This issue is related to the following question, which is of central importance in the NMF literature: given a collection of points  $\{r_1, \dots, r_m\} \subseteq \mathbb{R}^{K-1}$  presumed to lie within the convex hull of unobserved vertices  $\{v_1^*, \dots, v_K^*\}$ , when is recovery of these vertices possible? In the NMF context, separability means that each vertex coincides with a point in the observed point cloud, in which case we only need to identify which of the  $r_i$ 's correspond to simplex vertices. However, this is a very strong assumption and several efforts have been made to relax it. Javadi and Montanari (2020) show that vertex recovery is still possible under a uniqueness assumption that generalizes separability. Ge and Zou (2015) introduce the notion of *subset separability* which is also much weaker than separability. We note that many of the separability-based methods proposed in topic modeling, such as those of Arora et al. (2012), Bing et al. (2020a), and Bing et al. (2020b), have no obvious extension if the separability assumption is relaxed. This may not be important if the given corpus contains many specialized words and the topics are sufficiently distinct (an example is a collection of

research papers), but may matter more if the topics overlap significantly and the vocabulary is generic (for instance, a collection of high school English essays).

### 1.2.2 THE SVD-BASED APPROACH OF TRACY KE AND WANG (2024)

Tracy Ke and Wang (2024) are the first to establish the minimax-optimal rate of  $\sqrt{\frac{p}{nN}}$  for the  $\ell_1$ -loss  $\|\hat{A} - A\|_1 := \sum_{j=1}^p \sum_{k=1}^K |\hat{A}_{jk} - A_{jk}|$  where, for simplicity, all document lengths are assumed to be equal to  $N$ . Their procedure links topic estimation to the NMF setting discussed in the previous subsection and is summarized as follows. Let  $M := \text{diag}(n^{-1}D\mathbf{1}_n)$  where  $\mathbf{1}_n := (1, 1, \dots, 1)^T \in \mathbb{R}^n$ . Given  $K$ , the approach of Tracy Ke and Wang (2024) considers the first  $K$  left singular vectors  $\xi_1, \dots, \xi_K \in \mathbb{R}^p$  of  $M^{-1/2}D$ . Elementwise division of  $\xi_2, \dots, \xi_K$  by  $\xi_1$ , also known as *SCORE normalization* (Jin, 2015), yields a matrix  $\check{R} \in \mathbb{R}^{p \times (K-1)}$ , whose rows  $\check{r}_1, \dots, \check{r}_p \in \mathbb{R}^{K-1}$  can be shown to form a point cloud contained in a  $K$ -vertex simplex (up to stochastic errors). Since this corresponds precisely to the NMF setup discussed in the previous subsection, the simplex vertices can now easily be recovered using a suitable *vertex hunting* algorithm. Once these vertices are identified,  $A$  can then be estimated via a series of normalizations.

The work of Tracy Ke and Wang (2024) is an important contribution that motivates several other methods for topic modeling, including ours. However, this method was developed using strong assumptions on the parameter regimes and the behavior of word frequencies. More specifically, Corollary 3.1 of Tracy Ke and Wang (2024) states that the error upper bound  $\sqrt{\frac{p \log n}{nN}}$  is only applicable if we assume  $N > p^{4/3}$  or  $p \leq N < p^{4/3}$  and  $n \geq \max(Np^2, p^3, N^2p^5)$ . As the vocabulary size  $p$  is typically large, these are highly unrealistic assumptions on  $(n, N, p)$ . For example, the Associated Press (AP) dataset used by Tracy Ke and Wang (2024) (a corpus of news articles frequently used for topic model evaluation) has  $n = 2,134$  and  $p = 7,000$ . A typical AP article has between 300 and 700 words, so it is clear that none of the above assumptions holds. The error bound provided without these assumptions is  $\frac{p^2}{N\sqrt{N}} \sqrt{\frac{p \log n}{nN}}$ , which, when  $p$  is large and grows with  $n$ , may not necessarily converge to zero. Several other works that claim to establish minimax-optimal rates also do so by assuming  $N > p$ ; see Theorem 4.1 of Wu et al. (2023) and Remark 10 of Bing et al. (2020a).

In this paper, we do not seek to re-establish the rate  $\sqrt{\frac{p}{nN}}$ . Rather, we aim to provide a consistent error bound valid for all realistic parameter regimes (especially when  $p > \max(n, N)$ ). We propose to resolve some of the outstanding issues of the estimator of Tracy Ke and Wang (2024) by leveraging a sparsity structure that is often empirically observed in text documents, resulting in:

- Improved error bounds:** We observe that even the minimax-optimal rate  $\sqrt{\frac{p}{nN}}$  of Tracy Ke and Wang (2024) scales significantly with  $p$ . As the number of documents  $n$  increases, we can expect several previously unobserved words to be added to the corpus, whereas the average document length  $N$  may not change by much. However, many of these words may occur rarely, so the effective dimension of the parameter space may be quite small compared to the observed vocabulary size. This motivates us to restrict the parameter space by imposing a suitable column-wise sparsity as-

sumption on  $A$ , which enables an error bound that does not scale with  $p$  except for log factors.

2. **An increased signal-to-noise ratio:** The approach of Tracy Ke and Wang (2024) may not be suitable if many words in the corpus occur with low frequency. If for each word  $j$  we define  $h_j := \sum_{k=1}^K A_{jk}$ , the theoretical guarantees of Tracy Ke and Wang (2024) require  $\min_{1 \leq j \leq p} h_j \geq \frac{cK}{p}$  for some  $c \in (0, 1)$ . Note that since the columns of  $A$  sum up to 1, we always have  $\frac{1}{p} \sum_{j=1}^p h_j = \frac{K}{p}$ . Therefore, since  $h_j$  roughly indicates the frequency of word  $j$  in the corpus, this assumption restricts the frequencies of the least frequent words to be of the same order as the average frequency of all words.

Such a restrictive assumption is needed in the analysis of Tracy Ke and Wang (2024) because when many low-frequency words exist in the corpus, their procedure involves division by small and noisy numbers. This is a problem with their *pre-SVD normalization* step where  $D$  is pre-multiplied by the  $p \times p$  diagonal matrix  $M^{-1/2}$ , as the diagonal entries of  $M := \text{diag}(n^{-1}D\mathbf{1}_n)$  corresponding to infrequent words are usually small. This is also an issue with their elementwise division step, thus leading to higher errors from infrequent words in the point cloud obtained from their procedure (see Figures 8 and 19 for illustration). Although we also use SCORE normalization (Jin, 2015), our removal of infrequent words leads to a point cloud with a higher signal-to-noise ratio.

### 1.2.3 SPARSE TOPIC MODELING APPROACHES

To our knowledge, Bing et al. (2020b) and Wu et al. (2023) are the only two prior works that, like ours, impose additional sparsity constraints on  $A$ . However, the sparsity assumptions proposed in these papers are not appropriate for dealing with large  $p$ ; rather, they are more suitable for dealing with large  $K$ .

1. Bing et al. (2020b) assume that  $A$  is elementwise sparse, in the sense that the total number of nonzero entries of  $A$  (denoted as  $\|A\|_0$ ) is small. Their proposed procedure is then shown to satisfy the error upper bound

$$\|\hat{A} - A\|_1 \lesssim K \sqrt{\frac{\|A\|_0 \log(p \vee n)}{nN}}. \quad (3)$$

We note here that  $\|A\|_0$  can still be very large. Indeed, let  $\tilde{p}$  denote the number of words whose corresponding rows in  $A$  are *not* entirely zero. Technically we can have  $p > \tilde{p}$ , but words corresponding to zero rows of  $A$  are not observed with probability one, so  $\tilde{p}$  covers the entire set of all distinct words observed in the corpus. We have

$$\tilde{p} \leq \|A\|_0 \leq Kp. \quad (4)$$

In fact, one can see that their error bound depends on  $p$  from the error decomposition  $\|\hat{A} - A\|_1 \lesssim \text{I} + \text{II} + \text{III}$  in Theorem 2 of Bing et al. (2020b). For example,  $\text{I} = \frac{K}{\underline{\gamma}} \sqrt{\frac{p \log(n \vee p)}{nN}} + \frac{pK \log(p \vee n)}{\underline{\gamma} nN}$  for some constant  $\underline{\gamma}$ . This, together with (4), shows that the bound (3) is not very different from the rate  $\sqrt{\frac{p}{nN}}$  of Tracy Ke and Wang (2024),

except for possibly better dependence on  $K$ . Moreover, their theoretical results depend on several strong assumptions on the frequency of anchor words selected by their procedure. In contrast, our procedure is less affected by the frequency of anchor words, both in theory and in practice.

2. Wu et al. (2023) assume that each row of  $A$  has at most  $s_A$  nonzero entries. Since  $A$  has  $K$  columns, this sparsity assumption is only useful if  $K$  is large. Theorem 4.1 of Wu et al. (2023) then shows that their proposed estimator of  $A$  satisfies

$$\|\hat{A} - A\|_1 \lesssim K \sqrt{\frac{s_A \log n}{nN}}. \quad (5)$$

However, upon close examination of their proof, the  $\ell_1$  bound they achieve is actually  $K \sqrt{\frac{\|A\|_0 \log n}{nN}}$  (similar to (3)) so (5) is only possible by assuming that  $p = O(1)$  and using  $\|A\|_0 \leq ps_A$ . Furthermore, their result assumes  $N^{3/4} \geq p$  which, as we have noted in our discussion of Tracy Ke and Wang (2024), is highly restrictive.

In comparison with these two papers, our sparsity assumption is more compatible with the “large  $p$ ” setting, and we do not assume  $p = O(1)$  as in Wu et al. (2023).

### 1.3 Our Contributions

We summarize the main contributions of this paper below.

- We propose a new spectral procedure (Definition 8) for estimating  $A$ . This procedure takes into account the observation that, in most text datasets, the vocabulary size  $p$  is often large but many words occur very infrequently in the corpus. When  $K$  is unknown, a new estimator of  $K$  is also proposed (see Lemma 22).
- We introduce a new column-wise  $\ell_q$ -sparsity assumption (Assumption 5) for  $A$ . This assumption is motivated by Zipf’s law (Zipf, 1936) and links a word’s frequency of occurrence in a topic to its rank. Our proposed procedure is then shown to be adaptive to the unknown sparsity level  $s$  in the  $\ell_q$ -sparsity definition (15).
- We provide an error bound for our procedure using the  $\ell_1$  loss  $\|\hat{A} - A\|_1$  in Theorem 19. Under our sparsity assumption (15), our error bound is shown to be valid for all parameter regimes and only depends on  $p$  via weak factors. The common pre-processing step of removing infrequent words is incorporated into our procedure and accounted for in our analysis.
- Finally, in Section 2.4, we show that our theoretical results may still hold when the separability assumption is relaxed, provided that we choose a suitable vertex hunting procedure for non-separable point clouds in Definition 8.

Extensive experiments with synthetic datasets to confirm the effectiveness of our estimation procedure under a wide variety of parameter regimes are presented in Section 3. Furthermore, we also demonstrate the usefulness of our method for text analysis, as well as for other applications where the pLSI model is relevant, in Section 4.

### 1.4 Notations

For any set  $S$ , let  $|S|$  denote its cardinality, and let  $S^c$  denote its complement if it is clear in context which superset it refers to. For any  $k \in \mathbb{N}$ , let  $[k]$  denote the index set  $\{1, \dots, k\}$ . We use  $\mathbf{1}_d$  to denote the vector in  $\mathbb{R}^d$  with all entries equal to 1. For a general vector  $v \in \mathbb{R}^d$ , let  $\|v\|_r$  denote the vector  $\ell_r$  norm, for  $r = 0, 1, \dots, \infty$ , and let  $\text{diag}(v)$  denote the  $d \times d$  diagonal matrix with diagonal entries equal to entries of  $v$ . For any  $a, b \in \mathbb{R}$ , let  $a \vee b := \max(a, b)$  and  $a \wedge b := \min(a, b)$ .

Let  $I_m$  denote the  $m \times m$  identity matrix. For a general matrix  $Q \in \mathbb{R}^{m \times l}$  and  $r = 0, 1, \dots, \infty$ , let  $\|Q\|_r$  denote the vector  $\ell_r$ -norm of  $Q$  if one treats  $Q$  as a vector. Let  $\|Q\|_F$  and  $\|Q\|_{\text{op}}$  denote the Frobenius norm (that is,  $\|Q\|_F = \|Q\|_2$ ) and operator norm of  $Q$ , respectively. For any index  $j \in [m]$  and  $i \in [l]$ , let  $Q_{ji}$  or  $Q(j, i)$  denote the  $(j, i)$ -entry of  $Q$ . For index sets  $J \subseteq [m]$  and  $I \subseteq [l]$ , let  $Q_{JI}$  denote the submatrix of  $Q$  obtained by selecting only rows in  $J$  and columns in  $I$  (in particular, either  $J$  or  $I$  can be a single index). Also, let  $Q_{J*}$  denote the submatrix of  $Q$  obtained by selecting rows in  $J$  and *all* columns of  $Q$ ;  $Q_{*I}$  is similarly defined. This means  $Q_{j*}$  and  $Q_{*i}$  denote the  $j^{\text{th}}$  row and  $i^{\text{th}}$  column of  $Q$  respectively. For an integer  $k \leq m \wedge l$ , let  $\sigma_k(Q)$  denote the  $k^{\text{th}}$  largest singular value of  $Q$ , and if  $Q$  is a square matrix then, if applicable, let  $\lambda_k(Q)$  denote the  $k^{\text{th}}$  largest eigenvalue of  $Q$ . If  $m = l$ , then  $\text{tr}(Q)$  denotes the trace of  $Q$ .

Let  $C, c > 0$  denote absolute constants that may depend on  $K$  and  $q$ ; we assume that  $K$  and  $q$  are fixed, unobserved constants. Let  $C^*, c^* > 0$  denote numerical constants that do not depend on the unobserved quantities like  $K$  and  $q$  (this only matters when we discuss the estimation of  $K$ ). The constants  $C, c, C^*, c^*$  may change from line to line.

In our paper, for ease of presentation, we assume  $N_1 = \dots = N_n = N$ . Our results also hold if we assume the document lengths satisfy  $\max_{i \in [n]} N_i \leq C^* \min_{i \in [n]} N_i$  (that is, if  $N_1 \asymp \dots \asymp N_n$ ), in which case  $N = \frac{1}{n} \sum_{i=1}^n N_i$  denotes the average document length.

## 2. Our Procedure for Estimating $A$ and Its Theoretical Properties

For simplicity, we will first assume separability in order to explain our procedure. A discussion of possible relaxations of this condition will be deferred to Section 2.4.

### 2.1 The Oracle Procedure to Estimate $A$ Given $D_0$

Our oracle procedure concerns how  $A$  can be estimated if the non-stochastic matrix  $D_0$ , rather than  $D$ , is observed. Let  $J \subseteq [p]$  be an arbitrary collection of words in our vocabulary. We first need the definition of a vertex hunting procedure, which is relevant to the NMF setup discussed in Section 1.2.1.

**Definition 2 (Vertex hunting)** *Given  $K$ , a vertex hunting procedure is a function that takes a collection of points in  $\mathbb{R}^{K-1}$  and returns  $K$  points in  $\mathbb{R}^{K-1}$ .*

**Remark 3** A good vertex hunting procedure should return the vertices of the smallest simplex with  $K$  vertices containing the given point cloud. Throughout the paper, we will use  $\mathcal{V}(\cdot)$  to denote such a procedure.

The following definition of an ideal point cloud is based on the separability assumption. Any reasonable vertex hunting procedure should be able to successfully recover the simplex vertices from an ideal point cloud.

**Definition 4 (Ideal point cloud)** *Given  $K$ , an ideal point cloud is a collection of points in  $\mathbb{R}^{K-1}$  contained in the simplex defined by  $K$  vertices, such that the  $K$  vertices themselves belong to the point cloud.*

We are now ready to define the oracle procedure.

**Definition 5 (Oracle procedure)** *Given inputs  $K$ ,  $D_0$ , vertex hunting procedure  $\mathcal{V}(\cdot)$  and a set of words  $J \subseteq [p]$ , the oracle procedure returns  $\tilde{A} \in \mathbb{R}^{p \times K}$  defined as follows:*

1. (SVD) Perform SVD on  $[D_0]_{J^*}$  to obtain  $\Xi = [\xi_1, \dots, \xi_K] \in \mathbb{R}^{|J| \times K}$  containing the first  $K$  left singular vectors of  $[D_0]_{J^*}$ .
2. (Elementwise division) Divide  $\xi_2, \dots, \xi_K$  elementwise by  $\xi_1$  to obtain  $R \in \mathbb{R}^{|J| \times (K-1)}$ . This means  $R_{jk} = \xi_{k+1}(j)/\xi_1(j)$ , for  $k = 1, \dots, K-1$  and  $j \in J$ .
3. (Vertex hunting) Treat the rows  $\{r_j : j \in J\}$  of  $R$  as a point cloud in  $\mathbb{R}^{K-1}$ . Apply the vertex hunting procedure  $\mathcal{V}(\cdot)$  on this point cloud to obtain vertices  $v_1^*, \dots, v_K^*$ .
4. (Recovery of  $\Pi$ ) For each  $j \in J$ , solve for  $\pi_j \in \mathbb{R}^K$  from the linear equation

$$\begin{pmatrix} 1 & \dots & 1 \\ v_1^* & \dots & v_K^* \end{pmatrix} \pi_j = \begin{pmatrix} 1 \\ r_j \end{pmatrix}. \quad (6)$$

In other words,  $\pi_j$  satisfies  $\sum_{k=1}^K \pi_j(k) = 1$  and  $r_j = \sum_{k=1}^K \pi_j(k) v_k^*$ , for each  $j \in J$ . Let  $\Pi \in \mathbb{R}^{|J| \times K}$  be the matrix whose rows are  $\{\pi_j : j \in J\}$ .

5. (Normalization) Normalize the columns of  $\text{diag}(\xi_1) \cdot \Pi \in \mathbb{R}^{|J| \times K}$  so that the entries of each column sum up to 1. This yields  $\tilde{A}_{J^*}$ . Set  $\tilde{A}_{J^c} = 0$  to obtain  $\tilde{A}$ .

Our oracle procedure makes use of the SCORE normalization idea which was originally proposed for network data analysis (Jin, 2015). The elementwise division step (Step 2) is the most important step, as it provides a connection between singular vectors of  $D_0$  (or associated variables) and the NMF setup described in Section 1.2.1. The words in  $J$  are represented by the point cloud  $\{r_j : j \in J\}$ , which can be shown to be contained entirely in some  $K$ -vertex simplex. If the simplex vertices are identifiable and the vertex hunting procedure is successful in recovering them in Step 3, then (6) allows us to exactly recover the probabilistic weights  $\{\pi_j : j \in J\}$  associated with each word in  $J$ , which are connected to  $A$  via the relation

$$\text{diag}(\xi_1) \cdot \Pi = A_{J^*} \cdot \text{diag}(V_1), \quad (7)$$

for some vector  $V_1 \in \mathbb{R}^K$  containing only positive entries. This explains the column normalization step (Step 5), which essentially reverses the elementwise division step. For more details, we refer the reader to the proof of Lemma 27 in the appendix.

Based on the relation (7), we can show the following result.

**Lemma 6** *Suppose the set  $J$  contains at least one anchor word for each topic  $k \in [K]$ , and the vertex hunting procedure  $\mathcal{V}(\cdot)$  can successfully recover the simplex vertices from any ideal point cloud. The oracle procedure in Definition 5 then returns  $\tilde{A}$  satisfying  $\tilde{A}_{J^c*} = 0$  and*

$$\tilde{A}_{J*} = A_{J*} \cdot \text{diag}(\|A_{J1}\|_1^{-1}, \dots, \|A_{JK}\|_1^{-1}). \quad (8)$$

The proof of Lemma 6 is identical to that of Lemma 27 provided in Appendix B. The sole difference is that in Lemma 27, the set  $J$  is chosen as in (9) and we use Assumption 3.

**Remark 7** Our oracle procedure differs from that of Tracy Ke and Wang (2024) in two important ways. First, note that Step 1 only requires SVD to be performed on a submatrix of  $D_0$ . In general, we want the set  $J$  to contain words that occur with sufficiently high frequencies in the corpus so that the point cloud generated from our procedure has a higher signal-to-noise ratio. When  $p$  is large, we can often expect the corpus to contain many infrequently occurring words whose corresponding rows in  $A$  should be estimated as zero. Our oracle procedure yields  $\tilde{A}$  which is a good oracle approximation of  $A$  if  $\|A_{J^c*}\|_1$  is small, as in that case the diagonal matrix in (8) is close to the identity matrix.

Second, note that we consider the SVD of a submatrix of  $D_0$  and not  $M_0^{-1/2}D_0$  as in Tracy Ke and Wang (2024), where  $M_0 := \text{diag}(n^{-1}D_0\mathbf{1}_n)$ . This simplifies some parts of our theoretical analysis and allows us to obtain error bounds that depend less strongly on  $p$  (see Appendix C).

## 2.2 Estimation Procedure for $A$ Given $D$

Our procedure to estimate  $A$  below is designed to closely approximate the oracle procedure. Here we first assume  $K$  is known. The estimation of  $K$  is deferred to Section 2.5, and the choice of the vertex hunting procedure will be discussed in conjunction with identifiability assumptions on  $A$ .

**Definition 8 (Estimation procedure for  $A$ )** *Given inputs  $K$ , observation matrix  $D$  and vertex hunting procedure  $\mathcal{V}(\cdot)$ , our estimation procedure returns  $\hat{A}$  defined as follows:*

1. (Thresholding) Let  $M := \text{diag}(n^{-1}D\mathbf{1}_n)$  and  $p_n := p \vee n$ . Compute the set of words

$$J := \left\{ j \in [p] : M(j, j) \geq \alpha \sqrt{\frac{\log p_n}{nN}} \right\}. \quad (9)$$

Here,  $\alpha$  is a user-specified universal constant (see Remark 9).

2. (Spectral decomposition) Compute the first  $K$  eigenvectors  $\hat{\xi}_1, \dots, \hat{\xi}_K \in \mathbb{R}^{|J|}$  of the submatrix  $G_{JJ}$  of the  $p \times p$  matrix  $G$ , where

$$G := DD^T - \frac{n}{N}M. \quad (10)$$

Here, we assume all entries of  $\hat{\xi}_1$  are of the same sign, in which case we can choose  $\hat{\xi}_1$  to have all positive entries. If some entries of  $\hat{\xi}_1$  are negative, choose  $\hat{\xi}_1$  such that the majority of entries are positive, and apply Remark 10.

3. (*Elementwise division*) Divide  $\hat{\xi}_2, \dots, \hat{\xi}_K$  elementwise by  $\hat{\xi}_1$  to obtain  $\hat{R} \in \mathbb{R}^{|J| \times (K-1)}$ , with rows  $\{\hat{r}_j : j \in J\}$ . This means  $\hat{R}_{jk} = \hat{\xi}_{k+1}(j)/\hat{\xi}_1(j)$ , for  $k \in [K-1]$  and  $j \in J$ .
4. (*Vertex hunting*) Treat the rows of  $\hat{R}$  as a point cloud in  $\mathbb{R}^{K-1}$ . Apply the vertex hunting procedure  $\mathcal{V}(\cdot)$  to this point cloud to obtain vertices  $\hat{v}_1^*, \dots, \hat{v}_K^*$ .
5. (*Estimation of  $\Pi$* ) For each  $j \in J$ , solve for  $\hat{\pi}_j^\diamond \in \mathbb{R}^K$  from

$$\begin{pmatrix} 1 & \dots & 1 \\ \hat{v}_1^* & \dots & \hat{v}_K^* \end{pmatrix} \hat{\pi}_j^\diamond = \begin{pmatrix} 1 \\ \hat{r}_j \end{pmatrix}.$$

Obtain  $\hat{\pi}_j$  from  $\hat{\pi}_j^\diamond$  by first setting any negative entries to 0 and then normalizing so that the entries of  $\hat{\pi}_j$  sum up to 1. Let  $\hat{\Pi}$  be the matrix whose rows are  $\{\hat{\pi}_j : j \in J\}$ .

6. (*Normalization*) Normalize all columns of  $\text{diag}(\hat{\xi}_1) \cdot \hat{\Pi}$  so that they have unit  $\ell_1$ -norm. This yields  $\hat{A}_{J^*}$ . Set all entries of  $\hat{A}_{J^*}$  to zero to obtain  $\hat{A}$ .

As Steps 3-5 are also based on the SCORE normalization idea (Jin, 2015), we call this procedure the *Thresholded Topic-SCORE* (TTS). However, Step 1, Step 2 and Step 6 contain significant differences when compared with Topic-SCORE of Tracy Ke and Wang (2024).

**Remark 9 (Choice of  $\alpha$ )** The set  $J$  in (9) is chosen by examining the row sums of the observation matrix  $D$ , which indicate how frequently the words occur in the corpus. In (9),  $\alpha$  is meant to be a universal constant and thus does not affect our error rates, which are not optimized over constants. In our theoretical discussion, we choose  $\alpha = 8$  for convenience, but for most datasets this value of  $\alpha$  may result in too many words not meeting the threshold.

In practice, a good choice of  $\alpha$  is important for obtaining a good estimator of  $A$ . Based on our experiments, we recommend a smaller value of  $\alpha$ , such as  $\alpha = 0.005$ . The finite-sample choice  $\alpha = 0.005$  used in our experiments is a practical tuning choice and is not meant to be the constant used in the preceding asymptotic probability bound. This choice of  $\alpha$  should produce reasonable results for commonly observed values of  $(n, N, p)$ . Based on what we observe from experiments, if  $n \in [1000, 5000]$ ,  $N \in [300, 700]$ ,  $p \in [5000, 10000]$ , we can typically expect around 10-40% of words to be removed.

**Remark 10 (Signs of  $\hat{\xi}_1$ 's entries)** In the oracle procedure,  $\xi_1$  is the first left singular vector of  $[D_0]_{J^*}$  and so by Perron's theorem, the entries of  $\xi_1$  are all positive. In Step 2,  $\hat{\xi}_1$  is the first eigenvector of  $G_{JJ}$  which is not necessarily a Perron matrix, so  $\hat{\xi}_1$  technically may contain negative entries. Any word  $j$  for which  $\hat{\xi}_1(j)$  is negative should have the corresponding row of  $\hat{A}$  set to zero after Step 2, and then in Step 3 we form the point cloud by computing  $\hat{\xi}_{k+1}(j)/\hat{\xi}_1(j)$  for  $k \in [K-1]$  and  $j \in J$  with  $\hat{\xi}_1(j) > 0$  only.

In our theoretical analysis as well as in practice, however, this scenario will not happen with high probability. This is because  $G$  is chosen so that  $\max_{j \in J} |\hat{\xi}_1(j) - \xi_1(j)|$  is small. Since any word  $j$  that meets our threshold occurs with sufficiently high frequency,  $\xi_1(j)$  will also be sufficiently large for any  $j \in J$ , which implies  $|\hat{\xi}_1(j) - \xi_1(j)| \ll \xi_1(j)$  and thus  $\hat{\xi}_1(j) \geq \xi_1(j)/2$  for all  $j \in J$ . See Lemmas 40 and 41 in the appendix.

The set  $J$  as defined in (9) is data-dependent. It is quite useful to note that  $J$  can be approximated by the non-stochastic sets (11) with high probability. The proof of the lemma below can be found in Theorem 25(b) in the appendix.

**Lemma 11** *Let  $M_0 := \text{diag}(n^{-1}D_0\mathbf{1}_n)$ , and let*

$$J_{\pm} := \left\{ j \in [p] : M_0(j, j) > \alpha_{\pm} \alpha \sqrt{\frac{\log p_n}{nN}} \right\}, \quad (11)$$

where  $\alpha$  is from the definition of  $J$  in (9) and  $\alpha_- > 1$  and  $\alpha_+ \in (0, 1)$  are some suitably chosen constants depending on  $\alpha$  (for example if  $\alpha = 8$ , we can let  $\alpha_+ = \frac{1}{2}$ ,  $\alpha_- = 2$ ). Then the event  $\mathcal{E} := \{J_- \subseteq J \subseteq J_+\}$  occurs with probability at least  $1 - o(p_n^{-1})$ .

The following lemma bounds the size of  $J$ , and is obtained by bounding  $|J_+|$  and using Lemma 11.

**Lemma 12 (Size of  $J$ )** *With probability at least  $1 - o(p_n^{-1})$ ,*

$$|J| \leq \left( \frac{K}{\alpha\alpha_+} \sqrt{\frac{Nn}{\log p_n}} \right) \wedge p. \quad (12)$$

Our procedure requires the eigenvalue decomposition of a symmetric  $|J| \times |J|$  matrix. The bound (12) can be significantly smaller than  $\min(n, p)$  if  $nN \ll p^2$  and  $N \ll n$  (ignoring weak factors), which are reasonable assumptions for many text datasets. We can therefore expect the eigenvalue decomposition step in our procedure to be more computationally scalable than the SVD step (on a  $p \times n$  matrix) of Tracy Ke and Wang (2024).

### 2.3 Error Bounds for $\hat{A}$ Under Separability

We first discuss our theoretical results under separability, which is assumed in all of our proofs in the appendix. We begin by listing the assumptions underlying our analysis.

**Assumption 1 (A and W are well-conditioned)** *Let  $\Sigma_W := n^{-1}WW^T$ . For some constant  $c \in (0, 1)$ ,*

$$\sigma_K(A) \geq c\sqrt{K} \quad \text{and} \quad \sigma_K(\Sigma_W) \geq c. \quad (13)$$

**Assumption 2 (The topic-topic correlation matrix is regular)** *The entries of  $A^T A$  satisfy the following for some constant  $c > 0$ :*

$$\min_{1 \leq k, l \leq K} A^T A(k, l) \geq c. \quad (14)$$

**Assumption 3 (Separability)** *Each topic  $k \in [K]$  has at least one associated anchor word  $j$  belonging to the set  $J_-$  defined in (11).*

**Assumption 4 (Vertex hunting efficiency)** *Given  $K$  and an ideal point cloud defined in Definition 4, the vertex hunting function  $\mathcal{V}(\cdot)$  recovers the  $K$  vertices correctly. Furthermore, whenever  $\mathcal{V}(\cdot)$  is given as inputs two point clouds  $\{x_1, \dots, x_m\}$  and  $\{x'_1, \dots, x'_m\}$ , the outputs  $\{v_1, \dots, v_K\}$  and  $\{v'_1, \dots, v'_K\}$  satisfy for some absolute constant  $C > 0$  (up to a label permutation)*

$$\max_{k \in [K]} \|v_k - v'_k\|_2 \leq C \max_{j \in [m]} \|x_j - x'_j\|_2.$$

**Assumption 5 (Column-wise  $\ell_q$ -sparsity)** *Let the entries of each column  $A_{*k}$  of  $A$  be ordered as  $A_{(1)k} \geq \dots \geq A_{(p)k}$ . For some  $q \in (0, 1)$  and  $s > 0$ , the columns of  $A$  satisfy*

$$\max_{k \in [K]} \left( \max_{j \in [p]} j A_{(j)k}^q \right) \leq s. \quad (15)$$

Here, we assume that  $q$  is a fixed constant, whereas  $s$  is allowed to grow with  $n$ .

**Remark 13** We justify why Assumptions 1, 2 and 3 are reasonable below.

1. Equation (13) assumes the topic vectors in  $A$  are not too correlated. The assumption on  $W$  in (13) is necessary even when  $W$  is known, as its role is similar to that of the design matrix in the regression setting. Note that since the columns of  $A$  and  $W$  sum up to 1, we always have  $\sigma_1(A) \leq \sqrt{K}$  and  $\sigma_1(\Sigma_W) \leq 1$  (see Lemma 26(a) in the appendix).
2. The matrix  $A^T A \in \mathbb{R}^{K \times K}$  can be thought of as the topic-topic correlation matrix, since its entries are inner products of the columns of  $A$ . Therefore, (14) is especially true if the  $K$  topics are related to one another. However, even if the corpus covers unrelated topics, we expect all columns of  $A$  to assign significant weights to grammatical function words (such as ‘and’, ‘the’ in English) and filler words, which occur frequently in all documents regardless of the topics involved.
3. In light of Lemma 11, Assumption 3 requires that each topic has at least one anchor word that occurs in the corpus frequently enough so that it is included in  $J$ . Such an assumption on the frequency of anchor words is also commonly seen in other works that exploit the separability condition, and Assumption 3 is not strong since the threshold level of order  $\sqrt{\frac{\log pn}{nN}}$  in the definition of  $J_-$  is quite low. For comparison, Bing et al. (2020b) make the same assumption but with the threshold level of order  $\frac{\log pn}{N}$ , which may be higher than ours if the number of documents  $n$  far exceeds the average document length  $N$ .

**Remark 14 (Vertex hunting for separable point clouds)** Tracy Ke and Wang (2024) mention two vertex hunting algorithms which are suitable for separable point clouds, namely Successive Projection (SP) (Araújo et al., 2001) and Sketched Vertex Search (SVS) (Jin et al., 2017).

Given a point cloud  $r_1, \dots, r_m$ , SP starts by finding the point  $r_j$  whose Euclidean norm is the largest and sets this as the first estimated vertex  $\hat{v}_1$ . Then, for each  $2 \leq k \leq K$ , we can obtain  $\hat{v}_k$  from  $\hat{v}_1, \dots, \hat{v}_{k-1}$  by setting  $\hat{v}_k$  as the point  $r_j$  that maximizes  $\|(I - P_{k-1})r_j\|_2$ , where  $P_{k-1}$  denotes the projection matrix onto the linear span of  $\hat{v}_1, \dots, \hat{v}_{k-1}$ . SP can be shown to satisfy Assumption 4 when the volume of the true simplex is lower bounded by a constant (Gillis and Vavasis, 2013), which is a simple consequence of Lemma 26(f) in the appendix.

On the other hand, SVS starts by applying  $k$ -means clustering on the point cloud  $\{r_1, \dots, r_m\}$  to obtain cluster centers  $\hat{c}_1, \dots, \hat{c}_L$ , where  $L$  is a tuning parameter that is much larger than  $K$ . These clusters are meant to reduce the noise levels in the point cloud. Next, SVS exhaustively searches for all simplexes whose  $K$  vertices are located on these

cluster centers, in order to find the simplex  $S$  such that the maximum distance from any  $\hat{c}_l$  to  $S$  is minimized. In comparison to SP, SVS is more robust to noise in the point cloud but is computationally much slower if  $K$  is not small. SVS satisfies Assumption 4 under mild regularity conditions (Jin et al., 2017).

Note that these vertex hunting algorithms are only meant for separable point clouds, as the simplex vertices they produce are designed to belong to the convex hull of the point cloud. For more implementation details of SVS and SP, we refer the reader to Section A of Tracy Ke and Wang (2024).

**Remark 15 ( $\ell_q$ -sparsity)** To our knowledge, our work is the first to consider the  $\ell_q$ -sparsity assumption (15) in the topic modeling context, although similar assumptions have been adopted in other statistical settings such as sparse PCA and sparse covariance estimation (see, for example, Ma, 2013; Cai and Zhou, 2012). (15) imposes an assumption on the decay rate of the ordered entries of the columns of  $A$ , but does not restrict how small (or large) the smallest (or largest, assuming  $s \geq 1$ ) entries of  $A$ 's columns can be. Thus, our theoretical results are valid even in the presence of severe word frequency heterogeneity.

Note that if a column  $A_{*k}$  has  $s$  nonzero entries, then we always have  $\max_{j \in [p]} j A_{(j)k}^q \leq s$ . However, in light of (4) where we observe that  $\|A\|_0 \geq \tilde{p}$ , there exists at least one column of  $A$  with at least  $\lfloor \tilde{p}/K \rfloor$  nonzero entries, and so  $s$  in (15) cannot be much smaller than  $p$  if we impose hard sparsity ( $q = 0$ ) on *all* columns of  $A$ . Therefore, the  $\ell_q$ -sparsity assumption (15) gives us more flexibility as it allows for the possibility that most entries of  $A$  are small but nonzero. When  $q \approx 0$ , we can approximate the assumption of hard sparsity on all columns of  $A$ , whereas when  $q$  is close to 1, then (15) with  $s = O(1)$  corresponds to Zipf's law, which is the empirical observation that word frequency in text data is often inversely proportional to word rank.

The restriction that  $q \in (0, 1)$  is primarily due to the fact that we use the  $\ell_1$  loss  $\|\hat{A} - A\|_1$ . Since the columns of  $A$  sum up to 1, the columns of  $A$  already satisfy  $\ell_q$ -sparsity with  $q = s = 1$ , but this alone is not sufficient to control the error term  $\|A_{J^c}\|_1$  resulting from our thresholding step.

We are now ready to discuss our main theoretical results. Let  $\hat{\Xi} = [\hat{\xi}_1, \dots, \hat{\xi}_K] \in \mathbb{R}^{|J| \times K}$  contain the first  $K$  eigenvectors of  $G_{JJ}$  where  $G$  is defined as in (10). Recall its oracle counterpart  $\Xi = [\xi_1, \dots, \xi_K] \in \mathbb{R}^{|J| \times K}$  which contains the first  $K$  left singular vectors of  $[D_0]_{J*}$ . Let  $\{\Xi_j : j \in J\}$  and  $\{\hat{\Xi}_j : j \in J\}$  denote the rows of  $\Xi$  and  $\hat{\Xi}$  respectively.

**Lemma 16 (Row-wise error bounds for  $\hat{\Xi}$ )** *For all  $j \in [p]$ , let  $h_j := \sum_{k=1}^K A_{jk}$ . With probability  $1 - o(p_n^{-1})$ , there exist  $\omega \in \{\pm 1\}$  and a  $(K-1) \times (K-1)$  orthonormal matrix  $\Omega^*$  such that, if we define  $\Omega := \text{diag}(\omega, \Omega^*) \in \mathbb{R}^{K \times K}$ , we have*

$$\|\Omega \hat{\Xi}_j - \Xi_j\|_2 \leq C \sqrt{\frac{h_j \log p_n}{nN}} \quad \text{for all } j \in J. \quad (16)$$

The proof can be found in Lemma 40 and is an application of the well-known Davis-Kahan theorem (more specifically, we need to use the row-wise perturbation version of the theorem as proven in Lemma F.1 of Tracy Ke and Wang, 2024). We note here that the bound (16) depends on  $p$  only via the log term, and the  $h_j$ 's, which indicate how frequently one may encounter word  $j$  in the corpus, determine the magnitude of the bound (16).

As a consequence of the above lemma, one can provide error bounds for the point cloud obtained from our procedure. Again, recall that  $\{r_j : j \in J\}$  is the oracle point cloud from Step 3 of Definition 5, and  $\{\hat{r}_j : j \in J\}$  is the point cloud from Step 4 of Definition 8.

**Corollary 17 (Error bounds for the point cloud)** *With probability  $1 - o(p_n^{-1})$ , there exists a  $(K - 1) \times (K - 1)$  orthonormal matrix  $\Omega^*$  such that*

$$\max_{j \in J} \|\Omega^* \hat{r}_j - r_j\|_2 \leq C \left( \frac{\log p_n}{nN} \right)^{1/4}. \quad (17)$$

The proof can be found in Lemma 41. To elaborate further on (17), we can show that with high probability,

$$\|\Omega^* \hat{r}_j - r_j\|_2 \leq C \sqrt{\frac{\log p_n}{h_j n N}} \quad \text{for all } j \in J. \quad (18)$$

Observe that unlike (16), the bound (18) is inversely proportional to  $\sqrt{h_j}$  due to the fact that the point cloud is obtained from the elementwise division step. Since we do not restrict how small  $\min_{1 \leq j \leq p} h_j$  can be, the error bound (18) may be uncontrollable without appropriate thresholding of infrequent words. However, with the choice of  $J$  as in (9), one can show  $\min_{j \in J} h_j \geq c \sqrt{\frac{\log p_n}{nN}}$  with high probability, which when combined with (18) leads to (17).

From (18), we can also obtain bounds on how much the probabilistic weights  $\{\hat{\pi}_j : j \in J\}$  from Step 5 of Definition 8 deviate from the oracle weights  $\{\pi_j : j \in J\}$  from Step 4 of Definition 5. The proof of the following corollary can be found in Lemma 42 of the appendix.

**Corollary 18 (Error bounds for  $\hat{\Pi}$ )** *With probability  $1 - o(p_n^{-1})$ ,*

$$\max_{j \in J} \|\hat{\pi}_j - \pi_j\|_1 \leq C \left( \frac{\log p_n}{nN} \right)^{1/4}. \quad (19)$$

Note that while  $\{\Xi_j : j \in J\}$  and  $\{r_j : j \in J\}$  can be recovered only up to an orthonormal transformation  $\Omega^*$ , the bound (19) does not depend on  $\Omega^*$ . We also note that the bounds (16), (17) and (19) are derived without using the  $\ell_q$ -sparsity assumption (Assumption 5).

The next theorem is our main result, which provides the error rate for estimating  $A$  using the  $\ell_1$  loss  $\|\hat{A} - A\|_1$ . Recall the definition of  $\tilde{A}$  in Lemma 6.

**Theorem 19 (Estimation error for  $\hat{A}$ )** *Suppose Assumptions 1–4 are satisfied. Then with probability  $1 - o(p_n^{-1})$ ,*

$$\|\hat{A}_{J^*} - \tilde{A}_{J^*}\|_1 \leq C \left( \frac{\log p_n}{nN} \right)^{1/4}. \quad (20)$$

*If we further assume the  $\ell_q$ -sparsity assumption (Assumption 5) and  $s \left( \frac{\log p_n}{nN} \right)^{\frac{1-q}{2}} = o(1)$ , we also have with probability  $1 - o(p_n^{-1})$ ,*

$$\|\tilde{A}_{J^*} - A_{J^*}\|_1 = \|A_{J^*}\|_1 \leq C s \left( \frac{\log p_n}{nN} \right)^{\frac{1-q}{2}}, \quad (21)$$

and therefore with probability  $1 - o(p_n^{-1})$ ,

$$\|\hat{A} - A\|_1 \leq C \left[ \left( \frac{\log p_n}{nN} \right)^{1/4} + s \left( \frac{\log p_n}{nN} \right)^{\frac{1-q}{2}} \right], \quad (22)$$

for some constant  $C$  that may depend on  $K$  and  $q$ .

The proof of the above statements can be found in Appendix E.

**Remark 20** The bounds (20) and (21) can be interpreted as the estimation error and the approximation error respectively for using an estimator of  $A$  whose row support is contained in the set  $J$ . Note that the approximation error (21) is smaller if  $q$  is closer to 0; here we assume  $s$  does not grow too quickly relative to  $nN$ . In the most favorable setting where  $s = O(1)$  and  $0 < q < 1/2$  (strong sparsity regime), the aggregate error (22) is of the order  $\left( \frac{\log p_n}{nN} \right)^{1/4}$ , which clearly converges to zero as  $nN \rightarrow \infty$ . On the other hand, if  $s \geq 1$  and  $1/2 < q < 1$  (weak sparsity regime), the bound (22) is dominated by the term  $s \left( \frac{\log p_n}{nN} \right)^{\frac{1-q}{2}}$ .

**Remark 21** We note again that the bound (22), which does not depend on  $p$  except for log terms, is valid for all parameter regimes and in particular for the high-dimensional setting where  $p \gg \max(n, N)$ . This justifies the use of our method for many text datasets where the number of unique words observed across all documents is extremely large. Also, the bound (22) does not depend on  $\max_{j \in [p]} h_j$  or  $\min_{j \in [p]} h_j$  and is thus completely unaffected by variations in word frequencies. In these respects, our result improves upon the theoretical guarantees presented in prior work (Tracy Ke and Wang, 2024; Bing et al., 2020a; Arora et al., 2012; Wu et al., 2023).

## 2.4 Relaxation of the Separability Condition

Our main result (Theorem 19) may also hold under alternative identifiability assumptions on  $A$  if we use a suitable vertex hunting procedure that is effective even for non-separable point clouds. Recall  $v_1^*, \dots, v_K^*$  are the simplex vertices from the oracle point cloud  $\{r_j : j \in J\}$  in Definition 5 and  $\hat{v}_1^*, \dots, \hat{v}_K^*$  are the estimated vertices based on the point cloud  $\{\hat{r}_j : j \in J\}$  in Definition 8. The assumptions we made concerning separability and vertex hunting efficiency, namely Assumptions 3 and 4, are only useful in our analysis insofar as they allow the following bound to hold with high probability:

$$\max_{k \in [K]} \|\Omega^* \hat{v}_k^* - v_k^*\|_2 \leq C \max_{j \in J} \|\Omega^* \hat{r}_j - r_j\|_2, \quad (23)$$

for the same orthonormal matrix  $\Omega^*$  as in Corollary 17.

However, this bound may also hold if we adopt the identifiability assumption and the Archetype Analysis (AA) vertex hunting procedure proposed by Javadi and Montanari (2020). Figure 1 provides an example of a non-separable point cloud where AA recovers the simplex vertices much more effectively than SP and SVS, which only search for possible vertices within the point cloud itself or its convex hull. Appendix F summarizes important results from Javadi and Montanari (2020) that are relevant to our paper. In our estimation




Figure 1: Top left: a non-separable point cloud (blue) contained in a simplex (black) with 3 vertices. Top right: Estimated vertices from SVS (red). Bottom left: Estimated vertices from SP (red). Bottom right: Estimated vertices from AA (red).

procedure for  $A$ , once we obtain the matrix  $\hat{R}$  whose rows are  $\{\hat{r}_j : j \in J\}$  from Step 3 of Definition 8, the estimated simplex vertices can be obtained via AA by solving the following minimization problem:

$$\text{minimize } \mathcal{D}(V; \hat{R}) \text{ over } V \text{ s.t. } \mathcal{D}(\hat{r}_j; V) \leq \delta^2 \text{ for all } j \in J. \quad (24)$$

Here the rows of  $V$  represent the simplex vertices; see Appendix F for the definition of the distance function  $\mathcal{D}(\cdot, \cdot)$ . The main theoretical result of Javadi and Montanari (2020) (Theorem 48) is that the AA algorithm is robust to noise in the point cloud under certain conditions. In particular, if we replace Assumptions 3 and 4 by the following assumptions:

- (i) The matrix  $R$  from Step 2 of the oracle procedure (Definition 5) satisfies  $\alpha$ -uniqueness for some absolute constant  $\alpha > 0$ . Here,  $\alpha$ -uniqueness (described in Definition 47) is an identifiability assumption on the simplex vertices that is more general than separability.
- (ii) The convex hull of the rows of  $R$  contains a  $(K - 1)$ -dimensional ball of radius  $\mu > 0$
- (iii) The vertex hunting procedure  $\mathcal{V}(\cdot)$  is defined by (24) with  $\delta \asymp \left(\frac{\log p_n}{nN}\right)^{1/4}$ . This value of  $\delta$  is chosen based on Corollary 17 and Theorem 48.

then, in light of Theorem 48, (23) continues to hold and our main result, Theorem 19, remains valid. Alternatively, if we do not wish to use the  $\alpha$ -uniqueness condition for identi-

fiability, we can also assume that the distance from the oracle simplex vertices  $\{v_1^*, \dots, v_K^*\}$  to the convex hull of the oracle point cloud  $\{r_j : j \in J\}$  is not larger than  $\delta$ . In light of Theorem 49, this assumption can also be used to obtain (23).

Aside from Javadi and Montanari (2020), Ge and Zou (2015) also discuss an alternative identifiability assumption called *subset separability*. This notion can be illustrated by the point cloud in Figure 1 (top left), with  $K = 3$ . The point cloud (in blue) is contained in a triangle but is not separable as none of the triangle's vertices belongs to the point cloud. However, each edge of the triangle contains several blue points and thus can clearly be identified from the point cloud. The vertices can then be identified by taking intersections of the edges. Ge and Zou (2015) also provides a vertex hunting procedure which, under subset separability and additional regularity assumptions, can also be shown to be robust to noise in the point cloud, in the sense of (23).

In terms of computation, Javadi and Montanari (2020) describe two algorithms to solve the following Lagrangian variant of (24):

$$\hat{V}_\lambda = \arg \min_V [\mathcal{D}(\hat{R}; V) + \lambda \mathcal{D}(V; \hat{R})]. \quad (25)$$

Note that the objective function in (25) is non-convex and thus may have multiple minima. While AA may significantly reduce statistical error in the vertex hunting step when separability is not applicable, the trade-off is that its computational cost may be higher than that of the SP algorithm for separable point clouds.

## 2.5 Estimation of $K$

Our discussion so far assumes  $K$  is known. When  $K$  needs to be estimated, it is natural to examine the spectrum of any matrix that should be of rank  $K$  under the pLSI model.

Recall the definition of  $G$  in (10), and define  $G_0 := (1 - \frac{1}{N}) D_0 D_0^T$ . From Lemma 36, with probability  $1 - o(p_n^{-1})$  we have (here  $C^*$  is a numerical constant not dependent on unobserved constants but may depend on the choice of  $\alpha$ )

$$\|(G - G_0)_{JJ}\|_{\text{op}} \leq C^* K \sqrt{K} \sqrt{\frac{n \log p_n}{N}}. \quad (26)$$

Furthermore one can show  $[G_0]_{JJ}$  has rank  $K$  with high probability. By a simple application of Weyl's inequality, we then obtain the estimator (28) for  $K$ .

**Lemma 22** *Let  $g_n$  be a quantity satisfying*

$$c \sqrt{\frac{nN}{\log p_n}} \geq g_n \geq C^* K \sqrt{K}, \quad (27)$$

where  $C^*$  in (27) is the constant from (26) and  $c$  is another constant that may depend on  $K$ . If

$$\hat{K} := \max \left\{ k : \lambda_k(G_{JJ}) > g_n \sqrt{\frac{n \log p_n}{N}} \right\}, \quad (28)$$

then  $\hat{K} = K$  with probability  $1 - o(p_n^{-1})$ .




Figure 2: Scree plots of the eigenvalues of  $G_{JJ}$  for three synthetic datasets, with  $K \in \{3, 5, 10\}$ ,  $n = N = 500$  and  $p = 5,000$ . The x-axis is log-scaled. The largest eigenvalue is omitted from the plot. The red dots represent the next  $K - 1$  leading eigenvalues, while the blue dots represent the remaining eigenvalues.

The proof can be found in Corollary 37 of the appendix. In (27), the quantity  $g_n$  needs to be chosen to override the term  $C^*K\sqrt{K}$  but cannot converge to  $+\infty$  too quickly. Without any prior information on  $K$ , one can choose  $g_n$  to be a quantity that slowly converges to  $+\infty$ , such as  $g_n = 8 \log p_n$ . If one has prior knowledge on an upper bound for  $K$  (for example if  $K \leq 30$ ), the quantity  $g_n$  can be determined more specifically.

The estimator (28) is based on the bound (26), which depends on  $K$  and so we need to assume  $g_n \geq C^*K\sqrt{K}$ . However, one can also show that with probability at least  $1 - \frac{1}{n+p}$ ,

$$\|(D - D_0)_{J^*}\|_{\text{op}} \leq 4\sqrt{\frac{n \log(n+p)}{N}}. \quad (29)$$

(see Lemma 4 of Klopp et al., 2023). This bound does not depend on  $K$ . Under similar assumptions on  $\sigma_K(A)$  and  $\sigma_K(W)$ , we can consider the following estimator

$$\hat{K}' := \max \left\{ k : \sigma_k(D_{J^*}) > 4\sqrt{\frac{n \log(p+n)}{N}} \right\}, \quad (30)$$

and also show that, based on (29),  $\hat{K}' = K$  with high probability. The advantage of (28) over (30) is computational: both Step 2 of Definition 8 and (28) use the eigendecomposition of  $G_{JJ}$ , whereas (30) requires us to additionally perform SVD on  $D_{J^*}$ .

There are many choices of the quantity  $g_n$  that may satisfy (27) when  $nN$  is sufficiently large. In practice, the estimation of  $K$  may be sensitive to the choice of the eigenvalue cutoff, and moreover real datasets may not always adhere to our assumptions. As Lemma 22 suggests the spectrum of  $G_{JJ}$  is useful for estimating  $K$ , we note that it is often possible to determine the eigenvalue cutoff by inspecting the scree plot of  $G_{JJ}$ 's eigenvalues. Figure 2 displays the scree plots for several synthetic datasets with different values of  $K$ . In some situations, the top  $K$  eigenvalues of  $G_{JJ}$  are separated from the other eigenvalues by a

discernible gap, thus helping one to visually determine  $K$ . When such a gap is unavailable, one can use the Kneedle algorithm (Satopaa et al., 2011) to find the point of maximum curvature of the scree plot; this is a common technique to determine the number of principal components in principal component analysis.

While this method provides a simple way of selecting  $K$ , we note that the principled estimation of the number of factors (or matrix rank) has a long history. Information-criterion approaches balance fit and complexity (for example, Bai and Ng, 2002; Alessi et al., 2010). Spectral-gap rules detect sharp breaks in the spectrum, including the eigenvalue- and growth-ratio estimators of Ahn and Horenstein (2013). Thresholding or shrinkage of singular values offers another widely used route, with universal or data-adaptive cutoffs (Gavish and Donoho, 2014), and related thresholding ideas now underpin procedures for low-rank inference (for example, Chernozhukov et al., 2023; Choi et al., 2024). In principle, any of these selectors can be used as a plug-in in our setting by replacing the cutoff in (28) (or (30)) with the corresponding rule. We propose the spectral-thresholding rule in (28) for computational alignment with our estimator of  $G_{JJ}$ , but the most appropriate choice may depend on the noise structure, sample size, and modeling goals.

### 3. Experiments with Synthetic Data

In this section, we assess the empirical performance of our estimator through a series of synthetic experiments.<sup>1</sup> The controlled environment provided by these experiments allows us to better understand the behavior of our method in different parameter regimes.

Throughout this section, we benchmark our estimator’s performance against the following well-established methods: (a) Latent Dirichlet Allocation (Blei et al., 2003), the Bayesian counterpart to pLSI; (b) the anchor word recovery (AWR) approach of Arora et al. (2012), a procedure that identifies anchor words from the non-negative factorization of the second-order moment matrix  $DD^T$  and recovers topics by expressing each word’s weight as a nonnegative convex combination of the anchor words; (c) the Topic-SCORE procedure of Tracy Ke and Wang (2024); and (d) the Sparse Topic Model solver proposed by Bing et al. (2020b). This last method also recovers anchor words, then computes each word’s topic weights using non-negative convex regression against those anchors, with thresholding/regularization to enforce weak sparsity. We note the following regarding the procedure of Bing et al. (2020b):

- This procedure removes infrequently occurring words in the same manner as ours, but with the threshold  $\alpha\sqrt{\frac{\log p_n}{nN}}$  in (9) replaced by  $\frac{7\log p_n}{nN}$ . This threshold is lower than ours if  $\frac{\log p_n}{nN}$  is sufficiently small. In practice, however, the constant 7 used in their threshold is quite large and thus leads to excessive thresholding in some of our datasets, especially when the word frequencies decay according to Zipf’s law.
- This procedure requires a list of anchor words for each topic  $k \in [K]$  as input, rather than just the number of topics  $K$ . We therefore need to estimate a partition of anchor words using a special procedure that is included in their original implementation.

---

1. The code for our method and all the experiments presented in this section can be found on GitHub at the following link: <https://github.com/yatingliu2548/topic-modeling>.

Clearly, whether the anchor words are estimated and partitioned correctly has an impact on the overall estimation of  $A$ .

We therefore caution the reader that these factors put the Sparse Topic Model solver of Bing et al. (2020b) at a comparative disadvantage in our experiments.

*Data Generation Mechanism.* For simplicity, we ensure all documents are of the same length  $N$ . For each experiment, we create a document-to-topic matrix  $W \in \mathbb{R}^{K \times n}$  by independently drawing the columns  $W_{*i} \in \mathbb{R}^K, i = 1, \dots, n$  from the Dirichlet distribution with parameter  $\alpha_W = \mathbf{1}_K$ . We generate the matrix  $A \in \mathbb{R}^{p \times K}$  either without anchor words or with 5 anchor words per topic, in which case whenever word  $j$  is an anchor word for topic  $k$ , we set  $A_{jk} = \delta_{\text{anchor}}$  where  $\delta_{\text{anchor}} \in \{0.0001, 0.001, 0.01\}$ . In order to mimic the behavior of real text data, the entries of column  $k$  of  $A$  corresponding to non-anchor words are then chosen such that they decay according to Zipf’s law. This means for each column  $k$  of  $A$ , we ensure that the frequency  $f_{(j)}$  of the  $j^{\text{th}}$  most frequent non-anchor word follows the pattern

$$f_{(j)} \propto \frac{1}{(j + b_{\text{zipf}})^{a_{\text{zipf}}}}, \tag{31}$$

where  $a_{\text{zipf}} = 1, b_{\text{zipf}} = 2.7$ . Each column of  $A$  is subsequently normalized to unit  $\ell_1$ -norm. The pattern (31) has indeed been empirically shown to hold approximately for word frequencies in real datasets (Zipf, 1936; Piantadosi, 2014). Figure 12a in Appendix G illustrates the distribution of word frequencies generated under our data generation mechanism.

Having specified both  $A$  and  $W$ , the observation matrix  $D$  is then generated according to the pLSI model described in Section 1.1. We fit our method and the four benchmarks while varying the values of  $n, p, N$ , and  $K$ . In all of our experiments, unless otherwise specified, the constant  $\alpha$  in the threshold (9) is fixed at  $\alpha = 0.005$ . We evaluate the estimation error of all methods relative to the true underlying  $A$  by computing the  $\ell_1$  loss per topic

$$\mathcal{L}_1(\hat{A}, A) = \min_{\Pi \in \mathcal{P}} \frac{1}{K} \|\hat{A}\Pi - A\|_1,$$

where  $\mathcal{P}$  denotes the set of all  $K \times K$  permutation matrices.

*Varying  $(p, N, K)$ .* We first provide a snapshot of our method’s relative performance in different parameter regimes by fixing  $n = 500$  and varying  $(p, N, K)$ . Here we specify 5 anchor words per topic and set the anchor word frequency to  $\delta_{\text{anchor}} = 10^{-3}$ . The median  $\mathcal{L}_1(\hat{A}, A)$ -errors over 50 trials are plotted in Figure 3. As Figure 3 shows, our method (in blue) outperforms all other methods in most parameter regimes considered here. Interestingly, the estimation errors of AWR and LDA often appear constant as a function of document length  $N$ . As  $N$  increases, the errors from both Topic-SCORE and our method display a clearer pattern of consistency relative to AWR and LDA; this observation is also made by Tracy Ke and Wang (2024) in a similar experimental setup. However, our method’s errors decay to zero much faster than all other benchmarks when the vocabulary size is large ( $p \in \{5000, 10000\}$ ).

We note that in these experiments, the approach proposed by Bing et al. (2020b) does not perform very well. In particular, for small  $p$  and small  $N$ , the number of topics returned




Figure 3: Median  $\mathcal{L}_1(\hat{A}, A)$  errors for all methods based on 50 independent trials. Here the number of documents is fixed ( $n = 500$ ). In each panel, the errors are plotted as a function of document length  $N$  (log-scaled on the x-axis). The panels display results for different values of  $(p, K)$ , as specified by row and column labels.

by this method is smaller than the expected number of topics  $K$ , which prevents us from comparing its results with all four other methods. On inspection, we find that this is due to over-thresholding of the vocabulary, which leaves too few words to reliably estimate the matrix  $A$ . To provide a fair comparison with Bing et al. (2020b), we also compare all five methods using the data generation mechanism proposed by Bing et al. (2020b). This means that the non-anchor entries of each column of  $A$  no longer display the Zipf’s law pattern (31), but instead are generated from a Uniform distribution. We note that this data generation mechanism ensures all the non-anchor words for each topic are of roughly equal frequency and is thus also favorable to Topic-SCORE (Tracy Ke and Wang, 2024), which assumes  $\min_{j \in [p]} h_j \geq c\bar{h}$  where  $\bar{h} := \frac{1}{p} \sum_{j=1}^p h_j$ . The results are displayed in Figure 14 of Appendix G. Under this uniform data generation mechanism, our method (with  $\alpha = 0.005$ ) displays identical performance relative to Topic-SCORE, and both SCORE-based methods still perform well relative to other benchmarks in most parameter regimes. As expected, we also find that fewer words are removed by thresholding, in comparison with the Zipf’s law setting where our  $\ell_q$ -sparsity assumption (15) is more likely to hold with small  $s$  and




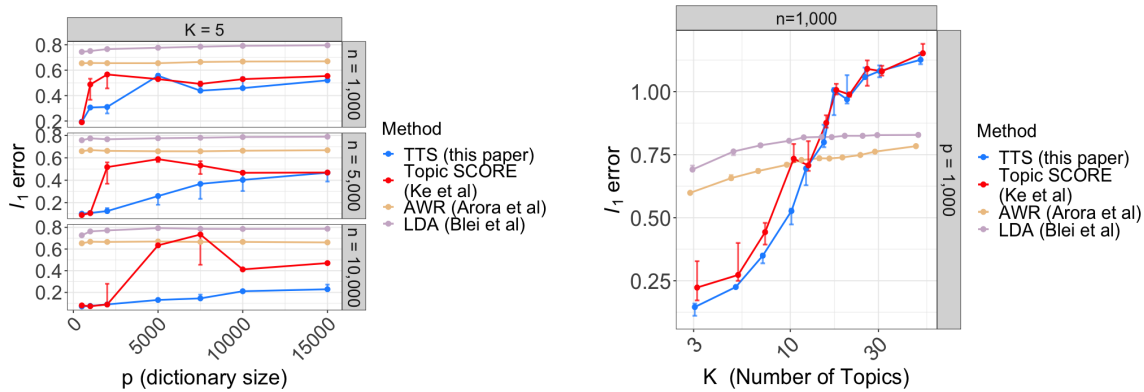
Figure 4:  $\mathcal{L}_1(\hat{A}, A)$ -errors from all methods as a function of  $n$ , for  $K \in \{5, 10\}$  with  $p$  and  $N$  fixed. Vertical error bars centered about the median errors indicate the errors’ interquartile ranges computed based on 50 independent trials.

many more words occur infrequently. These experiments empirically suggest that 1) TTS improves upon the performance of Topic-SCORE when the columns of  $A$  exhibit a Zipf’s law (or  $\ell_q$ -sparsity) decay pattern, and 2) our procedure’s performance remains reasonable and is similar to that of Topic-SCORE when the  $\ell_q$ -sparsity assumption (15) is violated.

*Varying the number of documents  $n$ .* We now focus on the effect of varying  $n$  on the estimation error. Fixing this time  $N = 500$  and  $p = 10,000$ , the  $\mathcal{L}_1(\hat{A}, A)$ -errors are presented in Figure 4 with  $K = 5$  and  $K = 10$ . Our method (in blue) consistently outperforms other methods and also displays a clear trend of consistency as  $n$  increases. When  $K$  increases, the estimation problem becomes more difficult due to the larger number of parameters, and so more documents are needed to achieve a reasonable performance. Nonetheless, our method still performs well when  $K = 10$  and  $n$  is reasonably large, whereas the error from Topic-SCORE decays to zero very slowly with this larger value of  $K$ .

*Varying the dictionary size  $p$ .* Figure 5a shows how the  $\mathcal{L}_1(\hat{A}, A)$ -errors vary as the vocabulary size  $p$  increases, with  $N = 500$ ,  $K = 5$  and  $n \in \{1000, 5000, 10000\}$ . We do not include the errors from the procedure of Bing et al. (2020b) as they are higher than those of LDA. As expected, the errors for all methods increase with the dictionary size  $p$ . However, our method mostly outperforms the other benchmarks, even in some high-dimensional parameter regimes where  $p > \max(n, N)$ . The performance of Topic-SCORE only converges to ours when  $p$  is too large relative to  $n$ , a setting which is challenging for all methods.

Additionally, our method also outperforms most other benchmarks in terms of computational runtime when  $p$  is large. We provide in Figure 18 of Appendix G a visualization of how the runtimes for all methods scale with  $p$ . Our method’s runtime is similar to that of AWR and is consistently better than that of Topic-SCORE, primarily due to our



(a)  $\mathcal{L}_1(\hat{A}, A)$ -errors as a function of  $p$ , with  $K = 5$  and  $N = 500$ . Results are obtained based on 15 independent trials.

(b)  $\mathcal{L}_1(\hat{A}, A)$ -errors as a function of  $K$ , with  $n = p = 10^3$  and  $N = 500$ . Results are obtained based on 50 independent trials.

Figure 5:  $\mathcal{L}_1(\hat{A}, A)$ -errors as a function of the dictionary size  $p$  (left) and the number of topics  $K$  (right). Vertical bars around median errors indicate interquartile ranges.

thresholding of infrequent words before performing eigendecomposition.

*Varying the number of topics  $K$ .* Figure 5b shows how the  $\mathcal{L}_1(\hat{A}, A)$ -errors vary as  $K$  increases, with  $n = p = 1000$  and  $N = 500$ . The main observation here is that LDA and AWR may be preferable to our method if  $K$  is *a priori* known to be large while the dataset we possess is relatively small. As Figure 5b illustrates, the SCORE-based methods perform worse than LDA and AWR when  $K > 15$ , but this is because the number of documents is quite small in this experiment ( $n = 1000$ ). If  $n$  and  $N$  are large enough, one can expect our method to accommodate a larger number of topics; see Figure 4 for an illustration.

*Relaxation of the separability assumption.* Section 2.4 suggests that the vertex hunting algorithm of Javadi and Montanari (2020) may reduce the vertex hunting error in some situations when separability fails to hold. Figure 6 compares the overall  $\mathcal{L}_1(\hat{A}, A)$ -errors as a function of  $n$  when we use Successive Projection (SP), Sketched Vertex Search (SVS) and Archetype Analysis (AA) in the vertex hunting step of TTS. As expected, when there are no anchor words, using the AA algorithm rather than SP/SVS can significantly improve the estimation of  $\hat{A}$ , especially when  $K$  is large. Again, this is because SP and SVS are not designed for non-separable point clouds and also perform better with small  $K$ . In fact, the AA algorithm also often works well under separability, since the  $\alpha$ -uniqueness condition used by Javadi and Montanari (2020) is satisfied. The main trade-off for this stronger statistical performance is the computational cost of solving the non-convex optimization problem required by AA. Nonetheless, the fact that our method accommodates non-separable datasets makes TTS more widely applicable compared to methods based on anchor word identification, such as those proposed by Bing et al. (2020b) and Arora et al. (2012).

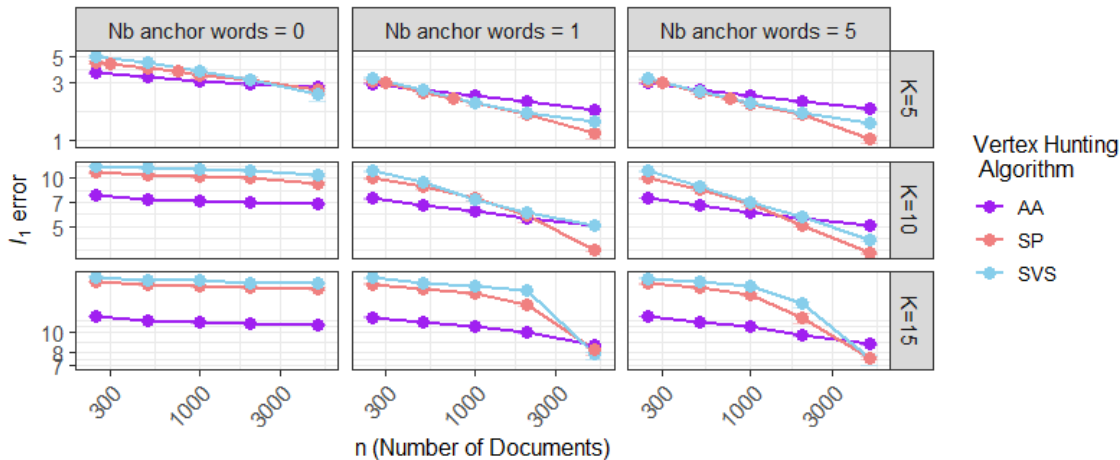



Figure 6:  $\mathcal{L}_1(\hat{A}, A)$ -errors as a function of  $n$  when we use three different vertex hunting algorithms in the vertex hunting step of TTS. Here,  $p = 10^4$  and  $N = 500$  are fixed, and  $K \in \{5, 10, 15\}$ . The number of anchor words per topic is either 0, 1, or 5. Results are averaged over 50 independent experiments.

*The importance of appropriate thresholding.* Figure 7a shows how the  $\mathcal{L}_1(\hat{A}, A)$ -error varies as the threshold level in (9) increases from zero, and Figure 7b shows the corresponding percentage of words removed. For this dataset, the performance of our method when  $\alpha = 0$  (no thresholding) is not too different from Topic-SCORE. As the threshold level increases, infrequent words that contribute noise to the point cloud are removed, thus leading to an improvement in the estimation of  $A_{J^*}$ . However, an excessively high threshold means we set too many rows of  $A$  to zero, and so the error from estimating  $A_{J^*}$  becomes higher. This explains the pattern observed in Figure 7a, which demonstrates the importance of choosing a balanced threshold in our procedure.

As we mentioned, the universal parameter  $\alpha$  should be independent of  $(p, n, N, K)$ . Our recommended value of  $\alpha = 0.005$  is obtained based on numerous such experiments with synthetic data where we vary the values of  $(p, n, N, K)$ . This choice of  $\alpha$  also works well in all real data applications of Section 4, where several parameter regimes are involved.

*Additional experiments and conclusion.* We also evaluate the impact of other aspects of the data generation mechanism on our estimator’s performance. We find that changing  $\delta_{\text{anchor}}$ , which controls the frequency of anchor words, does not significantly impact the overall performance of TTS. This is an advantage of SCORE-based methods over methods that rely on identification of anchor words, which are often affected by the frequency of anchor words both in theory and in practice. Additionally, when we increase the parameter  $a_{\text{zipf}}$  in (31), we find that our estimator’s performance improves significantly. This is not surprising as a larger  $a_{\text{zipf}}$  means the ordered entries of  $A$ ’s columns decay to zero faster, and our theoretical results also show that a strong sparsity regime (when  $q$  is close to 0 in Assumption 5) is favorable to our method. Further details about these experiments are



(a) Average  $\mathcal{L}_1(\hat{A}, A)$ -error as a function of the threshold parameter  $\alpha$ . (b) Corresponding percentage of the words discarded by thresholding as a function of  $\alpha$ .

Figure 7:  $\mathcal{L}_1(\hat{A}, A)$ -error averaged over 20 independent trials and the percentage of words removed as  $\alpha$  increases, for a synthetic dataset with  $p = 5000, n = N = 500$ .

deferred to Appendix G. Finally, we check the performance of our method on a set of semi-synthetic experiments based on the Associated Press dataset, which is included in the R package `tm` (Feinerer et al., 2015), thereby allowing us to test a different data generating mechanism. The results are also presented in Appendix G.

Overall, we have illustrated that our method (a) *performs well in a wide variety of parameter regimes, and notably in the high-dimensional setting where  $p$  is large*, and (b) *performs well even if our sparsity assumption is violated* (see the discussion on the uniform data generation mechanism, and also note that we use a weak sparsity regime with  $a_{\text{zipf}} \approx 1$  in most of our experiments). This makes our method applicable to the vast majority of real-world text datasets, which often are high-dimensional and exhibit Zipf’s law decay. However, alternative methods such as LDA and AWR may still be competitive in some settings, especially when the pLSI model fails to hold or if the number of documents  $n$  and the document length  $N$  are unusually small relative to the number of topics  $K$ .

#### 4. Practical Applications in Text Analysis and Beyond

In this section, we deploy our method on real-world datasets. Given the results of the previous section, we focus here on the comparison of our method with Topic-SCORE (Tracy Ke and Wang, 2024) and LDA (Blei et al., 2003).

Real datasets seldom have ground truth for  $A$ , and some may even lack an obvious choice for the number of topics  $K$ . Consequently, in this section we evaluate the estimators’ performance using, when appropriate, the following metrics:

- (a) *Topic Resolution* as a measure of topic consistency. We fit each estimator on two disjoint halves of the data and report the cosine similarity between estimated topics (after an appropriate permutation of the columns of  $A$ ). Mathematically, letting  $\hat{A}^{(i)}, i \in \{1, 2\}$  denote the estimated topic-word matrices obtained for each half of the data, we define the “average topic resolution”  $\eta$  as the mean cosine similarity (a classical similarity metric in natural language processing) between aligned topics:

$$\eta = \max_{\sigma \in \Pi_K} \frac{1}{K} \sum_{k=1}^K \frac{\hat{A}_{*k}^{(1)\top} \hat{A}_{*\sigma(k)}^{(2)}}{\|\hat{A}_{*k}^{(1)}\|_2 \|\hat{A}_{*\sigma(k)}^{(2)}\|_2}, \quad (32)$$

where  $\Pi_K$  denotes the set of all permutations of  $[K]$ . Thus, higher resolution indicates better-defined and more consistent topic vectors (although this does not necessarily mean better  $\ell_1$ -error).

- (b) *Multiscale Topic Refinement and Coherence* (Fukuyama et al., 2023): In the absence of an obvious number of topics  $K$ , we fit the method for multiple values of  $K$  and analyze the resulting topic hierarchy to check the stability of our estimator. We follow in particular the methodology of Fukuyama et al. (2023), which was developed to guide the choice of an appropriate number of topics  $K$  for LDA (Blei et al., 2003) by investigating the relationships among topics of increasing granularity. Given a hierarchy of topics, the method evaluates which topics consistently appear, constantly split, or are merely transient. We use these tools here, together with the associated package `alto` (Fukuyama et al., 2023), to analyze our estimator. The method of Fukuyama et al. (2023) starts by computing the alignment of topics across the hierarchy using the transport distance: for each  $K$ , this method computes how the mass of topic  $j \in \{1, \dots, K\}$  is split amongst the  $K + 1$  topics at the next level of the hierarchy. We refer the reader to the original work by Fukuyama et al. (2023) for a more detailed explanation of topic transport alignment. Once the relationships between consecutive topic models have been established, the method of Fukuyama et al. (2023) allows visualization of (a) topic refinement (that is, whether topics increase in granularity, as indicated by a small number of ancestors in the hierarchy; or conversely, whether topics are perpetually recombined from one level of the hierarchy to the next); and (b) topic coherence (whether a topic appears across multiple values of  $K$ ). We therefore favor methods with improved topic coherence and topic refinement, since these are markers of topic stability.

We explore the comparison between our method, LDA and Topic-SCORE under diverse parameter regimes (with varying  $n$ ,  $N$  and  $p$ ).

#### 4.1 Research Articles (High $p$ , High $n$ , Low $N$ )

For our first experiment, we consider a corpus of 20,140 research abstracts belonging to (at least) one of four categories: Computer Science, Mathematics, Physics and Statistics.<sup>2</sup> After pre-processing of the data (including the removal of standard stop words, numbers, and punctuation), our dataset involves a dictionary of size  $p = 81,649$  and  $n = 20,140$  documents with an average document size of  $N = 157$  words.

We first evaluate the topic consistency of all methods in estimating the topic-word matrix  $A$  using the mean topic resolution defined in equation (32). Table 1 displays the average topic resolution over 25 random splits of the data. As highlighted in the introductory

---

2. The data is available on Kaggle at <https://www.kaggle.com/datasets/blessondensil294/topic-modeling-for-research-articles/code?resource=download>. Although the original data set comprises six topics (with the addition of Quantitative Biology and Finance), due to the low representation of these last two topics (< 4% of the data), we drop them from our analysis.

Methods	Average Topic Resolution ( $\eta$ )	Interquartile range
LDA (Blei et al)	0.304	(0.270,0.330)
TTS (this paper)	0.332	(0.310,0.360)
Topic-SCORE (Ke et al)	0.145	(0.093,0.179)

Table 1: Average Topic Resolution on research article data. The interquartile range for the average topic resolution was computed over 25 random splits of the data.

	Top 10 most representative words per topic
Computer Science	learning, network, networks, model, can, neural deep, using, models, data
Physics	model, can, system, field, energy, systems magnetic, models, using, phase
Mathematics	problem, can, algorithm, show, method, paper results, also, time, using
Statistics	data, model, can, learning, using, models method, approach, based, paper

Table 2: Most common words found by our method

paragraph to this section, topic resolution can be taken as an indicator of the stability of the estimator of  $\hat{A}$  between two separate portions of the data. A method that produces higher topic resolution with a narrower interquartile range indicates a more stable estimation of the topic-word matrix  $A$ . As shown in Table 1, our approach consistently outperforms LDA and Topic-SCORE on this metric; it offers the highest average topic resolution score. Topic-SCORE’s performance exhibits more significant fluctuations, as indicated by its larger interquartile range.

Taking a closer look at the estimation of  $A$ , we consider the 10 most representative words generated by each of the three methods for every topic (obtained by selecting the top 10 largest entries in each column of  $\hat{A}$ ). The results are presented in Tables 2, 3, and 4. For the topics of Computer Science and Statistics, the top 10 most representative words produced by our method agree with 70% of LDA’s most representative words in the corresponding topics. There is much less agreement for the topic of Physics, but upon closer inspection we find that some of the words produced by our method in that category (such as ‘magnetic’, ‘energy’) are more indicative of the topic of Physics, whereas all of the top 10 words for Physics produced by LDA are generic words that can appear in other categories.

In contrast, the results of Topic-SCORE (Table 4) seem to diverge substantially from those of LDA and our method. It appears that the top 10 most representative words for Physics, Mathematics and Statistics from Topic-SCORE are dominated by infrequently occurring words and foreign words; the foreign words can be traced back to a few rare abstracts written in English and followed by a foreign language translation. This supports our hypothesis that Topic-SCORE amplifies the effects of infrequent words, unless significant *ad hoc* data pre-processing (removal or merger of rare words, and removal of documents with significant numbers of rare words) is applied.

	Top 10 most representative words per topic
Computer Science	data, network, learning, networks, can, model using, new, paper, based
Physics	show, data, analysis, two, can, problem results, field, system, performance
Mathematics	can, used, models, using, model, paper number, method, proposed, approach
Statistics	model, results, show, can, learning, method using, based, data, also

Table 3: Most common words found by LDA

	Top 10 most representative words per topic
Computer Science	data, can, model, using, learning, show results, method, paper, also
Physics	della, quantum, theory, del, year, teoria quantistica, per, nel, delle
Mathematics	die, der, collectors, problem, able, assumptions coupon, wir, based, one
Statistics	der, und, music, automatischen, learning, sheet die, musikverfolgung, deep, algorithms

Table 4: Most common words found by Topic-SCORE

In order to further investigate the performance gap between TTS and Topic-SCORE, we visualize the point cloud from both methods in Figure 8. As expected, we observe that the Topic-SCORE point cloud is severely stretched by a set of low-frequency words that include several foreign words. Again, with the presence of many rare words in the dataset, the lack of thresholding and the use of the pre-SVD multiplication step in Topic-SCORE contribute to a significant distortion of the point cloud. In comparison, the thresholding approach we adopt yields a more compact point cloud. As demonstrated in Figure 8(b), our method effectively recaptures the essential vertices of the point cloud simplex. A closer look at the words surrounding each vertex, as shown in Figure 8b, allows us to easily identify which simplex vertex belongs to which topic (Physics, Math, Computer Science and Statistics when moving in the anticlockwise direction). Under this “large  $p$ ” regime and in the presence of a myriad of rare words that may introduce significant noise, our method not only distinguishes words effectively but also clusters them into well-defined topics.

We note that this dataset comes with manually curated topic labels for each document. As a final verification, we analyze the performance of the different methods when used for recovering the ground truth labels for each document. Having estimated  $A$ , it is quite natural in light of the pLSI model to perform regression of  $D$  against  $\hat{A}$  in order to yield an estimator of  $W$ . To this end, we use the estimation procedure for  $W$  of Tracy Ke and Wang (2024), where the problem of estimating  $W$  given  $\hat{A}$  is reduced to a weighted constrained

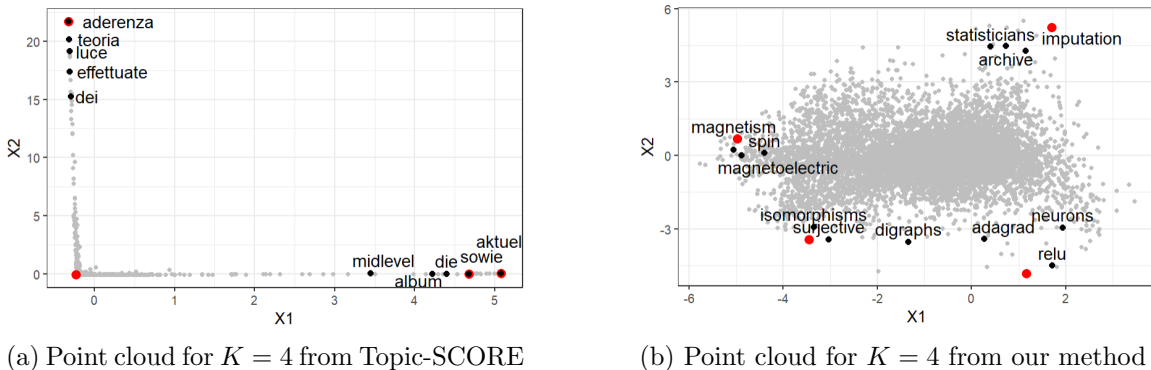


Figure 8: Comparison of the 3-dimensional point clouds from TTS (right) and Topic-SCORE (left), projected on the first two axes for visualization. Estimated vertices are colored red, and the point clouds are represented by gray dots. Most outlying words in Topic-SCORE’s point cloud are thresholded away by TTS, thus contributing to higher point cloud stability for our method.

linear regression problem:

$$\hat{W}_{*i} = \arg \min_{\omega \in \Delta_K} \frac{1}{p} \sum_{j: M_{jj} > 0} \frac{1}{M_{jj}} \left( D_{ji} - \sum_{k=1}^K \hat{A}_{jk} \omega_k \right)^2, \quad i \in [n], \tag{33}$$

where  $\Delta_K = \{\omega \in \mathbb{R}^K : \omega_k \geq 0, \sum_{k=1}^K \omega_k = 1\}$ . We strongly emphasize that the aim of this experiment is to evaluate the estimation of  $A$ , and we do not claim here that our method provides state-of-the-art results in the estimation of  $W$ . Other potentially better estimation procedures are available for  $W$ , many of which do not require estimating  $A$  first. Rather, as topic labels are available for this dataset, we use this simple estimation procedure for  $W$  via  $\hat{A}$  as another way of comparing the quality of  $\hat{A}$  obtained from TTS, Topic-SCORE and LDA. Since the  $\hat{W}$  obtained from (33) depends on  $\hat{A}$  as input, it stands to reason that a better estimation procedure for  $A$  may be reflected in a better agreement between  $\hat{W}$  and the provided topic labels for each document, if we use (33) to estimate  $W$ .

Let  $y_{ki} = 1$  if document  $i$  is labeled as belonging to topic  $k$  (and  $y_{ki} = 0$  otherwise). We compute the average  $\ell_1$  distance  $\mathcal{D}(\hat{W}, y)$  and cosine similarity  $S_k$  between the permuted matrix  $\hat{W}$  and the provided labels  $y$  for each topic  $k$  as follows:

$$\mathcal{D}(\hat{W}, y) := \min_{\sigma \in \Pi_K} \frac{1}{nK} \sum_{ki} |\hat{W}_{\sigma(k)i} - y_{ki}|, \quad S_k = \max_{\sigma \in \Pi_K} \frac{\sum_{i=1}^n \hat{W}_{\sigma(k)i} y_{ki}}{\|\hat{W}_{\sigma(k)*}\|_2 \|y_{k*}\|_2}.$$

Here, a smaller value of the  $\ell_1$  distance or a larger value of the cosine similarity score between  $y$  and  $\hat{W}$  indicate greater alignment with the provided topic labels. The results are displayed in Table 5.

Table 5 indicates that our method improves the estimation of  $W$  overall and provides the best topic alignment on average, when using (33) to estimate  $W$ . This suggests that our procedure yields a more accurate estimator of  $A$ .

Methods	$S_{CS}$	$S_{Phys}$	$S_{Math}$	$S_{Stat}$	$\bar{S}$	$\mathcal{D}(\hat{W}, y)$
LDA (Blei et al. (2003))	0.671	0.576	0.534	0.493	0.569	0.403
TTS (this paper)	0.610	0.748	0.636	0.494	0.622	0.305
Topic-SCORE (Tracy Ke and Wang (2024))	0.670	0.545	0.588	0.373	0.544	0.348

Table 5: The evaluation of  $\hat{W}$  obtained via estimating  $A$  first by using the three methods.  $\bar{S}$  is the average cosine similarity across all  $K$  topics

## 4.2 Single-Cell Analysis (Low $p$ , High $n$ , Low $N$ )

In this subsection, we consider a different application area for our methodology: the analysis of single-cell data. We revisit the mouse spleen dataset presented by Goltsev et al. (2018). This dataset consists of a set of images from both healthy and diseased mouse spleens. Each sample undergoes staining with 30 different antibodies via the CODEX process, as detailed by Goltsev et al. (2018). In Chen et al. (2020), each spleen sample is divided into a set of non-overlapping Voronoi bins, and the count of immune cell types is recorded in each bin. In this framework, each bin can be viewed as a document and cell types correspond to words. It is of interest to determine appropriate groupings of cell types (topics), as this may help one study the interactions between cells.

Since this dataset does not come with ground-truth labels, we sample two disjoint sets of size  $n = 10,000$  out of the 100,840 Voronoi tessellations across all spleen samples (where 10,000 is a number chosen to be large enough to ensure a “high  $n$ ” regime while still allowing all methods to have reasonable computational runtimes). By contrast, there are only 24 different cell types ( $p = 24$ ), while the average “document” length is  $N = 11.2$  with an interquartile range of (6, 16). While Chen et al. (2020) focus on evaluating estimators of the matrix  $W$ , here we repurpose the use of this dataset to study our estimator of  $A$ . In this dataset, the precise number of topics  $K$  is unknown. We thus apply the three methods for different values of  $K$  and use the metrics introduced at the beginning of this section (topic resolution, topic coherence and refinement) to compare the three methods. The results are presented in Figures 9 and 10.

*Discussion of the results.* Due to the structured nature of this dataset, all methods perform remarkably well, exhibiting an average topic similarity above 0.95. In more detail, we see that our method outperforms Topic-SCORE in terms of topic resolution. In particular, Topic-SCORE (in red) appears to have more variable performance, as reflected in its larger interquartile ranges and its jittery resolution as a function of  $K$ . Interestingly, in this specific instance, LDA seems to score higher on topic resolution (although we again emphasize that all methods perform very well on this metric). Additionally, Figure 10a shows the refinement and coherence of the topics for our method as  $K$  increases, in contrast to those of LDA in Figure 10b. In this data example, our method seems to provide topics with higher refinement (fewer ancestors per topic) and higher coherence (note in particular the stability of topic 1, 2, and 18) compared to LDA. In Figure 10b, it can be observed that topics 1, 2, and 18 are dispersed across different branches within the refinement plot as  $K$  varies.


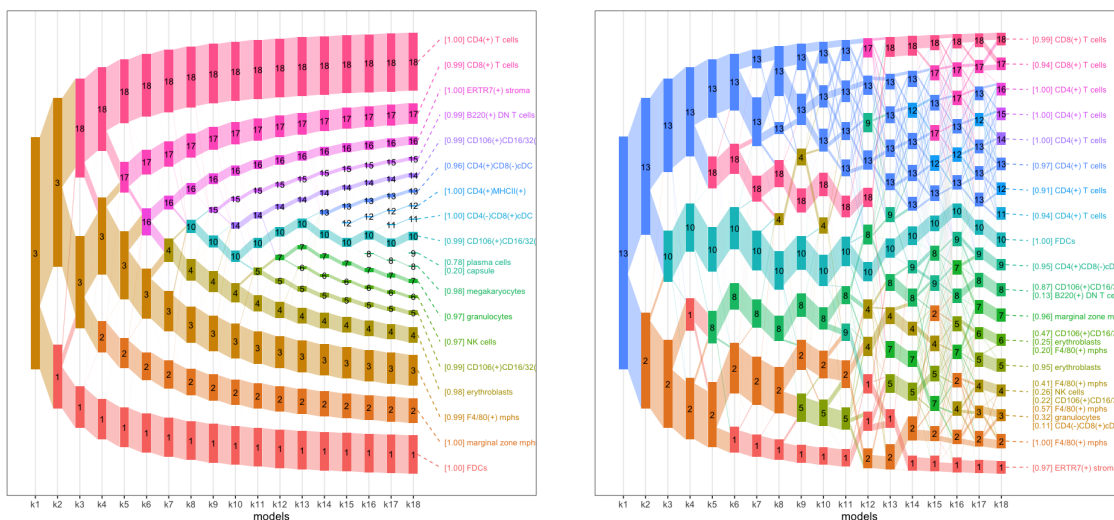


Figure 9: Median Topic Resolution as a function of  $K$  on the Mouse Spleen Data (Goltsev et al., 2018; Chen et al., 2020). Vertical error bars represent the interquartile range for the average topic resolution scores over 25 trials.



(a) Topic refinement for our method as  $K$  varies, provided by the package `alto` (Fukuyama et al., 2023).

(b) Topic refinement for LDA (Blei et al., 2003) as  $K$  varies, provided by the package `alto` (Fukuyama et al., 2023).

Figure 10: Comparison of the refinement and coherence of topics recovered using our method (left) and LDA (right).




Figure 11: Topic resolution (measured by the average cosine similarity between halves of the data) of our method (in blue) and Topic-SCORE (red) on the microbiome dataset of Yachida et al. (2019). Topic resolution is averaged over 25 random splits of the data.

### 4.3 Microbiome Examples (Low $p$ , Low $n$ , High $N$ )

We finish our discussion with an application of our method to microbiome data analysis. In particular, we reanalyze two datasets that have been previously analyzed through topic modeling: the colon dataset of Yachida et al. (2019) and the vaginal microbiome example of Callahan et al. (2017), which was reanalyzed by Fukuyama et al. (2023) using LDA. Microbiome data are represented in the form of a count matrix. In this matrix, each column corresponds to a different sample, while each row represents various taxa of bacteria. The entries within the matrix represent the abundance of each bacterial taxon in a given sample. Taking samples to be documents and bacteria as words, topic modeling offers an interesting way of exploring communities of bacteria (“topics”) (Sankaran and Holmes, 2019). For the sake of conciseness, we present the results here for the colon dataset of Yachida et al. (2019), and refer the reader to Appendix H for the results on the other dataset.

After pre-processing and eliminating species with a relative abundance below 0.001%, this dataset contains microbiome counts for  $p = 541$  distinct taxa from  $n = 503$  samples. In contrast, the length of each “document” is extremely high, with around  $N = 43$  million bacteria per sample. We test all three methods for different values of  $K$  and display the average topic resolution in Figure 11. On this metric, our method exhibits significantly better results than both LDA and Topic-SCORE for up to 15 topics. After 15 topics, LDA outperforms all SCORE-based methods in terms of topic resolutions. However, this comes at a much higher computational cost: while each of the SCORE methods in this example could be fitted in under a minute, each of the LDA fits took on the order of tens of minutes. Note that LDA’s high topic resolution could also be due to the higher weight of the prior in the estimation of the topic-word matrix  $A$ , which, due to the relatively small size of the dataset, could have a stabilizing effect on estimation. On the other hand, the performance of Topic-SCORE quickly drops to 0.65 as  $K$  increases, before reaching a plateau at around  $K \approx 10$ . By contrast, for small  $K$ , our method exhibits a resolution up to 40% higher than Topic-SCORE (for  $K = 10$ ) before also decreasing as the number of topics increases.

To understand the gap in performance between Topic-SCORE and our method, we again visualize the point clouds obtained by both methods. The visualization can be found in

Figure 19 in Appendix H. Similarly to our first example with text analysis, we observe that the point cloud of Topic-SCORE is heavily distorted; in contrast, ours is more compact.

## 5. Conclusion and Future Work

In this paper, we introduce *Thresholded Topic-SCORE* (TTS), a new estimation procedure for the topic-word matrix  $A$  that is based on eigenvalue decomposition and thresholding. Our procedure is shown to perform well under the column-wise  $\ell_q$ -sparsity assumption (15), which is exhibited by many real-world text datasets but to our knowledge has not been considered in prior works. TTS also accommodates non-separable data, simply by adopting a suitable vertex hunting algorithm such as Archetype Analysis from Javadi and Montanari (2020). Empirical results show that our method is competitive for a diverse range of parameter regimes, especially in the “large  $p$ ” setting where many words occur infrequently in the corpus. Overall, TTS is a compelling alternative to existing methods when  $p$  is large and the number of topics  $K$  is relatively small.

This paper suggests several potential research directions. First, the estimation of  $W$  is also of interest in applications, and the minimax-optimal  $\ell_1$ -error rate for estimating  $W$  has been established (see, for example, Klopp et al., 2023). However, the problem of estimating  $W$  essentially involves  $n$  independent sub-problems (one for each column of  $W$ ), and consequently the minimax-optimal error rate for  $W$  scales significantly with  $n$ . One may consider imposing additional structural assumptions on how the  $n$  documents are related to one another, in order to design an estimation procedure whose error decays to zero as  $n \rightarrow \infty$ .

Second, the minimax-optimal  $\ell_1$ -error rate given the  $\ell_q$ -assumption (15) remains an open problem. The minimax lower bound arguments of Tracy Ke and Wang (2024) are not directly applicable, as they are based on constructing hypotheses of  $A$  whose columns contain entries that are of roughly equal magnitudes (that is, the columns do not exhibit  $\ell_q$ -decay).

Third, our proposed method only makes use of the word counts for the documents, as the underlying pLSI model disregards the order of words in a document. Other language models, such as the multi-gram topic model, make use of word orders. An extension of our method using tensor factorization may be possible in this setting, as the corpus is stored using a multi-way tensor (Zheng et al., 2016).

## Acknowledgments

The authors would like to acknowledge support for this project from the National Science Foundation (NSF grant IIS-2238616). This work was completed in part with resources provided by the University of Chicago’s Research Computing Center.

All proofs in this appendix make use of notations described in Section 1.4. Assumptions 1–4 (which include separability) are assumed in Appendices A–D, whereas the sparsity assumption (Assumption 5) is further imposed in Appendix E.

## Appendix A. Properties of the Set $J$

**Lemma 23 (Weak sparsity of  $A$ )** *Order the  $\ell_2$  row norms of  $A$  so that*

$$\|A_{(1)*}\|_2 \geq \cdots \geq \|A_{(p)*}\|_2,$$

*Then the matrix  $A$  satisfies  $\max_{j \in [p]} j \|A_{(j)*}\|_2 \leq K$ .*

**Proof** Observe that for any  $j \in [p]$ ,

$$j \|A_{(j)*}\|_2 \leq \sum_{l=1}^p \|A_{l*}\|_2 \leq \sum_{l=1}^p \|A_{l*}\|_1 = K,$$

since  $A$  contains only non-negative entries and each column sums up to 1. ■

**Lemma 24** *If  $M_0 := \text{diag}(n^{-1}D_0\mathbf{1}_n)$  and  $h_j := \|A_{j*}\|_1$ , then for any  $j \in [p]$ ,*

$$\sigma_K(\Sigma_W)h_j \leq M_0(j, j) \leq h_j.$$

**Proof** Note that

$$M_0(j, j) = \frac{1}{n} \sum_{i=1}^n [D_0]_{ji} = \frac{1}{n} \sum_{i=1}^n \sum_{k=1}^K A_{jk} W_{ki} = \sum_{k=1}^K A_{jk} \left( \frac{1}{n} \sum_{i=1}^n W_{ki} \right),$$

and observe that  $h_j := \sum_{k=1}^K A_{jk}$  and for each  $k \in [K]$  (recall  $\Sigma_W := \frac{1}{n} W W^T$ ),

$$\sigma_K(\Sigma_W) \leq \Sigma_W(k, k) = \frac{1}{n} \sum_{i=1}^n W_{ki}^2 \leq \frac{1}{n} \sum_{i=1}^n W_{ki} \leq 1.$$

■

**Theorem 25** *Define  $p_n := p \vee n$ ,  $\tau_n := \sqrt{\frac{\log p_n}{Nn}}$ . Let*

$$J := \{j \in [p] : M(j, j) > \alpha \tau_n\}, \quad J_{\pm} := \{j \in [p] : M_0(j, j) > \alpha_{\pm} \alpha \tau_n\},$$

*for some suitably chosen  $\alpha > 0$  and  $0 < \alpha_+ < 1 < \alpha_-$ . The following statements hold:*

(a) *For a fixed  $j \in [p]$ , we have*

$$\mathbb{P}(|M(j, j) - M_0(j, j)| \geq t) \leq 2 \exp(-nNt^2/2).$$

(b) *The event*

$$\mathcal{E} := \{J_- \subseteq J \subseteq J_+\}$$

occurs with probability at least  $1 - o(p_n^{-1})$ . Here, we can select  $\alpha = 8, \alpha_+ = \frac{1}{2}, \alpha_- = 2$ . Note that this implies that on  $\mathcal{E}$ ,  $\min_{j \in J} h_j > \alpha_+ \alpha \tau_n$ .

(c) *We have*

$$\|A_{J_-^c}\|_F^2 \leq [2K(\beta\tau_n)^{-1} \wedge p] (\beta\tau_n)^2 = o(1)$$

where  $\beta := \frac{\alpha_- \alpha}{\sigma_K(\Sigma_W)} = \frac{16}{\sigma_K(\Sigma_W)} \leq C$ , and the same is true of  $\|A_{J^c}\|_F^2$  on event  $\mathcal{E}$ .

(d)  $|J_+| \leq \frac{K\tau_n^{-1}}{\alpha\alpha_+} \wedge p = \frac{K\tau_n^{-1}}{4} \wedge p$ , and the same is true of  $|J|$  on event  $\mathcal{E}$ .

(e)  $\sigma_K(A_{J^*}) > c\sqrt{K}$  for some absolute constant  $c > 0$  on event  $\mathcal{E}$ . This implies  $|J| \geq K$  and  $[D_0]_{J^*}$  have rank  $K$  on  $\mathcal{E}$ .

(f) There exists an absolute constant  $c \in (0, 1)$  such that the entries of  $A_{J^*}^T A_{J^*}$  are all greater than  $c$  on event  $\mathcal{E}$ .

## Proof

(a) Denote  $Z := D - D_0$ . We introduce the set of  $p$ -dimensional one-hot vectors

$$\{T_{im} : 1 \leq i \leq n, 1 \leq m \leq N\},$$

for each word in the dataset; note that  $T_{im} \sim \text{Multinomial}(1, [D_0]_{*i})$  and these one-hot vectors are mutually independent. It follows that each column of  $Z$  satisfies

$$[Z]_{*i} = \frac{1}{N} \sum_{m=1}^N (T_{im} - \mathbb{E}[T_{im}]).$$

Note that for a given  $j \in [p]$ :

$$M(j, j) - M_0(j, j) = \frac{1}{n} \sum_{i=1}^n Z_{ji} = \frac{1}{nN} \sum_{i=1}^n \sum_{m=1}^N (T_{im}(j) - \mathbb{E}[T_{im}(j)]), \quad (34)$$

and since  $|T_{im}(j) - \mathbb{E}[T_{im}(j)]| \leq 1$ , we can apply Hoeffding's inequality to conclude

$$\mathbb{P}(|M(j, j) - M_0(j, j)| \geq t) \leq 2 \exp(-nNt^2/2).$$

(b) Note that  $\alpha_- > 1$ . We have

$$\begin{aligned}
 \mathbb{P}(J_- \not\subseteq J) &= \mathbb{P}\left(\bigcup_{j \in J_-} \{M(j, j) \leq \alpha\tau_n\}\right) \\
 &\leq \sum_{j \in J_-} \mathbb{P}(M(j, j) \leq \alpha\tau_n) \\
 &\leq \sum_{j \in J_-} \mathbb{P}(M(j, j) - M_0(j, j) \leq \alpha\tau_n - \alpha_- \alpha\tau_n) \\
 &\leq \sum_{j \in J_-} \mathbb{P}(|M(j, j) - M_0(j, j)| \geq (\alpha_- - 1)\alpha\tau_n) \\
 &\leq \sum_{j \in J_-} 2 \exp(-Nn(\alpha_- - 1)^2 \alpha^2 \tau_n^2 / 2) \\
 &\leq 2p_n^{1 - (\alpha_- - 1)^2 \alpha^2 / 2},
 \end{aligned}$$

where in the last step we used  $|J_-| \leq p$ . We want to choose  $1 - \frac{(\alpha_- - 1)^2 \alpha^2}{2} < -1$  or equivalently  $(\alpha_- - 1)\alpha > 2$ .

Note that  $0 < \alpha_+ < 1$ . If  $\alpha > 2$ , we further have

$$\begin{aligned}
 \mathbb{P}(J \not\subseteq J_+) &= \mathbb{P}\left(\bigcup_{j \in J_+^c} \{M(j, j) > \alpha\tau_n\}\right) \\
 &\leq \sum_{j \in J_+^c} \mathbb{P}(M(j, j) - M_0(j, j) > (\alpha - \alpha\alpha_+)\tau_n) \\
 &\leq \sum_{j \in J_+^c} \mathbb{P}\left(|M(j, j) - M_0(j, j)| > (1 - \alpha_+)\alpha\sqrt{\frac{\log p_n}{nN}}\right) \\
 &\leq \sum_{j \in J_+^c} 2 \exp\left(-\frac{Nn(1 - \alpha_+)^2 \alpha^2 \log p_n}{2Nn}\right) \\
 &\leq 2p_n^{1 - (1 - \alpha_+)^2 \alpha^2 / 2}.
 \end{aligned}$$

Again, we want to choose  $1 - \frac{(1 - \alpha_+)^2 \alpha^2}{2} < -1$  or equivalently  $(1 - \alpha_+)\alpha > 2$ . A suitable choice is  $\alpha = 8, \alpha_+ = \frac{1}{2}, \alpha_- = 2$ .

When  $\alpha \leq 2$  (the practical range we use; see Remark 9), we need a tighter regime for  $\alpha$ . Instead of using Hoeffding's inequality, we replace Hoeffding by Bernstein to obtain a sharper tail bound:

$$\mathbb{P}(|M(j, j) - M_0(j, j)| \geq t) \leq 2 \exp\left(-\frac{nNt^2/2}{\max_{i,m} \text{Var}(T_{im}(j)) + t/3}\right).$$

For  $j \in J_+^c$ , we have  $M_0(j, j) \leq \alpha_+ \alpha\tau_n$ . Combined with Lemma 24 and Assumption 1, we have  $\text{Var}(T_{im}(j)) \leq [D_0]_{ji} = \sum_k A_{jk} W_{ki} \leq h_j \leq \sigma_K^{-1}(\Sigma_W) M_0(j, j) \leq c^{-1} \alpha_+ \alpha\tau_n$ .

Let  $\alpha'_+ = \max\{c^{-1}\alpha_+, (1 - \alpha_+)/3\}$ . Note that  $0 < \alpha_+ < 1$ . We further have

$$\begin{aligned}
 \mathbb{P}(J \not\subseteq J_+) &= \mathbb{P}\left(\bigcup_{j \in J_+^c} \{M(j, j) > \alpha\tau_n\}\right) \\
 &\leq \sum_{j \in J_+^c} \mathbb{P}(M(j, j) - M_0(j, j) > (\alpha - \alpha\alpha_+)\tau_n) \\
 &\leq \sum_{j \in J_+^c} \mathbb{P}\left(|M(j, j) - M_0(j, j)| > (1 - \alpha_+)\alpha\sqrt{\frac{\log p_n}{nN}}\right) \\
 &\leq \sum_{j \in J_+^c} 2 \exp\left(-\frac{Nn(1 - \alpha_+)^2\alpha^2\tau_n^2/2}{c^{-1}\alpha_+\alpha\tau_n + (1 - \alpha_+)\alpha\tau_n/3}\right) \\
 &\leq 2p_n \exp\left(-\frac{Nn(1 - \alpha_+)^2\alpha\tau_n/2}{\alpha'_+}\right).
 \end{aligned}$$

For any fixed  $c_\alpha > 1$ , if

$$\alpha \geq \frac{2\alpha'_+(1 + c_\alpha)}{(1 - \alpha_+)^2} \sqrt{\frac{\log p_n}{nN}},$$

then

$$\mathbb{P}(J \not\subseteq J_+) \leq 2p_n^{-c_\alpha} = o(p_n^{-1}).$$

If  $\alpha < \sqrt{\frac{\log p_n}{nN}}$ , then few or no entries are thresholded. In this case, many infrequent/noisy words remain and decrease the signal-to-noise ratio as shown in Figures 7 and 8.

(c) From Lemma 24, we have

$$M_0(j, j) \geq \sigma_K(\Sigma_W)\|A_{j*}\|_1 \geq \sigma_K(\Sigma_W)\|A_{j*}\|_2,$$

and so if we define

$$L := \{j \in [p] : \|A_{j*}\|_2 > \beta\tau_n\} \quad \text{where } \beta := \frac{\alpha - \alpha}{\sigma_K(\Sigma_W)} = \frac{16}{\sigma_K(\Sigma_W)}$$

then  $L \subseteq J_-$  by definition of  $J_-$ , and thus  $\|A_{J_-^c}\|_F \leq \|A_{L^c}\|_F$ . Now, if we order the  $\ell_2$  row norms  $\|A_{(1)*}\|_2 \geq \dots \geq \|A_{(p)*}\|_2$  and apply Lemma 23,

$$\begin{aligned}
 \|A_{L^c}\|_F^2 &= \sum_{j \notin L} \|A_{j*}\|_2^2 = \sum_{j \notin L} \min(\|A_{j*}\|_2^2, \beta^2\tau_n^2) \\
 &\leq \sum_{j=1}^p \min(\|A_{(j)*}\|_2^2, \beta^2\tau_n^2) \leq \sum_{j=1}^p \min\left(\frac{K^2}{j^2}, \beta^2\tau_n^2\right) \\
 &\leq \int_0^\infty \min(\beta^2\tau_n^2, K^2t^{-2})dt.
 \end{aligned}$$

Let  $t_0$  satisfy  $\beta^2\tau_n^2 = K^2t^{-2}$  or  $t_0 = \frac{K}{\beta\tau_n}$ . We continue:

$$\begin{aligned} \|A_{L^{c*}}\|_F^2 &\leq t_0\beta^2\tau_n^2 + K^2 \int_{t_0}^{\infty} t^{-2} dt \\ &= t_0\beta^2\tau_n^2 + K^2t_0^{-1} = 2t_0\beta^2\tau_n^2 \\ &= 2K\beta\tau_n = 2K\beta\sqrt{\frac{\log p_n}{Nn}} = o(1). \end{aligned}$$

given our assumption that  $\sigma_K(\Sigma_W) > c$  for some absolute constant  $c > 0$ . Moreover, it is also clear from the definition of  $L$  that

$$\|A_{L^{c*}}\|_F^2 \leq p(\beta\tau_n)^2.$$

- (d) For all  $j \in J_+ := \{j \in [p] : M_0(j, j) > \alpha\alpha_+\tau_n\}$ , note that  $h_j \geq M_0(j, j) > \alpha_+\alpha\tau_n$ . Then observe that

$$K = \sum_{j=1}^p h_j \geq \sum_{j \in J_+} h_j \geq |J_+|\alpha\alpha_+\tau_n = 4|J_+|\tau_n.$$

- (e) Here we use the assumption that  $\sigma_K(A) > c\sqrt{K}$  for some absolute constant  $c > 0$ . Observe that by Weyl's inequality for singular values, on event  $\mathcal{E}$  we have

$$\begin{aligned} \sigma_K(A_{J_*}) &\geq \sigma_K(A) - \|A_{J^{c*}}\|_{\text{op}} \\ &\geq \sigma_K(A) - \|A_{J^{c*}}\|_F \\ &\geq c\sqrt{K} - o(1) \geq c\sqrt{K}/2. \end{aligned}$$

when  $nN$  is sufficiently large, since in part (c) we have shown  $\|A_{J^{c*}}\|_F \leq C \left(\frac{\log p_n}{Nn}\right)^{1/4}$ . Hence,  $A_{J_*}$  has rank  $K$  on  $\mathcal{E}$ , and by Sylvester's rank inequality,

$$K = \text{rank}(A_{J_*}) + \text{rank}(W) - K \leq \text{rank}([D_0]_{J_*}) \leq \text{rank}(A_{J_*}) = K.$$

- (f) Here we use the assumption that the entries of  $A^T A$  are bounded below by an absolute constant. For any  $k, l \in [K]$ , since  $A^T A = A_{J_*}^T A_{J_*} + A_{J^{c*}}^T A_{J^{c*}}$ , on event  $\mathcal{E}$  the  $(k, l)$ -entry of  $A_{J_*}^T A_{J_*}$  satisfies

$$\begin{aligned} (A_{J_*}^T A_{J_*})(k, l) &= (A^T A)(k, l) - \sum_{j \notin J} A_{jk} A_{jl} \\ &\geq c - \|A_{J^{c*}}\|_2 \|A_{J^{c*}}\|_2 \\ &\geq c - \|A_{J^{c*}}\|_F^2 = c - o(1) \geq c/2, \end{aligned}$$

when  $nN$  is sufficiently large. ■

## Appendix B. Properties of Unobserved Quantities

**Lemma 26** *The following statements are true:*

- (a)  $\sigma_1(A) \leq \sqrt{K}$  and  $\sigma_1(\Sigma_W) \leq 1$ , where  $\Sigma_W := \frac{1}{n}WW^T$ .
- (b) If  $\Xi \in \mathbb{R}^{|J| \times K}$  contains the first  $K$  left singular vectors of  $[D_0]_{J^*}$ , then  $\Xi^T A_{J^*}$  is invertible. If  $V := (\Xi^T A_{J^*})^{-1} \in \mathbb{R}^{K \times K}$  then  $V$  satisfies the following:
- (i)  $\Xi = A_{J^*}V$
  - (ii) The singular values of  $V$  are the inverses of the singular values of  $A_{J^*}$
  - (iii) The columns  $V_1, \dots, V_K$  of  $V$  are eigenvectors of the matrix  $\Theta := \Sigma_W A_{J^*}^T A_{J^*}$ , associated with the eigenvalues

$$\lambda_k(\Theta) = \frac{\sigma_k^2([D_0]_{J^*})}{n} \quad \text{for } 1 \leq k \leq K.$$

- (c) The matrix  $\Theta_0 := \Sigma_W A^T A \in \mathbb{R}^{K \times K}$  satisfies the following:
- (i) The entries of  $\Theta_0$  are all positive and bounded below by an absolute constant  $c_1 > 0$ .
  - (ii) The gap between its first two eigenvalues is bounded below by an absolute constant  $c_2 > 0$ .
  - (iii) The entries of the unit-norm leading positive eigenvector of  $\Theta_0$  are all bounded below by an absolute constant  $c_3 > 0$ .
- (d) On event  $\mathcal{E}$ , the results of part (c) also apply to  $\Theta$ , possibly with smaller absolute constants  $c_1, c_2, c_3 > 0$ .
- (e) There exist absolute constants  $c, C > 0$  such that on  $\mathcal{E}$ , the entries of the first column of  $V$  satisfy

$$\frac{c}{\sqrt{K}} \leq \min_{k \in [K]} V_1(k) \leq \max_{k \in [K]} V_1(k) \leq \frac{C}{\sqrt{K}},$$

and if  $\xi_1, \dots, \xi_K$  are the columns of  $\Xi$ , then for any  $j \in J$ , its first column satisfies

$$\frac{ch_j}{\sqrt{K}} \leq \xi_1(j) \leq \frac{Ch_j}{\sqrt{K}}.$$

- (f) Let  $Q \in \mathbb{R}^{K \times K}$  be defined by  $Q^T = [\text{diag}(V_1)]^{-1}V$ , and note that the entries of the first row of  $Q$  are all equal to 1. If  $v_1^*, \dots, v_K^* \in \mathbb{R}^{K-1}$  are defined by the relation

$$Q = \begin{pmatrix} 1 & \dots & 1 \\ v_1^* & \dots & v_K^* \end{pmatrix},$$

then we have

$$c \leq \sigma_K(Q) \leq \sigma_1(Q) \leq C.$$

Consequently,  $v_1^*, \dots, v_K^*$  are affinely independent (which means the simplex defined by their convex hull is non-degenerate) and  $\max_{k \in [K]} \|v_k^*\|_2 \leq C$ .

**Proof**

- (a) The  $k^{\text{th}}$  diagonal entry of  $A^T A$  is  $\|A_{*k}\|_2^2 \leq \|A_{*k}\|_1 = 1$ , so  $\text{tr}(A^T A) \leq K$  which implies  $\sigma_1(A) \leq \sqrt{K}$ . Similarly,

$$\sigma_1(\Sigma_W) = \frac{\sigma_1(W^T W)}{n} \leq \frac{\text{tr}(W^T W)}{n} = \frac{\sum_{i=1}^n \|W_{*i}\|_2^2}{n} \leq \frac{\sum_{i=1}^n \|W_{*i}\|_1}{n} = 1.$$

- (b) By singular value decomposition, we have

$$[D_0]_{J_*} = \Xi \Lambda B^T,$$

where  $\Lambda = \text{diag}(\sigma_1, \dots, \sigma_K) \in \mathbb{R}^{K \times K}$  contains the singular values of  $[D_0]_{J_*}$  and  $B \in \mathbb{R}^{n \times K}$  contains its right singular vectors. Here  $\Xi^T \Xi = B^T B = I_K$ . Then

$$\Xi = \Xi \Lambda B^T B \Lambda^{-1} = [D_0]_{J_*} B \Lambda^{-1} = A_{J_*} W B \Lambda^{-1},$$

If we let  $V = W B \Lambda^{-1} \in \mathbb{R}^{K \times K}$ , then  $\Xi = A_{J_*} V$ . Furthermore, since

$$\Xi^T \Xi = \Xi^T A_{J_*} V = I_K,$$

we can see that  $V$  can be defined as the inverse of  $\Xi^T A_{J_*}$ , thus proving (i). Also, since

$$\Xi^T \Xi = V^T A_{J_*}^T A_{J_*} V = I_K,$$

we have  $V V^T A_{J_*}^T A_{J_*} V V^T = V V^T$ , which implies  $V V^T = (A_{J_*}^T A_{J_*})^{-1}$  and (ii) follows.

Now let  $\xi_1, \dots, \xi_K$  be the columns of  $\Xi$ , and let  $V_1, \dots, V_K$  be the columns of  $V$ . Note that  $\xi_k = A_{J_*} V_k$ . Since  $[D_0]_{J_*} [D_0]_{J_*}^T \xi_k = \sigma_k^2 \xi_k$  and  $\Sigma_W := \frac{1}{n} W W^T$ , we have

$$A_{J_*} \Sigma_W A_{J_*}^T A_{J_*} V_k = A_{J_*} \Sigma_W A_{J_*}^T \xi_k = \frac{1}{n} [D_0]_{J_*} [D_0]_{J_*}^T \xi_k = \frac{\sigma_k^2}{n} \xi_k = \frac{\sigma_k^2}{n} A_{J_*} V_k.$$

Multiplying both sides by  $(A_{J_*}^T A_{J_*})^{-1} A_{J_*}^T$  on the left, we have

$$\Sigma_W A_{J_*}^T A_{J_*} V_k = \frac{\sigma_k^2}{n} V_k,$$

and so  $V_1, \dots, V_K$  are eigenvectors (not necessarily orthonormal) of  $\Theta := \Sigma_W A_{J_*}^T A_{J_*}$ , associated with eigenvalues  $\sigma_k^2/n$  for  $k = 1, \dots, K$ . This proves (iii).

- (c) For any  $1 \leq k, l \leq K$ , (i) follows from our assumptions:

$$\begin{aligned} \Theta_0(k, l) &= \sum_{s=1}^K \Sigma_W(k, s) \cdot (A^T A)(s, l) \\ &\geq \min_{t, u \in [K]} (A^T A)(t, u) \cdot \sum_{s=1}^K \Sigma_W(k, s) \\ &\geq \min_{t, u \in [K]} (A^T A)(t, u) \cdot \Sigma_W(k, k) \\ &\geq \min_{t, u \in [K]} (A^T A)(t, u) \cdot \sigma_K(\Sigma_W) > c. \end{aligned}$$

Let  $\gamma(\Theta_0) := \lambda_1(\Theta_0) - \lambda_2(\Theta_0) \geq 0$  denote the gap between the first two eigenvalues of  $\Theta_0$ . The proof of (ii) is an asymptotic argument. If we consider a sequence  $\{\Theta_0^{(n)}\}$  that varies with  $n$  as  $n \rightarrow \infty$ , then (ii) follows if we can establish that

$$\liminf_{n \rightarrow \infty} \gamma(\Theta_0^{(n)}) > 0.$$

Assume to the contrary that  $\liminf_{n \rightarrow \infty} \gamma(\Theta_0^{(n)}) = 0$ . Then there exists a subsequence  $\{\Theta_0^{(n_m)}\}_{m=1}^\infty$  such that the gap between the first two eigenvalues decays to zero. Since

$$\|\Theta_0^{(n)}\|_{\text{op}} \leq \|\Sigma_W^{(n)}\|_{\text{op}} \|A^{(n)}\|_{\text{op}}^2 \leq K$$

and  $K$  is fixed as  $n$  varies, there must exist a further subsequence that converges to some matrix  $\Theta_0^{(\infty)}$ . By part (i), this matrix  $\Theta_0^{(\infty)}$  has entries that are bounded below by some absolute constant  $c > 0$ , and yet its first two eigenvalues are equal (by eigenvalue continuity). By Perron's theorem (see Section 8.2 of Horn and Johnson, 2013 for a reference), such a matrix  $\Theta_0^{(\infty)}$  cannot exist.

(iii) is also proven in a similar manner. Let  $\eta_0^{(n)} \in \mathbb{R}^K$  denote the leading unit-norm positive eigenvector of  $\Theta_0^{(n)}$ ; its entries are all positive by Perron's theorem. Suppose there exists some  $k \in [K]$  such that

$$\liminf_{n \rightarrow \infty} \eta_0^{(n)}(k) = 0.$$

Note that the mapping from a matrix in  $\mathbb{R}^{K \times K}$  with strictly positive entries to its leading unit-norm positive eigenvector is continuous (this will be further elaborated in part (d)). Again, this implies that there exists a subsequence  $\{\Theta_0^{(n_m)}\}$  that converges to some  $\Theta_0^{(\infty)}$  having strictly positive entries, and yet its leading eigenvector contains a zero entry. This contradicts Perron's theorem.

- (d) In light of Theorem 25(f) which shows  $A_{J_*}^T A_{J_*}$  has entries bounded below by  $c > 0$  on  $\mathcal{E}$ , (i) is proven similarly as in part (c).

We will first show (iii). Note that we refrain from applying the asymptotic arguments of part (c) directly to  $\Theta$  since, unlike  $\Theta_0$ ,  $\Theta$  depends on  $J$  which is random. Also, the  $\sin \theta$  theorem is not applicable to eigenvectors of  $\Theta$  and  $\Theta_0$  as these matrices are not symmetric. Hence, we opt for the approach presented below.

Define the open domain

$$E = \{\Psi \in \mathbb{R}^{K \times K} : \Psi(k, l) > 0 \text{ for all } k, l \in [K]\},$$

and define  $\mathbf{f} : E \rightarrow \mathbb{R}^K$  as the function mapping a matrix in  $E$  to its leading unit-norm positive eigenvector. Also, fix  $\Psi_0 \in E$  and  $1 \leq k, l \leq K$ . For any real-valued  $t$  in a neighborhood of zero, consider the function

$$f_{kl}^{\Psi_0}(t) := \mathbf{f}(\Psi_0 + t\mathbf{e}_k\mathbf{e}_l^T),$$

where  $\mathbf{e}_k$  and  $\mathbf{e}_l$  are the  $k^{\text{th}}$  and  $l^{\text{th}}$  canonical basis vectors of  $\mathbb{R}^k$  respectively.

Since the algebraic multiplicity of the first eigenvalue of any matrix in  $E$  is 1 (Perron's theorem), by Theorem 2 of Greenbaum et al. (2020), for any  $\Psi_0 \in E$  and any  $k, l \in [K]$ , the function  $f_{kl}^{\Psi_0}(\cdot)$  is continuously differentiable around 0 (more specifically, one can write  $f_{kl}^{\Psi_0}(t) = \frac{x(t)}{\|x(t)\|_2}$  for some eigenvector function  $x(t)$  that is analytic in a neighborhood of 0). Therefore, the function  $\mathbf{f}$  itself is continuously differentiable on  $E$ , and we can define its derivative  $\mathbf{f}'(\Psi)$  as a matrix in  $\mathbb{R}^{K^2 \times K}$  containing all the partial derivatives of  $\mathbf{f}$  at  $\Psi \in E$ . Since these partial derivatives are all continuous,  $\mathbf{f}' : E \rightarrow \mathbb{R}^{K^2 \times K}$  is a continuous function.

Now if  $c > 0$  is an absolute constant such that all the entries of  $\Theta_0$  and  $\Theta$  are greater than  $c$  (the latter on event  $\mathcal{E}$ ), then  $\Theta$  and  $\Theta_0$  belong to the set

$$E' = \{\Psi \in \mathbb{R}^{K \times K} : \Psi(k, l) \geq c \text{ for all } k, l \in [K] \text{ and } \|\Psi\|_{\text{op}} \leq K\},$$

which is a compact subset of  $E$ . Let  $\eta$  and  $\eta_0$  be the unit-norm positive first eigenvectors of  $\Theta$  and  $\Theta_0$  respectively. On event  $\mathcal{E}$ , by Theorem 9.19 of Rudin (1976),

$$\begin{aligned} \|\eta - \eta_0\|_2 &= \|\mathbf{f}(\Theta) - \mathbf{f}(\Theta_0)\|_2 \leq \max_{\Psi \in E'} \|\mathbf{f}'(\Psi)\|_{\text{op}} \|\Theta - \Theta_0\|_F \\ &\leq C \|\Sigma_W\|_{\text{op}} \|A^T A - A_{J_*}^T A_{J_*}\|_F \\ &= C \|\Sigma_W\|_{\text{op}} \|A_{J^{c_*}}^T A_{J^{c_*}}\|_F \leq C \|A_{J^{c_*}}\|_F^2 = o(1), \end{aligned}$$

where we note that  $A^T A = A_{J_*}^T A_{J_*} + A_{J^{c_*}}^T A_{J^{c_*}}$ . Hence, for any  $k \in [K]$ ,

$$\eta(k) \geq \eta_0(k) - \|\eta - \eta_0\|_2 \geq c - o(1) > c/2,$$

if  $nN$  is sufficiently large. We have shown  $\min_k \eta(k) \geq c/2 > 0$  on  $\mathcal{E}$ .

As for (ii), we have shown in (b)(iii) for  $\Theta$  (and the proof is similar for  $\Theta_0$ ) that

$$\lambda_k(\Theta) = \frac{\sigma_k^2([D_0]_{J_*})}{n}, \quad \lambda_k(\Theta_0) = \frac{\sigma_k^2(D_0)}{n}.$$

Note that since  $\|A\|_{\text{op}} \leq \sqrt{K}$  and  $\|W\|_{\text{op}} \leq \sqrt{n}$ ,

$$\max[\sigma_k([D_0]_{J_*}), \sigma_k(D_0)] \leq \|A\|_{\text{op}} \|W\|_{\text{op}} \leq \sqrt{Kn},$$

and by Weyl's inequality for singular values (which can be applied after appending zero rows to the matrix  $A_{J_*}$  so as to match the dimension of  $A$ ),

$$\begin{aligned} |\lambda_k(\Theta) - \lambda_k(\Theta_0)| &\leq \frac{|\sigma_k([D_0]_{J_*}) - \sigma_k(D_0)| |\sigma_k([D_0]_{J_*}) + \sigma_k(D_0)|}{n} \\ &\leq \frac{\|A_{J^{c_*}}\|_{\text{op}} \|W\|_{\text{op}} (2\sqrt{Kn})}{n} \\ &\leq 2\sqrt{K} \|A_{J^{c_*}}\|_{\text{op}} = o(1). \end{aligned}$$

on event  $\mathcal{E}$ , so

$$\begin{aligned} |\lambda_1(\Theta) - \lambda_2(\Theta)| &\geq |\lambda_1(\Theta_0) - \lambda_2(\Theta_0)| - o(1) \\ &\geq c - o(1) \geq c/2, \end{aligned}$$

if  $nN$  is sufficiently large, for some absolute constant  $c > 0$ .

(e) By Theorem 25(e), on the event  $\mathcal{E}$  we have

$$\sigma_K(A_{J_*}) \geq c\sqrt{K}.$$

Therefore,

$$\max_{k \in [K]} V_1(k) \leq \|V_1\|_2 \leq \sigma_1(V) = \sigma_K^{-1}(A_{J_*}) \leq \frac{C}{\sqrt{K}},$$

and since

$$\|V_1\|_2 \geq \sigma_K(V) = \sigma_1^{-1}(A_{J_*}) \geq \sigma_1^{-1}(A) \geq \frac{1}{\sqrt{K}}$$

and  $\frac{1}{\|V_1\|_2} V_1$  is the unit-norm leading positive eigenvector of  $\Theta$ , on event  $\mathcal{E}$  we have

$$\min_{k \in [K]} V_1(k) = \|V_1\|_2 \min_{k \in [K]} \left\{ \frac{V_1(k)}{\|V_1\|_2} \right\} \geq \frac{c}{\sqrt{K}},$$

Since  $\xi_1 = A_{J_*} V_1$ , it follows that on event  $\mathcal{E}$ , for any  $j \in J$ ,

$$\frac{ch_j}{\sqrt{K}} \leq \xi_1(j) \leq \frac{Ch_j}{\sqrt{K}}.$$

(f) It can be seen by the definition of  $Q$  that

$$\sigma_K(Q) \geq \frac{\sigma_K(V)}{\max_{k \in [K]} V_1(k)} \geq c > 0,$$

and

$$\sigma_1(Q) \leq \frac{\sigma_1(V)}{\min_{k \in [K]} V_1(k)} \leq C,$$

for some  $c, C > 0$ . Thus,  $\max_{k \in [K]} \|v_k^*\|_2 \leq C$  and  $Q$  has independent columns, which implies  $v_1^*, \dots, v_K^*$  are affinely independent. ■

**Lemma 27** *Let  $A_{J_1}, \dots, A_{J_K}$  be the columns of  $A_{J_*}$ . Under Assumption 4 on the vertex hunting function  $\mathcal{V}(\cdot)$ , the oracle procedure in Definition 5 returns*

$$\tilde{A}_{J_*} = A_{J_*} \cdot \text{diag}(\|A_{J_1}\|_1^{-1}, \dots, \|A_{J_K}\|_1^{-1}). \quad (35)$$

on event  $\mathcal{E}$ .

**Proof** Note that from Lemma 26(b)(i), we have  $\Xi = A_{J_*} V$ . Let  $\mathbf{1}_J$  be the vector of size  $|J|$  with entries all equal to 1. Now, by the definition of  $R$ ,

$$[\mathbf{1}_J, R] = [\text{diag}(\xi_1)]^{-1} \Xi = [\text{diag}(\xi_1)]^{-1} A_{J_*} V.$$

Recall from Lemma 26(f) the definition  $Q^T := [\text{diag}(V_1)]^{-1} V = \begin{pmatrix} 1 & \cdots & 1 \\ v_1^* & \cdots & v_K^* \end{pmatrix}^T$ . Then

$$[\mathbf{1}_J, R] = [\text{diag}(\xi_1)]^{-1} A_{J_*} \cdot \text{diag}(V_1) Q^T = \Pi \begin{pmatrix} 1 & \cdots & 1 \\ v_1^* & \cdots & v_K^* \end{pmatrix}^T, \quad (36)$$

where  $\Pi$  is defined as follows:

$$\Pi := [\text{diag}(\xi_1)]^{-1} A_{J^*} \cdot \text{diag}(V_1) \in \mathbb{R}^{|J| \times K}. \quad (37)$$

From (36) and (37), we can see that  $\Pi$  contains only non-negative entries and the rows of  $\Pi$  sum up to 1. This means the rows of  $R$  (the point cloud) lie inside the convex hull of simplex vertices  $\{v_1^*, \dots, v_K^*\} \subseteq \mathbb{R}^{K-1}$ .

By Assumption 3, for each topic there exists at least an anchor word for that topic in the set  $J$  on event  $\mathcal{E}$ . This means that the point cloud contains at least one point on each vertex  $v_1^*, \dots, v_K^*$ . By Assumption 4, the vertex hunting procedure  $\mathcal{V}(\cdot)$  returns precisely the vertices  $v_1^*, \dots, v_K^*$ . Now let  $\{\pi_j : j \in J\} \subseteq \mathbb{R}^K$  denote the rows of  $\Pi$ . From taking the transpose of (36),  $\Pi$  is then estimated correctly by solving

$$\begin{pmatrix} 1 & \dots & 1 \\ v_1^* & \dots & v_K^* \end{pmatrix} \pi_j = \begin{pmatrix} 1 \\ r_j \end{pmatrix}.$$

Now, by the definition of  $\Pi$  in (37),

$$\text{diag}(\xi_1) \cdot \Pi = A_{J^*} \cdot \text{diag}(V_1),$$

and thus (35) follows if we normalize  $\text{diag}(\xi_1) \cdot \Pi$  to ensure its columns sum up to 1. ■

## Appendix C. Concentration Inequalities Involving $Z = D - D_0$

**Remark 28** This section contains all the concentration inequalities necessary for our analysis, and is comparable to Section E in the appendix of Tracy Ke and Wang (2024).

Lemma 30 and Lemma 31 are similar to Lemmas E.1 and E.2 of Tracy Ke and Wang (2024) in that they are simple applications of Bernstein's inequality. However, it is crucial to note that our results are applicable even when  $\min_{j \in [p]} h_j$  is extremely small, as we only restrict our attention to  $j \in J_+$  (where  $J_+$  is defined in Appendix A). In contrast, Lemmas E.1 and E.2 of Tracy Ke and Wang (2024) require  $\min_{j \in [p]} h_j \geq cK/p$  (or at least  $\min_{j \in [p]} h_j \gg (Nn)^{-1} \log n$ ).

Lemma 32 in our paper is based on standard techniques for deriving concentration inequalities for U-statistics. Our results here can be compared to Lemmas E.3-E.6 of Tracy Ke and Wang (2024), which use a truncation argument and the fact that the product of two sub-Gaussian variables is sub-exponential. Our bounds do not depend on  $p$  except for log factors and are applicable to all parameter regimes (in particular when  $p \gg n \vee N$ ), whereas the bounds in Lemmas E.3-E.6 of Tracy Ke and Wang (2024) depend heavily on  $p$  and  $\min_{j \in [p]} h_j$ .

**Lemma 29 (Bernstein's inequality)** *Let  $X_1, \dots, X_n$  be independent random variables with  $\mathbb{E}(X_i) = 0$ ,  $\text{Var}(X_i) \leq \sigma_i^2$ , and  $|X_i| \leq b$  almost surely for all  $i$ . Let  $\sigma^2 := n^{-1} \sum_{i=1}^n \sigma_i^2$ . Then for any  $t > 0$ ,*

$$\mathbb{P} \left( n^{-1} \left| \sum_{i=1}^n X_i \right| \geq t \right) \leq 2 \exp \left( -\frac{nt^2/2}{\sigma^2 + bt/3} \right).$$

**Lemma 30** Denote  $\tilde{h}_j := h_j \wedge 1$ . With probability at least  $1 - o(p_n^{-1})$ ,

$$|M(j, j) - M_0(j, j)| \leq C^* \sqrt{\frac{\tilde{h}_j \log p_n}{nN}} \quad \text{for all } j \in J_+. \quad (38)$$

**Proof** Similar to (34), for a fixed  $j \in J_+$  we have

$$M(j, j) - M_0(j, j) = \frac{1}{n} \sum_{i=1}^n Z_{ji} = \frac{1}{nN} \sum_{i=1}^n \sum_{m=1}^N (T_{im}(j) - \mathbb{E}[T_{im}(j)]).$$

Note that since  $T_{im}(j) \sim \text{Bernoulli}(D_0(j, i))$ ,  $|T_{im}(j) - \mathbb{E}[T_{im}(j)]| \leq 1$  and

$$\text{Var}(T_{im}(j)) \leq D_0(j, i) = \sum_{k=1}^K A_{jk} W_{ki} \leq \sum_{k=1}^K A_{jk} = h_j.$$

(and also  $\text{Var}(T_{im}(j)) \leq 1$ ). We apply Bernstein's inequality to conclude for any  $t > 0$ :

$$\mathbb{P}(|M(j, j) - M_0(j, j)| \geq t) \leq 2 \exp\left(-\frac{nNt^2/2}{\tilde{h}_j + t/3}\right).$$

One can choose  $t = C^* \sqrt{\frac{\tilde{h}_j \log p_n}{nN}}$  or  $t = \frac{C^* \log p_n}{nN}$  depending on whether  $\tilde{h}_j \geq \frac{\log p_n}{nN}$  holds. Thus with probability at least  $1 - o(p_n^{-2})$ ,

$$\begin{aligned} |M(j, j) - M_0(j, j)| &\leq C^* \max\left(\sqrt{\frac{\tilde{h}_j \log p_n}{nN}}, \frac{\log p_n}{nN}\right) \\ &\leq C^* \sqrt{\frac{\tilde{h}_j \log p_n}{nN}}, \end{aligned}$$

since if  $j \in J_+$ , then  $\tilde{h}_j > \alpha_+ \alpha \sqrt{\frac{\log p_n}{nN}} \geq \frac{c^* \log p_n}{nN}$  when  $nN$  is sufficiently large so that  $\frac{\log p_n}{nN} \leq 1$ . We then take union bound over  $j \in J_+$ .  $\blacksquare$

**Lemma 31** Denote  $\{Z_j : j \in J_+\} \subseteq \mathbb{R}^n$  as the rows of  $Z$  in  $J_+$ , and  $\{W_k : k \in [K]\} \subseteq \mathbb{R}^n$  as the rows of  $W$ . With probability at least  $1 - o(p_n^{-1})$ ,

$$|Z_j^T W_k| \leq C^* \sqrt{\frac{n\tilde{h}_j \log p_n}{N}} \quad \text{for all } j \in J_+ \text{ and } k \in [K]. \quad (39)$$

**Proof** Note that

$$Z_{ji} = \frac{1}{N} \sum_{m=1}^N (T_{im}(j) - \mathbb{E}[T_{im}(j)]), \quad (40)$$

and so for any  $j \in J_+$  and  $k \in [K]$ ,

$$Z_j^T W_k = \sum_{i=1}^n Z_{ji} W_{ki} = \frac{1}{nN} \sum_{i=1}^n \sum_{m=1}^N n W_{ki}(T_{im}(j) - \mathbb{E}[T_{im}(j)]).$$

We note that  $|nW_{ki}(T_{im}(j) - \mathbb{E}[T_{im}(j)])| \leq n$  and  $\text{Var}[nW_{ki}(T_{im}(j) - \mathbb{E}[T_{im}(j)])] \leq n^2 \tilde{h}_j$ , so by Bernstein inequality, for any  $t > 0$ ,

$$\mathbb{P}(|Z_j^T W_k| > t) \leq 2 \exp\left(-\frac{nNt^2/2}{n^2 \tilde{h}_j + nt/3}\right).$$

We can let  $t = C^* \sqrt{\frac{n \tilde{h}_j \log p_n}{N}}$  and again by noting that  $\tilde{h}_j \geq \alpha_+ \alpha \sqrt{\frac{\log p_n}{nN}} \geq c^* \frac{\log p_n}{nN}$  if  $j \in J_+$ , we obtain (39). ■

**Lemma 32** *With probability at least  $1 - o(p_n^{-1})$ ,*

$$|Z_j^T Z_l - \mathbb{E}(Z_j^T Z_l)| \leq C^* \sqrt{\frac{n \tilde{h}_j \tilde{h}_l \log p_n}{N}} \quad \text{for all } j, l \in J_+ \text{ with } j \neq l. \quad (41)$$

$$|Z_j^T Z_j - \mathbb{E}(Z_j^T Z_j)| \leq C^* \sqrt{\frac{n \tilde{h}_j^2 \log p_n}{N}} + \frac{C^*}{N} \sqrt{\frac{n \tilde{h}_j \log p_n}{N}} \quad \text{for all } j \in J_+. \quad (42)$$

**Proof** Denote  $X_{im}(j) := T_{im}(j) - \mathbb{E}[T_{im}(j)]$ . Fix  $j, l \in J_+$ . By (40), note that

$$\begin{aligned} Z_j^T Z_l &= \sum_{i=1}^n Z_{ji} Z_{li} = \frac{1}{N^2} \sum_{i=1}^n \sum_{m=1}^N \sum_{s=1}^N X_{im}(j) X_{is}(l) \\ &= \frac{1}{N^2} \sum_{i=1}^n \sum_{m=1}^N X_{im}(j) X_{im}(l) + \frac{1}{N^2} \sum_{i=1}^n \sum_{\substack{1 \leq m, s \leq N \\ m \neq s}} X_{im}(j) X_{is}(l) \\ &= \frac{n}{N} V_1 + \frac{N-1}{N} V_2, \end{aligned}$$

where we define

$$\begin{aligned} V_1 &:= \frac{1}{nN} \sum_{i=1}^n \sum_{m=1}^N X_{im}(j) X_{im}(l). \\ V_2 &:= \frac{1}{N(N-1)} \sum_{i=1}^n \sum_{\substack{1 \leq m, s \leq N \\ m \neq s}} X_{im}(j) X_{is}(l). \end{aligned}$$

Note that  $\mathbb{E}(V_2) = 0$ , and we need an upper bound on  $|V_1 - \mathbb{E}(V_1)|$  and  $|V_2|$ . We will deal with  $V_2$  first. Define  $S_N$  as the set of permutations on  $\{1, \dots, N\}$  and  $N' := \lfloor N/2 \rfloor$ . Also define

$$W_i(X_{i1}, \dots, X_{iN}) := \frac{1}{N'} \sum_{m=1}^{N'} X_{i,2m-1}(j) X_{i,2m}(l),$$

Then by symmetry (note that the inner sum over  $m, s$  in the definition of  $V_2$  has  $N(N-1)$  summands),

$$V_2 = \frac{\sum_{i=1}^n \sum_{\pi \in S_N} W_i(X_{i,\pi(1)}, \dots, X_{i,\pi(N)})}{N!}.$$

Define, for a given  $\pi \in S_N$ ,

$$Q_\pi := \sum_{i=1}^n N' W_i(X_{\pi(1)}, \dots, X_{\pi(N)}),$$

so that  $N'V_2 = \frac{1}{N!} \sum_{\pi \in S_N} Q_\pi$ . For arbitrary  $t, s > 0$ , by Markov's inequality and the convexity of the exponential function,

$$\mathbb{P}(N'V_2 \geq t) \leq e^{-st} \mathbb{E}(e^{sN'V_2}) \leq e^{-st} \frac{\sum_{\pi \in S_N} \mathbb{E}(e^{sQ_\pi})}{N!}.$$

Also, define  $Q = Q_\pi$  for  $\pi$  being the identity permutation. Observe that

$$Q = \sum_{i=1}^n \sum_{m=1}^{N'} Q_{im} \quad \text{where } Q_{im} = X_{i,2m-1}(j) X_{i,2m}(l)$$

so  $Q$  is a (double) summation of mutually independent variables. We have  $|Q_{im}| \leq 1$ ,  $\mathbb{E}(Q_{im}) = 0$  and  $\mathbb{E}(Q_{im}^2) \leq \tilde{h}_j \tilde{h}_l$ . The rest of the proof for  $V_2$  is similar to the standard proof for the usual Bernstein's inequality and one can skip to the conclusion (43).

If we denote  $G(x) = \frac{e^x - 1 - x}{x^2}$ , observe  $G(x)$  is increasing. Hence,

$$\begin{aligned} \mathbb{E}(e^{sQ_{im}}) &= \mathbb{E}\left(1 + sQ_{im} + \frac{s^2 Q_{im}^2}{2} + \dots\right) \\ &= \mathbb{E}[1 + s^2 Q_{im}^2 G(sQ_{im})] \\ &\leq \mathbb{E}[1 + s^2 Q_{im}^2 G(s)] \\ &\leq 1 + s^2 \tilde{h}_j \tilde{h}_l G(s) \leq e^{s^2 \tilde{h}_j \tilde{h}_l G(s)}. \end{aligned}$$

Hence,

$$e^{-st} \mathbb{E}(e^{sQ}) = \exp(-st + N' n \tilde{h}_j \tilde{h}_l s^2 G(s)),$$

Since this bound is applicable to all the other  $Q_\pi$  and not just  $\pi$  being equal to the identity permutation, we have

$$\mathbb{P}(N'V_2 \geq t) \leq \exp(-st + N' n \tilde{h}_j \tilde{h}_l s^2 G(s)) = \exp\left(-st + N' n \tilde{h}_j \tilde{h}_l (e^s - 1 - s)\right).$$

Now we choose  $s = \log\left(1 + \frac{t}{N' n \tilde{h}_j \tilde{h}_l}\right) > 0$ . Then

$$\begin{aligned} \mathbb{P}(N'V_2 \geq t) &\leq \exp\left[-t \log\left(1 + \frac{t}{N' n \tilde{h}_j \tilde{h}_l}\right) + N' n \tilde{h}_j \tilde{h}_l \left(\frac{t}{N' n \tilde{h}_j \tilde{h}_l} - \log\left(1 + \frac{t}{N' n \tilde{h}_j \tilde{h}_l}\right)\right)\right] \\ &= \exp\left[-N' n \tilde{h}_j \tilde{h}_l H\left(\frac{t}{N' n \tilde{h}_j \tilde{h}_l}\right)\right], \end{aligned}$$

where we define the function  $H(x) = (1+x)\log(1+x) - x$ . Note that we have the inequality

$$H(x) \geq \frac{3x^2}{6+2x},$$

for all  $x > 0$ . Hence,

$$\mathbb{P}(N'V_2 \geq t) \leq \exp\left(-\frac{t^2/2}{N'n\tilde{h}_j\tilde{h}_l + t/3}\right),$$

or by rescaling,

$$\mathbb{P}(N'V_2 \geq N'nt) \leq \exp\left(-\frac{N'nt^2/2}{\tilde{h}_j\tilde{h}_l + t/3}\right). \quad (43)$$

We can choose  $t^2 = \frac{C^*\tilde{h}_j\tilde{h}_l}{N'n} \log p_n$  and note that  $\tilde{h}_j\tilde{h}_l \geq (\alpha_+)^2 \frac{\log p_n}{nN}$  if  $j, l \in J_+$ . Hence, with probability  $1 - o(p_n^{-1})$  (even after taking union bound over  $j, l \in J_+$ ),

$$|V_2| \leq C^* \sqrt{\frac{n\tilde{h}_j\tilde{h}_l \log p_n}{N}},$$

As for  $V_1$ , we can just apply the usual Bernstein's inequality. Let  $\mu_{ij} = \mathbb{E}[T_{im}(j)] = [D_0]_{ji}$  and define  $\mu_{il}$  similarly; note  $\mu_{ij} \leq \tilde{h}_j$ . Since  $X_{im}(j) = T_{im}(j) - \mu_{ij}$ ,

$$X_{im}(j)X_{im}(l) = T_{im}(j)T_{im}(l) - \mu_{ij}T_{im}(l) - \mu_{il}T_{im}(j) + \mu_{ij}\mu_{il}, \quad (44)$$

If  $j \neq l$  then  $T_{im}(j)T_{im}(l) = 0$  and so

$$\begin{aligned} \text{Var}[X_{im}(j)X_{im}(l)] &= \text{Var}[\mu_{ij}T_{im}(l) + \mu_{il}T_{im}(j)] \\ &\leq \mathbb{E}[\mu_{ij}T_{im}(l) + \mu_{il}T_{im}(j)]^2 \\ &= \mu_{ij}^2\mu_{il} + \mu_{il}^2\mu_{ij} = \mu_{ij}\mu_{il}(\mu_{ij} + \mu_{il}) \\ &\leq \mu_{ij}\mu_{il} \leq \tilde{h}_j\tilde{h}_l, \end{aligned}$$

since  $\mu_{ij} + \mu_{il} \leq 1$ . Hence, by Bernstein's inequality,

$$\mathbb{P}(|V_1 - \mathbb{E}(V_1)| \geq t) \leq 2 \exp\left(-\frac{nNt^2/2}{\tilde{h}_j\tilde{h}_l + t/3}\right),$$

which is similar to (43), so we obtain with probability  $1 - o(p_n^{-1})$  that

$$\frac{n}{N}|V_1 - \mathbb{E}(V_1)| \leq \frac{C^*}{N} \sqrt{\frac{n\tilde{h}_j\tilde{h}_l \log p_n}{N}} \leq C^* \sqrt{\frac{n\tilde{h}_j\tilde{h}_l \log p_n}{N}},$$

and (41) is proven.

If  $j = l$  then since  $T_{im}^2(j) = T_{im}(j)$ , (44) leads to

$$X_{im}^2(j) = T_{im}(j)(1 - 2\mu_{ij}) + \mu_{ij}^2,$$

and since  $|1 - 2\mu_{ij}| \leq 1$  and  $\text{Var}(T_{im}(j)) = \mu_{ij}(1 - \mu_{ij})$ ,

$$\text{Var}[X_{im}^2(j)] \leq \mu_{ij} \leq \tilde{h}_j,$$

and so we obtain (42) since with probability  $1 - o(p_n^{-1})$

$$\frac{n}{N} |V_1 - \mathbb{E}(V_1)| \leq \frac{C^*}{N} \sqrt{\frac{n\tilde{h}_j \log p_n}{N}}.$$

■

**Corollary 33** *With probability  $1 - o(p_n^{-1})$ , the following statements hold:*

$$\|[ZW_k]_{J_+}\|_2 \leq C^* \sqrt{\frac{nK \log p_n}{N}} \text{ for all } 1 \leq k \leq K. \quad (45)$$

$$\|[ZZ^T - \mathbb{E}(ZZ^T)]_{jJ_+}\|_2 \leq C^* \sqrt{\frac{n\tilde{h}_j K \log p_n}{N}} \text{ for all } j \in J_+. \quad (46)$$

$$\|[ZZ^T - \mathbb{E}(ZZ^T)]_{J_+J_+}\|_F \leq C^* K \sqrt{\frac{n \log p_n}{N}}. \quad (47)$$

**Proof** This follows from (39), (41) and (42) after squaring the error bounds and summing them up. We note that  $\sum_{j \in J_+} \tilde{h}_j \leq \sum_{j=1}^p h_j = K$ . ■

## Appendix D. Estimation Errors for Singular Vectors and the Point Cloud

We will use the following theorem (a row-wise perturbation bound for eigenvectors) from Tracy Ke and Wang (2024).

**Lemma 34 (Lemma F.1 of Tracy Ke and Wang, 2024)** *Let  $B_0$  and  $B$  be  $m \times m$  symmetric matrices with  $\text{rank}(B_0) = K$ , and assume  $B_0$  is positive semi-definite. For  $1 \leq k \leq K$ , let  $\delta_k^0$  and  $\delta_k$  be the  $k^{\text{th}}$  largest eigenvalues of  $B_0$  and  $B$  respectively, and let  $u_k^0$  and  $u_k$  be the  $k^{\text{th}}$  eigenvectors of  $B_0$  and  $B$ . Fix  $1 \leq s \leq k \leq K$ . If for some  $c \in (0, 1)$ , suppose (by default, if  $s = 1$  then  $\delta_{s-1}^0 - \delta_s^0 = \infty$ )*

$$\min(\delta_{s-1}^0 - \delta_s^0, \delta_k^0 - \delta_{k+1}^0, \min_{l \in [K]} \delta_l^0) \geq c \|B_0\|_{op}, \quad \|B - B_0\|_{op} \leq (c/3) \|B_0\|_{op}.$$

Write  $U_0 = [u_s^0, \dots, u_k^0]$ ,  $U = [u_s, \dots, u_k]$  and  $\Xi = [u_1^0, \dots, u_K^0]$ . There exists an orthonormal matrix  $O$  such that for all  $1 \leq j \leq p$ ,

$$\|(UO - U_0)_{j*}\|_2 \leq \frac{5}{c \|B_0\|_{op}} (\|B - B_0\|_{op} \|\Xi_{j*}\|_2 + \sqrt{K} \|(B - B_0)_{j*}\|_2).$$

If we define

$$G := DD^T - \frac{n}{N}M.$$

$$G_0 := \left(1 - \frac{1}{N}\right) D_0 D_0^T,$$

then the above lemma can be applied to the submatrices  $G_{J,J}$  and  $[G_0]_{J,J}$  (see Lemma 40).

**Lemma 35** *With probability  $1 - o(p_n^{-1})$ , we have*

$$\|(G - G_0)_{J_+ J_+}\|_{op} \leq C^* K \sqrt{\frac{nK \log p_n}{N}},$$

and for any  $j \in J_+$ , row  $j$  of  $(G - G_0)_{J_+ J_+}$  has  $\ell_2$  norm satisfying

$$\|(G - G_0)_{j J_+}\|_2 \leq C^* K \sqrt{\frac{nh_j \log p_n}{N}}.$$

**Proof** From basic properties of the multinomial distribution, we can show that

$$\mathbb{E}(ZZ^T) = \sum_{i=1}^n \text{Cov}(Z_{*i}) = \sum_{i=1}^n \text{Cov}(D_{*i}) = \frac{n}{N}M_0 - \frac{1}{N}D_0 D_0^T,$$

and therefore

$$\begin{aligned} G - G_0 &= DD^T - \frac{n}{N}M - \left(1 - \frac{1}{N}\right) D_0 D_0^T \\ &= (D_0 + Z)(D_0 + Z)^T - \frac{n}{N}M - \left(1 - \frac{1}{N}\right) D_0 D_0^T \\ &= ZD_0^T + D_0 Z^T + ZZ^T - \frac{n}{N}M + \frac{1}{N}D_0 D_0^T \\ &= ZD_0^T + D_0 Z^T + (ZZ^T - \mathbb{E}[ZZ^T]) + \frac{n}{N}(M_0 - M), \end{aligned}$$

and so we can write  $(G - G_0)_{J_+ J_+} = E_1 + E_2 + E_3$  where

$$E_1 := (ZD_0^T + D_0 Z^T)_{J_+ J_+}.$$

$$E_2 := (ZZ^T - \mathbb{E}[ZZ^T])_{J_+ J_+}.$$

$$E_3 := \frac{n}{N}(M_0 - M)_{J_+ J_+}.$$

We can deal with  $E_3$  first. From (38), with probability  $1 - o(p_n^{-1})$  we have

$$\|E_3\|_{op} \leq \frac{C^* n}{N} \sqrt{\frac{(\max_{j \in J_+} \tilde{h}_j) \log p_n}{nN}} \leq \frac{C^*}{N} \sqrt{\frac{n \log p_n}{N}},$$

and for any  $j \in J_+$ ,

$$\|[E_3]_{j*}\|_2 = \frac{n}{N} |M(j, j) - M_0(j, j)| \leq \frac{C^*}{N} \sqrt{\frac{n\tilde{h}_j \log p_n}{N}}.$$

From (46) and (47), with probability  $1 - o(p_n^{-1})$

$$\|E_2\|_{\text{op}} \leq \|E_2\|_F \leq C^* K \sqrt{\frac{n \log p_n}{N}}.$$

$$\|[E_2]_{j*}\|_2 \leq C^* \sqrt{\frac{n\tilde{h}_j K \log p_n}{N}},$$

If we denote  $A_1, \dots, A_K$  as the columns of  $A$  and  $W_1, \dots, W_k$  as the rows of  $W$ ,

$$D_0 Z^T = \sum_{k=1}^K A_k (Z W_k)^T,$$

and so from (45) and the fact that  $\sum_{k=1}^K \|A_k\|_2 \leq \sum_{k=1}^K \|A_k\|_1 \leq K$ ,

$$\|E_1\|_{\text{op}} \leq 2 \|[D_0 Z^T]_{J_+ J_+}\|_{\text{op}} \leq 2 \sum_{k=1}^K \|A_k\|_2 \|Z_{J_+*} W_k\|_2 \leq C^* K \sqrt{\frac{nK \log p_n}{N}}.$$

Let  $Z_1, \dots, Z_p$  denote the rows of  $Z$ . From (39), (45) and the fact that  $\sum_{k=1}^K A_k(j) = h_j$  and  $h_j \leq K$ , for any  $j \in J_+$ :

$$\begin{aligned} \|[E_1]_{j*}\|_2 &\leq \sum_{k=1}^K A_k(j) \|Z_{J_+*} W_k\|_2 + \sum_{k=1}^K |Z_j^T W_k| \|A_k\|_2 \\ &\leq C^* h_j \sqrt{\frac{nK \log p_n}{N}} + C^* K \sqrt{\frac{n\tilde{h}_j \log p_n}{N}} \\ &\leq C^* K \sqrt{\frac{nh_j \log p_n}{N}}, \end{aligned}$$

Since the bounds for  $\|E_1\|_{\text{op}}$  and  $\|[E_1]_{j*}\|_2$  dominate those for  $E_2$  and  $E_3$ , our result follows.  $\blacksquare$

**Lemma 36** *With probability  $1 - o(p_n^{-1})$ , we also have*

$$\|(G - G_0)_{JJ}\|_{\text{op}} \leq C^* K \sqrt{\frac{nK \log p_n}{N}}, \quad (48)$$

and for any  $j \in J$ ,

$$\|(G - G_0)_{jJ}\|_2 \leq C^* K \sqrt{\frac{nh_j \log p_n}{N}}.$$

**Proof** This is simply a consequence of the previous lemma and the fact that  $J \subseteq J_+$  with probability  $1 - o(p_n^{-1})$ , which implies that  $(G - G_0)_{JJ}$  is a submatrix of  $(G - G_0)_{J_+J_+}$ . Note that we refrain from applying the argument of the previous lemma directly to  $(G - G_0)_{JJ}$ , since  $J$  and  $Z$  are not independent (whereas  $J_+$  is a non-random index set).  $\blacksquare$

**Corollary 37** *Let  $g_n$  be a quantity satisfying*

$$c\sqrt{\frac{nN}{\log p_n}} \geq g_n \geq C^*K\sqrt{K}, \quad (49)$$

where  $C^*$  in (49) is the constant from (48) and  $c$  is another constant to be determined. If

$$\hat{K} := \max \left\{ k : \lambda_k(G_{JJ}) > g_n \sqrt{\frac{n \log p_n}{N}} \right\},$$

then  $\hat{K} = K$  with probability  $1 - o(p_n^{-1})$ .

**Proof** We have shown in Theorem 25(e) that  $[G_0]_{JJ}$  has rank  $K$  on  $\mathcal{E}$ . By Weyl's inequality,

$$\lambda_{K+1}(G_{JJ}) \leq \|(G - G_0)_{JJ}\|_{\text{op}} \leq C^*K\sqrt{K}\sqrt{\frac{n \log p_n}{N}} \leq g_n\sqrt{\frac{n \log p_n}{N}}.$$

This implies  $\hat{K} \leq K$ . On the other hand, again by Weyl's inequality,

$$|\lambda_K(G_{JJ}) - \lambda_K([G_0]_{JJ})| \leq C^*K\sqrt{K}\sqrt{\frac{n \log p_n}{N}} \leq g_n\sqrt{\frac{n \log p_n}{N}},$$

and since  $G_0 := (1 - \frac{1}{N})D_0D_0^T$ , by Theorem 25(e) and Assumption 1,

$$\lambda_K([G_0]_{JJ}) \geq \left(1 - \frac{1}{N}\right) n[\sigma_K(A_{J_*})]^2 \sigma_K(\Sigma_W) \geq cKn > 2g_n\sqrt{\frac{n \log p_n}{N}},$$

when  $nN$  is sufficiently large and  $c$  in (49) is chosen appropriately. Hence,

$$\lambda_K(G_{JJ}) \geq \lambda_K([G_0]_{JJ}) - |\lambda_K(G_{JJ}) - \lambda_K([G_0]_{JJ})| > g_n\sqrt{\frac{n \log p_n}{N}},$$

and thus  $K \leq \hat{K}$  with probability  $1 - o(p_n^{-1})$ .  $\blacksquare$

Recall that  $\hat{\Xi}$  contains the first  $K$  eigenvectors of  $G_{JJ}$  and  $\Xi$  contains the first  $K$  left singular vectors of  $[D_0]_{J_*}$ , or equivalently the first  $K$  eigenvectors of  $[G_0]_{JJ}$ . We will provide a coordinate-wise error bound for  $\hat{\Xi}$  in Lemma 40. First we need a few lemmas.

**Lemma 38** *For any  $j \in J$ ,  $\|\Xi_{j_*}\|_2 \leq Ch_j$ .*

**Proof** Note  $\Xi = A_{J^*}V$  so  $\Xi_{j^*} = A_{j^*}V$ . By Theorem 25(e), on the event  $\mathcal{E}$  we have

$$\sigma_K(A_{J^*}) \geq c\sqrt{K}.$$

Therefore,

$$\|\Xi_{j^*}\|_2 \leq \|V\|_{\text{op}}\|A_{j^*}\|_2 = \sigma_K(A_{J^*})^{-1}\|A_{j^*}\|_2 \leq Ch_j,$$

where the equality uses Lemma 26(b)(ii), and the last inequality uses that  $K$  is fixed.  $\blacksquare$

Note that on event  $\mathcal{E}$ ,  $[D_0]_{J^*}$  and hence  $[G_0]_{JJ}$  has rank  $K$ .

**Lemma 39** *On event  $\mathcal{E}$ ,*

$$cnK \leq \lambda_k([G_0]_{JJ}) \leq nK \text{ for all } k \in [K] \text{ and } \lambda_1([G_0]_{JJ}) \geq cn + \max_{2 \leq k \leq K} \lambda_k([G_0]_{JJ}).$$

**Proof** We note that  $[D_0D_0^T]_{JJ} = A_{J^*}WW^T A_{J^*}^T = nA_{J^*}\Sigma_W A_{J^*}^T$ . Hence,

$$\lambda_1([G_0]_{JJ}) \leq n\|A\|_{\text{op}}^2\|\Sigma_W\|_{\text{op}} \leq nK,$$

$$\lambda_K([G_0]_{JJ}) \geq n[\sigma_K(A_{J^*})]^2\sigma_K(\Sigma_W) \geq cnK,$$

where the last inequality uses Theorem 25(e) and Assumption 1. We also note that for any two matrices  $P$  and  $Q$ , the nonzero eigenvalues of  $PQ$  are the same as those of  $QP$ . Thus the nonzero eigenvalues of  $[D_0D_0^T]_{JJ}$  are the same as the nonzero eigenvalues of  $WW^T A_{J^*}^T A_{J^*} =: n\Theta$ . We have already shown in Lemma 26(d) that the gap between the first two eigenvalues of  $\Theta$  are at least an absolute constant on  $\mathcal{E}$ . Hence, our result follows.  $\blacksquare$

**Lemma 40 (Row-wise estimation error for  $\hat{\Xi}$ )** *Denote  $\{\Xi_j : j \in J\}$  as the rows of  $\Xi$  and  $\{\hat{\Xi}_j : j \in J\}$  as the rows of  $\hat{\Xi}$ . With probability  $1 - o(p_n^{-1})$ , there exist  $\omega \in \{\pm 1\}$  and an orthonormal matrix  $\Omega^* \in \mathbb{R}^{(K-1) \times (K-1)}$  such that, if  $\Omega := \text{diag}(\omega, \Omega^*) \in \mathbb{R}^{K \times K}$ , we have*

$$\|\Omega \hat{\Xi}_j - \Xi_j\|_2 \leq C \sqrt{\frac{h_j \log p_n}{nN}} \text{ for all } j \in J.$$

**Proof** Let  $\hat{\xi}_1$  and  $\xi_1$  be the first eigenvectors of  $[G]_{JJ}$  and  $[G_0]_{JJ}$  respectively. The gap between the first two eigenvalues of  $[G_0]_{JJ}$  is at least  $cn$ , which is much greater than  $C^*K\sqrt{\frac{nK \log p_n}{N}}$  (the high-probability bound on  $\|(G - G_0)_{JJ}\|_{\text{op}}$ ). By applying Lemma 34, there exists  $\omega \in \{\pm 1\}$  such that with probability  $1 - o(p_n^{-1})$ , for all  $j \in J$ ,

$$\begin{aligned} |\omega \hat{\xi}_1(j) - \xi_1(j)| &\leq C \frac{h_j \|(G - G_0)_{JJ}\|_{\text{op}} + \sqrt{K} \|(G - G_0)_{jJ}\|_2}{n} \\ &\leq C \frac{h_j \sqrt{\frac{n \log p_n}{N}} + \sqrt{\frac{nh_j \log p_n}{N}}}{n} \\ &\leq C \sqrt{\frac{h_j \log p_n}{nN}}, \end{aligned}$$

where we applied  $h_j \leq K$ .

Let  $\Xi^* = [\xi_2, \dots, \xi_K]$  contain the other  $(K-1)$  eigenvectors of  $[G_0]_{JJ}$ , and define  $\hat{\Xi}^*$  similarly. Again, since the smallest nonzero eigenvalue of  $[G_0]_{JJ}$  is at least  $cnK$ , there exists an orthonormal matrix  $\Omega^* \in \mathbb{R}^{(K-1) \times (K-1)}$  such that for all  $j \in J$ ,

$$\|(\hat{\Xi}^* \Omega^* - \Xi^*)_{j*}\|_2 \leq C \sqrt{\frac{h_j \log p_n}{nN}}.$$

We then define  $\Omega = \text{diag}(\omega, \Omega^*)$  and combine the above results.  $\blacksquare$

**Lemma 41 (Estimation error for the point cloud)** *With probability  $1 - o(p_n^{-1})$ , all entries of  $\hat{\xi}_1$  have the same sign and there exists an orthonormal matrix  $\Omega^* \in \mathbb{R}^{(K-1) \times (K-1)}$  such that*

$$\max_{j \in J} \|\Omega^* \hat{r}_j - r_j\|_2 \leq C \left( \frac{\log p_n}{nN} \right)^{1/4}.$$

**Proof** First, we note that WLOG, we can assume  $\omega = 1$ . This is because from the previous lemma, for any  $j \in J$ , since  $h_j \geq c \sqrt{\frac{\log p_n}{nN}}$ ,

$$|\omega \hat{\xi}_1(j) - \xi_1(j)| \leq C \sqrt{\frac{h_j \log p_n}{nN}} \leq Ch_j \left( \frac{\log p_n}{nN} \right)^{1/4},$$

whereas we also know from Lemma 26(e) that

$$\xi_1(j) > ch_j > 0.$$

We can see that  $|\omega \hat{\xi}_1(j) - \xi_1(j)| \ll \xi_1(j)$  with high probability as  $nN$  is sufficiently large, and this implies  $\omega \hat{\xi}_1(j) \geq \xi_1(j)/2$ . If  $\hat{\xi}_1$  is defined such that the majority of its entries are positive (and in fact its entries are all of the same sign with high probability), we can simply assume  $\omega = 1$  from now on.

Denote  $\{\Xi_j : j \in J\}$  as the rows of  $\Xi$  and  $\{\hat{\Xi}_j : j \in J\}$  as the rows of  $\hat{\Xi}$ . Now, since by definition,

$$\begin{pmatrix} 1 \\ r_j \end{pmatrix} = [\xi_1(j)]^{-1} \Xi_j, \quad \begin{pmatrix} 1 \\ \Omega^* \hat{r}_j \end{pmatrix} = [\hat{\xi}_1(j)]^{-1} \Omega \hat{\Xi}_j.$$

it follows that

$$\begin{aligned} \|\Omega^* \hat{r}_j - r_j\|_2 &= \left\| \frac{1}{\hat{\xi}_1(j)} \Omega \hat{\Xi}_j - \frac{1}{\xi_1(j)} \Xi_j \right\|_2 \\ &= \left\| \frac{1}{\hat{\xi}_1(j)} (\Omega \hat{\Xi} - \Xi) - \frac{\hat{\xi}_1(j) - \xi_1(j)}{\hat{\xi}_1(j)} r_j \right\|_2 \\ &\leq |\hat{\xi}_1(j)|^{-1} (\|\Omega \hat{\Xi} - \Xi\|_2 + \|r_j\|_2 |\hat{\xi}_1(j) - \xi_1(j)|). \end{aligned}$$

We have noted in Lemma 27 that the point cloud  $\{r_j : j \in J\}$  lies entirely in the convex hull of  $v_1^*, \dots, v_K^*$ , and Lemma 26(f) shows that  $\max_{k \in [K]} \|v_k^*\|_2 \leq C$ , so we also have  $\max_{j \in J} \|r_j\|_2 \leq C$ . We have also noted before that

$$\hat{\xi}_1(j) \geq \frac{\xi_1(j)}{2} > ch_j$$

with high probability. Therefore, with probability  $1 - o(p_n^{-1})$ , for all  $j \in J$ :

$$\begin{aligned} \|\Omega^* \hat{r}_j - r_j\|_2 &\leq C \sqrt{\frac{\log p_n}{h_j n N}} \\ &\leq C \left( \frac{\log p_n}{n N} \right)^{1/4}, \end{aligned}$$

since  $\min_{j \in J} h_j \geq c \sqrt{\frac{\log p_n}{n N}}$  with probability  $1 - o(p_n^{-1})$ . ■

**Lemma 42** *If we denote the rows of  $\hat{\Pi}$  from our proposed procedure as  $\{\hat{\pi}_j : j \in J\}$  and the rows of  $\Pi$  from the oracle procedure as  $\{\pi_j : j \in J\}$ , then with probability  $1 - o(p_n^{-1})$ ,*

$$\max_{j \in J} \|\hat{\pi}_j - \pi_j\|_1 \leq C \left( \frac{\log p_n}{n N} \right)^{1/4}.$$

**Proof** Recall that  $\hat{\pi}_j^\diamond \in \mathbb{R}^K$  is the unnormalized vector solving

$$\begin{pmatrix} 1 & \dots & 1 \\ \hat{v}_1^* & \dots & \hat{v}_K^* \end{pmatrix} \hat{\pi}_j^\diamond = \begin{pmatrix} 1 \\ \hat{r}_j \end{pmatrix} \iff \begin{pmatrix} 1 & \dots & 1 \\ \Omega^* \hat{v}_1^* & \dots & \Omega^* \hat{v}_K^* \end{pmatrix} \hat{\pi}_j^\diamond = \begin{pmatrix} 1 \\ \Omega^* \hat{r}_j \end{pmatrix}.$$

Therefore,

$$\hat{\pi}_j^\diamond = \hat{Q}^{-1} \begin{pmatrix} 1 \\ \Omega^* \hat{r}_j \end{pmatrix} \quad \text{where} \quad \hat{Q} := \begin{pmatrix} 1 & \dots & 1 \\ \Omega^* \hat{v}_1^* & \dots & \Omega^* \hat{v}_K^* \end{pmatrix}.$$

We also have

$$\pi_j = Q^{-1} \begin{pmatrix} 1 \\ r_j \end{pmatrix} \quad \text{where} \quad Q = \begin{pmatrix} 1 & \dots & 1 \\ v_1^* & \dots & v_K^* \end{pmatrix}$$

Consequently,

$$\|\hat{\pi}_j^\diamond - \pi_j\|_2 \leq \|\hat{Q}^{-1}\|_{\text{op}} \|\Omega^* \hat{r}_j - r_j\|_2 + \|\hat{Q}^{-1} - Q^{-1}\|_{\text{op}} \sqrt{\|r_j\|_2^2 + 1}.$$

Note that  $\max_{j \in J} \|r_j\|_2 \leq C$  since the  $r_j$ 's are in the convex hull of  $v_1^*, \dots, v_K^*$ . Also, since  $Q^T = [\text{diag}(V_1)]^{-1} V$ , we have  $\|Q^{-1}\|_{\text{op}} \leq C$  since

$$\max_{k \in [K]} V_1(k) \leq \frac{C}{\sqrt{K}} \quad \text{and} \quad \|V^{-1}\|_{\text{op}} = \sigma_1(A) \leq \sqrt{K}.$$

Now, we note that with probability  $1 - o(p_n^{-1})$ ,

$$\begin{aligned} \|\hat{Q} - Q\|_{\text{op}} &\leq \|\hat{Q} - Q\|_F \leq \sqrt{K} \max_{k \in [K]} \|\Omega^* \hat{v}_k^* - v_k^*\|_2 \\ &\leq \sqrt{K} \max_{j \in J} \|\Omega^* \hat{r}_j - r_j\|_2 \leq C \left( \frac{\log p_n}{nN} \right)^{1/4} = o(1), \end{aligned}$$

where we used Assumption 4, since in the oracle procedure the vertex hunting algorithm correctly returns  $v_1^*, \dots, v_K^*$ . Therefore,

$$\|\hat{Q}^{-1} - Q^{-1}\|_{\text{op}} = \|\hat{Q}^{-1}(Q - \hat{Q})Q^{-1}\|_{\text{op}} \leq \|\hat{Q}^{-1}\|_{\text{op}} \|\hat{Q} - Q\|_{\text{op}} \|Q\|_{\text{op}}^{-1} \leq C \left( \frac{\log p_n}{nN} \right)^{1/4}.$$

Here we note  $\sigma_K(\hat{Q}) \geq \sigma_K(Q) - \|\hat{Q} - Q\|_{\text{op}} \geq c - o(1) \geq c/2$  if  $nN$  is large enough, so  $\|\hat{Q}^{-1}\|_{\text{op}} \leq C$ . Therefore, we obtain

$$\|\hat{\pi}_j^\diamond - \pi_j\|_2 \leq C \left( \frac{\log p_n}{nN} \right)^{1/4}.$$

Now if we define  $\hat{\pi}_j = \frac{\tilde{\pi}_j^\diamond}{\|\tilde{\pi}_j^\diamond\|_1}$  where  $\tilde{\pi}_j^\diamond(k) = \max(\hat{\pi}_j^\diamond(k), 0)$ , then since  $\|\hat{\pi}_j\|_1 = \|\pi_j\|_1 = 1$ ,

$$\begin{aligned} \|\hat{\pi}_j - \pi_j\|_1 &\leq \|\hat{\pi}_j - \tilde{\pi}_j^\diamond\|_1 + \|\tilde{\pi}_j^\diamond - \pi_j\|_1 \\ &= |1 - \|\tilde{\pi}_j^\diamond\|_1| \|\hat{\pi}_j\|_1 + \|\tilde{\pi}_j^\diamond - \pi_j\|_1 \\ &= |\|\pi_j\|_1 - \|\tilde{\pi}_j^\diamond\|_1| + \|\tilde{\pi}_j^\diamond - \pi_j\|_1 \\ &\leq 2\|\pi_j - \tilde{\pi}_j^\diamond\|_1 \\ &\leq 2\|\pi_j - \hat{\pi}_j^\diamond\|_1 \leq 2\sqrt{K} \|\hat{\pi}_j^\diamond - \pi_j\|_2 \\ &\leq C \left( \frac{\log p_n}{nN} \right)^{1/4}. \end{aligned}$$

■

## Appendix E. Estimation Error of $\hat{A}$

In this section, we will additionally impose the  $\ell_q$ -sparsity assumption (15) for  $q \in (0, 1)$ .

**Lemma 43** *Under Assumption 5, if  $\beta := \frac{\alpha - \alpha}{\sigma_K(\Sigma_W)}$  and  $\tau_n := \sqrt{\frac{\log p_n}{nN}}$ , on event  $\mathcal{E}$*

$$\|A_{J^c*}\|_1 \leq \frac{K}{1-q} s(\beta\tau_n)^{1-q}. \quad (50)$$

**Remark 44** We assume from now on that  $s$  does not grow too quickly relative to  $nN$  so that the RHS of (50) is  $o(1)$ .

**Proof** On event  $\mathcal{E}$  we have  $J_- \subseteq J$ , so  $j \notin J$  implies  $M_0(j, j) \leq \alpha_- \alpha \tau_n$  where  $\tau_n := \sqrt{\frac{\log p_n}{nN}}$ . Since  $\sigma_K(\Sigma_W) h_j \leq M_0(j, j)$ ,  $j \notin J$  implies  $A_{jk} \leq h_j \leq \beta \tau_n$  for any  $k \in [K]$  on  $\mathcal{E}$ . Then with probability  $1 - o(p_n^{-1})$ , for any  $k \in [K]$ ,

$$\begin{aligned} \|A_{J^c k}\|_1 &= \sum_{j \notin J} \min(A_{jk}, \beta \tau_n) \leq \sum_{j=1}^p \min(A_{(j)k}, \beta \tau_n) \\ &\leq \sum_{j=1}^p \min(s^{1/q} j^{-1/q}, \beta \tau_n) \leq \int_0^\infty \min(s^{1/q} t^{-1/q}, \beta \tau_n) dt. \end{aligned}$$

Now, let  $t_0 := s(\beta \tau_n)^{-q}$  so that  $s^{1/q} t_0^{-1/q} = \beta \tau_n$ . Then continuing from the above display,

$$\begin{aligned} \|A_{J^c k}\|_1 &\leq t_0 \beta \tau_n + s^{1/q} \int_{t_0}^\infty t^{-1/q} dt \\ &= t_0 \beta \tau_n + \frac{q}{1-q} s^{1/q} t_0^{1-1/q} = \frac{1}{1-q} t_0 \beta \tau_n \\ &= \frac{1}{1-q} s(\beta \tau_n)^{1-q}, \end{aligned}$$

and the result follows by summing up this bound across  $k \in [K]$ . Note that the assumption  $q \in (0, 1)$  ensures the integrals above converge.  $\blacksquare$

**Lemma 45** *On event  $\mathcal{E}$ , if  $\tilde{A}_{J^*}$  is defined as in (35),*

$$\|\tilde{A}_{J^*} - A_{J^*}\|_1 = \|A_{J^c*}\|_1 \leq \frac{K}{1-q} s(\beta \tau_n)^{1-q}. \quad (51)$$

**Proof** We note that the columns of  $\tilde{A}_{J^*}$  sum up to 1, the columns of  $A$  sum up to 1, and as a result of the definition of  $\tilde{A}_{J^*}$  in (35),

$$\tilde{A}_{J^*} - A_{J^*} = \tilde{A}_{J^*} \cdot \text{diag}(\|A_{J^c 1}\|_1, \dots, \|A_{J^c K}\|_1),$$

Then  $\|\tilde{A}_{J^*} - A_{J^*}\|_1 = \|A_{J^c*}\|_1$  and our result follows from the previous lemma.  $\blacksquare$

**Theorem 46** *With probability  $1 - o(p_n^{-1})$ , for some constant  $C$  that may depend on  $K$  and  $q$ , we have*

$$\|\hat{A} - A\|_1 \leq C \left[ \left( \frac{\log p_n}{nN} \right)^{1/4} + s \left( \frac{\log p_n}{nN} \right)^{\frac{1-q}{2}} \right].$$

**Proof** Consider the unnormalized matrices

$$\hat{A}_{J^*}^\diamond := \text{diag}(\hat{\xi}_1) \hat{\Pi} \quad \text{and} \quad \tilde{A}_{J^*}^\diamond := \text{diag}(\xi_1) \Pi$$

Then with probability  $1 - o(p_n^{-1})$ , for any  $j \in J$ ,

$$\begin{aligned}
 \|(\hat{A}^\diamond - \tilde{A}^\diamond)_{j*}\|_1 &= \|\hat{\xi}_1(j)\hat{\pi}_j - \xi_1(j)\pi_j\|_1 \\
 &\leq |\hat{\xi}_1(j)|\|\hat{\pi}_j - \pi_j\|_1 + |\hat{\xi}_1(j) - \xi_1(j)|\|\pi_j\|_1 \\
 &\leq C \left[ h_j \left( \frac{\log p_n}{nN} \right)^{1/4} + \sqrt{\frac{h_j \log p_n}{nN}} \right] \\
 &\leq Ch_j \left( \frac{\log p_n}{nN} \right)^{1/4}, \tag{52}
 \end{aligned}$$

where again we note that on event  $\mathcal{E}$ ,  $h_j > \alpha_+ \alpha \sqrt{\frac{\log p_n}{nN}}$  if  $j \in J$ . Since  $\sum_{j=1}^p h_j = K$ , with probability  $1 - o(p_n^{-1})$ ,

$$\|\hat{A}_{J*}^\diamond - \tilde{A}_{J*}^\diamond\|_1 \leq C \left( \frac{\log p_n}{nN} \right)^{1/4} = o(1). \tag{53}$$

Now,  $\hat{A}_{J*}$  and  $\tilde{A}_{J*}$  are defined by normalizing the columns of  $\hat{A}_{J*}^\diamond$  and  $\tilde{A}_{J*}^\diamond$ , so we have for each  $j \in J$  and  $k \in [K]$

$$\hat{A}_{jk} = \frac{\hat{A}_{jk}^\diamond}{\|\hat{A}_{Jk}^\diamond\|_1} \quad \text{and} \quad \tilde{A}_{jk} = \frac{\tilde{A}_{jk}^\diamond}{\|\tilde{A}_{Jk}^\diamond\|_1} = \frac{A_{jk}}{\|A_{Jk}\|_1}.$$

Therefore, for each  $j \in J$  and  $k \in [K]$ ,

$$\begin{aligned}
 |\hat{A}_{jk} - \tilde{A}_{jk}| &= \left| \frac{\hat{A}_{jk}^\diamond}{\|\hat{A}_{Jk}^\diamond\|_1} - \frac{\tilde{A}_{jk}^\diamond}{\|\tilde{A}_{Jk}^\diamond\|_1} \right| \\
 &\leq \frac{|\hat{A}_{jk}^\diamond - \tilde{A}_{jk}^\diamond|}{\|\hat{A}_{Jk}^\diamond\|_1} + \tilde{A}_{jk}^\diamond \left| \frac{1}{\|\hat{A}_{Jk}^\diamond\|_1} - \frac{1}{\|\tilde{A}_{Jk}^\diamond\|_1} \right| \\
 &\leq \frac{|\hat{A}_{jk}^\diamond - \tilde{A}_{jk}^\diamond| + \tilde{A}_{jk}^\diamond \|\hat{A}_{Jk}^\diamond - \tilde{A}_{Jk}^\diamond\|_1}{\|\hat{A}_{Jk}^\diamond\|_1} \\
 &= \frac{|\hat{A}_{jk}^\diamond - \tilde{A}_{jk}^\diamond|}{\|\hat{A}_{Jk}^\diamond\|_1} + \frac{A_{jk} \|\hat{A}_{Jk}^\diamond - \tilde{A}_{Jk}^\diamond\|_1}{\|A_{Jk}\|_1 \|\hat{A}_{Jk}^\diamond\|_1}. \tag{54}
 \end{aligned}$$

Now,

$$\|A_{Jk}\|_1 = 1 - \|A_{Jc_k}\|_1 \geq 1 - \frac{1}{1-q} s(\beta\tau_n)^{1-q} \geq c$$

for some absolute constant  $c \in (0, 1)$  as  $nN$  becomes sufficiently large. Furthermore, since by definition of  $\Pi$  we have  $\hat{A}_{J*}^\diamond = \text{diag}(\xi_1)\Pi = A_{J*}\text{diag}(V_1)$ , so  $\hat{A}_{Jk}^\diamond = V_1(k)A_{Jk}$ , and  $\min_{k \in [K]} V_1(k) \geq \frac{c}{\sqrt{K}}$ , so

$$\|\tilde{A}_{Jk}^\diamond\|_1 = V_1(k)\|A_{Jk}\|_1 \geq c,$$

and thus

$$\|\hat{A}_{Jk}^\diamond\|_1 \geq \|\tilde{A}_{Jk}^\diamond\|_1 - \|\tilde{A}_{Jk}^\diamond - \hat{A}_{Jk}^\diamond\|_1 \geq c - C \left( \frac{\log p_n}{nN} \right)^{1/4} \geq c/2$$

as  $nN$  becomes sufficiently large. Hence, we have from (54), (52) and (53) that

$$|\hat{A}_{jk} - \tilde{A}_{jk}| \leq Ch_j \left( \frac{\log p_n}{nN} \right)^{1/4},$$

and so for any  $j \in J$ ,

$$\|\hat{A}_{j*} - \tilde{A}_{j*}\|_1 \leq Ch_j \left( \frac{\log p_n}{nN} \right)^{1/4},$$

which, since  $\sum_{j=1}^p h_j = K$ , implies

$$\|\hat{A}_{J*} - \tilde{A}_{J*}\|_1 \leq C \left( \frac{\log p_n}{nN} \right)^{1/4}. \quad (55)$$

We combine (50), (51) and (55) to obtain what we need to prove.  $\blacksquare$

## Appendix F. Results on Archetype Analysis (Javadi and Montanari, 2020)

To facilitate our discussion on relaxing the separability assumption, we summarize the results of Javadi and Montanari (2020) in this section.

We first introduce the notations in this paper. For a point  $u \in \mathbb{R}^d$  and a matrix  $V \in \mathbb{R}^{m \times d}$ , let

$$\mathcal{D}(u; V) := \min\{\|u - V^T \pi\|_2^2 : \pi \in \Delta^m\},$$

where

$$\Delta^m := \{x \in \mathbb{R}^m : x^T \mathbf{1}_m = 1 \text{ and } x_j \geq 0 \text{ for all } j \in [m]\}.$$

In words,  $\mathcal{D}(u; V)$  is the square of the distance between  $u$  and  $\text{conv}(V)$ , where  $\text{conv}(V)$  denotes the convex hull of the rows of  $V$ . If  $U \in \mathbb{R}^{p \times d}$  is a matrix with rows  $u_1, \dots, u_p \in \mathbb{R}^d$ , we generalize the above definition by letting

$$\mathcal{D}(U; V) := \sum_{l=1}^p \mathcal{D}(u_l; V).$$

Now, consider a factorization of the form  $X_0 = W_0 H_0$ , where the rows of  $X_0 \in \mathbb{R}^{m \times (K-1)}$  form a point cloud,  $W_0 \in \mathbb{R}^{m \times K}$  is a matrix of weights whose rows are in  $\Delta^K$ , and the rows of  $H_0 \in \mathbb{R}^{K \times (K-1)}$  are the  $K$  simplex vertices.

**Definition 47 ( $\alpha$ -uniqueness)** *We say that the point cloud  $X_0 = W_0 H_0$  satisfies uniqueness with parameter  $\alpha > 0$  (or  $\alpha$ -uniqueness) if for all  $H \in \mathbb{R}^{K \times (K-1)}$  with  $\text{conv}(X_0) \subseteq \text{conv}(H)$ , we have*

$$\mathcal{D}(H; X_0)^{1/2} \geq \mathcal{D}(H_0; X_0)^{1/2} + \alpha[\mathcal{D}(H; H_0)^{1/2} + \mathcal{D}(H_0; H)^{1/2}]. \quad (56)$$

The motivation behind this assumption is quite clear. Any  $H$  with  $\text{conv}(X_0) \subseteq \text{conv}(H)$  is a plausible explanation of the data. For  $H_0$  to be identifiable, we want  $\mathcal{D}(H; X_0) > \mathcal{D}(H_0; X_0)$  if  $H \neq H_0$ , and so (56) is a quantitative formulation of this requirement. Note that if  $X_0 = W_0 H_0$  is a separable factorization, then it always satisfies uniqueness with  $\alpha = 1$ . Indeed, whenever  $\text{conv}(H_0) = \text{conv}(X_0)$ , one has  $\mathcal{D}(H; X_0) = \mathcal{D}(H; H_0)$  and  $\mathcal{D}(H_0; X_0) = \mathcal{D}(H_0; H) = 0$ .

The vertex hunting procedure considered by Javadi and Montanari (2020) is as follows. Suppose we observe  $X$  which is a noisy version of  $X_0$ :

$$X = X_0 + Z = W_0 H_0 + Z. \quad (57)$$

Let  $x_1, \dots, x_m$  be the rows of  $X$ . We can obtain an estimator  $\hat{H}$  of  $H_0$  by solving the following optimization problem (Archetype Analysis):

$$\text{minimize } \mathcal{D}(H; X) \text{ s.t. } \mathcal{D}(x_i; H) \leq \delta^2 \text{ for all } i \in [m], \quad (58)$$

where  $\delta \geq \max_{i \in [m]} \|Z_{i*}\|_2$ . In light of Corollary 17, we want to choose  $\delta \geq C \left(\frac{\log p_n}{nN}\right)^{1/4}$  in our context, where  $C$  is the constant in (17) (replace  $X_0$  in (57) with the point cloud matrix  $R$  from our oracle procedure, and  $X$  with the point cloud matrix  $\hat{R}$  from Definition 8).

The main theoretical result of Javadi and Montanari (2020) is that their vertex hunting procedure is robust to noise in the point cloud.

**Theorem 48 (Theorem 1 of Javadi and Montanari, 2020)** *Suppose  $X_0$  satisfies the  $\alpha$ -uniqueness assumption, and  $\text{conv}(X_0)$  contains a  $(K-1)$ -dimensional ball of radius  $\mu > 0$ . Consider the vertex hunting procedure defined by (58), with  $\delta = \max_{i \in [m]} \|Z_{i*}\|_2$ . If*

$$\max_{i \in [m]} \|Z_{i*}\|_2 \leq \frac{\alpha\mu}{30K^{3/2}},$$

then

$$\|\hat{H} - H_0\|_F^2 \leq \frac{C^2 K^5}{\alpha^2} \delta^2.$$

Here, the constant  $C$  may depend on  $\mu$  and the maximum/minimum singular values of  $H_0$ , and we ignore the vertex label permutation (by redefining  $\hat{H}$  if necessary).

Using similar proof techniques as in the above theorem, we can also show the following robustness result for Archetype Analysis without using the  $\alpha$ -uniqueness condition (the proof is omitted for brevity). In (59), we do not need to assume separability (in which case one has  $\mathcal{D}(H_0; X_0) = 0$ ), but we want the distance from the vertices in  $H_0$  to the convex hull of the point cloud  $X_0$  to be no larger than  $\delta$ . Again,  $\delta \asymp \left(\frac{\log p_n}{nN}\right)^{1/4}$  when applied to our topic modeling setup.

**Theorem 49** *Using the same assumptions as in Theorem 48 except the  $\alpha$ -uniqueness condition, if  $\max_{i \in [m]} \|Z_{i*}\|_2 \leq \delta \leq \frac{\mu}{2K+2}$ , the vertex hunting procedure (58) satisfies for some constants  $C_1, C_2 > 0$ :*

$$\|\hat{H} - H_0\|_F^2 \leq C_1 \mathcal{D}(H_0; X_0) + C_2 \delta^2. \quad (59)$$

In practice, the vertex hunting procedure defined by (58) is difficult to use. When applied to real datasets, one may prefer to work with the Lagrangian form of (58):

$$\hat{H}_\lambda = \arg \min_H [\mathcal{D}(X; H) + \lambda \mathcal{D}(H; X)].$$

Algorithms to solve this non-convex optimization problem are available in Section 4 of Javadi and Montanari (2020).

## Appendix G. Synthetic Experiments: Additional Results

This appendix supplements the synthetic experiments presented in the main text. In particular, we discuss and illustrate here in more detail the Zipf law used in our experiments and provide further plots to assess the impact of additional model parameters (including the intensity of the anchor words, value of the Zipf law coefficients, etc).

### G.1 Zipf’s Law: Illustration and Comparison

Most of the synthetic experiments we use in this paper rely on the generation of documents where word frequencies follow a Zipf’s law distribution (see (31)). Figure 12a illustrates instances of such frequency distributions for a dictionary of size  $p = 10,000$  words as we vary the parameters of this distribution (namely, the values of  $\alpha_{zipf}$  and  $\beta_{zipf}$ ). Figure 12b compares the frequency heterogeneity resulting from sampling frequencies  $f_{(j)}$ ’s from a Zipf law distribution (31) to frequencies sampled from a Uniform distribution:

$$f_{(j)} \propto \text{Uniform}(0, 1).$$

These two figures illustrate in particular the fast decay in word frequencies under Zipf’s law. In fact, for the reference Zipf law ( $a_{zipf} = 1$  and  $b_{zipf} = 2.7$ ), only 10% of words have frequencies above 0.001—making the rest of the words extremely rare. As shown in Figure 12a, this decay increases rapidly as the parameter  $\alpha_{zipf}$  increases. By comparison, for the uniform distribution, all word frequencies are of the same order of magnitude. One of the main assumptions in this paper is the weak sparsity of the row sums of the topic matrix  $A$ , described in Assumption 5 of the main text:  $\max_{k \in [K]} \max_{j \in [p]} j A_{(j)k}^q \leq s$ , which, as argued in the main text, is best reflected by Zipf’s law.

### G.2 Synthetic Experiment: Uniform Distribution of Non-Anchor Words

In this paragraph, we describe the results of our synthetic experiments using a Uniform distribution for the generation of non-anchor words.

*Data Generation Mechanism.* In this setting, the data is generated as follows:

$$\begin{aligned} \forall i \in \{1, \dots, n\}, \quad W_i &\sim \text{Dirichlet}(\mathbf{1}_K) \\ \forall k \in [K] \text{ and } \forall j \in \{1, \dots, 5\}, \quad A_{j+5(k-1),k} &= \delta \\ \forall k \in [K] \text{ and } \forall j \in \{5K + 1, \dots, p\}, \quad A_{jk} &\sim \text{Uniform}(0, 1) \end{aligned} \tag{60}$$

This is thus an identical generation setup as in the main text, but using a Uniform distribution rather than a Zipf law to generate the frequencies. As noted in the previous section, this corresponds to a setting in which all frequencies have roughly the same amplitude (homogeneous frequencies), which does not agree (a) with our Assumption 5 (weak-sparsity); and (b) observations or realistic models of word frequencies in real documents (Corral et al., 2015). Nonetheless, we run our simulation pipeline in this setting and report the results in Figure 13a. Results are averaged over 50 experiments, using  $K = 5$ ,  $\alpha_{dirichlet} = 1$ , and 5 anchor words with intensity  $\delta_{anchor} = 0.001$ .

*Analysis of the results.* We note that Topic-SCORE (Tracy Ke and Wang, 2024) and our method perform similarly in these conditions. As shown in Figure 13b, our method does

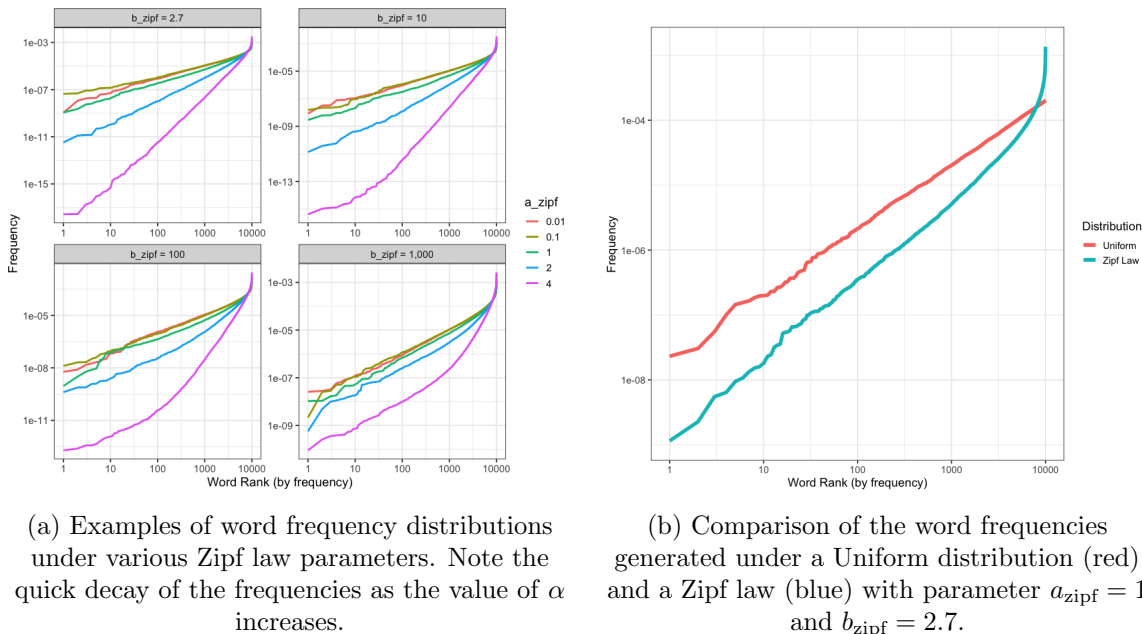


Figure 12: Comparison of word frequencies in a dictionary of size  $p = 10,000$  under various generation mechanisms. The  $x$  and  $y$  axis are presented on a log-scale.

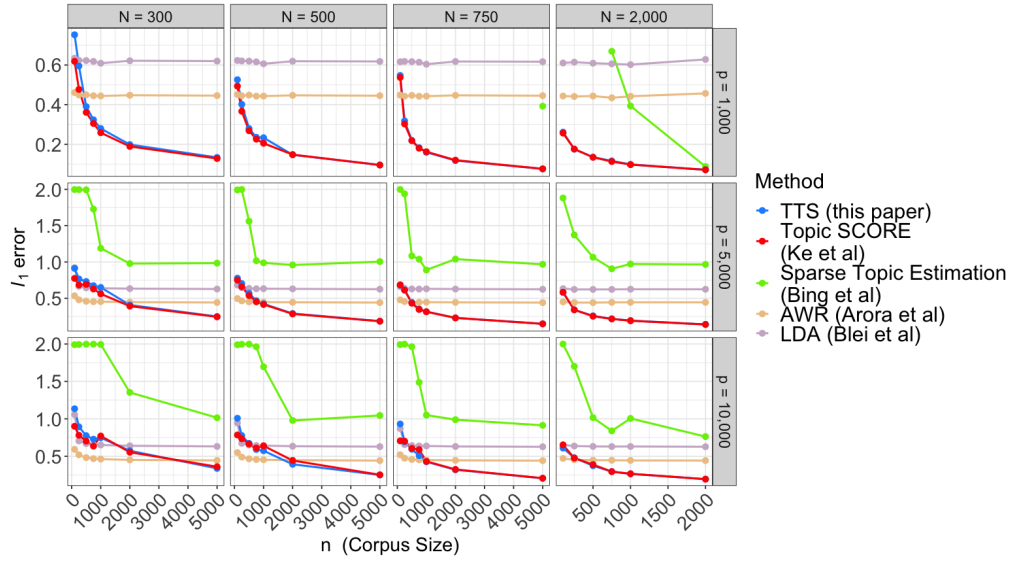
not threshold much in this regime — as is probably to be expected, since the frequencies of all the words are of the same order. As highlighted in the main text, this suggests that, whilst designed to leverage the weak sparsity of the matrix  $A$ , our method is robust to various frequency regimes. Our method is, in fact, able to perform satisfactorily even in homogeneous (and less sparse) settings. Interestingly, in these regimes, AWR becomes an alternative approach (in particular, as the number of words  $p$  increases).

### G.3 Varying Additional Parameters


In this subsection, we evaluate the effect of other parameters in the data generation mechanism detailed in Section 3. In particular, we assess here (a) the effect of the anchor word frequency  $\delta_{anchor}$  (Figure 15); and (b) the impact of the Zipf law coefficient  $a_{zipf}$ .

*Discussion: Effect of the frequency of the anchor words.* As observed in Figure 15, the frequency of the anchor words does not appear to have a great impact on the results of the SCORE-based methods. Increasing the frequency of anchor words does seem to improve the performance of the Sparse Topic Estimation Method of Bing et al. (2020a) and LDA (Blei et al., 2003).

*Discussion: Effect of the parameter  $a_{zipf}$ .* As observed in Figure 16, our method offers significant improvement over others as the word frequency heterogeneity increases. In fact, both the Sparse Topic Estimation of Bing et al. (2020a) and our method’s  $\ell_1$  errors decrease significantly as the heterogeneity of the word frequency (high  $a_{zipf}$ ) increases. This is reassuring, since these are the only two of the five methods that actively leverage sparsity of the topic matrix to improve estimation.



(a) Median  $\ell_1$  error  $\mathcal{L}_1(\hat{A}, A) = \min_{\Pi \in \mathcal{P}} \frac{1}{K} \|\hat{A}\Pi - A\|_1$  for the different methods. Rows indicate the size of the vocabulary  $p$ , while columns indicate the document length  $N$ .



(b) Percentage of thresholded words as a function of corpus length, dictionary size and document length.

Figure 13: Comparison of the different methods under the uniform frequency generation mechanism detailed in (60). Results are shown for a fixed number of topics  $K = 5$ , averaged over 50 independent trials, and plotted as a function of corpus size  $n$ . For low values of  $p$ , the “Sparse Topic Estimation” method of Bing et al. (2020b) does not appear because the number of topics it estimated was smaller than the true value  $K = 5$ ; therefore we were unable to evaluate it and do not report it in the plots.

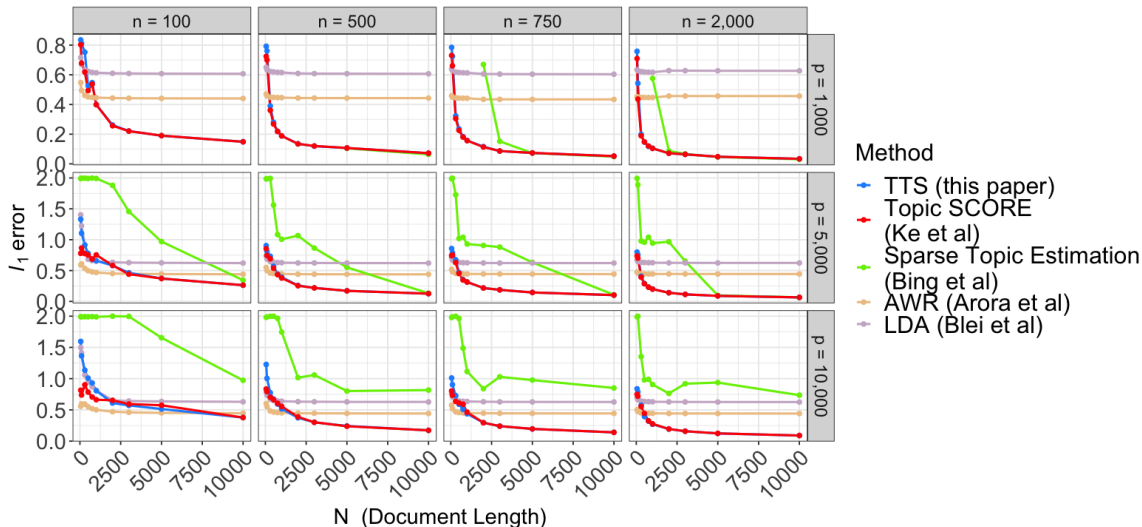


Figure 14: Same experiments as depicted in Figure 13. Results are displayed as a function of document length  $N$ .

### G.4 Semi-Synthetic Experiments

*Data Generation Mechanism.* To ensure the realism of our experiments, we now propose to evaluate the performance of our method on a semi-synthetic dataset. To this end, we consider the Associated Press Dataset, a dataset of 2,246 documents with 10,473 terms. We fit an LDA model on the induced document-term matrix and vary the number of topics to obtain an underlying  $A_0$  and  $W_0$ . For each experiment, we select a subset of documents by sampling  $i \in \{1, \dots, n\}$  without replacement, and concatenate the selected columns of  $W_i$  to construct the topic-to-document matrix  $W$ . Having obtained  $A$  and  $W$ , we then sample each column of the matrix  $D$  from a multinomial distribution as before. We note that: (a) this data generation mechanism matches the one assumed by LDA, so it comes as no surprise that LDA performs better in this set of experiments than in the previous subsection; and (b) this data generation mechanism does not impose as much sparsity as the Zipf law. Indeed, the  $A$  matrices imputed by LDA, which serve as ground truth here, are less heterogeneous than in our other synthetic experiments. Nonetheless, our method improves slightly over Topic-SCORE for corpora of moderate size, across most data regimes. We also note that both SCORE methods outperform AWR across all data regimes. Interestingly, LDA exhibits impressive accuracy in this particular regime.




Figure 15: Effect of the anchor word frequency  $\delta_{anchor}$  on the results. Here the number of topics is fixed to  $K = 3$  and the dictionary has size  $p = 5,000$ . Consistently with the experiments presented in the main text, each topic has 5 anchor words.




Figure 16: Effect of the parameter  $a_{zipf}$  (see (31)) on the results. The number of topics is fixed to  $K = 5$  and the dictionary has size  $p = 10,000$ . Consistently with the experiments presented in the main text, each topic has 5 anchor words with frequency  $\delta = 0.001$ .




Figure 17: Semi-synthetic experiment: Comparison of the  $\ell_1$  error in the reconstruction of  $A$  averaged over 50 independent experiments. The estimation error is plotted as a function of corpus size  $n$  (on the x-axis), as a function of document length  $N$  (one per column) and number of topics  $K$  (one per row).

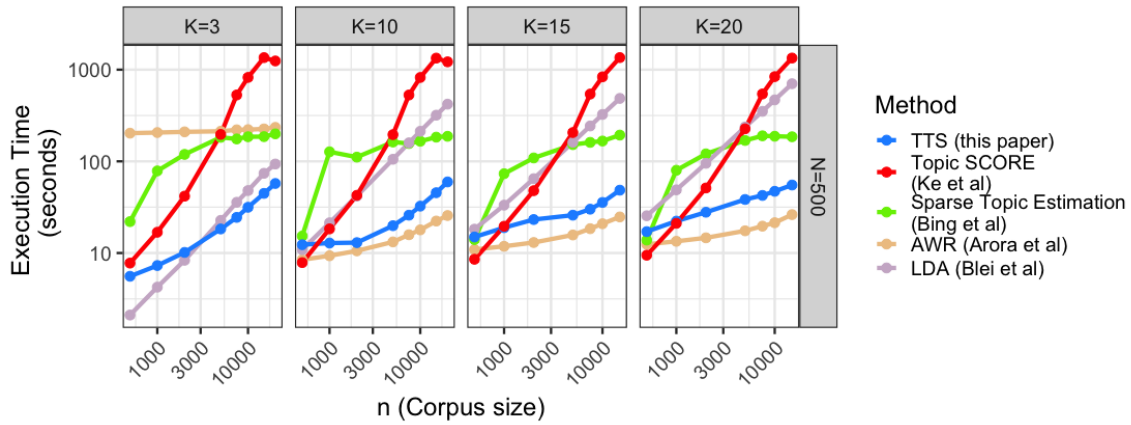


Figure 18: Median computational time for each method as a function of corpus size  $n$  over 50 independent synthetic experiments.

*Computational Speed.*

Finally, we report here the computational speed measured in our experiments. All experiments here were performed on the University of Chicago cluster. We observe a substantial increase in the time required by Topic-SCORE as the number of documents  $n$  increases.

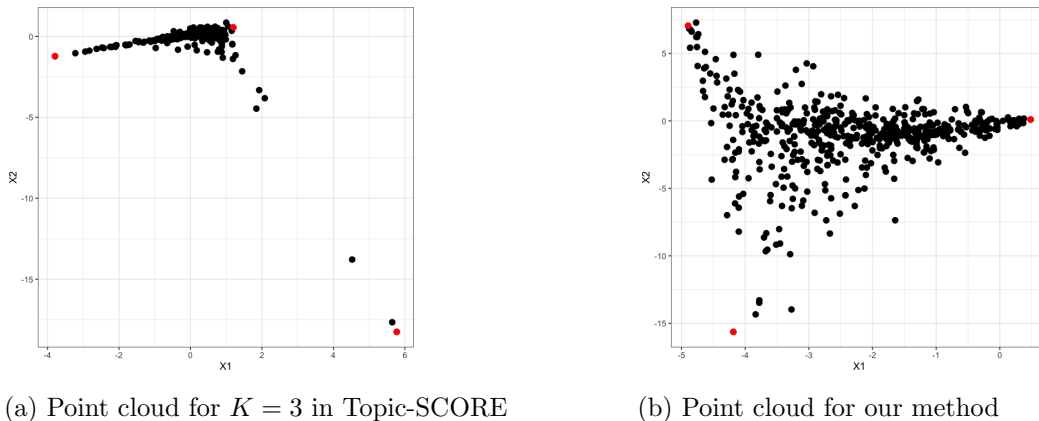


Figure 19: Comparison of the point clouds  $R$  obtained by our method (right) and Topic-SCORE (left). Vertices are colored red. Note in particular how the point cloud on the left is stretched by a few outliers.

## Appendix H. Real-World Experiments: Additional Results

In this section, we provide additional visualization plots for microbiome data analysis to evaluate the estimated topic representatives obtained by LDA, Topic-SCORE, and our method. We focus on two microbiome datasets: the colon dataset of Yachida et al. (2019), as before, with low  $p$ , low  $n$ , and high  $N$ ; and the vaginal microbiome data of Callahan et al. (2017), with moderate  $p$ , moderate  $n$ , and high  $N$ .

### H.1 Microbiome Dataset from Yachida et al. (2019) (Low $p$ , Low $n$ , High $N$ )

As depicted in Figure 11, our approach consistently demonstrates a higher average topic resolution compared to Topic-SCORE. To understand the reasons for the gap in performance between Topic-SCORE (Tracy Ke and Wang, 2024) and our method, we compare their respective point clouds in Figure 19. The left plot in Figure 19a shows the adverse impact of pre-SVD normalization on Topic-SCORE’s point cloud. It is evident that the outliers from small  $h_j$  skew the point cloud, thereby compromising the quality of its subsequent vertex hunting. In comparison, our point cloud in Figure 19b shows an equilateral triangle formed by the three estimated vertices. Our method combines the thresholding phase with spectral decomposition and guarantees a higher signal-to-noise ratio.






Figure 20: Topic resolution (measured by the average cosine similarity between halves of the data) of our method (in blue), Topic-SCORE (red), and LDA (purple).

## H.2 Microbiome Dataset from Callahan et al. (2017) (High $N$ , Mid $p$ , Mid $n$ )


We also reanalyze the dataset of Callahan et al. (2017), which serves as an example in Fukuyama et al. (2023) to exemplify their topic refinement procedure. This dataset comprises ASV counts for 2,699 different bacterial species from 2,179 longitudinal samples collected throughout pregnancy in 135 individuals (Callahan et al., 2017). In this case, the average sample length is around  $N = 157,500$ . Based on the refinement results of LDA, Fukuyama et al. (2023) conclude that the topic analysis should be done using  $K = 7$  topics, or with up to  $K = 12$  if one allows for the possibility of spurious topics. We thus fit up to 12 topics and plot the average resolution (Figure 20) and refinement of the methods in Figure 21. We find that our method compares favorably to Topic-SCORE in terms of its average resolution and is broadly similar to LDA. For a low number of topics ( $K \leq 7$ ), our method seems even preferable to LDA in terms of topic resolution, achieving better resolution at much greater speed. The topic coherence seems slightly higher for LDA at  $K = 7$ , the recommended choice of  $K$  by Fukuyama et al. (2023), than for the other methods. For  $K \geq 8$ , our method seems to yield topics that are more recombined from one level of the hierarchy to the next than those observed in LDA. This concurs with the choice of  $K = 7$  by the authors, but seems to highlight the fact that for high values of  $K$ , in datasets of moderate sizes, LDA appears a preferable choice. However, we note that Topic-SCORE seems to (1) exhibit more recombination of topics early on in the hierarchy (see bottom topics in Figure 21a); and (2) put very little mass on topics as  $K$  increases (indicated by the small size of the rectangles in Figure 21a): as  $K$  increases, most of the dataset’s mass is distributed along roughly 5 or 6 topics. This could in particular mean that the method does not really identify more than 6 topics.



(a) Coherence between topics for Topic-SCORE.



(b) Coherence between topics for TTS (our method).



(c) Coherence between topics for LDA.

Figure 21: Topic coherence and refinement (as computed by the method of Fukuyama et al., 2023) for the different methods in the vaginal microbiome dataset of Callahan et al. (2017). Topics are colored by coherence.

## References

- Seung C. Ahn and Alex R. Horenstein. Eigenvalue ratio test for the number of factors. *Econometrica*, 81(3):1203–1227, 2013. ISSN 0012-9682,1468-0262. doi: 10.3982/ECTA8968. URL <https://doi.org/10.3982/ECTA8968>.
- Lucia Alessi, Matteo Barigozzi, and Marco Capasso. Improved penalization for determining the number of factors in approximate factor models. *Statist. Probab. Lett.*, 80(23-24):1806–1813, 2010. ISSN 0167-7152,1879-2103. doi: 10.1016/j.spl.2010.08.005. URL <https://doi.org/10.1016/j.spl.2010.08.005>.
- Mário César Ugulino Araújo, Teresa Cristina Bezerra Saldanha, Roberto Kawakami Harrop Galvao, Takashi Yoneyama, Henrique Caldas Chame, and Valeria Visani. The successive projections algorithm for variable selection in spectroscopic multicomponent analysis. *Chemometrics and Intelligent Laboratory Systems*, 57(2):65–73, 2001.
- Sanjeev Arora, Rong Ge, and Ankur Moitra. Learning topic models—going beyond SVD. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science—FOCS 2012*, pages 1–10. IEEE Computer Soc., Los Alamitos, CA, 2012. ISBN 978-0-7685-4874-6.
- Jushan Bai and Serena Ng. Determining the number of factors in approximate factor models. *Econometrica*, 70(1):191–221, 2002. ISSN 0012-9682,1468-0262. doi: 10.1111/1468-0262.00273. URL <https://doi.org/10.1111/1468-0262.00273>.
- Manuele Bicego, Pietro Lovato, Alessandro Perina, Marianna Fasoli, Massimo Delledonne, Mario Pezzotti, Annalisa Polverari, and Vittorio Murino. Investigating topic models’ capabilities in expression microarray data classification. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 9(6):1831–1836, 2012.
- Xin Bing, Florentina Bunea, and Marten Wegkamp. A fast algorithm with minimax optimal guarantees for topic models with an unknown number of topics. *Bernoulli*, 26(3):1765–1796, 2020a. ISSN 1350-7265,1573-9759. doi: 10.3150/19-BEJ1166. URL <https://doi.org/10.3150/19-BEJ1166>.
- Xin Bing, Florentina Bunea, and Marten Wegkamp. Optimal estimation of sparse topic models. *J. Mach. Learn. Res.*, 21:Paper No. 177, 45, 2020b. ISSN 1532-4435,1533-7928.
- David M Blei, Andrew Y Ng, and Michael I Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022, 2003.
- T. Tony Cai and Harrison H. Zhou. Optimal rates of convergence for sparse covariance matrix estimation. *Ann. Statist.*, 40(5):2389–2420, 2012. ISSN 0090-5364,2168-8966. doi: 10.1214/12-AOS998. URL <https://doi.org/10.1214/12-AOS998>.
- Benjamin J Callahan, Daniel B DiGiulio, Daniela S Aliaga Goltsman, Christine L Sun, Elizabeth K Costello, Pratheepa Jegannathan, Joseph R Biggio, Ronald J Wong, Maurice L Druzin, Gary M Shaw, et al. Replication and refinement of a vaginal microbial signature of preterm birth in two racially distinct cohorts of us women. *Proceedings of the National Academy of Sciences*, 114(37):9966–9971, 2017.

- Zhenghao Chen, Ilya Soifer, Hugo Hilton, Leeat Keren, and Vladimir Jojic. Modeling multiplexed images with spatial-lda reveals novel tissue microenvironments. *Journal of Computational Biology*, 27(8):1204–1218, 2020.
- Victor Chernozhukov, Christian Hansen, Yuan Liao, and Yinchu Zhu. Inference for low-rank models. *Ann. Statist.*, 51(3):1309–1330, 2023. ISSN 0090-5364,2168-8966. doi: 10.1214/23-aos2293. URL <https://doi.org/10.1214/23-aos2293>.
- Jungjun Choi, Hyukjun Kwon, and Yuan Liao. Inference for low-rank completion without sample splitting with application to treatment effect estimation. *J. Econometrics*, 240(1):Paper No. 105682, 23, 2024. ISSN 0304-4076,1872-6895. doi: 10.1016/j.jeconom.2024.105682. URL <https://doi.org/10.1016/j.jeconom.2024.105682>.
- Álvaro Corral, Gemma Boleda, and Ramon Ferrer-i Cancho. Zipf’s law for word frequencies: Word forms versus lemmas in long texts. *PLOS ONE*, 10(7):e0129031, 2015.
- Stephan A Curiskis, Barry Drake, Thomas R Osborn, and Paul J Kennedy. An evaluation of document clustering and topic modelling in two online social networks: Twitter and reddit. *Information Processing & Management*, 57(2):102034, 2020.
- David Donoho and Victoria Stodden. When does non-negative matrix factorization give a correct decomposition into parts? *Advances in Neural Information Processing Systems*, 16, 2003.
- Ingo Feinerer, Kurt Hornik, and Maintainer Ingo Feinerer. Package ‘tm’. *Corpus*, 10(1), 2015.
- Julia Fukuyama, Kris Sankaran, and Laura Symul. Multiscale analysis of count data through topic alignment. *Biostatistics*, 24(4):1045–1065, 2023. ISSN 1465-4644,1468-4357. doi: 10.1093/biostatistics/kxac018. URL <https://doi.org/10.1093/biostatistics/kxac018>.
- Matan Gavish and David L. Donoho. The optimal hard threshold for singular values is  $4/\sqrt{3}$ . *IEEE Trans. Inform. Theory*, 60(8):5040–5053, 2014. ISSN 0018-9448,1557-9654. doi: 10.1109/TIT.2014.2323359. URL <https://doi.org/10.1109/TIT.2014.2323359>.
- Rong Ge and James Zou. Intersecting faces: Non-negative matrix factorization with new guarantees. In *International Conference on Machine Learning*, pages 2295–2303. PMLR, 2015.
- Nicolas Gillis and Stephen A Vavasis. Fast and robust recursive algorithms for separable nonnegative matrix factorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(4):698–714, 2013.
- Yury Goltsev, Nikolay Samusik, Julia Kennedy-Darling, Salil Bhate, Matthew Hale, Gustavo Vazquez, Sarah Black, and Garry P Nolan. Deep profiling of mouse splenic architecture with CODEX multiplexed imaging. *Cell*, 174(4):968–981, 2018.
- Anne Greenbaum, Ren-Cang Li, and Michael L. Overton. First-order perturbation theory for eigenvalues and eigenvectors. *SIAM Rev.*, 62(2):463–482, 2020. ISSN 1095-7200,0036-1445. doi: 10.1137/19M124784X. URL <https://doi.org/10.1137/19M124784X>.

- Thomas Hofmann. Probabilistic latent semantic indexing. In *Proceedings of the 22nd annual international ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 50–57, 1999.
- Roger A. Horn and Charles R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, second edition, 2013. ISBN 978-0-521-54823-6.
- Hamid Javadi and Andrea Montanari. Nonnegative matrix factorization via archetypal analysis. *J. Amer. Statist. Assoc.*, 115(530):896–907, 2020. ISSN 0162-1459,1537-274X. doi: 10.1080/01621459.2019.1594832. URL <https://doi.org/10.1080/01621459.2019.1594832>.
- Jiashun Jin. Fast community detection by SCORE. *Ann. Statist.*, 43(1):57–89, 2015. ISSN 0090-5364,2168-8966. doi: 10.1214/14-AOS1265. URL <https://doi.org/10.1214/14-AOS1265>.
- Jiashun Jin, Zheng Tracy Ke, and Shengming Luo. Estimating network memberships by simplex vertex hunting. *arXiv preprint arXiv:1708.07852*, 12, 2017.
- Olga Klopp, Maxim Panov, Suzanne Sigalla, and Alexandre B. Tsybakov. Assigning topics to documents by successive projections. *Ann. Statist.*, 51(5):1989–2014, 2023. ISSN 0090-5364,2168-8966. doi: 10.1214/23-aos2316. URL <https://doi.org/10.1214/23-aos2316>.
- Li-Jia Li, Chong Wang, Yongwhan Lim, David M Blei, and Li Fei-Fei. Building and using a semantivisual image hierarchy. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3336–3343. IEEE, 2010.
- Zongming Ma. Sparse principal component analysis and iterative thresholding. *Ann. Statist.*, 41(2):772–801, 2013. ISSN 0090-5364,2168-8966. doi: 10.1214/13-AOS1097. URL <https://doi.org/10.1214/13-AOS1097>.
- Steven T Piantadosi. Zipf’s word frequency law in natural language: A critical review and future directions. *Psychonomic Bulletin & Review*, 21:1112–1130, 2014.
- Jonathan K Pritchard, Matthew Stephens, and Peter Donnelly. Inference of population structure using multilocus genotype data. *Genetics*, 155(2):945–959, 2000.
- Walter Rudin. *Principles of mathematical analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill Book Co., New York-Auckland-Düsseldorf, third edition, 1976.
- Kris Sankaran and Susan P. Holmes. Latent variable modeling for the microbiome. *Biostatistics*, 20(4):599–614, 2019. ISSN 1465-4644,1468-4357. doi: 10.1093/biostatistics/kxy018. URL <https://doi.org/10.1093/biostatistics/kxy018>.
- Ville Satopaa, Jeannie Albrecht, David Irwin, and Barath Raghavan. Finding a “kneedle” in a haystack: detecting knee points in system behavior. In *2011 31st International Conference on Distributed Computing Systems Workshops*, pages 166–171. IEEE, 2011.

- Zheng Tracy Ke and Minzhe Wang. Using SVD for topic modeling. *J. Amer. Statist. Assoc.*, 119(545):434–449, 2024. ISSN 0162-1459,1537-274X. doi: 10.1080/01621459.2022.2123813. URL <https://doi.org/10.1080/01621459.2022.2123813>.
- Ruijia Wu, Linjun Zhang, and T. Tony Cai. Sparse topic modeling: computational efficiency, near-optimal algorithms, and statistical inference. *J. Amer. Statist. Assoc.*, 118(543): 1849–1861, 2023. ISSN 0162-1459,1537-274X. doi: 10.1080/01621459.2021.2018329. URL <https://doi.org/10.1080/01621459.2021.2018329>.
- Shinichi Yachida, Sayaka Mizutani, Hirotsugu Shiroma, Satoshi Shiba, Takeshi Nakajima, Taku Sakamoto, Hikaru Watanabe, Keigo Masuda, Yuichiro Nishimoto, Masaru Kubo, et al. Metagenomic and metabolomic analyses reveal distinct stage-specific phenotypes of the gut microbiota in colorectal cancer. *Nature Medicine*, 25(6):968–976, 2019.
- Xiaolin Zheng, Weifeng Ding, Zhen Lin, and Chaochao Chen. Topic tensor factorization for recommender system. *Information Sciences*, 372:276–293, 2016.
- George Kingsley Zipf. *The Psycho-biology of Language: An Introduction to Dynamic Philology*. Houghton-Mifflin, 1936.