
Finding k in Latent k -Polytope

Chiranjib Bhattacharyya^{*1} Ravi Kannan^{*2} Amit Kumar^{*3}

Abstract

The recently introduced Latent k -Polytope(LkP) encompasses several stochastic Mixed Membership models including Topic Models. The problem of finding k , the number of extreme points of LkP, is a fundamental challenge and includes several important open problems such as determination of number of components in Ad-mixtures. This paper addresses this challenge by introducing Interpolative Convex Rank(ICR) of a matrix defined as the minimum number of its columns whose convex hull is within Hausdorff distance ε of the convex hull of all columns. The first important contribution of this paper is to show that under *standard assumptions* k equals the ICR of a *subset smoothed data matrix* defined from Data generated from an LkP. The second important contribution of the paper is a polynomial time algorithm for finding k under standard assumptions. An immediate corollary is the first polynomial time algorithm for finding the *inner dimension* in Non-negative matrix factorisation(NMF) with assumptions which are qualitatively different than existing ones such as *Separability*.

1. Introduction

Latent variable models have found large number of applications in the real world. Such models specify a generative process between the observed variables and a set of underlying un-observed variables, often called Latent variables. Examples of latent variable models include Clustering like k -means, Finite Mixture models, Topic models (LDA), and Stochastic block models. In such models learning the parameters of the generative process is often intractable and remains an active area of research.

¹Department of Computer Science and Automation, IISc Bangalore, India ²Microsoft Research India Lab., Bangalore, India ³Department of Computer Science and Engineering, IIT Delhi, India. Correspondence to: Chiranjib Bhattacharyya <chiru@iisc.ac.in>.

(Bhattacharyya & Kannan, 2020) suggested Latent k polytope(LkP), an unifying model which includes all of the above mentioned instances of latent variable models as special cases. LkP defines a latent polytope K on k vertices and data is generated by first picking *latent points* from this polytope, denoted by columns of a matrix \mathbf{P} , and then perturbing them in an adversarial manner to produce data-points. The perturbations can be quite large in magnitude, and the data points can lie quite far from the actual polytope K . The problem of recovering (vertices of) K from data generated in this manner is in general hard, unless we make assumptions about K and the data matrix \mathbf{A} , each of whose columns represents one data point. They crystallized a set of assumptions, which will be referred in the sequel as *standard assumptions*, and showed that the vertices of the polytope K can be provably recovered to an ε approximation from a data Matrix \mathbf{A} satisfying the assumptions.

The algorithm generally applies to all special cases and results in Model estimation with guaranteed approximation from a finite sample of data. However, as is usual practice in Latent Variable models the value of k is assumed known. For example, most algorithms for k -Means or LDA with k topics would require the value of k .

In this paper we rely on the geometry of the problem to find k from the data Matrix \mathbf{A} . Straightforward attempts such as examining the rank etc of \mathbf{A} are fruitless because most of the points in \mathbf{A} can be very far from the actual polytope K . We use two key insights here: (i) Suppose first we have access to the latent points matrix \mathbf{P} instead of \mathbf{A} . In this case, the convex hull of \mathbf{P} , $\text{CH}(\mathbf{P})$, would be a close approximation of K . Still it is not clear how to identify k from the $\text{CH}(\mathbf{P})$. (ii) Such a strategy is still not practical as \mathbf{P} is not available. However the eigen structure of \mathbf{A} can provide important clues.

Contributions: This paper addresses some of these challenges and a summary of contributions are listed below.

- The paper introduces the notion of Interpolative Convex Rank(ICR) of a matrix, and shows that $k = \text{ICR}$ of a *subset smoothed data matrix* where k is the number of vertices in LkP (see details in Theorem 1). The notion of ICR should be of independent interest.
- The paper introduces new techniques based on the hy-

perplane separator theorem for proving lower bounds on the ICR of a matrix (see details in Section 3.4).

- Under standard assumptions, the paper gives a polynomial time algorithm for finding the correct value of the number of vertices of the polytope K (see Theorem 18). We do this by using the familiar singular value thresholding method, but from a different perspective than commonly used in numerical analysis.
- In this paper we show that Non-negative Matrix factorization(NMF) is a special case of Latent- k -polytope. In Theorem 20 we show that the algorithm recovers the correct inner dimension for NMF under standard assumptions.

2. Preliminaries and Problem Definition

Notation: For any natural number n we will denote $[n] = \{1, \dots, n\}$. The points will lie in the d -dimensional Euclidean space \mathbb{R}^d . For a matrix \mathbf{B} and index ℓ , we shall use $B_{\cdot, \ell}$ to denote the ℓ^{th} column of \mathbf{B} . Given a subset S of columns of \mathbf{B} , the notation $\mathbf{B}_{\cdot, S}$ refers to $\frac{1}{|S|} \sum_{\ell \in S} B_{\cdot, \ell}$. The spectral norm of \mathbf{B} will be denoted by $\|\mathbf{B}\|$. The 2-norm (i.e., Euclidean norm) of a vector v is denoted by $|v|$. The r^{th} singular value of a matrix \mathbf{B} is denoted by $s_r(\mathbf{B})$. The convex hull of a set of points $\mathbf{x}_i \in \mathbb{R}^d, i \in [m]$, is denoted by $\text{CH}(\{\mathbf{x}_1, \dots, \mathbf{x}_m\}) = \{\sum_{i=1}^m \alpha_i \mathbf{x}_i | 0 \leq \alpha_i, \sum_{i=1}^m \alpha_i = 1\}$. We abuse notation and use $\text{CH}(\mathbf{B})$, where \mathbf{B} is a matrix, to denote the convex hull of the columns of \mathbf{B} . An extreme point of a convex set, \mathbf{C} , is any point, $\mathbf{x} \in \mathbf{C}$ such that it cannot be expressed as $\mathbf{x} = \beta \mathbf{x}_1 + (1 - \beta) \mathbf{x}_2, 0 \leq \beta \leq 1, \mathbf{x}_1 \neq \mathbf{x}_2, \mathbf{x}_1, \mathbf{x}_2 \in \mathbf{C}$.

Definition 1. Given two sets $S, T \subseteq \mathbb{R}^d$, the Hausdorff distance $\text{dist}(S, T)$ between them is defined as

$$\text{dist}(S, T) := \max \left(\max_{x \in S} d(x, T), \max_{y \in T} d(y, S) \right),$$

where $d(z, A)$, for a point z and a set $A \subseteq \mathbb{R}^d$, is defined as $\min_{t \in A} |z - t|$.

2.1. Review of LkP and Problem statement

In this subsection we introduce relevant definitions and present the problem statement.

2.1.1. REVIEW OF LATENT-K-POLYTOPE

In this subsection we briefly review the Latent- k -polytope introduced in (Bhattacharyya & Kannan, 2020).

Definition 2. Let the columns of $\mathbf{M} \in \mathbb{R}^{d \times k}$ be the extreme points of the Latent k Polytope, $K = \text{CH}(\mathbf{M})$. The data matrix $\mathbf{A} \in \mathbb{R}^{d \times n}$ is said to be generated from LkP with parameters $\mathbf{M}, \sigma, \sigma > 0$, if there exists a $d \times n$ matrix \mathbf{P} , whose columns lie in K , such that $\sigma = \frac{\|\mathbf{A} - \mathbf{P}\|}{\sqrt{n}}$.

For each $j \in [n]$, $A_{\cdot, j}$, a column of the data-matrix as the perturbation of the latent data point $P_{\cdot, j}$. The parameter σ , referred in the sequel as *perturbation parameter*, can be thought of as a measure of the maximum mean squared perturbation in any direction. In previous work (Bhattacharyya & Kannan, 2020) the problem of learning \mathbf{M} from \mathbf{A} was addressed. In general this is an intractable problem but under certain standard assumptions stated in the next subsection, a polynomial time algorithm was given for recovering \mathbf{M} to an additive error. The dimension of \mathbf{M} , i.e. the value of k , was assumed to be known.

2.1.2. STANDARD ASSUMPTIONS

(Bhattacharyya & Kannan, 2020) defined a set of assumptions on LkP which will be referred in the sequel as *Standard Assumptions*. A LkP model with parameters \mathbf{M}, σ , and the data matrix \mathbf{A} is said to satisfy the standard assumptions if the following are true.

- **Separation:** Every vertex of K is far from the affine or the convex hull of the remaining vertices of K .
- **Proximity:** There is a parameter $\delta \in (0, 1)$ such that for every vertex of K , there are at least δn points (columns of \mathbf{A}) whose corresponding latent points (columns in \mathbf{P}) are close to the vertex.
- **Spectrally Bounded Perturbations:** The parameter $\sigma (= \|\mathbf{A} - \mathbf{P}\|/\sqrt{n})$ is much smaller than the size of K .

The first assumption is on the model, \mathbf{M} , the second is on observed data, \mathbf{A} , and last assumption is on the perturbation parameter, σ . In the sequel these assumptions will be stated more precisely. We prove that the set of standard assumptions (see Theorem 1) imply that there are k subsets of data each of cardinality δn whose averages nearly contain in their convex hull the averages of all δn sized subsets of data. This observation will be leveraged to give a data driven characterization of k . It is to be noted that the oft used notion of *separability* in Non-negative matrix Factorization (see (Gillis & Luce, 2018)) assumes that there are k actual data points whose hull nearly contains all data points. This is restrictive and does not hold even for simple cases of LkP.

2.1.3. PROBLEM DEFINITION

The objective of the paper is to address the problem of finding k under the above mentioned assumptions.

Definition 3. Given \mathbf{A} generated from LkP and δ , find the number of extreme points of K .

Note that σ, k, \mathbf{P} and \mathbf{M} are not given as input parameters. Only data, \mathbf{A} , is available. In (Bhattacharyya & Kannan,

2020) the emphasis was on computing \mathbf{M} given k . In contrast this paper attempts to address the problem of finding k from \mathbf{A} under standard assumptions.

2.2. Non-negative Matrix factorization(NMF)

LkP is an interesting geometric alternative to several Unsupervised learning paradigms such as Topic Models, Clustering and Community Detection (Bhattacharyya & Kannan, 2020). An important aim of this paper is to explore NMF as a special case of LkP.

Over the last two decades NMF have attracted significant research interest. In this section we do not review the vast literature on NMF but point the reader to many readable surveys(e.g. (Gillis, 2014)). However, to aid further discussion the definition of NMF is briefly reviewed.

Exact NMF seeks a factorization of the data matrix \mathbf{A} , $\mathbf{A}_{d \times n} = \mathbf{M}_{d \times k} \mathbf{W}_{k \times n}$ where n is the number of datapoints, d is the ambient data dimension and all entries of \mathbf{M} , \mathbf{W} , \mathbf{A} are assumed to be positive. Each column of \mathbf{M} represents a feature and each column of \mathbf{W} denotes the corresponding weights on the individual features. The parameter k is often called *inner-dimension* (Vavasis, 2009). It is possible that Exact NMF may not exist and bulk of the research in NMF has been directed towards understanding the case of approximate NMF defined as follows

$$\mathbf{A}_{d \times n} = \mathbf{M}_{d \times k} \mathbf{W}_{k \times n} + \mathbf{N}. \quad (\text{Approx-NMF})$$

Here again \mathbf{W} , \mathbf{M} are assumed to have *non-negative* entries. Each column of \mathbf{N} can be thought of as perturbation or noise.

The Inexact NMF is a special case of Latent- k -Polytope with the requirement that each column of \mathbf{W} should sum to 1. Each datapoint, $\mathbf{A}_{\cdot, j}$ is a perturbation of a point $\mathbf{P}_{\cdot, j} = \sum_{l=1}^r w_{jl} \mathbf{M}_{\cdot, l}$, lying on the polytope $\mathbf{K} = \text{CH}(\mathbf{M})$. The parameter σ depends on the norm of the matrix \mathbf{N} . It is straightforward to see that if we are give data-matrix \mathbf{A} , which satisfies the standard assumptions, and the inner dimension r , the algorithm mentioned in (Bhattacharyya & Kannan, 2020) directly applies and one can recover the features.

Our goal in this paper is to recover the *inner-dimension* for NMF.

2.3. Related Literature: finding the number of vertices of a Latent Polytope

To the best of our knowledge the problem of finding the number of vertices of the latent polytope has not been studied. The problem of enumerating all the vertices of a polytope, described by linear inequalities, is known to be hard (Khachiyan et al., 2008). The problem described

here is much more complicated as we have access only to some perturbed points. Since LkP is a geometric construct it maybe difficult to adopt probabilistic approaches such as Hierarchical Dirichlet Process (Teh et al., 2006) which have shown empirical success in many applications. However, to the best of our understanding it does not provide a guarantee for finding the correct k , the number of extreme points of \mathbf{K} , from a finite number of samples.

3. Finding the number of vertices of LkP from data: An interpolative approach

Finding the number of vertices of \mathbf{K} , the Latent k Polytope, defined in Definition 2 is in general a hard problem. Many open problems of active research such as finding the number of topics in a topic model are important special cases of this problem. The challenges in addressing the issue of determining the number of vertices can be distilled into two following issues.

Firstly, one needs to establish a suitable quantity which relates the observed data \mathbf{A} to the columns of \mathbf{M} . Existing attempts in special cases (e.g. use of Hierarchical Dirichlet processes for Topic models(Teh et al., 2006)) do not translate into a finite-time algorithm which guarantees the recovery of the correct number of vertices. A well-formulated quantity which can yield tractable algorithmic insights under realistic assumptions is important aim to have. Secondly, even if such a quantity is devised, the overall goal of finding a tractable algorithm is a harder challenge.

In this paper we address both these issues by appealing to convex geometry. It involves conceptualizing a new notion of Matrix rank (see Subsection 3.1), and following which we develop a polynomial time algorithm(see Subsection 3.2).

3.1. Interpolative Convex Rank (ICR) of a Matrix

In this section the concept of Interpolative Convex Rank (ICR) of a matrix is introduced.

Definition 4. Given a parameter $\varepsilon > 0$, the *Interpolative Convex Rank (ICR) of a matrix $\mathbf{C} \in \mathbb{R}^{d \times m}$ denoted $\text{ICR}(\varepsilon, \mathbf{C}) = l$, where l is the minimum number such that there exists columns $\mathbf{C}_{\cdot, i_1}, \dots, \mathbf{C}_{\cdot, i_l}$ of \mathbf{C} such that*

$$\text{dist}(\text{CH}(\mathbf{C}), \text{CH}(\mathbf{C}_{\cdot, i_1}, \dots, \mathbf{C}_{\cdot, i_l})) \leq \varepsilon$$

In words it is the minimum number of columns of \mathbf{C} whose convex hull is ε close the convex hull of all columns of \mathbf{C} . A geometric insight behind the definition can be obtained by observing that if $\varepsilon = 0$ then ICR of \mathbf{C} is equal to the number of *extreme points* of $\text{CH}(\mathbf{C})$. For any $\varepsilon > 0$, $\text{ICR}(\varepsilon, \mathbf{C})$ is the minimum number of vertices drawn from the columns of \mathbf{C} whose convex hull approximates the convex hull of

all the columns of \mathbf{C} with error at most ε . Such approximations have been studied with different notions of distance in Convex Geometry (Barvinok, 2014; Gruber, 1993).

For any non zero ε , the quantity ICR rank is the smallest number of extreme points of a polytope whose convex hull approximates the polytope.

One of the goals of this paper is to relate NMF to LkP. We argue that Non-Negative Rank (NNR), a often used measure for number of factors, may not be suited and ICR is a better alternative. To support this claim we first recall the definition of NNR (Vavasis, 2009); we define it through the *inner dimension*.

Definition 5. *The non-negative rank (NNR) of a matrix is the minimum inner dimension of a non-negative factorization of the matrix.*

There is no known provable method for identifying the correct *inner dimension*; it is mostly determined by heuristics (see Section 3 in (Gillis, 2014)). It is folklore that (under some conditions) NNR reflects or is equal to the correct inner dimension. However, in many cases NNR can be very different than inner-dimension. In the Supplement this is illustrated with several examples where ICR recovers the true inner-dimension while NNR could be much lower. The intuitive reason for the contrast between NNR and ICR in these examples is that the interpolative property forces ICR to restrict attention to the columns of the given data matrix, but NNR does not have such restriction and hence can choose suitable but arbitrary points not from the data-matrix to yield a low value of NNR.

The use of interpolative assumption for NMF is not new and is mentioned in (Recht et al., 2012). However the algorithms assume that the value of inner-dimension is known and using ICR to characterize the inner-dimension is new.

3.2. Recovering the number of vertices of LkP from

ICR

In this section, we prove the result that the parameter k can be recovered from the data matrix \mathbf{A} as ICR of a suitably smoothed data matrix $\tilde{\mathbf{A}}$ obtained from \mathbf{A} . As pointed out earlier, the raw data matrix \mathbf{A} does not yield k since $\text{CH}(\mathbf{A})$ can be much larger than \mathbf{K} .

We now define the subset smoothed version of \mathbf{A} .

Definition 6. (Subset smoothed data matrix) *Let \mathcal{R}_δ denote the collection of all subsets $R \subseteq [n]$, $|R| = \delta n$. Then the subset smoothed data matrix $\tilde{\mathbf{A}}$ has $|\mathcal{R}_\delta|$ columns indexed by the elements in \mathcal{R}_δ . For each $R \in \mathcal{R}_\delta$, the column $\tilde{\mathbf{A}}_{\cdot,R}$ of $\tilde{\mathbf{A}}$ is same as $\mathbf{A}_{\cdot,R}$, i.e., $\frac{1}{|\mathcal{R}_\delta|} \sum_{\ell \in R} \mathbf{A}_{\cdot,\ell}$.*

We now state the main result of this section.

Theorem 1. *Assume that the input data matrix \mathbf{A} is gener-*

ated by LkP (as in Definition 2). Let $\rho > 0$, $\delta \in (0, 1)$ and suppose the following three assumptions are satisfied

1. *Every vertex $M_{\cdot,\ell}$ of \mathbf{K} satisfies the following condition:*

$$d(M_{\cdot,\ell}, \text{CH}(\{M_{\cdot,\ell'} : \ell' \neq \ell\})) \geq 7\rho. \quad (1)$$

2. *For every $\ell \in [k]$, there exists a subset $S_\ell \subseteq [n]$ of size δn such that*

$$|M_{\cdot,\ell} - P_{\cdot,S_\ell}| \leq \frac{\rho}{30} \quad (2)$$

3. *The perturbation parameter σ is defined as follows*

$$\frac{\sigma}{\sqrt{\delta}} \leq \frac{\rho}{36}. \quad (3)$$

Then the $\text{ICR}(\rho/3, \tilde{\mathbf{A}}) = k$.

Proof. The proof consists of computing an upper-bound (see Lemma 3) and a lower bound (see Corollary 8). \square

Remarks: Before proceeding to give formal details we wish to make some remarks to convey the main intuitions. The assumptions are re-statements of the three standard assumptions defined in Subsection 2.1.2. We would like to note that the first condition (1) above is much weaker than requiring that $M_{\cdot,\ell}$ is far from the affine hull of the other columns of $M_{\cdot,\ell}$, made in (Bhattacharyya & Kannan, 2020). We now give some intuitions behind the proof. It proceeds in two parts. In the upper bound proof, we show that $\text{ICR}(\rho/3, \tilde{\mathbf{A}})$ is at most k (see Lemma 3). The idea of the proof is as follows. For $R \in \mathcal{R}_\delta$, $\mathbf{A}_{\cdot,R}$ and $P_{\cdot,R}$ are close to each other (this is where the subset smoothing operation helps). Now, $P_{\cdot,R}$ lies inside \mathbf{K} , and so can be expressed as a convex combination of the columns $M_{\cdot,\ell}$ of \mathbf{M} , each of which is close to the corresponding point A_{\cdot,S_ℓ} , where S_ℓ is the subset defined in (2). It follows that $\mathbf{A}_{\cdot,R}$ is close to the convex hull of the points $\{A_{\cdot,S_\ell}, \ell \in [k]\}$. This exhibits a convex polytope on k vertices which are close to $\text{CH}(\tilde{\mathbf{A}})$, proving the upper bound.

The lower bound part of the proof (see Corollary 8) is harder and involves new tools from Convex Geometry. Suppose there is a subset of points W whose convex hull is close to $\text{CH}(\tilde{\mathbf{A}})$. It is not hard to show that the convex hull of W is also close to $\text{CH}(\mathbf{M})$ (again using subset smoothing). Now we show that we can partition (a slightly larger set than) $\text{CH}(\mathbf{M})$ into k parts by cutting it with suitable hyperplanes obtained from the Separating Hyperplane Theorem, and W must have a non-empty intersection with each of these parts. This shows that $|W|$ must be at least k .

3.3. Upper Bound Proof

Let H denote $\text{CH}(\tilde{\mathbf{A}})$. Let H' denote the convex hull of $\{A_{\cdot, S_\ell}, \ell \in [k]\}$, where $S_\ell \in \mathcal{R}_\delta$ are as defined in (2). We show that H is within Hausdorff distance $\rho/3$ of H' . Since $H' \subseteq H$, it suffices to show that for every point in H , there is a point in H' within distance at most $\rho/3$. H is a convex set and $d(x, H')$ is a convex function of x , so the maximum of $d(x, H')$ over H is attained at a vertex of H . Hence it suffices to show that for each of the points $A_{\cdot, R}, R \in \mathcal{R}_\delta$, we have $d(A_{\cdot, R}, H') \leq \rho/3$.

The proof proceeds by first showing that for any $R \in \mathcal{R}_\delta$, the vector $A_{\cdot, R}$ is close to $P_{\cdot, R}$, in fact within distance at most $\sigma/\sqrt{\delta}$. Now, any vector of the form $P_{\cdot, R}$ can be written as a convex combination of the vertices of K , i.e., the vectors $M_{\cdot, \ell}, \ell \in [k]$. The same argument also shows that each of the vectors $M_{\cdot, \ell}$ is close to the corresponding vector A_{\cdot, S_ℓ} . Thus $A_{\cdot, R}$ is close to a convex combination of the vectors $A_{\cdot, S_\ell}, \ell \in [k]$. We now prove these claims formally.

Claim 2. For all $S \subseteq [n], |S| \geq \delta n$, $|A_{\cdot, S} - P_{\cdot, S}| \leq \frac{\sigma}{\sqrt{\delta}}$.

Proof. We first note that $|A_{\cdot, S} - P_{\cdot, S}| \leq \|\mathbf{A} - \mathbf{P}\|/\sqrt{|S|}$. Indeed, let u be the unit vector which is $1/\sqrt{|S|}$ on coordinates corresponding to S , 0 otherwise. Then

$$|A_{\cdot, S} - P_{\cdot, S}| = \frac{1}{\sqrt{|S|}} \cdot |(\mathbf{A} - \mathbf{P})u| \leq \frac{\|\mathbf{A} - \mathbf{P}\|}{\sqrt{|S|}}.$$

The claim now follows by using the definition of σ , i.e., $\|\mathbf{A} - \mathbf{P}\| = \sigma\sqrt{n}$. \square

We are now ready to show the upper bound result.

Lemma 3. Under the conditions of Theorem 1, $\text{ICR}(\rho/3, \tilde{\mathbf{A}})$ is at most k .

Proof. Fix an $R \in \mathcal{R}_\delta$. Since $P_{\cdot, R} \in K$, it can be written as a convex combination $\sum_{\ell \in [k]} \lambda_\ell M_{\cdot, \ell}$ of the vertices of K . Then, we have

$$\left| A_{\cdot, R} - \sum_{\ell \in [k]} \lambda_\ell A_{\cdot, S_\ell} \right| \leq |A_{\cdot, R} - P_{\cdot, R}| + \sum_{\ell} \lambda_\ell \left(|M_{\cdot, \ell} - P_{\cdot, S_\ell}| + |A_{\cdot, S_\ell} - P_{\cdot, S_\ell}| \right) \leq \frac{2\sigma}{\sqrt{\delta}} + \frac{\rho}{30} \leq \rho/3,$$

since the first and third term (in the RHS of the first inequality) are each at most $\sigma/\sqrt{\delta}$ by the Claim 2 and the second term is at most $\rho/30$ by (2). This shows that that the $\text{ICR}(\rho/3, \tilde{\mathbf{A}})$ is at most k . \square




Figure 1. Illustration of the proof of Theorem 5: the matrix \mathbf{R} has four points. The lightly shaded half-space around $R_{\cdot, \ell}$ denotes V_ℓ , whereas the darker shaded region around the same vertex denotes Q_ℓ .

3.4. Lower Bound Proof

We now prove the more difficult part of Theorem 1, i.e., $\text{ICR}(\rho/3, \tilde{\mathbf{A}})$ cannot be strictly less than k . We will prove a stronger assertion, namely, that there is no set of $k-1$ points in \mathfrak{R}^d whose convex hull is within distance $\rho/3$ of H , where $H = \text{CH}(\tilde{\mathbf{A}})$.

Assume, for sake of contradiction, that there exist points $w_1, \dots, w_{k-1} \in \mathfrak{R}^d$ such that

$$\text{dist}(H, \text{CH}(\{w_1, \dots, w_{k-1}\})) \leq \rho/3. \quad (4)$$

Recall that $K = \text{CH}(\mathbf{M})$. Let W_0 denote $\text{CH}(\{w_1, \dots, w_{k-1}\})$. We first begin by showing that $\text{dist}(K, W_0)$ is at most ρ . The intuition is that the Hausdorff distance between H and K is small – the argument follows along the same lines as in the upper bound proof. It follows from (4) that $\text{dist}(K, W_0)$ is also small. We give the complete proof of the following result in the supplementary material.

Claim 4. Let W_0 denote the convex hull of the points w_1, \dots, w_{k-1} . Then, $\text{dist}(K, W_0) \leq \rho$.

The desired contradiction follows easily from the above claim and the following result:

Theorem 5. Let W be a set of points such that the Hausdorff distance between the the convex hull of W and that of the columns of \mathbf{M} is at most ρ , i.e., $\text{dist}(\text{CH}(W), \text{CH}(\mathbf{M})) \leq \rho$. Then $|W| \geq k$.

The intuition behind the proof is as follows: each vertex $M_{\cdot, \ell}$ of $\text{CH}(\mathbf{M})$ is far from the convex hull formed by rest of the columns of \mathbf{M} , and hence we can find a separating hyperplane which separates $M_{\cdot, \ell}$ from $\text{CH}(\mathbf{M}^{-\ell})$ by some large enough margin (see the region V_ℓ in Figure 1). We define a further refinement of this region, shown as Q_ℓ in Figure 1. The proof consists of two conceptual steps: (i) we

show that the regions Q_ℓ are disjoint for distinct columns $M_{\cdot,\ell}$ of \mathbf{M} , and (ii) we argue that each of the regions Q_ℓ must contain at least one vertex of W . Since there are k columns in \mathbf{M} , this would imply that $|W|$ must be at least k .

Proof. For two subsets A and B of points in \mathbb{R}^d , define their Minkowski sum, $A + B$, as $\{a + b : a \in A, b \in B\}$. For an index $\ell \in [k]$, let $\mathbf{M}^{-\ell}$ denote the matrix \mathbf{M} with column ℓ removed. Since the columns of \mathbf{M} satisfy condition (1), this implies that (here \mathbf{B} denotes the unit ball)

$$[M_{\cdot,\ell} + 7\rho\mathbf{B}] \cap \text{CH}(\mathbf{M}^{-\ell}) = \emptyset.$$

Therefore, the Separating Hyperplane Theorem from Convex Geometry implies that there is a unit vector $v^{(\ell)}$ such that for all $\ell' \neq \ell$,

$$v^{(\ell)} \cdot M_{\cdot,\ell} > v^{(\ell)} \cdot M_{\cdot,\ell'} + 7\rho \quad (5)$$

For an index $\ell \in [k]$, let V_ℓ denote the set

$$V_\ell := \{x : v^{(\ell)} \cdot x > v^{(\ell)} \cdot M_{\cdot,\ell} - 2\rho\},$$

and define $Q_\ell := (\text{CH}(\mathbf{M}) + \rho\mathbf{B}) \cap V_\ell$.

Claim 6. *For every $\ell \in [k]$, there is a point $w \in W$ such that $w \in Q_\ell$.*

Proof. Assume for the sake of contradiction that there is an index $\ell \in [k]$ such that $Q_\ell \cap W = \emptyset$. Since $\text{dist}(\text{CH}(W), \text{CH}(\mathbf{M})) \leq \rho$, $W \subseteq \text{CH}(\mathbf{M}) + \rho\mathbf{B}$, and so it must be the case that $V_\ell \cap W = \emptyset$. The definition of V_ℓ implies that for every $w \in W$,

$$w \cdot v^{(\ell)} \leq M_{\cdot,\ell} \cdot v^{(\ell)} - 2\rho.$$

Consequently, for every point $y \in \text{CH}(W)$,

$$v^{(\ell)} \cdot y \leq v^{(\ell)} \cdot M_{\cdot,\ell} - 2\rho. \quad (6)$$

Since $\text{dist}(\text{CH}(\mathbf{M}), \text{CH}(W)) \leq \rho$, there is a point $y \in \text{CH}(W)$ with $|y - M_{\cdot,\ell}| \leq \rho$. But this contradicts (6). This proves the claim. \square

So, we must have that for each $\ell \in [k]$, $Q_\ell \cap W$ is non-empty. We now show that the sets Q_ℓ are disjoint, which will yield the desired result.

Lemma 7. *The sets $Q_\ell, \ell \in [k]$ are mutually disjoint.*

Proof. Suppose, for the sake of contradiction, that there is a point $z \in Q_\ell \cap Q_{\ell'}$ for distinct indices ℓ, ℓ' . Since $z \in \text{CH}(\mathbf{M}) + \rho\mathbf{B}$, there is a point $y \in \text{CH}(\mathbf{M})$, with $|y - z| \leq \rho$.

So y can be written as a convex combination $\sum_{\ell'' \in [k]} \alpha_{\ell''} M_{\cdot,\ell''}$. We now also write y as

$$y = \alpha_\ell M_{\cdot,\ell} + (1 - \alpha_\ell)x, \quad x \in \text{CH}(M_{\cdot,\ell''} : \ell'' \neq \ell) \quad (7)$$

Inequality (5) implies that $v^{(\ell)} \cdot x < v^{(\ell)} \cdot M_{\cdot,\ell} - 7\rho$. So,

$$v^{(\ell)} \cdot y < v^{(\ell)} \cdot M_{\cdot,\ell} - (1 - \alpha_\ell)7\rho.$$

Since $v^{(\ell)}$ is unit vector and $|y - z| \leq \rho$,

$$v^{(\ell)} \cdot z \leq v^{(\ell)} \cdot M_{\cdot,\ell} - (1 - \alpha_\ell)7\rho + \rho.$$

But since $z \in V_{\ell'}$, we have by the definition of $V_{\ell'}$

$$v^{(\ell)} \cdot z > v^{(\ell)} \cdot M_{\cdot,\ell} - 2\rho.$$

Thus by the last two inequalities, $\alpha_\ell > 4/7 > 0.5$. A similar argument shows that $\alpha_{\ell'} > 0.5$. But this is not possible because the quantities $\alpha_{\ell''}$ are non-negative and add to 1. This proves the result. \square

Combining Lemma 7 and Claim 6, we see that $|W| \geq k$. This proves the desired theorem. \square

As a corollary, we get the lower bound result:

Corollary 8. *Under the assumptions of Theorem 1, $\text{ICR}(\rho/3, \tilde{\mathbf{A}}) \geq k$.*

Proof. Combining Theorem 5 and Claim 4, we see that if W is any set of points for $\text{dist}(\text{CH}(\tilde{\mathbf{A}}), \text{CH}(W)) \leq \rho/3$, then $|W| \geq k$. By definition, if $\text{ICR}(\rho/3, \tilde{\mathbf{A}}) = k'$, then there is a set of k' points whose convex hull is within Hausdorff distance $\rho/3$ of $\text{CH}(\tilde{\mathbf{A}})$. Therefore, $k' \geq k$. \square

It is worth noting that the matrix $\tilde{\mathbf{A}}$ has exponential (in n) number of columns, and so cannot be used directly by an efficient algorithm. We have given an algebraic characterization of the number of vertices of K in terms of ICR , but no polynomial time algorithm is known to find ICR . In the next section, we give a polynomial time algorithm to find the number of vertices of K .

4. Algorithm for finding the number of vertices of the latent polytope

In this section, we consider a slightly stronger set of standard assumptions than those in Theorem 1. We prove that if the input points satisfy these conditions, then we can recover the correct value of k in polynomial time. In fact, we will show that there is a polynomial time computable threshold τ such that

$$k = \text{Max}_\tau : s_\tau(\mathbf{A}) \geq \tau.$$

Singular value thresholding is a well known procedure which yields provable guarantees in many applications (Cai

et al., 2010; Chatterjee, 2015; Donoho & Gavish, 2014). However the techniques considered in the literature do not apply to the problem at hand and hence new methods are warranted.

The main difficulty in choosing such a threshold τ is that it cannot be kept at some suitable multiple of the largest singular value of the data matrix \mathbf{A} : the matrix \mathbf{A} may be ill-conditioned, and so the largest singular value may not give us any information about the parameter k . Instead, we write down a convex program whose output is used for deciding this threshold τ . Once we have the correct value of k , we can use ideas based on ICR to recover the actual vertices of \mathbf{M} (see for example Algorithm 2 and Theorem 20).

Algorithm 1 Algorithm for finding the number of extreme points of K .

Input: $d \times n$ data matrix \mathbf{A} and a parameter $\delta \in (0, 1)$.
Let opt be the optimal value of the convex program (12).

Let k^* be the maximum value of $k' \in [d]$ such that the singular value $s_{k'}(\mathbf{A}) \geq \delta^2 \text{opt}/8$.

Output k^* .

We now state the main result of this section. We give some definitions. For a subspace V and a point x , let $\mathbf{proj}(x, V)$ denote the orthogonal projection of x on V . For a subset S of points, let $\mathbf{Null}(S)$ denote the subspace orthogonal to the linear span of S .

Theorem 9. *Suppose that the input data satisfies the following assumptions. There exists a $\delta \in (0, 1)$ such that*

Every vertex $M_{\cdot,\ell}$ of K satisfies the following condition:

$$|\mathbf{proj}(M_{\cdot,\ell}, \mathbf{Null}(\mathbf{M} \setminus M_{\cdot,\ell}))| \geq \delta |M_{\cdot,\ell}|. \quad (8)$$

The LHS above is same as the distance of $M_{\cdot,\ell}$ from the affine hull of the other columns of \mathbf{M} .

$$\text{For } S_\ell := \{j : |M_{\cdot,\ell} - P_{\cdot,j}| \leq \frac{4\sigma}{\sqrt{\delta}}\}, \quad |S_\ell| \geq \delta n \quad (9)$$

$$\sigma \leq \delta^3 \min_\ell |M_{\cdot,\ell}|/20. \quad (10)$$

Further, assume that all the entries of \mathbf{M} are non-negative. Then, given δ and the data matrix \mathbf{A} , Algorithm 1 outputs the correct value of k .

Note that the LHS of condition (8) is smaller than the LHS of the corresponding constraint (1).

The following result follows as corollary of the above result. It shows that the ICR of the subset smoothed data matrix $\tilde{\mathbf{A}}$ is equal to the number of vertices k of K . Details are given in the supplementary material.

Lemma 10. *Suppose the conditions (8), (9) and (10) stated in Theorem 9 are satisfied by the input data. Then $\text{ICR}(6\sigma/\sqrt{\delta}, \tilde{\mathbf{A}}) = k$.*

4.1. Proof of Theorem 9

We now proceed to prove Theorem 9. There are two phases in Algorithm 1. In the first phase, we estimate $\min_\ell |M_{\cdot,\ell}|$ by writing a suitable convex program. We use this estimate as a threshold for the singular values of \mathbf{A} to determine the value of k .

The convex program is as follows (recall that \mathbf{A} is the data matrix generated from LkP defined in Definition 2):

$$\min \text{opt} := |\mathbf{A} \cdot x| \quad (11)$$

$$\sum_{j \in [n]} x_j = 1, \quad 0 \leq x_j \leq 1/(\delta n), \forall j \in [n]. \quad (12)$$

We now show that the optimal value of the convex program can be used to approximate $\min_\ell |M_{\cdot,\ell}|$ within a constant factor.

Lemma 11. *Let $\rho_0 := \min_\ell |M_{\cdot,\ell}|$, and opt be defined in (11). Then*

$$\rho_0/(2k) \leq \text{opt} \leq 2\rho_0.$$

Proof. The proof is done by two subclaims, each of which prove one of the above inequalities. Let x^* denote the vector which satisfies $\text{opt} = |\mathbf{A} \cdot x^*|$.

Claim 12. $|\mathbf{A} \cdot x^*| \geq \rho_0/(2k)$.

Proof. The vector $\mathbf{P} \cdot x^*$ is a convex combination of the columns of \mathbf{P} . Each column of \mathbf{P} can be expressed as a convex combination of the columns of \mathbf{M} . It follows that $\mathbf{P} \cdot x^*$ is a convex combination of the columns of \mathbf{M} , i.e., it is of the form $\sum_{j \in [k]} \alpha_j M_{\cdot,j}$, where α_j 's are non-negative and sum to 1. It follows that there is an index j with $\alpha_j \geq 1/k$. Since all the columns of \mathbf{M} are non-negative, $\mathbf{P} \cdot x$ is component-wise at least $M_{\cdot,j}/k$, and so, has length at least ρ_0/k . Finally,

$$|\mathbf{A} \cdot x| \geq |\mathbf{P} \cdot x| - \|\mathbf{P} - \mathbf{A}\| \cdot |x| \geq \rho_0/k - \sigma/\delta \stackrel{(10)}{\geq} \rho_0/(2k),$$

where the second last inequality follows from the fact that $x_j \leq 1/(\delta n)$ and (10). \square

Claim 13. $|\mathbf{A} \cdot x^*| \leq 2\rho_0$

Proof. Let ρ_0 be $|M_{\cdot,\ell}|$. Define a feasible solution x to the convex program \mathcal{P} as follows: $x_j = 1/(\delta n)$ if $j \in S_\ell$, 0 otherwise. Observe that $|x| = \frac{1}{\sqrt{\delta} \sqrt{n}}$. Now,

$$|\mathbf{A} \cdot x| = |A_{\cdot,S_\ell}| \leq |P_{\cdot,S_\ell}| + \frac{\sigma}{\sqrt{\delta}} \leq |M_{\cdot,\ell}| + \frac{4\sigma}{\sqrt{\delta}} + \frac{\sigma}{\sqrt{\delta}} \leq 2\rho_0,$$

where the first inequality is by Claim 4 and the second by (9). \square

The above two results show that $\rho_0/(2k) \leq \text{opt} \leq 2\rho_0$. \square

We now use opt as a threshold for the singular values of \mathbf{A} , and show that k can be recovered from this threshold. This will also prove correctness of Algorithm 1, and hence yield a provable data determined threshold for finding the ICR of \mathbf{A} .

Theorem 14. Define $k^* = \max_{k' \in [d]} : s_{k'}(\mathbf{A})/\sqrt{n} \geq \delta^2 \text{opt}/8$. Then $k = k^*$.

Proof. In order to prove this result, we need to estimate $s_k(\mathbf{A})$ and $s_{k+1}(\mathbf{A})$. We shall do this by estimating the singular values of some related matrices. We begin by bounding the singular values of \mathbf{M} .

Claim 15. $s_k(\mathbf{M}) \geq \frac{\delta \rho_0}{\sqrt{k}}$, whereas $s_{k+1}(\mathbf{M}) = 0$.

Proof. The second statement follows trivially since \mathbf{M} has only k columns, and so is a rank k matrix. For the first statement, observe that

$$s_k(\mathbf{M}) = \min_{x:|x|=1} |\mathbf{M}x|.$$

Let x be any unit vector in \mathfrak{R}^k . There must be a coordinate $\ell \in [k]$ such that $|x_\ell| \geq \frac{1}{\sqrt{k}}$. Now, $|\mathbf{M}x| = \left| \sum_{j \in [k]} x_j M_{.,j} \right| \geq |x_\ell| \cdot |\text{proj}(M_{.,\ell}, \text{Null}(\mathbf{M} \setminus M_{.,\ell}))| \geq \frac{\delta \rho_0}{\sqrt{k}}$, where the last inequality follows from condition (8). \square

Let S_1, \dots, S_k be subsets of $[n]$ guaranteed by condition (9). We assume that each of these subsets have size exactly δn (otherwise we remove some elements from them). Note that these sets are mutually disjoint, since (8) and (10) imply that for any two distinct $\ell, \ell' \in [k]$, $|M_{.,\ell} - M_{.,\ell'}| > \frac{8\sigma}{\sqrt{\delta}}$.

Let $\mathbf{B} = (P_{.,S_1} \mid P_{.,S_2} \mid \dots \mid P_{.,S_k})$. We relate $s_k(\mathbf{B})$ and $s_k(\mathbf{M})$.

Claim 16. $s_k(\mathbf{B}) \geq s_k(\mathbf{M}) - \frac{4\sigma\sqrt{k}}{\sqrt{\delta}} \geq \frac{\delta \rho_0}{2\sqrt{k}}$.

Proof. Condition (9) implies that $\|\mathbf{B} - \mathbf{M}\| \leq \frac{4\sigma\sqrt{k}}{\sqrt{\delta}}$. This implies the first inequality in the claim, since we have (from Linear Algebra), $s_k(\mathbf{B}) \geq s_k(\mathbf{M}) - \|\mathbf{M} - \mathbf{B}\|$. The second inequality follows from Claim 15 and (10). \square

We now relate $s_k(\mathbf{B})$ and $s_k(\mathbf{P})$.

Claim 17. $s_k(\mathbf{P}) \geq \sqrt{\delta n} \cdot s_k(\mathbf{B}) \geq \frac{\delta^{1.5} \rho_0}{2\sqrt{k}} \cdot \sqrt{n}$.

Proof. We shall exhibit a subspace of dimension k such that for any unit vector x in this subspace, $|\mathbf{P} \cdot x| \geq \sqrt{\delta n} \cdot s_k(\mathbf{P})$. This will prove the desired result by the Max-Min theorem of Linear Algebra which states:

$$s_k(\mathbf{B}) = \max_{V: \dim(V)=k} \min_{x \in V, |x|=1} |\mathbf{P}x|.$$

For each index $\ell \in [k]$ defined a unit vector $e_\ell \in \mathfrak{R}^n$ as follows:

$$(e_\ell)_j = \begin{cases} \frac{1}{\sqrt{\delta n}} & \text{if } j \in S_\ell \\ 0 & \text{otherwise} \end{cases}$$

Since the sets S_1, \dots, S_k are pair-wise disjoint, the vectors $e_\ell, \ell \in [k]$, span a subspace U of dimension k (in fact these vectors are form an orthonormal family of k vectors).

Now consider any unit vector $v = \sum_{j \in [k]} \beta_j e_j$. Since e_j 's are orthonormal, $\sum_j \beta_j^2 = 1$. Now,

$$\mathbf{P} \cdot v = \sum_{j \in [k]} \beta_j \sqrt{\delta n} \cdot P_{.,S_j} = \sqrt{\delta n} \cdot \mathbf{P} \cdot y,$$

where y is the vector whose j^{th} coordinate is β_j . Since y is a unit vector, $|\mathbf{P} \cdot y| \geq s_k(\mathbf{B})$. This implies the desired result (the last inequality in the claim follows from Claim 16). \square

Finally, we give bounds on the singular values of the data matrix \mathbf{A} .

Lemma 18. $s_k(\mathbf{A}) \geq \delta^2 \text{opt}/8 > s_{k+1}(\mathbf{A})$.

Proof. Since $\|\mathbf{P} - \mathbf{A}\| \leq \sigma\sqrt{n}$, $|s_r(\mathbf{P})/\sqrt{n} - s_r(\mathbf{A})/\sqrt{n}|$ is at most σ for any $r \in [d]$. Using this fact and Claim 17, we see that

$$s_k(\mathbf{A})/\sqrt{n} \geq \frac{\delta^{1.5} \rho_0}{4\sqrt{k}} \geq \frac{\delta^{1.5} \text{opt}}{8\sqrt{k}} \geq \frac{\delta^2 \text{opt}}{8},$$

where the second inequality follows from Claim 13, and the last inequality follows from the fact that $\delta \leq 1/k$.

Similarly, (since $s_{k+1}(\mathbf{P}) = 0$)

$$\begin{aligned} s_{k+1}(\mathbf{A})/\sqrt{n} &\leq \sigma \stackrel{(10)}{\leq} \delta^3 \rho_0/20 \leq k\delta^3 \text{opt}/10 \\ &\leq \delta^2 \text{opt}/10 < \delta^2 \text{opt}/8, \end{aligned}$$

where the third inequality follows from Claim 12. \square

This proves that $k^* = k$. \square

Thus we have shown that $k^* = k$. This proves that we can recover the number of vertices in the latent polytope if the conditions (8), (9), (10) are satisfied, hence proving Theorem 9.

4.2. Finding a non-negative matrix factorization without the knowledge of inner-dimension

Given \mathbf{A} finding an approximate non-negative factorization is intractable. In a long line of work starting from (Arora et al., 2012) several ingenious algorithms have emerged (Recht et al., 2012; Gillis & Vavasis, 2014; Rong & Zou, 2015) which using *Separability*, first introduced by (Donoho & Stodden, 2003), provably recover an approximate factorization. These procedures assume that the inner-dimension k is known. By imposing the assumption that *Every data point is approximately a non-negative combination of k data points*, several algorithms (e.g. (Ambikapathi et al., 2013; Fu et al., 2015; Gillis & Luce, 2018)) can find minimum such k by solving Linear or Convex programs. However the assumption is too strong to hold for LkP and can be shown to fail in many simple instances. To the best of our knowledge there are no known algorithms which can provably recover the inner-dimension and the factorization under the conditions of LkP.

In this section we state polynomial time procedures which provably recovers the factorization (the matrix \mathbf{M}) along with k , the inner-dimension (see equation (Approx-NMF)). Noting that Approx-NMF is a special case of LkP, the application of Algorithm 1 recovers k .

It is important to note that the *standard assumptions* stated earlier are different than separability. In fact conditions (9) and (10) of Theorem 9, are analogous to assumptions which have been used for NMF, namely, *nearly pure records* and *noise* assumptions of (Bhattacharyya et al., 2016). Instead of the *Well-Separatedness* condition (8) of Theorem 9, we make the following ‘‘Dominant Features’’ assumptions which is similar to, but, weaker than their assumption with the same name.

Dominant Features This stipulates that for $\ell \in [k]$, there is a set T_ℓ of ‘‘dominant features’’ of column ℓ of \mathbf{M} . More precisely,

$$\exists T_1, T_2, \dots, T_k \text{ disjoint} : \forall \ell, \sum_{i \in T_\ell} M_{i,\ell} \geq 4 \sum_{i \in T_\ell} \sum_{\ell' \neq \ell} M_{i,\ell'}.$$

Lemma 19. *The Dominant Features assumption above implies condition (8) of Theorem 9.*

The proof is based on diagonal dominance and is given in the Supplementary Material. Finally we state the main theorem, which gives a procedure, described in Algorithm 2, for finding close approximations to the vertices of \mathbf{M} .

The algorithm is a simplified version of (Bhattacharyya & Kannan, 2020), requiring new proofs. At the outset, the algorithm finds the left k -singular subspace V of \mathbf{A} and projects data points onto V . It runs in k iterations and in each iteration, finds an approximation to a new vertex

of K . This approximation is of the form $A_{\cdot,R}$ for some $R \subseteq [n]$, $|R| = \delta n$. R is obtained as follows: We pick a random vector u_r in $V \cap$ the null space of the vectors found so far and set: $R \leftarrow \arg \max_R |u \cdot A_{\cdot,R}|$, noting that the optimization can be carried out in polynomial time. Indeed, the optimization to find R is a 1-dimension problem. We project all the columns $A_{\cdot,j}$ along u_r (note that this is the matrix \mathbf{A} which has only n columns). Now the optimal R would be either the δn columns of \mathbf{A} with the highest dot product with u_r or the smallest, and so, can be done in polynomial time.

Theorem 20. *Assuming conditions (9) and (10) and also the Dominant Features assumption, Algorithm 1 finds k in polynomial time. Using this value of k , Algorithm 2 finds k points, each of which is within $O(k^4 \sigma / \delta^{1.5})$ distance of a unique vertex of \mathbf{M} .*

The proof is given in the supplementary material. The above theorem yields an NMF without knowing the value of *inner-dimension*. A detailed discussion of the relative strengths and weakness of various assumptions is beyond the scope of this paper and will be discussed elsewhere. It suffices to stay that the assumptions, techniques and the results presented in this section are novel contributions to NMF.

Algorithm 2 Algorithm for finding the approximations to the vertices of K .

Input: $d \times n$ data matrix \mathbf{A} and a parameter $\delta \in (0, 1)$, k .

Let V be the subspace of \mathfrak{R}^d spanned by the top k left singular vectors of \mathbf{A} .

Let $\hat{A}_{\cdot,j}$ be the projection of $A_{\cdot,j}$ on V , for $j = 1, \dots, n$.

for $r = 0, \dots, k - 1$ **do**

$u_r \leftarrow$ unit vector chosen uniformly at random from the subspace $U_r := V \cap \mathbf{Null}(\hat{A}_{\cdot,R_1}, \dots, \hat{A}_{\cdot,R_r})$.

$R_{r+1} \leftarrow \arg \max_{R \in \mathcal{R}_\delta} |u \cdot \hat{A}_{\cdot,R}|$.

end for

Output $\hat{A}_{\cdot,R_1}, \dots, \hat{A}_{\cdot,R_k}$.

5. Conclusions

This paper uses Convex geometry in finding the number of extreme points of the Latent- k -polytope. This settles several interesting problems related to finding the number of components in Mixed-membership models.

Acknowledgement The authors thank the reviewers for their insightful comments which greatly helped in preparing the camera ready version. The author CB gratefully acknowledges the support of Microsoft Research.

References

- Ambikapathi, A., Chan, T., Chi, C., and Keizer, K. Hyperspectral data geometry-based estimation of number of endmembers using p -norm-based pure pixel identification algorithm. *IEEE Trans. Geosci. Remote. Sens.*, 51(5-1):2753–2769, 2013. URL <https://doi.org/10.1109/TGRS.2012.2213261>.
- Arora, S., Ge, R., Kannan, R., and Moitra, A. Computing a nonnegative matrix factorization—provably. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pp. 145–162. ACM, 2012.
- Barvinok, A. Thrifty approximations of convex bodies by polytopes. *International Mathematics Research Notices*, 2014(16):4341–4356, 2014. doi: 10.1093/imrn/rnt078.
- Bhattacharyya, C. and Kannan, R. Near-optimal sample complexity bounds for learning latent k -polytopes and applications to ad-mixtures. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 854–863. PMLR, 2020.
- Bhattacharyya, C., Goyal, N., Kannan, R., and Pani, J. Non-negative matrix factorization under heavy noise. In Balcan, M. and Weinberger, K. Q. (eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 1426–1434. JMLR.org, 2016. URL <http://proceedings.mlr.press/v48/bhattacharya16.html>.
- Cai, J., Candès, E. J., and Shen, Z. A singular value thresholding algorithm for matrix completion. *SIAM J. Optim.*, 20(4):1956–1982, 2010. doi: 10.1137/080738970. URL <https://doi.org/10.1137/080738970>.
- Chatterjee, S. Matrix estimation by universal singular value thresholding. *Annals of Statistics*, 43(1):177–214, 2015.
- Donoho, D. and Gavish, M. Minimax risk of matrix denoising by singular value thresholding. *Annals of Statistics*, 42(6):2413–2440, 2014.
- Donoho, D. and Stodden, V. When does non-negative matrix factorization give a correct decomposition into parts? In *Advances in Neural Information Processing Systems*, pp. 1141–1148, 2003.
- Fu, X., Ma, W., Chan, T., and Bioucas-Dias, J. M. Self-dictionary sparse regression for hyperspectral unmixing: Greedy pursuit and pure pixel search are related. *IEEE J. Sel. Top. Signal Process.*, 9(6):1128–1141, 2015.
- Gillis, N. The why and how of nonnegative matrix factorization. *Regularization, Optimization, Kernels, and Support Vector Machines*, 12, 2014.
- Gillis, N. and Luce, R. A fast gradient method for non-negative sparse regression with self-dictionary. *IEEE Trans. Image Process.*, 27(1):24–37, 2018. URL <https://doi.org/10.1109/TIP.2017.2753400>.
- Gillis, N. and Vavasis, S. Fast and robust recursive algorithms for separable nonnegative matrix factorization. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 36(4):698–714, April 2014. ISSN 0162-8828. doi: 10.1109/TPAMI.2013.226.
- Gruber, P. M. *Aspects of approximation of convex bodies, Handbook of Convex Geometry*, volume A, pp. 319–345. North-Holland, Amsterdam, 1993.
- Khachiyan, L., Boros, E., Borys, K., Elbassioni, K., and Gurvich, V. Generating all vertices of a polyhedron is hard. *Discrete & Computational Geometry*, 39(1):174–190, 2008. doi: 10.1007/s00454-008-9050-5. URL <https://doi.org/10.1007/s00454-008-9050-5>.
- Recht, B., Re, C., Tropp, J., and Bittorf, V. Factoring non-negative matrices with linear programs. In *Advances in Neural Information Processing Systems*, pp. 1214–1222, 2012.
- Rong, G. and Zou, J. Intersecting faces: Non-negative matrix factorization with new guarantees. In *ICML*, pp. 2295–2303, 2015.
- Teh, Y. W., Jordan, M. I., Beal, M. J., and Blei, D. M. Hierarchical dirichlet processes. *Journal of the American Statistical Association*, 101(476):1566–1581, 2006. ISSN 01621459. URL <http://www.jstor.org/stable/27639773>.
- Vavasis, S. A. On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization*, 20(3):1364–1377, 2009.