

PERTURBATION OF MATRICES AND NON-NEGATIVE RANK WITH A VIEW TOWARD STATISTICAL MODELS

CRISTIANO BOCCI*, ENRICO CARLINI†, AND FABIO RAPALLO‡

Abstract. In this paper we study how perturbing a matrix changes its non-negative rank. We prove that the non-negative rank is upper-semicontinuous and we describe some special families of perturbations. We show how our results relate to Statistics in terms of the study of Maximum Likelihood Estimation for mixture models.

Key words. Frobenius norm, independence of random variables, Jacobian matrix, mixture models.

AMS subject classifications. 15B51, 62H17

1. Introduction. The rank of a matrix gives the least number of rank one matrices, also known as dyadic products, needed to write the matrix as a sum of dyads. More precisely an $n \times m$ matrix P such that $\text{rk}(P) = k$ can be written as

$$P = c_1(r_1)^t + \dots + c_k(r_k)^t, \quad (1.1)$$

where the column vectors c_h and r_h have the proper sizes. Even if P has non-negative entries, the vectors c_h and r_h , are allowed to have negative entries. If we require the vectors to have non-negative entries, then the least number of summands is called the non-negative rank of P , namely $\text{rk}_+(P)$. The non-negativity constraints make the situation more complex, and the non-negative rank of a matrix is harder to study than the ordinary rank, see e.g. [5]. From this description it is clear that $\text{rk}_+(P) \geq \text{rk}(P)$. Therefore, it could be impossible to decompose a rank k matrix into the sum of exactly k dyadic products $c_h(r_h)^t$, where c_h and r_h are non-negative vectors. The relations between the ordinary rank and the non-negative rank have received an increasing attention in the last years, both from a theoretical and an applied point of view. Some recent references are [3], [6], [16], [19] and [4].

Computing the non-negative rank of a matrix P is related to compute a non-negative factorization of P . There are many recently proposed algorithms to deal with the problem of non-negative matrix factorization, e.g. see [15] or [14] for an application to stochastic matrices. However, the non-negative factorization problem is known to be NP-hard ([22]). Roughly speaking, we can say that there is no efficient way to compute the non-negative rank.

In this paper we study how the non-negative rank of a matrix is affected by small perturbations of the matrix. This is of particular interest when the matrix arises in Probability and Statistics. In fact, when the data entries of the matrices are determined by experimental data, small perturbations must be taken into account.

Here, a perturbation is intended in the following topological sense. Given a matrix P we consider a neighborhood of P in the topology induced by the Frobenius norm on matrices. We call any matrix in the neighborhood a perturbation of P . Clearly this

*Dipartimento di Scienze Matematiche e Informatiche “R. Magari”, Università di Siena, Pian dei Mantellini 44, 53100 Siena, Italy (cristiano.bocci@unisi.it)

†Dipartimento di Matematica, Politecnico di Torino, Corso Duca degli Abruzzi, 24, 10129 Turin, Italy (enrico.carlini@polito.it).

‡Dipartimento di Scienze e Tecnologie Avanzate, Università del Piemonte Orientale, Viale Teresa Michel 11, 15121 Alessandria, Italy (fabio.rapallo@mf.unipmn.it)

notion is more meaningful and interesting when a small neighborhood is considered and hence matrices close to P are studied.

We show that the non-negative rank is upper-semicontinuous with respect to the Frobenius norm, see Theorem 3.1, and hence it cannot be decreased by small perturbations of the matrix. We also produce examples of perturbations preserving the non-negative rank, see Proposition 3.2. Using a Jacobian analytic approach we show that, under some mild conditions, perturbing a matrix leaving the ordinary rank fixed also leaves the non-negative rank unchanged, see Proposition 4.1.

The notion of non-negative rank has also relevant applications in Probability and Statistics. In fact, a probability matrix with dyadic expansion as in Equation (1.1) belongs to the mixture of k independence models for categorical data (in the case that all the involved vectors are non-negative). Mixture models play a central role in applied probability, as they are the key tool in modelling partially observed phenomena, see [1] for more details. Three major topics in Probability and Statistics where mixture models are used as a key ingredient are: (a) the study of sequence alignment, with special attention to DNA sequences and phylogenetic trees, see e.g. [18, 2], and the book [21] for a detailed construction of the underlying mathematical models; (b) the cluster analysis for categorical multidimensional data, see e.g. [11]; (c) multivariate methods for text mining, see e.g. [24].

There are many unsolved problems concerning mixture models. Among these, one of the most important is the determination of the maximum likelihood estimators. Despite the fact that Maximum Likelihood Estimation (MLE) is a largely investigated topic, and many numerical solutions are available, a complete theoretical solution is not available yet. Therefore, any advance in the geometric description of such models can be useful to address the maximization problem from the theoretical viewpoint.

Recently, mixture models for categorical data have been considered also in the framework of Algebraic Statistics, a branch of Statistics which uses notions and techniques from Computational Algebra and Algebraic Geometry, see [18, 7, 9]. In this paper, we show how the geometric description of the set of matrices with fixed non-negative rank leads to a better understanding of MLE for mixture models.

The paper is structured as follows. In Section 2 we recall some basic notions. In Section 3 and Section 4 we use a topological and analytic approach to study perturbations. In Section 5 we use our results to work out some significant examples. Finally, in Section 6 we show how our results relate to the study of MLE in Statistics.

2. Basic facts. In this section, we recall some known facts about the non-negative rank. The definitions and the results presented below will be used throughout the paper.

Non-negative matrices. A non-negative $n \times m$ matrix is a point in $\mathbb{R}_{\geq 0}^{nm}$ where

$$\mathbb{R}_{\geq 0}^{nm} = \{(p_{i,j}) : p_{i,j} \in \mathbb{R}, p_{i,j} \geq 0\}.$$

Stochastic matrices. A stochastic matrix is a non-negative matrix having column sums equal to one. To each non-negative matrix without zero columns, we can associate a stochastic matrix. Denote by $P = [c_1, \dots, c_m]$ the set of columns of a non-negative matrix P , where $c_j \neq 0$ for all j . Define the scaling factor $\sigma(P)$ by

$$\sigma(P) := \text{diag}\{\|c_1\|_1, \dots, \|c_m\|_1\}$$

where $\|\cdot\|_1$ is the 1-norm in \mathbb{R}^n . Then the pullback map θ defined by

$$\theta(A) = A\sigma(A)^{-1}$$

produces stochastic matrices.

Remark. In Probability, stochastic matrices are defined as the non-negative matrices having row sums equal to one. Here we adopt the convention of normalizing the columns. The rank and the non-negative rank are clearly invariant under matrix transposition. Thus this convention does not affect our results.

Simplex. The n -simplex in \mathbb{R}^n is

$$\Delta^n = \left\{ (x_1, \dots, x_n) \in \mathbb{R}^n : x_i \geq 0, \sum_{i=1}^n x_i \leq 1 \right\}.$$

Note that an $n \times m$ stochastic matrix P can be seen as a collection of m points in Δ^n . More precisely we consider the map π_n assigning to a matrix the set of its columns, that is $\pi_n(P) = \{c_1, \dots, c_m\} \subset \Delta^n$. All the points c_j lie on the same face of the n -simplex, which is a $(n-1)$ -simplex. Hence, by dropping the last component of each c_j , we have a map π_{n-1} sending P into a collection of m points in Δ^{n-1} . Following [16] we will use this geometric interpretation to visualize small size matrices and their ranks, see Section 5.

Non-negative rank. Given a $n \times m$ non-negative matrix P , the non-negative rank of P is the smallest integer k such that

$$P = c_1(r_1)^t + \dots + c_k(r_k)^t$$

where the vectors $c_h \in \mathbb{R}^n$ and the vectors $r_h \in \mathbb{R}^m$ have non-negative entries. The non-negative rank of the matrix P is denoted with $\text{rk}_+(P)$. For different description of the non-negative rank we refer the reader to [5].

Remark. For a non-negative matrix P without zero columns, one has that $\text{rk}_+(P) = \text{rk}_+(\theta(P))$. Hence, the study of the non-negative rank of stochastic matrices coincides with the study of the non-negative rank of non-negative matrices without zero columns (see [16]). Thus, from now on, we will often restrict our attention to stochastic matrices.

Combinatorics, Geometry and non-negative rank. The Nested Polytopes Problem (NPP) is introduced in [10] inspired by the intermediate simplex problem of [22]. We can state a simplified version of NPP directly related to the study of the non-negative rank. Let \mathcal{P} be a polytope. Given a set of r points Z in \mathcal{P} is it possible to find k points ($k < r$) in \mathcal{P} having convex hull \mathcal{P}_k such that

$$Z \subset \mathcal{P}_k \subset \mathcal{P}?$$

The following result is contained in [13] and it will be of crucial importance in this paper. Thus we provide a proof here for the convenience of the reader.

LEMMA 2.1. *Let P be an $n \times m$ stochastic matrix. If we set $Z = \pi_{n-1}(P)$, then*

$$\text{rk}_+(P) = \min_t \{t : Z \subset \mathcal{P}_t \subset \Delta^{n-1}\}$$

where \mathcal{P}_t is the convex hull of t points.

Proof. If $\text{rk}_+(P) = k$, then $P = c_1(r_1)^t + \dots + c_k(r_k)^t$, where the vectors c_h and r_h have non-negative entries. As P is a stochastic matrix, we can consider the dyads $\frac{c_h}{\|c_h\|_1}(\|c_h\|_1 r_h)^t$. Hence, without loss of generality, we can assume that the following hold:

- $\|c_h\|_1 = 1$, for all $1 \leq h \leq k$;

- $(r_h)^t = (r_h^{(1)}, \dots, r_h^{(m)})$ and $\sum_h r_h^{(j)} = 1$ for all $1 \leq j \leq m$.
- Let $(c_h)^t = (c_h^{(1)}, \dots, c_h^{(n)})$ and consider the points

$$Q_h = (c_h^{(1)}, \dots, c_h^{(n-1)}) \in \Delta^{n-1}.$$

It is now straightforward to check that $\pi_{n-1}(P)$ is contained in the convex hull of the points Q_1, \dots, Q_k . Hence, $\text{rk}_+(P) = k \geq \min_t \{t : Z \subset \mathcal{P}_t \subset \Delta^{n-1}\}$.

Conversely, if $k = \min_t \{t : Z \subset \mathcal{P}_t \subset \Delta^{n-1}\}$ then $\pi_{n-1}(P) = \{c_1, \dots, c_m\}$ is in the convex hull of points $Q_1, \dots, Q_k \in \Delta^{n-1}$. Namely

$$c_j = \sum_{h=1}^k \alpha_h^{(j)} Q_h$$

and $\sum_h \alpha_h^{(j)} = 1$. If we let $Q_h = (q_h^{(1)}, \dots, q_h^{(n-1)})$ then

$$P = \sum_{h=1}^k \begin{pmatrix} q_h^{(1)} \\ \vdots \\ q_h^{(n-1)} \\ 1 - \sum_{i=1}^{n-1} q_h^{(i)} \end{pmatrix} \begin{pmatrix} \alpha_h^{(1)} & \dots & \alpha_h^{(m)} \end{pmatrix}.$$

Hence $\min_t \{t : Z \subset \mathcal{P}_t \subset \Delta^{n-1}\} = k \geq \text{rk}_+(P)$ and the result follows. \square

Remark. One can also consider the NPP for $Z \subset \mathcal{P}_k \subset \Delta^{n-1} \cap H$ where H is the linear span of Z . This amounts to study the restricted non-negative rank as described in [13] and [10].

3. Upper-semicontinuity of non-negative rank. In this section we will use the ideas recalled in Section 2 to show that the non-negative rank is upper-semicontinuous in the topology given by the Frobenius norm.

Given a non-negative matrix $P = (p_{i,j}) \in \mathbb{R}_{\geq 0}^{nm}$ and $\epsilon > 0$ define the ball of center P and radius ϵ

$$B(P, \epsilon) = \left\{ N = (n_{i,j}) \in \mathbb{R}_{\geq 0}^{nm} : \sqrt{\sum (p_{i,j} - n_{i,j})^2} < \epsilon \right\}.$$

THEOREM 3.1. *Let P be an $n \times m$ non-negative matrix, without zero columns, such that $\text{rk}_+(P) = k$, then there exists a ball $B(P, \epsilon)$ such that $\text{rk}_+(N) \geq k$, for all $N \in B(P, \epsilon)$.*

Proof. We give a proof by contradiction. Suppose that for all natural numbers r there exists $N(r) \in B(P, \frac{1}{r})$ such that $\text{rk}_+(N(r)) = t < k$. Clearly, the limit of the sequence $N(r)$ is P . By hypothesis we know that there exist convex polytopes $\mathcal{P}(r) \subset \Delta^{n-1}$ such that

$$\pi_{n-1} \circ \theta(N(r)) \subset \mathcal{P}(r),$$

where each $\mathcal{P}(r)$ is the convex hull of the points

$$q_1(r), \dots, q_t(r) \in \Delta^{n-1}.$$

We now *claim* that there exists a limit polytope $\bar{\mathcal{P}} \subset \Delta^{n-1}$ which is the convex hull of points $\bar{q}_1, \dots, \bar{q}_t$ (possibly not distinct) obtained by the sequences $q_h(r)$. As $t < k$ it is enough to show that $\pi_{n-1} \circ \theta(P) \subset \bar{\mathcal{P}}$ to get a contradiction using Lemma 2.1.

Let

$$\pi_{n-1} \circ \theta(N(r)) = \{c_1(r), \dots, c_m(r)\}$$

and

$$\pi_{n-1} \circ \theta(P) = \{c_1, \dots, c_m\}$$

and notice that (possibly after reordering) the limit of $c_j(r)$ is c_j . Also notice that for each j we have

$$c_j(r) = \alpha_{j,1}(r)q_1(r) + \dots + \alpha_{j,t}(r)q_t(r)$$

where the coefficients $\alpha_{j,h}(r)$ vary in the compact set $[0, 1]$, i.e. c_j belongs to the convex hull of the points $q_1(r), \dots, q_t(r)$. Taking subsequences and passing to the limit (limits exist as our sequences have values in compact sets) for each j we get

$$c_j = \bar{\alpha}_{j,1}\bar{q}_1 + \dots + \bar{\alpha}_{j,t}\bar{q}_t$$

and hence $c_j \in \bar{\mathcal{P}}$, for $j = 1, \dots, m$. Thus a contradiction and the statement is proved.

Proof of the claim The sequences $q_h(r)$ have values in the compact set Δ^{n-1} and hence they each have converging subsequences. To show that a limit polytope exists, we proceed as follows. Take a subsequence of $q_1(r)$ and let \bar{q}_1 be its limit. Then, either $q_2(r)$ has a subsequence with limit $\bar{q}_2 \neq \bar{q}_1$ or it does not and, in this case, we set $\bar{q}_2 = \bar{q}_1$. In the latter case the limit polytope will be the convex hull of strictly less than t distinct points. Iterating the process we obtain points $\bar{q}_1, \dots, \bar{q}_t$ and their convex hull is the limit polytope $\bar{\mathcal{P}}$. \square

Thus, given a matrix P , we know that in a suitable neighborhood of P the non-negative rank can only increase, i.e. the non-negative rank is upper-semicontinuous.

Clearly each neighborhood of a matrix P contains matrices having the same non-negative rank of P . Consider, for example, the matrices λP for $\lambda \in \mathbb{R}$ close to one. But even more is true as shown in the following statement.

PROPOSITION 3.2 (Barycentric perturbation). *Let P be a non-negative $n \times m$ matrix, without zero columns, such that $\text{rk}(P) > 1$. For any $\epsilon > 0$ there exists $N_\epsilon \in B(P, \epsilon)$ such that*

$$\text{rk}_+(N_\epsilon) = \text{rk}_+(P)$$

and $N_\epsilon \neq \lambda P$ for any $\lambda \in \mathbb{R}$.

Proof. Clearly, it is enough to prove the result for small ϵ . Thus, we can assume that each matrix in $B(P, \epsilon)$ has non-negative rank at least $\text{rk}_+(P)$, i.e. ϵ is small enough for Theorem 3.1 to apply.

Let P have columns c_j and consider the vector $b = \frac{1}{m} \sum_{j=1}^m c_j$. Roughly speaking we consider the barycenter of the points $\pi_{n-1} \circ \theta(P)$. Then we consider the $n \times m$ matrix N_δ having the j -th column defined as

$$c_j + \delta(b - c_j),$$

for $\delta \in [0, 1]$. When δ moves from zero to one, the points $\pi_{n-1} \circ \theta(N_\delta)$ approach the barycenter b . Thus, by Lemma 2.1, we have that $\text{rk}_+(N_\delta) \leq \text{rk}_+(P)$ for $\delta \in [0, 1]$.

By choosing δ small enough we get $N_\delta \in B(P, \epsilon)$. Hence we have that $\text{rk}_+(N_\delta) = \text{rk}_+(P)$. Letting $N_\epsilon = N_\delta$ the existence part of the proof is done.

To complete the proof, we only have to show that N_ϵ and P are not proportional. If $N_\epsilon = \lambda P$ for some $\lambda \in \mathbb{R}$, then, by the construction of N_ϵ , either all the columns c_j are proportional or $b = 0$. As $\text{rk}(P) > 1$ the matrix cannot have proportional columns. Moreover, P is non-negative, thus b cannot be the zero vector. Hence, if $N_\epsilon = \lambda P$, we get a contradiction and this completes the proof. \square

4. Jacobian approach. Throughout this section we assume $k \leq \min\{n, m\}$ and we let $X_{n \times m, k} \subset \mathbb{R}^{mn}$ be the variety of $n \times m$ matrices of rank at most k . It is well-known that $\dim(X_{n \times m, k}) = k(n + m - k)$, e.g. see [12, Proposition 12.2].

Consider the map $f : \mathbb{R}^{k(n+m)} \rightarrow X_{n \times m, k} \subseteq \mathbb{R}^{mn}$ which sends the point

$$p = (x_{1,1}, \dots, x_{1,n}, y_{1,1}, \dots, y_{1,m}, \dots, x_{k,1}, \dots, x_{k,n}, y_{k,1}, \dots, y_{k,m})$$

to the matrix

$$f(p) = \sum_{h=1}^k \begin{pmatrix} x_{h,1} \\ \vdots \\ x_{h,n} \end{pmatrix} (y_{h,1} \quad \dots \quad y_{h,m}). \quad (4.1)$$

Let f_+ be the restriction of f to the non-negative orthant $\mathbb{R}_{\geq 0}^{k(n+m)}$. The image of f_+ is the set $X_{n \times m, k}^+$ of $n \times m$ matrices of non-negative rank at most k . It is clear that $X_{n \times m, k}^+ \subseteq X_{n \times m, k}$.

Remark. We let $f_+^*(p)$ be the Jacobian matrix of f_+ at p and we say that $f_+^*(p)$ has maximal rank if its rank is $k(n + m - k)$. Thus, if $f_+^*(p)$ has maximal rank, then the map f_+ is locally surjective at p , e.g. see [23, page 25 Corollary (d)].

We can use this Jacobian approach to investigate properties of the non-negative rank under perturbations preserving the rank.

PROPOSITION 4.1 (Isorank perturbation). *Let P be an $n \times m$ non-negative matrix such that $\text{rk}_+(P) = k$ and consider the map f as defined in (4.1). If $P = f_+(p)$ is such that $f_+^*(p)$ has maximal rank and p has positive coordinates, then there exists a ball $B(P, \epsilon)$ such that for each $N \in B(P, \epsilon)$ we have:*

$$\text{if } \text{rk}(N) = \text{rk}(P), \text{ then } \text{rk}_+(N) = \text{rk}_+(P).$$

Proof. By the hypothesis we get that f_+ is locally surjective at p . Hence, there exist balls $B(P, \epsilon)$ and $B(p, \delta)$ such that each $N \in B(P, \epsilon) \cap X_{n \times m, k}$ has a preimage in $B(p, \delta)$ using the map f_+ . Moreover, if p has positive coordinates we can find, possibly smaller, ϵ and δ such that $B(p, \delta)$ is in the positive orthant. Thus, given $N \in B(P, \epsilon) \cap X_{n \times m, k}$ there exists q with positive coordinates such that $f(q) = N$. Hence $\text{rk}_+(N) \leq k = \text{rk}_+(P)$. The conclusion follows by Theorem 3.1 by taking an ϵ small enough. \square

Remark. The proof above also shows that there exists a neighborhood U of P such that each matrix in U of rank at most k has non-negative rank at most k . In other words, if $P = f_+(p)$, $f_+^*(p)$ has maximal rank and p has non-negative coordinates, then $X_{n \times m, k}^+ \cap U = X_{n \times m, k} \cap U$ for U a suitable neighborhood of P .

Now we describe sufficient conditions on p granting that $f_+^*(p)$ has maximal rank.

THEOREM 4.2. *If $p \in \mathbb{R}^{k(n+m)}$ is a point with coordinates*

$$p = (x_{1,1}, \dots, x_{1,n}, y_{1,1}, \dots, y_{1,m}, \dots, x_{k,1}, \dots, x_{k,n}, y_{k,1}, \dots, y_{k,m}),$$

such that

- $(x_{h,1}, \dots, x_{h,n})$, $h = 1, \dots, k$ are linearly independent vectors of \mathbb{R}^n ;
 - $(y_{h,1}, \dots, y_{h,m})$, $h = 1, \dots, k$ are linearly independent vectors of \mathbb{R}^m ;
- then $\text{rk}(f_+^*(p)) = k(n + m - k)$.

Proof. We may assume that $n \leq m$. Since the Jacobian is given by all possible derivatives with respect to $x_{h,i}$ and $y_{h,j}$, it is enough to show that exactly $k(m + n - k)$ of them are linearly independent. First of all we notice that the derivative with respect to $x_{h,i}$ is a matrix of the form

$$f_{x_{h,i}} = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} (y_{h,1} \quad \dots \quad y_{h,m}),$$

where the one is in position i . Similarly, the derivative with respect to $y_{h,j}$ is a matrix of the form

$$f_{y_{h,j}} = \begin{pmatrix} x_{h,1} \\ \vdots \\ x_{h,n} \end{pmatrix} (0 \quad \dots \quad 0 \quad 1 \quad 0 \quad \dots \quad 0),$$

where the one is in position j . That is, the derivative with respect to $x_{h,i}$, $i = 1, \dots, n$ is a matrix with all zeros but the i -th row consisting of the vector $(y_{h,1}, \dots, y_{h,m})$. Similarly the derivative with respect to $y_{h,j}$, $j = 1, \dots, m$ is a matrix with all zeros but the j -th column consisting of the vector $(x_{h,1}, \dots, x_{h,n})$.

We now build a set consisting of $k(m + n - k)$ linearly independent derivatives and hence we prove the statement.

The derivatives $f_{x_{h,i}}$, $i = 1, \dots, n$ are clearly linearly independent and we let \mathcal{S} be the kn -dimensional vector space that they span. Thus, if $m = k$ we are done.

If $m < k$ we proceed as follows. Let $V = \langle (y_{h,1}, \dots, y_{h,m}), h = 1, \dots, k \rangle$ where $\dim V = k$ by hypothesis. Now consider all the vectors in \mathbb{R}^m with at most one non-zero component and notice that they span a vector space of dimension $m > k$. Hence, V can not contain all these vectors and we may assume that $(1, 0, \dots, 0) \notin V$. Thus it is easy to see that the linear span

$$\mathcal{S}_1 = \langle \mathcal{S}, f_{y_{h,1}} \text{ such that } 1 \leq h \leq k \rangle$$

is such that $\dim \mathcal{S}_1 = \dim \mathcal{S} + k = kn + k$.

If $m - k = 1$ we are done. If $m - k > 1$ we argue as above. Namely, as $\dim V = k$, V can not contain all vectors with first component one and at most one more non-vanishing component. In particular, we may assume that $(1, *, 0, \dots, 0) \notin V$, where $*$ is any non-zero real number. Then we consider

$$\mathcal{S}_2 = \langle \mathcal{S}_1, f_{y_{h,2}} \text{ such that } 1 \leq h \leq k \rangle$$

and we readily see that $\dim \mathcal{S}_2 = \dim \mathcal{S}_1 + k = kn + 2k$.

For each $j \leq m - k$ we can repeat the process above increasing the dimension of \mathcal{S}_j by k each time. Hence, we can construct \mathcal{S}_{m-k} such that it is spanned by derivatives and $\dim \mathcal{S}_{m-k} = k(n + m - k)$. The statement is now proved. \square

5. Examples. In this section we will present some interesting examples. Some of these examples were inspired to us by the matrix

$$\begin{pmatrix} 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}$$

which is the most well-known example of a matrix with rank three and non-negative rank four (see [5]).

Small cases. Let P be an $n \times m$ matrix and assume $n \leq m$. We want to describe how the non-negative rank of P changes under perturbations for small values of n . If $n \leq 3$, then it is easy to show that $\text{rk}(P) = \text{rk}_+(P)$, see [5]. Thus the first interesting cases are for $n = 4$. If $\text{rk}_+(P) = 4$ then, by Theorem 3.1, any small perturbation will not change the non-negative rank. Thus, let us assume that $\text{rk}_+(P) = 3$. Using Proposition 3.2, we know that there are small perturbations preserving the non-negative rank. Of course, there are small perturbations not preserving it: it is enough to increase the ordinary rank. Hence, we ask: are there small perturbations of P , say P_ϵ , such that $\text{rk}(P_\epsilon) = \text{rk}(P)$ and $\text{rk}_+(P_\epsilon) = 4$? Not surprisingly, the answer depends on the choice of P . It is easy to construct a matrix P with the required ranks and satisfying the hypothesis of Proposition 4.1. Thus, in this case, the answer to our question is no. But, for a different choice of P , the answer can be yes. Consider, for example, P_ϵ defined as follows:

$$P_\epsilon = \frac{1}{4} \begin{pmatrix} 2 & 0 & 2 & 1 - \epsilon \\ 0 & 2 & 0 & 1 + \epsilon \\ 0 & 0 & 2 & 1 + \epsilon \\ 2 & 2 & 0 & 1 - \epsilon \end{pmatrix},$$

and let $P = P_0$. It is easy to see that $\text{rk}(P_\epsilon) = 3$ for all ϵ while $\text{rk}_+(P_0) = 3$ and $\text{rk}_+(P_\epsilon) = 4$ for small positive values of ϵ . To see this we use the graphical presentation in Figure 5.1 where we denote with c_1, \dots, c_4 the points corresponding to the columns of P while $c_4(\epsilon)$ corresponds to the fourth column of P_ϵ . In Figure 5.1, and in the following figures, we use the graphic representation described in [16] and related to the map π_3 . More precisely, a 4×4 matrix will be presented as a set of four points in a tetrahedron. This presentation allows for an easy visualization of rank related properties. We notice, for example, that a rank three matrix will correspond to four coplanar points.

Failing of upper-semicontinuity. The upper-semicontinuity of the non-negative rank is of course a local property as shown by the following example. Consider the matrix

$$M_\epsilon = \frac{1}{2(1+2\epsilon)} \begin{pmatrix} 1+\epsilon & \epsilon & 1+\epsilon & \epsilon \\ 1+\epsilon & \epsilon & \epsilon & 1+\epsilon \\ \epsilon & 1+\epsilon & 1+\epsilon & \epsilon \\ \epsilon & 1+\epsilon & \epsilon & 1+\epsilon \end{pmatrix}$$

and let $c_1(\epsilon), \dots, c_4(\epsilon)$ be the four column vectors where we set $c_j = c_j(0)$, $i = 1, \dots, 4$. When $\epsilon = 0$ the matrix has non-negative rank equal to four. We use the map π_3 to represent the columns in the simplex Δ^3 , which is a tetrahedron in \mathbb{R}^3 . To simplify the drawings, we have dropped the first coordinate of each column instead of the last

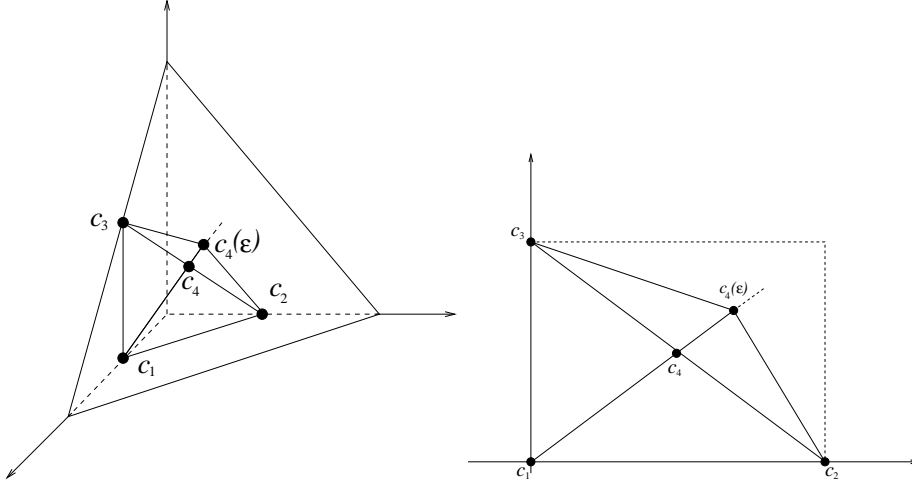


FIG. 5.1. The matrices P_ϵ for $\epsilon = 0$ and a small positive value of ϵ represented in the tetrahedron and in the plane.

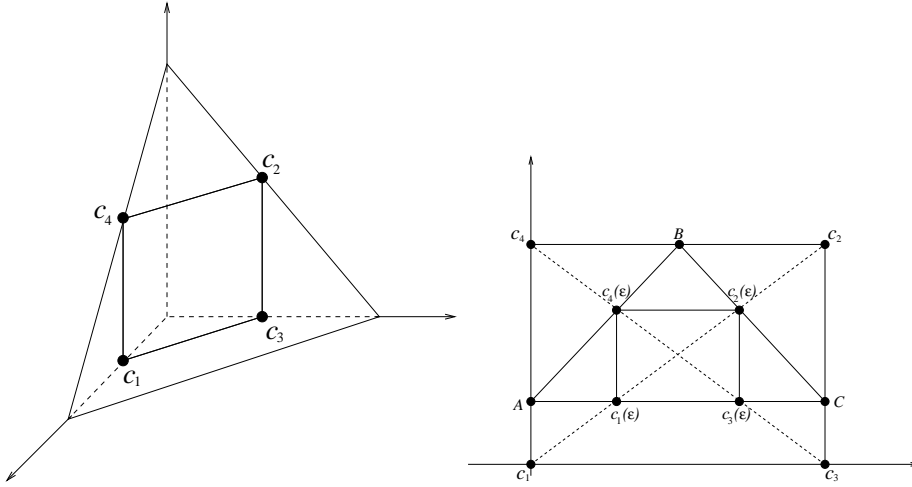


FIG. 5.2. The matrix M_0 in the simplex Δ^3 (left) and the points $c_1(\epsilon), \dots, c_4(\epsilon)$ in the critical configuration for $\epsilon = \sqrt{2}/2$ (right).

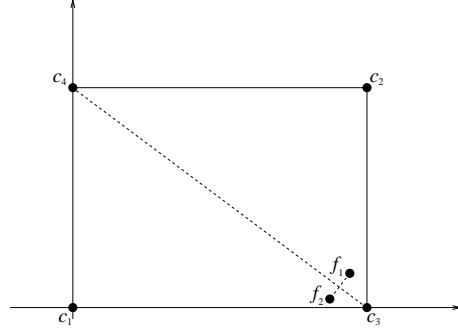
one, but of course this does not affect our analysis. The four points for the matrix M_0 are plotted in Figure 5.2 (left).

The points $c_1(\epsilon), \dots, c_4(\epsilon)$ are the vertices of a rectangle R_ϵ which we can draw in the plane. As $\epsilon > 0$ increases, the four points move along the main diagonals, as in Figure 5.2 (right), and R_ϵ will eventually be contained in the triangle ABC where

$$A = (0, \sqrt{2}/2 - 1/2) \quad B = (\sqrt{2}/4, 1/2) \quad C = (\sqrt{2}/2, \sqrt{2}/2 - 1/2).$$

It is not hard to show that, for $\epsilon < \sqrt{2}/2$, we have $\text{rk}_+(M_\epsilon) = 4$ while $\text{rk}_+(M_{\sqrt{2}/2}) = 3$. Hence, moving far enough from M_0 the non-negative rank can decrease.

Non-convexity of $X_{4 \times 4, 3}^+$. In the 4×4 case, the properties of the non-negative rank imply that the unique non-trivial case is the case of rank 3. The matrices in $X_{4 \times 4, 3}$

FIG. 5.3. The points $c_1, \dots, c_4, f_1, f_2$ defining the matrices A_1 and A_2 .

can belong to $X_{4 \times 4, 3}^+$ or to $X_{4 \times 4, 4}^+ \setminus X_{4 \times 4, 3}^+$. With the same graphical approach as above, we can show that the set $X_{4 \times 4, 3}^+$ is not convex (even if the ordinary rank is constant). To do this, it is enough to consider the two matrices $A_1 = [c_4, c_2, c_3, f_1]$ and $A_2 = [c_4, c_3, c_1, f_2]$ where the columns $c_1, c_2, c_3, c_4, f_1, f_2$ are displayed in Figure 5.3 in the same plane as in Figure 5.2 (right). It is immediate to see that both A_1 and A_2 have rank 3 and non-negative rank 3, but the matrix $A = (A_1 + A_2)/2$ has rank 3 (its 4 points are coplanar) but non-negative rank 4. With the same technique, one can also see that the set $X_{4 \times 4, 3} \setminus X_{4 \times 4, 3}^+$ is not convex.

$$B_1 = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix} \quad B_2 = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \end{pmatrix}.$$

B_1 and B_2 have rank 3 and non-negative rank 4, as they are obtained from M_0 possibly with permutation of columns, but the matrix $B = (B_1 + B_2)/2$ has non-negative rank 3.

6. Relations with the analysis of statistical mixture models. The results about the non-negative rank presented above have a useful counterpart in Probability and Mathematical Statistics. In particular the notion of nonnegative rank is useful in the study of mixture of independence models for discrete distributions. We now recall some basic definitions.

Distribution. The distribution (or density) of a random variable X on a set of n possible outcomes $\{1, \dots, n\}$ is a vector of n non-negative numbers (p_1, \dots, p_n) such that

$$p_i \geq 0 \text{ for all } i \text{ and } \sum_i p_i = 1,$$

where $p_i = \mathbb{P}(X = i)$ is the probability that X assumes the value i .

Joint distribution. If we consider a pair (X, Y) of random variables on $\{1, \dots, n\}$ and $\{1, \dots, m\}$ respectively, the joint distribution of X and Y is a probability matrix, i.e. a non-negative matrix $P = (p_{i,j})$ such that

$$p_{i,j} \geq 0 \text{ for all } i, j \text{ and } \sum_{i,j} p_{i,j} = 1, \quad (6.1)$$

where $p_{i,j} = \mathbb{P}(X = i, Y = j)$ is the probability that $(X = i)$ and $(Y = j)$.

Probability models. A matrix P satisfying the constraints in Equation (6.1) is also called a two-way table. The set

$$\Delta = \left\{ P \in \mathbb{R}^{nm} : p_{i,j} \geq 0 \text{ for all } i, j \text{ and } \sum_{i,j} p_{i,j} = 1 \right\}$$

is the $n \times m$ (closed) standard simplex and each probability distribution for a pair (X, Y) belongs to Δ . A probability model \mathcal{M} is a subset of Δ . In many cases \mathcal{M} is defined through a set of polynomial equations, and in such case we call \mathcal{M} an algebraic model.

The independence model. For two-way tables, one among the most simple models is the independence model. The construction of the independence model is described for instance in [1]. Under independence of X and Y we have

$$\mathbb{P}(X = i, Y = j) = \mathbb{P}(X = i)\mathbb{P}(Y = j)$$

for all $i = 1, \dots, n$ and for all $j = 1, \dots, m$, and therefore P is a rank one matrix, i.e., there exist vectors r and c such that $P = c(r)^t$. Thus, the independence model for $n \times m$ tables is the set:

$$\mathcal{M}_I = \{P : \text{rank}(P) = 1\} \cap \Delta.$$

Remark. It is a well known fact in Linear Algebra that a non-zero matrix P has rank 1 if and only if all 2×2 minors of P vanish. This shows that the independence model is an algebraic model. Thus, an equivalent definition of the independence model is as follows. The independence model is the set:

$$\mathcal{M}_I = \{P : p_{i,j}p_{k,h} - p_{i,h}p_{k,j} = 0 \text{ for all } 1 \leq i < k \leq n, 1 \leq j < h \leq m\} \cap \Delta.$$

Notice that the model is defined through pure binomials and that the set of all the 2×2 minors of a matrix are a system of generator of a toric ideal. This is a general fact in the analysis of algebraic statistical models and the models of this form are called toric models. The reader can refer to [7] and [20] for further details.

Mixture models. The mixture of two independence models is defined through the following procedure:

- Take two distributions $P_1, P_2 \in \mathcal{M}_I$;
- Toss a (biased) coin and choose P_1 with probability α and P_2 with probability $(1 - \alpha)$.

It is clear that the resulting distribution is a convex combination of P_1 and P_2 , i.e., a matrix of the form $\alpha P_1 + (1 - \alpha)P_2$. This process can be generalized. We can consider k distributions $P_1, \dots, P_k \in \mathcal{M}_I$ and define the mixture of k independence models as follows. The mixture of k independence models is the set

$$\mathcal{M}_{kI} = \{P : P = \alpha_1 c_1(r_1)^t + \dots + \alpha_k c_k(r_k)^t\}, \quad (6.2)$$

where the vectors r_h , the vectors c_h and $\alpha = (\alpha_1, \dots, \alpha_k)$ are probability distributions, i.e. all the components are non-negative and each vector has sum one. Some results and examples about this type of statistical models are presented in [8].

Notice that in the decomposition in Equation (6.2), the components must be non-negative, and therefore the model coincides with the set of $n \times m$ matrices with non-negative rank at most k and with sum equal to one, i.e., $\mathcal{M}_{kI} = X_{n \times m, k}^+$.

Maximum Likelihood Estimation. The problem of MLE consists in finding the global maxima of a suitable function, called likelihood, $L : \mathcal{M}_{kI} \rightarrow \mathbb{R}$. This problem is of great relevance and it has stimulated a lot of research, both from the theoretical and the numerical point of view. The interested reader can find a summary in [17]. In the case of mixture models, MLE is quite difficult mainly for two reasons: (a) while for a large class of statistical models the likelihood is a concave function, for mixture model this is not true; (b) the natural parametrization of the model as in Eq. (6.2) is redundant, see [4] for more on this. The papers [8] and [11] present some *ad hoc* solutions for mixture models with special interest in applications. Hence, a precise investigation of the geometric structure of the statistical model is essential to handle the maximization problem.

From the geometric investigation carried out in the previous section, we get a negative result. It is not possible to approximate a probability matrix P , with $\text{rk}_+(P) \neq \text{rk}(P)$, using matrices with ordinary rank and non-negative rank which coincide. This fact is summarized in the following corollary.

COROLLARY 6.1. *Given n, m and $2 < k < \min\{n, m\}$, the set \mathcal{M}_{kI} is not dense in $X_{n \times m, k}$.*

Proof. This result is a straightforward application of Proposition 3.1. Let P be a non-negative matrix such that $\text{rk}(P) = k$ and $\text{rk}_+(P) > k$. Then there is no sequence of matrices P_n whose limit is P such that $\text{rk}(P_n) = \text{rk}_+(P_n) = k$. \square

We consider Corollary 6.1 a negative result in the following sense: to study MLE on mixture models one must consider matrices with non-negative rank different from the ordinary rank. In particular, it is necessary to investigate the (not clear and not trivial) geometry of the set $X_{n \times m, k}^+$ and of its boundary. This study is necessary in order to be able to exploit optimization techniques and to avoid redundant variables.

Acknowledgments. The authors thank the two anonymous referees and the editor for their suggestions and comments which led to a better paper.

REFERENCES

- [1] ALAN AGRESTI, *Categorical Data Analysis*, Wiley, New York, 2 ed., 2002.
- [2] ELIZABETH ALLMAN AND JOHN A. RHODES, *Molecular phylogenetics from an algebraic viewpoint*, Statist. Sinica, 17 (2007), pp. 1299–1316.
- [3] LEROY B. BEASLEY AND THOMAS J. LAFFEY, *Real rank versus nonnegative rank*, Linear Algebra Appl., 431 (2009), pp. 2330–2335.
- [4] ENRICO CARLINI AND FABIO RAPALLO, *Probability matrices, non-negative rank, and parametrization of mixture models*, Linear Algebra Appl., 433 (2010), pp. 424–432.
- [5] JOEL E. COHEN AND URIEL G. ROTHBLUM, *Nonnegative ranks, decompositions, and factorizations of nonnegative matrices*, Linear Algebra Appl., 190 (1993), pp. 149–168.
- [6] BO DONG, MATTHEW M. LIN, AND MOODY T. CHU, *Nonnegative rank factorization via rank reduction*. Preprint, 2008.
- [7] MATHIAS DRTON, BERND STURMFELS, AND SETH SULLIVANT, *Lectures on Algebraic Statistics*, Birkhauser, Basel, 2009.
- [8] STEPHEN E. FIENBERG, PATRICIA HERSH, ALESSANDRO RINALDO, AND YI ZHOU, *Maximum likelihood estimation in latent class models for contingency table data*, in Algebraic and Geometric Methods in Statistics, Paolo Gibilisco, Eva Riccomagno, Maria Piera Rogantin, and Henry P. Wynn, eds., Cambridge University Press, Cambridge, 2010, pp. 27–62.
- [9] PAOLO GIBILISCO, EVA RICCOMAGNO, MARIA PIERA ROGANTIN, AND HENRY P. WYNN, *Algebraic and Geometric Methods in Statistics*, Cambridge University Press, Cambridge, 2010.
- [10] NICOLAS GILLIS AND FRANÇOIS GLINEUR, *On the geometric interpretation of the nonnegative rank*. arXiv:1009.0880, 2010.
- [11] GÉRARD GOVAERT AND MOHAMED NADIF, *Latent block model for contingency table*, Comm. Statist. Th. Methods, 39 (2010), pp. 416–425.
- [12] JOE HARRIS, *Algebraic Geometry: A first course*, Springer, New York, 1992.

- [13] MICHEL HAZEWINDEL, *On positive vectors, positive matrices and the specialization ordering*, Tech. Report PM-R8407, Centre for Mathematics and Computer Science, Amsterdam, 1984.
- [14] NGOC-DIEP HO AND PAUL VAN DOOREN, *Non-negative matrix factorization with fixed row and column sums*, Linear Algebra Appl., 429 (2008), pp. 1020–1025.
- [15] DANIEL D. LEE AND H. SEBASTIAN SEUNG, *Algorithms for non-negative matrix factorization*, Adv. Neural Inform. Process. Systems, 13 (2001), pp. 556–562.
- [16] MATTHEW M. LIN AND MOODY T. CHU, *On the nonnegative rank of euclidean distance matrices*, Linear Algebra Appl., 433 (2010), pp. 681–689.
- [17] GEOFFREY MCLACHLAN AND DAVID PEEL, *Finite Mixture Models*, John Wiley and Sons, New York, 2000.
- [18] LIOR PACTER AND BERND STURMFELS, *Algebraic statistics for computational biology*, Cambridge University Press, New York, 2005.
- [19] V. PAÚL PAUCA, J. PIPER, AND ROBERT J. PLEMMONS, *Nonnegative matrix factorization for spectral data analysis*, Linear Algebra Appl., 416 (2006), pp. 29–47.
- [20] FABIO RAPALLO, *Toric statistical models: Parametric and binomial representations*, Ann. Inst. Statist. Math., 59 (2007), pp. 727–740.
- [21] RONALD W. SHONKWILER AND JAMES HEROD, *Mathematical Biology. An Introduction with Maple and Matlab*, Springer, New York, 2 ed., 2009.
- [22] STEPHEN A. VAVASIS, *On the complexity of nonnegative matrix factorization*, SIAM J. Optim., 20 (2009), pp. 1364–1377.
- [23] FRANK W. WARNER, *Foundations of differentiable manifolds and Lie groups*, Springer, New York, 1971.
- [24] CHENGXIANG ZHAI, ATULYA VELIVELLI, AND BEI YU, *A cross-collection mixture model for comparative text mining*, in Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, New York, 2004, pp. 27–62.