The Nonnegative Rank of a Matrix: Hard Problems, Easy Solutions*

Yaroslav Shitov†

Abstract. Using elementary linear algebra, we develop a technique that leads to solutions of two widely known problems on nonnegative matrices. First, we give a short proof of the result by Vavasis stating that the nonnegative rank of a matrix is NP-hard to compute. This proof is essentially contained in the paper by Jiang and Ravikumar, who discussed this topic in different terms fifteen years before the work of Vavasis. Second, we present a solution of the Cohen–Rothblum problem on rational nonnegative factorizations, which was posed in 1993 and remained open until now.

Key words. nonnegative matrix factorization, computational complexity

AMS subject classifications. 15A03, 15A23, 68Q17

DOI. 10.1137/16M1080999

1. Introduction. Let A be a matrix with nonnegative real entries. The *nonnegative rank* of A is the smallest k for which there exist k nonnegative rank-one matrices that sum to A. This concept is related to the *nonnegative matrix factorization* problem, which is the task of finding the optimal approximation of a given matrix with a matrix of given nonnegative rank. This problem has important applications in different branches of modern science, including data mining [13], combinatorial optimization [21], statistics [12], quantum mechanics [4], and many others.

This paper presents a combinatorial approach that leads to very short solutions of two widely known problems on nonnegative matrices, and one of these accomplishments is a proof that the nonnegative rank is NP-hard to compute. The paper [19] by Vavasis has become a standard reference to this result in applied mathematics, and the proof it contains is based on ingenious geometric considerations. The proof we present is short and does not require any special knowledge. When this article was about to be completed, the author learned that a similar proof is contained in the paper [10] by Jiang and Ravikumar. Their Lemma 3.3 is stated in different terms, but it is essentially the oldest reference to the NP-hardness of nonnegative rank we are aware of.

Another result that we obtain is a solution of the *Cohen–Rothblum problem*: Assume that A is a rational matrix with nonnegative rank k; do there exist k *rational* nonnegative rank-one matrices that sum to A? This problem was posed in the foundational paper [4] more than 20 years ago, and it has been widely discussed in the literature. Vavasis [19] demonstrated the connection between the Cohen–Rothblum problem and algorithmic complexity of nonnegative rank. Another notable applica-

tion of nonnegative ranks is the theory of extended formulations of polytopes [21], and the possible lack of optimal rational factorizations is a difficulty in this theory [15]. Kubjas, Robeva, and Sturmfels [12] consider this problem in the context of modern statistics; they also give a partial solution. We note the paper [6], which contains a solution of a similar problem for *positive semidefinite* rank. Our article contains an explicit example of a 21×21 matrix with integral entries that can be written as a sum of 19 nonnegative rank-one matrices but not as a sum of 19 rational nonnegative rank-one matrices. This gives a solution of the Cohen–Rothblum problem, which has been open until now.[1]

2. Our Technique. Let A be a nonnegative matrix with entries in a subfield $\mathbb{F} \subset \mathbb{R}$. The *nonnegative rank* of A *with respect to* $\mathbb{F}$ (denoted $\mathrm{rank}_{\mathbb{F}+} A$) is the smallest k for which A is a sum of k nonnegative rank-one matrices over $\mathbb{F}$. In particular, the quantity $\mathrm{rank}_{\mathbb{R}+}(A)$ is the usual nonnegative rank as defined in the introduction. Combinatorial methods of studying the nonnegative rank have a long history; see the recent survey [8], the older papers [7, 14], and references therein. Methods that employ zero patterns of matrices are still being developed and used in modern investigations [1, 9]. As we will see later, the combinatorial analysis of zero patterns plays a crucial role in our approach.

Let $\alpha_1, \ldots, \alpha_n$ be nonnegative real numbers. The $5 \times (n+4)$ matrix

$$\mathcal{B}(\alpha_1, \ldots, \alpha_n) = \begin{pmatrix} \alpha_1 & \ldots & \alpha_n & 1 & 1 & 1 & 1 \\ 1 & \ldots & 1 & 1 & 1 & 0 & 0 \\ 0 & \ldots & 0 & 0 & 1 & 1 & 0 \\ 0 & \ldots & 0 & 0 & 0 & 1 & 1 \\ 0 & \ldots & 0 & 1 & 0 & 0 & 1 \end{pmatrix}$$

will be particularly important in our consideration. Its lower-right 4×4 submatrix

$$\mathcal{B}_0 = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 1 \end{pmatrix}$$

appears in the note [18], which is one of the oldest references on the topic of nonnegative matrix factorizations. The reason why $\mathcal{B}_0$ is so important is that it has different conventional and nonnegative ranks, $\mathrm{rank}\,\mathcal{B}_0 = 3$ and $\mathrm{rank}_+\,\mathcal{B}_0 = 4$. In particular, the row of ones can be represented as a linear combination of the rows of $\mathcal{B}_0$ in different ways: one can take the sum of the first and third rows, or the sum of the second and fourth rows, or any convex combination of these two sums. This leads us to the following easy but important observation.

OBSERVATION 1. $\mathrm{Rank}_{\mathbb{F}+}\,\mathcal{B}(\alpha_1, \ldots, \alpha_n) = 4$ *if and only if* $\alpha_1 = \cdots = \alpha_n \in [0, 1] \cap \mathbb{F}$.

In other words, this matrix has the following interesting property. For any $\varepsilon \in (0, 1)$, the 5×5 matrices obtained from $\mathcal{B}(\varepsilon)$ by a small perturbation of the $(1, 1)$ entry have the same nonnegative rank as the $(1, 1)$ cofactor of $\mathcal{B}(\varepsilon)$. Basic results of

---

[1]Our solution of the Cohen–Rothblum problem appears on arXiv as the manuscript [17], which was uploaded concurrently with another solution of the problem (https://arxiv.org/abs/1605.06848; now available as [3]).

linear algebra show that such a property cannot hold for a conventional rank function of matrices over a field. In fact, the usual rank of a matrix equals the order of the largest nonsingular submatrix, so a small perturbation of an entry would lead to a matrix with rank greater than the rank of its cofactor. We use this distinction between the conventional and nonnegative rank functions in the following matrix completion problem: Given nonnegative matrices $A \in \mathbb{F}^{m \times n}$, $B \in \mathbb{F}^{m \times k}$, $c \in \mathbb{F}^{1 \times n}$ and a subset $I \subset \mathbb{F}$, what is the smallest nonnegative rank of

$$(2.1) \qquad \mathcal{A}(x) = \left(\begin{array}{c|c} A & B \\ \hline c & x \ldots x \end{array} \right)$$

provided that $x \in I$? We can reduce this completion version to the standard formulation of the nonnegative rank problem as follows.

THEOREM 2. *Let $A \in \mathbb{F}^{m \times n}$, $B \in \mathbb{F}^{m \times k}$, $c \in \mathbb{F}^{1 \times n}$ be nonnegative matrices, and let r, s be positive integers. We assume that either $\mathrm{rank}_{\mathbb{F}+} A \geqslant r$ or $k = 1$. We define*

$$\mathcal{G} = \left(\begin{array}{cc|c|cccc} & & & 0 & 0 & 0 & 0 \\ & A & B & \vdots & \vdots & \vdots & \vdots \\ & & & 0 & 0 & 0 & 0 \\ \hline & c & s \ldots s & 1 & 1 & 1 & 1 \\ \hline 0 \ldots 0 & 1 \ldots 1 & 1 & 1 & 0 & 0 \\ 0 \ldots 0 & 0 \ldots 0 & 0 & 1 & 1 & 0 \\ 0 \ldots 0 & 0 \ldots 0 & 0 & 0 & 1 & 1 \\ 0 \ldots 0 & 0 \ldots 0 & 1 & 0 & 0 & 1 \end{array} \right)$$

and $\mathcal{A}(x)$ as in (2.1). Then the inequality $\mathrm{rank}_{\mathbb{F}+} \mathcal{G} \leqslant r + 4$ holds if and only if there is a $\xi \in [s - 1, s] \cap \mathbb{F}$ such that $\mathrm{rank}_{\mathbb{F}+} \mathcal{A}(\xi) \leqslant r$.

Proof. Note that $\mathcal{G}$ is the sum of $\mathcal{A}(\xi)$ and $\mathcal{B}(s - \xi, \ldots, s - \xi)$ up to adding zero rows and columns to the latter matrices. Since $\mathrm{rank}_{\mathbb{F}+} \mathcal{B}(s - \xi, \ldots, s - \xi) = 4$ by Observation 1, the "if" direction follows immediately.

Let us prove the "only if" direction. First, we note that the inequalities

$$(2.2) \qquad \mathrm{rank}_{\mathbb{F}+} (A \mid B) \geqslant r \quad \text{and} \quad \mathrm{rank}_{\mathbb{F}+} \left(\begin{array}{c} A \\ c \end{array} \right) \geqslant r$$

hold trivially if $\mathrm{rank}_{\mathbb{F}+} A \geqslant r$. (We note that the block matrices involved in (2.2) have sizes $m \times (n + k)$ and $(m + 1) \times n$, respectively.) If $k = 1$, then $\mathcal{A}(x)$ can be obtained by adding a row (a column, respectively) to the matrix in the left (right, respectively) inequality of (2.2). Therefore, if (2.2) were not true, we would obtain $\mathrm{rank}_{\mathbb{F}+} \mathcal{A}(x) \leqslant r$ and complete the proof. So we can assume that the inequalities (2.2) are true.

Now we assume that $G_1, \ldots, G_{r+4}$ are nonnegative rank-one matrices that sum to $\mathcal{G}$. If there were a G_i with two nonzero red entries, then $\mathcal{G}$ would have a 2×2 submatrix with positive entries, two of which are red—but this is not the case. Therefore, any G_i contains at most one nonzero red entry, so there are at least four G_i's with nonzero red entries. We call these G_i's red, denote their sum by R, and observe that they have zero black entries. Using the first inequality in (2.2), we find that there are r nonred G_i's, and every one of them has a nonzero entry either in the block corresponding to

A or in the block corresponding to B. Using the second inequality in (2.2), we see that every nonred G_i contains a nonzero entry either in the block corresponding to A or in the block corresponding to c. The two previous sentences allow us to conclude that the nonred G_i's do not contribute to the green entries of $\mathcal{G}$. Taking into account Observation 1, we find that, for some $a \in [0,1]$, the matrix R equals $\mathcal{B}(a, \ldots, a)$ up to the zeros at black positions. Therefore, the top-left $(m+1) \times (n+k)$ submatrix of $\mathcal{G} - R$ is the matrix $\mathcal{A}(s-a)$ as in (2.1), which completes the proof. $\quad\square$

COROLLARY 3. *Let $\mathcal{A}(x)$, $\mathcal{G}$ be as in Theorem 2. If $k = 1$, then we have*

$$\operatorname{rank}_{\mathbb{F}+} \mathcal{G} = \min \operatorname{rank}_{\mathbb{F}+} \mathcal{A}(\xi) + 4,$$

where the minimum is taken over all $\xi \in [s-1, s] \cap \mathbb{F}$.

3. On the Algorithmic Complexity of the Nonnegative Rank. Corollary 3 allows us to construct a polynomial time reduction to nonnegative rank from the following problem, which is known (see [20]) as the *nonnegative rank of a partial* 0-1 *matrix*. Given a matrix X whose (i,j) entry can be $0, 1$, or x_{ij}, what is the smallest nonnegative rank of matrices that are obtained from X by assigning numbers in $[0,1]$ to the variables x_{ij}? We are going to prove that this problem is NP-hard.

Let G be a simple graph with vertex set V and edge set E. A *clique* in G is a subset $U \subset V$ such that $\{u_1, u_2\} \in E$ for all distinct $u_1, u_2 \in U$. The *clique covering number* of G is the smallest number of cliques needed to cover V. Let $X = X(G)$ be the matrix (whose rows and columns are indexed with vertices of G) defined by $X_{vv} = 1$, $X_{uv} = x_{uv}$ if u, v are adjacent, and $X_{uv} = 0$, otherwise.

Let $U_1, \ldots, U_c$ be cliques whose union is V, and assume without loss of generality that these cliques are disjoint. For all $i \in \{1, \ldots, c\}$, we define the matrix H^i as $H^i_{\alpha\beta} = 1$ if $\alpha, \beta \in U_i$ and $H^i_{\alpha\beta} = 0$, otherwise. Clearly, the matrix $H_1 + \cdots + H_c$ has nonnegative rank at most c and can be obtained from X by replacing the variables with zeros and ones.

Conversely, let $M^1, \ldots, M^r$ be nonnegative rank-one matrices whose sum can be obtained from X by replacing the variables with numbers in $[0,1]$. If one of the sets $V^l = \{v \in V \mid M^l_{vv} > 0\}$ was not a clique, we would find nonadjacent distinct vertices u, v in V^l. This would imply $X_{uv} = 0$ and contradict $M^l_{uv} > 0$. So we see that $V^1, \ldots, V^r$ are cliques, and their union is V. This proves that the clique covering number of G equals the smallest nonnegative rank of matrices that are obtained from $X(G)$ by replacing the variables with numbers in $[0,1]$. Therefore, since the clique covering number is NP-hard to compute (see [11]), so is the nonnegative rank for partial 0-1 matrices. The observation at the beginning of this section shows that the standard version of the nonnegative rank problem is NP-hard as well.

As is pointed out in the introduction, Lemma 3.3 in [10] contains essentially the same result: Jiang and Ravikumar prove the NP-hardness of the *normal set basis problem*, which is a reformulation of the more commonly known *biclique partition number* problem. We refer the reader to [2] for a discussion of these two formulations and further complexity results. There is also a straightforward correspondence between the biclique partition number of a bipartite graph and the nonnegative rank of its adjacency matrix; see Remark 6.4 in [14] and Lemma 2.2 in [7]. We hope that these references can help an interested reader to translate the proof of Lemma 3.3 in [10] into the language of matrix theory and come up with the NP-hardness proof for nonnegative rank.

4. Nonnegative Rank Depends on the Field. We proceed with a solution of the Cohen–Rothblum problem. To this end, we consider the matrix

$$\begin{pmatrix}
2 & 2 & 2 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\
1 & 2 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 2 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 1 & 2 & 2 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1
\end{pmatrix}$$

and denote it by $\mathcal{M}$. We are going to show that the rational and real nonnegative ranks of $\mathcal{M}$ are different. Our strategy is to apply Theorem 2 to $\mathcal{M}$, and we assume that $\mathbb{F}$ is a field as in section 2. Let us remove the last twelve rows and columns from $\mathcal{M}$, replace the $(1,1)$, $(3,4)$, $(4,5)$ entries by the variables a, b, c, and denote the resulting matrix by $\mathcal{M}_1(a, b, c)$. The threefold application of Corollary 3 implies

$$(4.1) \qquad \operatorname{rank}_{\mathbb{F}+} \mathcal{M} = \min_{a,b,c \in [1,2] \cap \mathbb{F}} \operatorname{rank}_{\mathbb{F}+} \mathcal{M}_1(a, b, c) + 12.$$

Observing that the 3×3 submatrix defined as the intersection of the second, third, and fourth rows and the first three columns of $\mathcal{M}$ has rank three, we apply Theorem 2 with $r = 3$ to $\mathcal{M}_1(a, b, c)$. Taking into account (4.1), we find that the inequality $\operatorname{rank}_{\mathbb{F}+} \mathcal{M} \leqslant 19$ holds if and only if there are $a, b, c, d \in [1, 2] \cap \mathbb{F}$ such that the matrix

$$\mathcal{C}(a, b, c, d) = \begin{pmatrix}
a & 2 & 2 & 1 & 0 \\
1 & 2 & 1 & 0 & 1 \\
0 & 0 & 1 & b & 0 \\
0 & 1 & 0 & 0 & c \\
0 & 1 & 1 & d & d
\end{pmatrix}$$

has nonnegative rank at most three with respect to $\mathbb{F}$. Basic tools of linear algebra allow us to show that the conventional rank of $\mathcal{C}(a, b, c, d)$ exceeds three unless $b = c = d = 1 \pm \sqrt{0.5}$ and $a = 2 - 1/b$. In particular, the rational nonnegative rank of $\mathcal{C}$ cannot be less than four, which implies $\operatorname{rank}_{\mathbb{Q}+} \mathcal{M} \geqslant 20$. On the other hand, for

$\alpha = 1 + \sqrt{0.5}$,

$$
\begin{pmatrix}
0 & 0 & 0 & 0 & 0 \\
0 & \alpha^{-1} & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & \alpha \\
0 & 1 & 0 & 0 & \alpha
\end{pmatrix}
+
\begin{pmatrix}
0 & 0 & \alpha^{-1} & 1 & 0 \\
0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & \alpha & 0 \\
0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & \alpha & 0
\end{pmatrix}
+
\begin{pmatrix}
\sqrt{2} & 2 & \sqrt{2} & 0 & 0 \\
1 & \sqrt{2} & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0
\end{pmatrix}
$$

is a representation of $\mathcal{C}(\sqrt{2}, \alpha, \alpha, \alpha)$ as a sum of three nonnegative rank-one matrices. This shows that $\operatorname{rank}_{\mathbb{R}_+} \mathcal{M} \leqslant 19$ and completes the proof that the rational and real nonnegative ranks of $\mathcal{M}$ are different.

Note Added in Proof. The idea of my construction comes from the earlier paper [16], which contains a similar NP-hardness proof but in the setting of Boolean matrices. As I have subsequently learned, Theorem 3.3 in [5] has the same formulation and proof as the result in [16] but in the language of graph theory. I feel sorry that I did not know about [5] when I was preparing [16] for publication, but such occurrences seem to be ubiquitous in modern mathematics and hard to avoid. (As said above, a similar situation arose with Lemma 3.3 in [10], which states the NP-hardness of nonnegative rank but not in the terms common in applied mathematics.)

REFERENCES

[1] M. CONFORTI, G. CORNUÉJOLS, AND G. ZAMBELLI, *Extended formulations in combinatorial optimization*, 4OR, 8 (2010), pp. 1–48. (Cited on p. 795)

[2] P. CHALERMSOOK, S. HEYDRICH, E. HOLM, AND A. KARRENBAUER, *Nearly tight approximability results for minimum biclique cover and partition*, in Proceedings of the 22nd Annual European Symposium on Algorithms (Wroclaw, 2014), 2014, pp. 235–246. (Cited on p. 797)

[3] D. CHISTIKOV, S. KIEFER, I. MARUŠIĆ, M. SHIRMOHAMMADI, AND J. WORRELL, *Nonnegative matrix factorization requires irrationality*, SIAM J. Appl. Alg. Geom., 1 (2017), pp. 285–307, https://doi.org/10.1137/16M1078835. (Cited on p. 795)

[4] J. E. COHEN AND U. G. ROTHBLUM, *Nonnegative ranks, decompositions, and factorizations of nonnegative matrices*, Linear Algebra Appl., 190 (1993), pp. 149–168. (Cited on p. 794)

[5] M. C. GOLUMBIC, T. HIRST, AND M. LEWENSTEIN, *Uniquely restricted matchings*, Algorithmica, 31 (2001), pp. 139–154. (Cited on p. 799)

[6] J. GOUVEIA, H. FAWZI, AND R. Z. ROBINSON, *Rational and real positive semidefinite rank can be different*, Oper. Res. Lett., 44 (2016), pp. 59–60. (Cited on p. 795)

[7] D. A. GREGORY, N. J. PULLMAN, K. F. JONES, AND J. R. LUNDGREN, *Biclique coverings of regular bigraphs and minimum semiring ranks of regular matrices*, J. Combin. Theory Ser. B, 51 (1991), pp. 73–89. (Cited on pp. 795, 797)

[8] S. FIORINI, V. KAIBEL, K. PASHKOVICH, AND D. O. THEIS, *Combinatorial bounds on nonnegative rank and extended formulations*, Discrete Math., 313 (2013), pp. 67–83. (Cited on p. 795)

[9] S. FIORINI, S. MASSAR, S. POKUTTA, H. R. TIWARY, AND R. DE WOLF, *Linear vs. semidefinite extended formulations: Exponential separation and strong lower bounds*, in Proceedings of

the Forty-Fourth Annual ACM Symposium on Theory of Computing, ACM, New York, 2012. (Cited on p. 795)

[10] T. Jiang and B. Ravikumar, *Minimal NFA problems are hard*, SIAM J. Comput., 22 (1993), pp. 1117–1141, https://doi.org/10.1137/0222067. (Cited on pp. 794, 797, 799)

[11] R. Karp, *Reducibility among combinatorial problems*, in Complexity of Computer Computations (Proc. Sympos., IBM Thomas J. Watson Res. Center, Yorktown Heights, NY, 1972), Plenum, New York, 1972, pp. 85–103. (Cited on p. 797)

[12] K. Kubjas, E. Robeva, and B. Sturmfels, *Fixed points of the EM algorithm and nonnegative rank boundaries*, Ann. Statist., 43 (2015), pp. 422–461. (Cited on pp. 794, 795)

[13] D. D. Lee and H. S. Seung, *Learning the parts of objects by non-negative matrix factorization*, Nature, 401 (1999), pp. 788–791. (Cited on p. 794)

[14] J. Orlin, *Contentment in graph theory: Covering graphs with cliques*, Indag. Math., 80 (1977), pp. 406–424. (Cited on pp. 795, 797)

[15] T. Rothvoss, *Some 0/1 polytopes need exponential size extended formulations*, Math. Program., 142 (2013), pp. 255–268. (Cited on p. 795)

[16] Ya. Shitov, *On the complexity of Boolean matrix ranks*, Linear Algebra Appl., 439 (2013), pp. 2500–2502. (Cited on p. 799)

[17] Ya. Shitov, *Nonnegative Rank Depends on the Field* II, preprint, https://arxiv.org/abs/1605.07173, 2016. (Cited on pp. 795, 799)

[18] L. B. Thomas, *Rank factorization of nonnegative matrices*, SIAM Rev., 16 (1973), pp. 393–394, https://doi.org/10.1137/1016064. (Cited on p. 795)

[19] S. A. Vavasis, *On the complexity of nonnegative matrix factorization*, SIAM J. Optim., 20 (2009), pp. 1364–1377, https://doi.org/10.1137/070709967. (Cited on p. 794)

[20] T. Watson, *Nonnegative rank vs. binary rank*, Chicago J. Theoret. Comput. Sci., 2 (2016), pp. 1–13. (Cited on p. 797)

[21] M. Yannakakis, *Expressing combinatorial optimization problems by linear programs*, J. Comput. System Sci., 43 (1991), pp. 441–466. (Cited on pp. 794, 795)