Lower bound computations for the nonnegative rank

Fiorini, Samuel; Guo, Krystal; MacChia, Marco; Walter, Matthias

Publication date
2019

Document Version
Final published version

Published in
Proceedings of the 17th Cologne-Twente Workshop on Graphs and Combinatorial Optimization

Lower Bound Computations for the Nonnegative Rank

Samuel Fiorini[1], Krystal Guo[1], Marco Macchia[1], and Matthias Walter[2]

[1]Université libre de Bruxelles, Brussels, Belgium
[2]University of Twente, Enschede, Netherlands

Abstract

The nonnegative rank of a matrix is the smallest inner dimension when writing it as a product of two nonnegative matrices. Such nonnegative matrix factorizations have numerous applications in machine learning and data mining, where one usually allows inexact factorizations. The exact counterpart is related to the so-called extension complexity of a polytope, an important parameter in combinatorial optimization.

We implemented different algorithms for computing lower bounds on the nonnegative rank of a matrix. In this extended abstract we focus on results that relate our best algorithms' performance to the size of the matrix.

1 Introduction

Let $M \in \mathbb{R}^{m \times n}_{\geq 0}$ be a nonnegative matrix. A *nonnegative factorization* of M is a product $M = X \cdot Y$ with $X \in \mathbb{R}^{m \times r}_{\geq 0}$ and $Y \in \mathbb{R}^{r \times n}_{\geq 0}$, where we call r the *rank of the factorization*. Expressing a given data matrix as $M \approx X \cdot Y$ with a small r is the (approximate) nonnegative matrix factorization problem. The exact version of the problem is to compute the *nonnegative rank* $\text{rk}_+(M)$ of M, defined as the smallest such r with $M = X \cdot Y$. This quantity has an important interpretation in polyhedral combinatorics [1], which we briefly outline.

The *extension complexity* of a polytope P is the smallest number of facets of any (typically higher-dimensional) polytope Q that affinely projects to P. The number of facets is equal to the (minimum) number of inequalities that are necessary to describe Q in a linear programming (LP) formulation. Suppose we have an inner description of $P \subseteq \mathbb{R}^d$ in terms of its vertices $x_1, x_2, \dots, x_n \in \mathbb{R}^d$ as well as an outer description in terms of linear inequalities $Ax \leq b$ with $A \in \mathbb{R}^{m \times d}$. Then the corresponding slack matrix $M \in \mathbb{R}^{m \times n}_{\geq 0}$ has entries $M_{i,j} \coloneqq b_i - A_{i,*} x_j$. Yannakakis [8] proved that $\text{rk}_+(M)$ is equal to the extension complexity of P. Since its computation is notoriously hard, in our paper we deal with the computation of lower bounds.

2 Lower bounds

Combinatorial bounds. Let $S \coloneqq \text{supp}(M) \coloneqq \{(i,j) \in [m] \times [n] : M_{i,j} > 0\}$ denote the support of M. We call an index set $R \subseteq [m] \times [n]$ a *rectangle* if $R \subseteq S$ and if it is of the form $R = I \times J$ for $I \subseteq [m]$ and $J \subseteq [n]$. We denote by $\mathcal{R}$ the set of all rectangles. Every nonnegative factorization with inner dimension r induces a covering of S with r rectangles (see [8, 2]), which establishes the *rectangle covering bound* $\text{rc}(M)$. It can be computed using the integer program (IP)

$$\min_{\lambda \in \mathbb{Z}^{\mathcal{R}}_{\geq 0}} \left\{ \sum_{R \in \mathcal{R}} \lambda_R : \sum_{R \in \mathcal{R} :\, s \in R} \lambda_R \geq 1 \ \forall s \in S \right\}. \tag{rc}$$

Since this is a minimization problem, the optimum value of its LP relaxation, called the *fractional rectangle covering bound* and denoted by $\mathrm{frc}(M)$, is also a lower bound on $\mathrm{rk}_+(M)$.

Both bounds can be refined slightly by considering entry pairs $s_1, s_2 \in S$ with $s_1 = (i_1, j_1)$ and $s_2 = (i_2, j_2)$. We call $\{s_1, s_2\}$ a *fooling pair* if $M_{i_1,j_1} \cdot M_{i_2,j_2} > M_{i_1,j_2} \cdot M_{i_2,j_1} > 0$ holds and denote by $\mathcal{P}$ the set of all fooling pairs. Although $\{i_1, i_2\} \times \{j_1, j_2\}$ is a rectangle, its induced submatrix of M has rank 2. Thus, at least one of the entries must be covered by two rectangles in any rectangle covering induced by a nonnegative factorization of M. This observation was made in [6] and establishes the *refined rectangle covering bound* $\mathrm{rrc}(M)$, defined via the IP

$$\min_{\lambda \in \mathbb{Z}^{\mathcal{R}}_{\geq 0}} \left\{ \sum_{R \in \mathcal{R}} \lambda_R : \sum_{R \in \mathcal{R} \colon s \in R} \lambda_R \geq 1 \ \forall s \in S, \sum_{R \in \mathcal{R} \colon s_1 \in R \vee s_2 \in R} \lambda_R \geq 2 \ \forall \{s_1, s_2\} \in \mathcal{P} \right\}. \tag{rrc}$$

Once again we define the bound based on the value of the its LP relaxation $\mathrm{frrc}(M)$, called the *fractional refined rectangle covering bound.*

Another combinatorial bound is the *fooling set bound* $\mathrm{fs}(M)$ that is derived from the packing counterpart of the rectangle covering problem It is defined as

$$\max \{|F| : F \subseteq S \text{ such that } |F \cap R| \leq 1 \text{ for all } R \in \mathcal{R}\}, \tag{fs}$$

and we call the feasible sets $F \subseteq S$ *fooling sets*. It is equal to the maximum size of a clique in an auxiliary graph in which nodes correspond to entries of S and two nodes are connected by an edge if and only if their entries are not contained in a rectangle.

Hyperplane separation bounds. In contrast to the combinatorial bounds the *hyperplane separation bound* also exploits the actual entries of M. It was suggested by Fiorini and used to establish an exponential lower bound on the extension complexity of the perfect matching polytope [7]. For its correctness we refer to Lemma 1 therein. It reads

$$\mathrm{hsb}(M) \coloneqq \max \left\{ \frac{\langle W, M \rangle}{||M||_\infty} : W \in \mathbb{R}^{m \times n}, \ \langle W, \chi(R) \rangle \leq 1 \text{ for all } R \in \mathcal{R} \right\}, \tag{hsb}$$

where $\chi(R) \in \{0,1\}^{m \times n}$ is the characteristic vector of R. Note that the entries $W_{i,j}$ with $(i,j) \notin S$ do not play a role in the model. By restricting $W_{i,j}$ to be nonnegative for all $(i,j) \in S$, we obtain the weaker version, called the *nonnegative hyperplane separation bound* and denoted by $\mathrm{nhsb}(M)$.

3 Computational study

In our software tool `nonnegrank` [3] we implemented all lower bounds from Section 2 (see Table 1). Our code is written in `C++` and relies on different external libraries. Several bound computations can be made more efficient in two ways. First, if the input matrix M has some symmetry, then the computational effort can often be reduced by reducing the dimension of the problem. Second, if the set $\mathcal{R}^{\max} \subseteq \mathcal{R}$ of inclusion-wise maximal rectangles is available explicitly, then several subproblems can be solved more quickly. Our software can compute $\mathcal{R}^{\max}$ in time polynomial in $|S|$ and $|\mathcal{R}^{\max}|$ using Ganter's Next Closure Algorithm [4]. Finally, in case this enumeration shall be avoided, we provide an implementation of (rc) and (rrc) using branch-and-price.

Table 1: Implemented bounds with their used external libraries, depending on availability of $\mathcal{R}^{\max}$ (column **Enum**) and information whether symmetry can be exploited (column **Symmetry**).

Bound	Enum	Implementation [libraries]	Symmetry
fs	No	Graph model for (fs) [`Cliquer`] or IP model [`Gurobi or SCIP`]	No
rc	Yes	Model (rc) over $\mathcal{R}^{\max}$ [`Gurobi or SCIP`]	No
	No	Branch-and-price for Model (rc) [`SCIP`]	No
rrc	Yes	Model (rrc) over $\mathcal{R}^{\max}$ [`Gurobi or SCIP`]	No
	No	Branch-and-price for Model (rrc) [`SCIP`]	No
frc	Yes	LP of Model (rc) over $\mathcal{R}^{\max}$ [`Gurobi or SoPlex`]	Yes
	No	LP of Model (rc) [`Gurobi or SCIP`]	Yes
frrc	Yes	LP of Model (rrc) over $\mathcal{R}^{\max}$ [`Gurobi or SoPlex`]	Yes
	No	LP of Model (rrc) [`Gurobi or SCIP`]	Yes
hsb	Yes/No	Model (hsb) [`Gurobi or SCIP`]	Yes
nhsb	Yes	Model (hsb) with $W_s \geq 0 \ \forall s \in S$ over $\mathcal{R}^{\max}$ [`Gurobi or SoPlex`]	Yes
	No	Model (hsb) with $W_s \geq 0 \ \forall s \in S$ [`Gurobi or SCIP`]	Yes

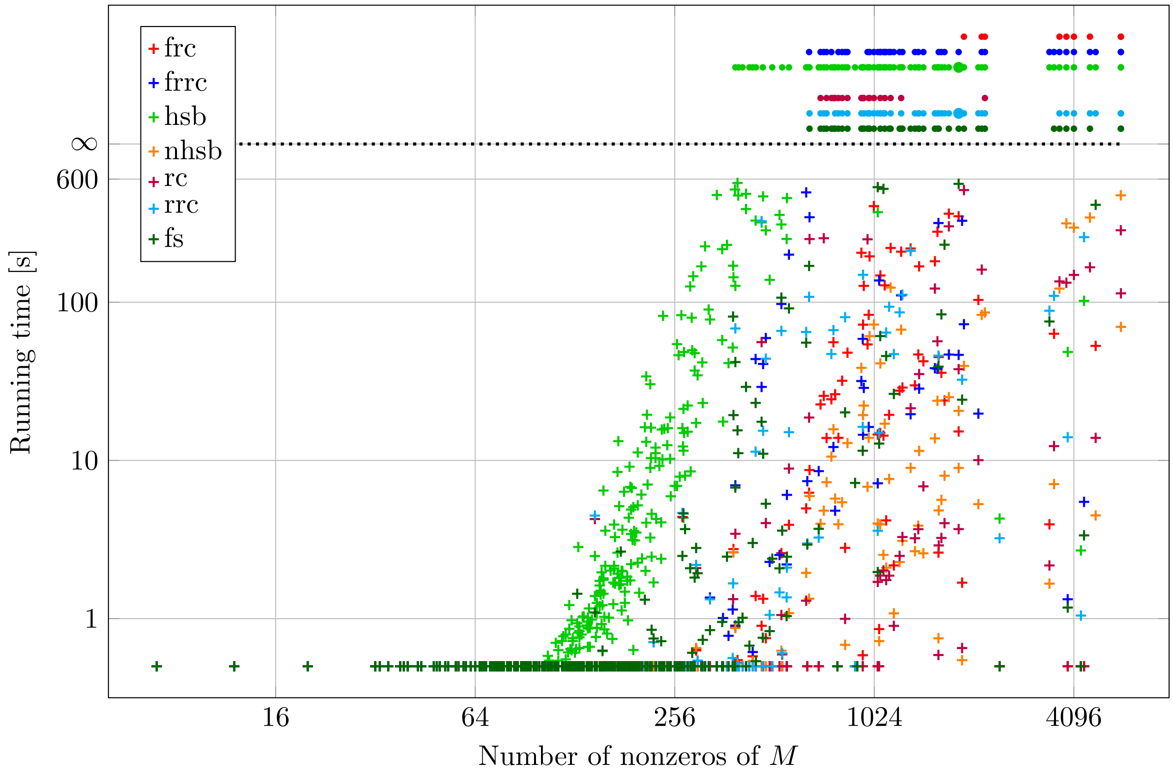

Figure 1: Running times for various bound computations. Time limit: 600 s. Balls above dotted line indicate number of timeouts. IP/LP solver was `Gurobi`, rectangle enumeration turned on, symmetry detection turned off.

Setup. In this extended abstract we report results only for two types of instances, both derived from slack matrices of polytopes with dimensions 6, 7 or 8. The first are random 0/1-polytopes with prescribed ambient dimension and expected number of vertices. The second set is derived from 2-level polytopes randomly chosen from the 2-level polytope database from [5]. Our test were run single threaded on a machine with 2.97 GHz and 8 GB memory.

Experimental results. For the sake of this extended abstract, we restrict ourselves to considering computational results corresponding to the aggregated dataset. We set a time limit of 600 s and use `Gurobi` as the undelying solver, since `SCIP` turned out to be generally slower on our instances. Moreover, we do not report detailed results for symmetry breaking algorithms since they were generally faster in case symmetry was present. On the one hand, we notice that slack matrices of 2-level polytopes exhibit considerable symmetries. On the other hand slack matrices of 0/1-polytopes did not present symmetries. When possible, we use the complete enumeration of rectangles in $\mathcal{R}^{\max}$ and provide the corresponding reduced running times. Figure 1 shows the resulting computation times as a function of the size of the support of the input nonnegative matrix. Moreover, it is easy to see that the computation of hsb is the most time-consuming, the reason being that a cutting plane method has to be adopted (even when $\mathcal{R}^{\max}$ is computed). It is interesting to notice that rc can be computed within the time limit of 600 s for many instances with support of size greater than 1000, even though it requires to solve an IP. The slightly more complicated bound rrc is harder to compute, involving contraints corresponding to fooling pairs.

Future work. In the full version of this abstract, we plan to have a detailed comparison of the lower bounds in Section 2, in order to provide a complete evaluation of the quality of the bounds, and compare it with the running times.

Acknowledgements. The first, the second and the third authors are supported by ERC grant *FOREFRONT* (grant agreement no. 615640) funded by the European Research Council under the EU's 7th Framework Programme (FP7/2007-2013).

References

[1] Michele Conforti, Gérard Cornuéjols, and Giacomo Zambelli. Extended formulations in combinatorial optimization. *Annals of Operations Research*, 204(1):97–143, 2013.

[2] Samuel Fiorini, Volker Kaibel, Kanstantsin Pashkovich, and Dirk O. Theis. Combinatorial bounds on nonnegative rank and extended formulations. *Discrete Mathematics*, 313(1):67–83, 2013.

[3] Samuel Fiorini, Marco Macchia, Krystal Guo, and Matthias Walter. nonnegrank – a tool to compute lower bounds on the nonnegative rank, 2019. Software page: `https://bitbucket.org/matthias-walter/nonnegrank/`.

[4] Bernhard Ganter and Klaus Reuter. Finding all closed sets: A general approach. *Order*, 8(3):283–290, 1991.

[5] Marco Macchia. *Two level polytopes: geometry and optimization.* PhD thesis, Université libre de Bruxelles, 2018.

[6] Michael Oelze, Arnaud Vandaele, and Stefan Weltge. Computing the extension complexities of all 4-dimensional 0/1-polytopes. arXiv:1406.4895, June 2014.

[7] Thomas Rothvoss. The matching polytope has exponential extension complexity. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pages 263–272, New York, NY, USA, 2014. ACM.

[8] Mihalis Yannakakis. Expressing combinatorial optimization problems by linear programs. *Journal of Computer and System Sciences*, 43(3):441–466, 1991.