

Computing approximate PSD factorizations

Amitabh Basu^{1,*}, Michael Dinitz^{2,**}, and Xin Li²

¹ Dept. of Applied Mathematics and Statistics, Johns Hopkins University

² Dept. of Computer Science, Johns Hopkins University

Abstract. We give an algorithm for computing approximate PSD factorizations of nonnegative matrices. The running time of the algorithm is polynomial in the dimensions of the input matrix, but exponential in the PSD rank and the approximation error.

1 Introduction

Matrix factorization is a fundamental operation that has importance for diverse areas of mathematics and engineering such as machine learning, communication complexity, polyhedral combinatorics, statistical inference, and probability theory, to name a few. The problem can be stated quite simply as follows:

Given two sequences of sets $\mathcal{K} = \{K_d\}_{d \in \mathbb{N}}$ and $\mathcal{K}' = \{K'_d\}_{d \in \mathbb{N}}$ where K_d, K'_d are subsets of \mathbb{R}^d for all $d \in \mathbb{N}$, and a matrix $M \in \mathbb{R}^{n \times m}$, find a factorization $M = UV$ where $U \in \mathbb{R}^{n \times d}$ and $V \in \mathbb{R}^{d \times m}$, and each row of U is in K_d and each column of V is in K'_d .

Such a factorization is called a $\mathcal{K}, \mathcal{K}'$ factorization. The smallest $d \in \mathbb{N}$ such that such a factorization exists is called the $\mathcal{K}, \mathcal{K}'$ rank. Most of the literature on this problem focuses on the case when the matrix M is nonnegative. In this context, when $K_d = K'_d = \mathbb{R}_+^d$, the factorization is called *nonnegative factorization*, and the corresponding rank is called *nonnegative rank*. When $K_d = K'_d$ are the cone of $d \times d$ PSD matrices, the factorization is known as a *PSD factorization* and the corresponding rank is called the *PSD rank*. These notions will be the object of study in this paper. A more general notion is that of *cone factorizations*, where \mathcal{K} is a family of cones and \mathcal{K}' is the family of corresponding dual cones; see [12].

One of the most elegant applications of such factorizations arises in combinatorial optimization. A very common technique in approaching combinatorial optimization problems is to formulate the problem as a linear programming problem. However, a naive formulation of a problem may result in a polytope (the feasible region of the LP) with a large number of facets (exponentially many in the size of the problem), making it intractable to actually solve. One way around this is to try to express the polytope as the projection of a higher dimensional convex set. In particular, suppose that it can be expressed as the

* Supported in part by the NSF grant CMMI1452820.

** Supported in part by NSF awards CCF-1464239 and CCF-1535887.

projection of either a higher dimensional polytope (LP), the feasible region of an SDP, or the feasible region of a more general convex optimization problem. Furthermore suppose that the number of “extra” dimensions is polynomial in the size of the original problem, and the description of the higher-dimensional convex optimization problem is also polynomial in the size of the original problem (i.e. there are not an exponential number of facets). Then we can efficiently solve the higher-dimensional problem, which means we can efficiently solve the original LP, even if its size makes solving it directly intractable.

It turns out that the smallest size of such a reformulation is a direct function of the nonnegative rank (for LP reformulations), the PSD rank (for SDP reformulations), or more general cone factorization ranks of the so-called *slack matrix* of the original LP formulation. The actual factorization can be used to explicitly find the smallest reformulation. This line of research started with a seminal paper by Yannakakis [22], and has recently seen a flurry of research activity – see the surveys [14, 8] and [11, 10, 19, 7, 15, 20, 16, 6] for some of the most recent breakthroughs.

In machine learning applications the actual factorization is perhaps more important than the value of the rank, as this factorization is key to certain text mining, clustering, imaging and bioinformatics applications. A key algorithmic question is compute such a factorization. Unfortunately, this question is computationally challenging – even computing the nonnegative rank was proved to be NP-hard by Vavasis [21].

A recent algorithmic breakthrough was achieved by Arora et al [1], where they showed that computing nonnegative factorizations can be done in polynomial time (in the dimensions of the input matrix) for the family of matrices with fixed (constant) nonnegative rank. The run time of their algorithm was doubly exponential in the nonnegative rank, and this was later improved to a singly exponential algorithm by Moitra [17], which he showed to be nearly optimal under the Exponential Time Hypothesis. The analogous question for PSD factorizations is largely open (the question is also posed in the survey [9]):

Question 1. Let $r \in \mathbb{N}$ be a constant. Does there exist an algorithm which, given any $n \times m$ nonnegative matrix M with PSD rank r , computes a PSD factorization of rank r in time polynomial in n, m ?

In the above question, and in the rest of the paper, we assume the real arithmetic model of computation developed by Blum, Shub and Smale [4], thus ignoring questions of approximating irrational computations by rational arithmetic. Our main result is a polynomial time algorithm to compute *approximate* factorizations of matrices with fixed PSD rank. We consider the space \mathcal{S}^r of $r \times r$ symmetric matrices, and the cone of $r \times r$ PSD matrices in this space, denoted by \mathcal{S}_+^r . Given any matrix $M \in \mathbb{R}^{n \times m}$, we use the notation $\|M\|_\infty := \max_{i,j} |M_{ij}|$.

More precisely, we prove:

Theorem 1. *Let $r \in \mathbb{N}$ be fixed. Then there exists an algorithm which, given any $\epsilon > 0$ and any $n \times m$ nonnegative matrix M with PSD rank r , computes a*

factorization $M = UV$ such that each row of U and each column of V are in \mathcal{S}_+^r such that

$$\|M - UV\|_\infty \leq \epsilon \|M\|_\infty$$

and has runtime polynomial in the dimensions of M .

Approximate PSD factorizations can be useful for reformulation questions in combinatorial optimization, where one seeks approximations of the original polyhedron using SDPs, as opposed to an exact reformulation – see [13] for results along this direction.

Technical overview. Our algorithm for Theorem 1 is inspired by ideas behind the algorithm in Arora et al [1]. However, there are some important differences. Arora et al’s algorithm uses properties of the nonnegative orthant that do not hold for the cone of PSD matrices. To overcome this difficulty, we need to use an approximation of the PSD cone by a polyhedral cone obtained by intersecting enough tangent halfspaces. We then generalize Arora et al’s techniques to compute factorizations inside a general polyhedron, as opposed to the nonnegative orthant. The nonnegative orthant is a very special polyhedron, and many of its special properties are utilized in the algorithm of Arora et al. We have to use interesting techniques from polyhedral theory (such as Fourier-Motzkin elimination) to extend these ideas to handle general polyhedra (see Theorem 4). Finally, to bound the errors in the approximate factorization, we use some technical results on rescaling PSD factorizations due to Briet et al [7] (Theorem 5).

2 Preliminaries

For any normed space $(V, \|\cdot\|)$, we denote the distance between two subsets $X, Y \subseteq V$ by $\text{dist}(X, Y) := \inf\{\|x - y\| : x \in X, y \in Y\}$. A closed subset P of a normed space V is called a *closed cone* if it is convex and $\lambda P \subseteq P$ for all $\lambda \geq 0$. A cone is called a *polyhedral cone* if it is the intersection of finitely many halfspaces. For any closed cone P in an inner product space $(V, \langle \cdot, \cdot \rangle)$ (closed with respect to the norm obtained from the inner product), the dual cone will be denoted by

$$P^* = \{v \in V : \langle v, y \rangle \geq 0 \ \forall y \in P\}.$$

We recall a standard fact about dual cones:

Fact 1 *Let $(V, \langle \cdot, \cdot \rangle)$ be an inner product space with $\|\cdot\|$ denoting the norm on V induced by the inner product. For any closed cone $P \subseteq V$, if $x \in V$ such that $\text{dist}(x, P) = \delta$, then there exists a vector $a \in P^*$ with $\|a\| = 1$ such that the distance of x from the hyperplane $\{y \in V : \langle a, y \rangle = 0\}$ is δ , i.e., $\langle a, x \rangle = -\delta$.*

On the space \mathcal{S}^r of $r \times r$ symmetric matrices, we consider the inner product $\langle A, B \rangle = \sum_{i,j} A_{ij} B_{ij}$.

Fact 2 *The PSD cone \mathcal{S}_+^r is self-dual, i.e., $(\mathcal{S}_+^r)^* = \mathcal{S}_+^r$.*

Definition 1. Let C be a subset of a normed space $(V, \|\cdot\|)$. For $\epsilon > 0$, $\mathcal{X}_\epsilon \subseteq C$ is called an ϵ -covering for C with respect to the norm $\|\cdot\|$ if for every $a \in C$, there exists $a' \in \mathcal{X}_\epsilon$ such that $\|a - a'\| < \epsilon$.

Definition 2. For any closed cones $P_1 \subseteq P_2$ in a normed space $(V, \|\cdot\|)$, we say P_2 is an ϵ -approximation of P_1 with respect to $\|\cdot\|$ for some $\epsilon > 0$, if for every $p_2 \in P_2$, there exists a point $p_1 \in P_1$ such that $\|p_2 - p_1\| \leq \epsilon\|p_2\|$.

Theorem 2. Let $C = \{x \in \mathcal{S}_+^r : \|x\|_2 = 1\}$ be the spherical cap on the PSD cone. Let $\epsilon > 0$ and let $\mathcal{X}_\epsilon \subseteq C$ be any finite ϵ -covering for C with respect to some norm $\|\cdot\|$. Then the polyhedral cone

$$P := \{x \in \mathcal{S}^r : \langle a', x \rangle \geq 0 \quad \forall a' \in \mathcal{X}_\epsilon\}$$

is an ϵ -approximation for \mathcal{S}_+^r with respect to $\|\cdot\|$.

Proof. It suffices to prove that for any $x \in \mathcal{S}^r$ such that $\text{dist}(x, \mathcal{S}_+^r) > \epsilon\|x\|$, then $x \notin P$. By Fact 1 (2) and Fact 2, there exists $a \in C$ such that $\langle a, x \rangle < -\epsilon\|x\|$. By definition of ϵ -covering, there exists $a' \in \mathcal{X}_\epsilon$ such that $\|a - a'\| < \epsilon$. By Cauchy-Schwartz, we have $|\langle a', x \rangle - \langle a, x \rangle| \leq \|a - a'\|\|x\| < \epsilon\|x\|$. Combined with $\langle a, x \rangle < -\epsilon\|x\|$, this implies $\langle a', x \rangle < 0$. Thus, by definition of P , $x \notin P$. \square

Remark 1. Since C is a compact set, there always exists a finite ϵ -covering of C for any $\epsilon > 0$. Rabani and Shpilka [18] give explicit constructions of small ϵ -coverings of the sphere $\mathcal{S}^{d-1} = \{x \in \mathbb{R}^d : \|x\|_2 = 1\}$, which will prove useful for us.

The following fact from linear algebra is useful.

Proposition 1 Any linear transformation $T : \mathbb{R}^d \rightarrow \mathbb{R}^m$ can be expressed as $T = A \circ \phi$ where $\phi : \mathbb{R}^d \rightarrow \ker(T)^\perp$ is the projection of \mathbb{R}^d onto $\ker(T)^\perp$ and $A : \ker(T)^\perp \rightarrow \text{Im}(T)$ is an invertible linear transformation.

This leads to the following observation about linear transformation of polyhedra.

Proposition 2 Let $P \subseteq \mathbb{R}^d$ be a polyhedron defined by p inequalities. Let $T : \mathbb{R}^d \rightarrow \mathbb{R}^m$ be any linear transformation. Then $T(P)$ is a polyhedron defined by at most $O(p^{2^d})$ inequalities.

Proof. Let us make a change of coordinates such that $\ker(T)^\perp = \mathbb{R}^{d'}$ with $d \geq d' \geq 0$ - this does not change the number of inequalities required to describe P or $T(P)$. By Proposition 1, T can be expressed as $A \circ \phi$ where ϕ is the projection from $\mathbb{R}^d \rightarrow \mathbb{R}^{d'}$, and A is an invertible transformation from $\mathbb{R}^{d'} \rightarrow \text{Im}(T)$. So we just need to analyze the effect of ϕ and A on the number of inequalities.

To analyze $\phi(P)$, we note that the Fourier-Motzkin elimination process [23] implies that projecting out a single variable can be done by squaring the number of inequalities. By repeatedly applying this, we get that $\phi(P)$ has at most $p^{2^{d-d'}}$ inequalities. Since A is an invertible linear transformation, $A(\phi(P))$ has the same number of inequalities as $\phi(P)$. The result follows. \square

We list one final linear algebraic observation. Let $\dim(W)$ denote the dimension of an affine subspace W , and let $\text{aff}(X)$ denote the affine hull of the columns of a matrix X (or just a finite set of vectors X).

Proposition 3 *Let $\{m_1, \dots, m_t\} \subseteq \mathbb{R}^m$ and $\{b_1, \dots, b_t\} \subseteq \mathbb{R}^d$ such that there exists a linear transformation $A : \mathbb{R}^d \rightarrow \mathbb{R}^m$ such that $m_i = A(b_i)$ for all $i = 1, \dots, t$. Further suppose that $\dim(\text{aff}(\{m_1, \dots, m_t\})) = \dim(\text{aff}(\{b_1, \dots, b_t\})) = k$ and that m_1, \dots, m_{k+1} and b_1, \dots, b_{k+1} are maximal affinely independent subsets, respectively.*

Then, for every $i > k + 1$, $m_i = \lambda_1 m_1 + \dots + \lambda_{k+1} m_{k+1}$ implies that $b_i = \lambda_1 b_1 + \dots + \lambda_{k+1} b_{k+1}$.

We also use the following deep result from real algebraic geometry and quantifier elimination.

Theorem 3. [2] *There is an algorithm that tests the feasibility of any system of s polynomial equalities involving N variables with d as the maximum degree of any polynomial, that runs in time $(sd)^{O(N)}$.*

3 Factorizations from a polyhedron

Our main tool for proving Theorem 1 will be the following generalization of the algorithm of Arora et al. [1], who proved it for the special case of P being the nonnegative cone. We generalize this to an arbitrary polyhedron P .

Theorem 4. *Let M be an $n \times m$ matrix with nonnegative entries, and let P be some polyhedron in \mathbb{R}^d described by p inequalities. If there exists a factorization $M = UV$ such that each row of U and each column of V is in P , then one can compute such a factorization in time polynomial in n and m (assuming d and p to be constants).*

In order to prove this theorem, we first need a few useful lemmas. Let X be a $p \times q$ matrix. For any subset $C \subseteq \{1, \dots, q\}$, let X^C denote the matrix formed by the subset of columns indexed by C . Similarly, for any subset $R \subseteq \{1, \dots, p\}$, let X_R denote the matrix formed by the rows indexed by R .

Lemma 1. *Let M be an $n \times m$ matrix with nonnegative entries. Let P be some polyhedron in \mathbb{R}^d described by p inequalities. Suppose there exists a factorization $M = UV$ such that each row of U and each column of V is in P . Then there exists a partition $C_1 \uplus C_2 \uplus \dots \uplus C_k = \{1, \dots, m\}$, a partition $R_1 \uplus R_2 \uplus \dots \uplus R_\ell = \{1, \dots, n\}$, and matrices $\bar{U} \in \mathbb{R}^{n \times d}$, $\bar{V} \in \mathbb{R}^{d \times m}$ such that the following properties all hold:*

1. $M = \bar{U}\bar{V}$.
2. Each row of \bar{U} and each column of \bar{V} is in P .
3. $\dim(\text{aff}(M^{C_j})) = \dim(\text{aff}(\bar{V}^{C_j}))$ for all $j = 1, \dots, k$.
4. $\dim(\text{aff}((M_{R_i})^T)) = \dim(\text{aff}((\bar{U}_{R_i})^T))$ for all $j = 1, \dots, \ell$.

5. $k, \ell \leq p^d$.

Proof. We use an idea from Arora et al. [1] to produce \bar{U}, \bar{V} and the partitions with the stated properties. Starting from U, V , we will first construct \bar{V} , and then use this to construct \bar{U} . Slightly more formally, we will first construct a partition $C_1 \uplus C_2 \uplus \dots \uplus C_k = \{1, \dots, m\}$ and \bar{V} such that $M = U\bar{V}$, all columns of \bar{V} are in P , and condition 3 in the statement is satisfied. We will then keep \bar{V} fixed and will construct a partition $R_1 \uplus R_2 \uplus \dots \uplus R_\ell = \{1, \dots, n\}$ such that each row of \bar{U} is in P , and condition 4 from the statement is satisfied. Condition 5 will then be straightforward.

For any $p \in P$, let F_p be the face of P of minimum dimension containing p . This induces a partial ordering \succ on the points in P , where $p_1 \succ p_2$ if $F_{p_1} \supsetneq F_{p_2}$.

For every column v^j of V , consider the set $(v^j + \ker(U)) \cap P$ and define \bar{v}^j to be a minimal element in this set according to this partial order. Note that for any $p \in P$, if there exists $u \in \ker(U) \setminus \{0\}$ such that the line $p + \lambda u$, $\lambda \in \mathbb{R}$ lies in the affine hull of F_p , then one can choose λ such that $p + \lambda u$ is in a strict face of F_p . Thus, by the minimal choice of \bar{v}^j , we have that $(\bar{v}^j + \ker(U)) \cap \text{aff}(F_{\bar{v}^j}) = \bar{v}^j$ for every $j \in \{1, \dots, m\}$. We set \bar{V} to be the matrix with columns \bar{v}^j . Note that $M = U\bar{V}$ as desired, since $U\bar{v}^j = U(v^j + x^j) = Uv^j$ for every $j \in \{1, \dots, m\}$, where x^j is some vector in $\ker(U)$.

The partition $C_1 \uplus C_2 \uplus \dots \uplus C_k$ of the columns of \bar{V} is obtained by grouping the columns together based on the face of minimum dimension that they lie on. Thus, $k \leq p^d$ which is an upper bound on the number of faces of P . We now need to verify that $\dim(\text{aff}(M^{C_j})) = \dim(\text{aff}(\bar{V}^{C_j}))$ for all $j = 1, \dots, k$. Fix some j and let the columns of M^{C_j} be $\{m_0, m_1, \dots, m_h\}$ and let the columns of \bar{V}^{C_j} be v_0, v_1, \dots, v_h . Since $M^{C_j} = U\bar{V}^{C_j}$, we know that $\dim(\text{aff}(M^{C_j})) \leq \dim(\text{aff}(\bar{V}^{C_j}))$. If the inequality is strict, then there exists a $v \in \ker(U) \setminus \{0\}$ such that $v = \lambda_0 v_0 + \lambda_1 v_1 + \dots + \lambda_h v_h$ and $\lambda_0 + \lambda_1 + \dots + \lambda_h = 0$. But then, if F is the face of minimum dimension containing v_0, v_1, \dots, v_h , we find that $v_0 + \lambda v$, $\lambda \in \mathbb{R}$ lies in the affine hull of F . This would contradict the construction of the columns of \bar{V} . Therefore, $\dim(\text{aff}(M^{C_j})) = \dim(\text{aff}(\bar{V}^{C_j}))$.

In a similar manner, we can change the rows of U (keeping \bar{V} fixed) to obtain \bar{U} so that condition 1 still holds and condition 2 is now satisfied. We can also construct the partition $R_1 \uplus R_2 \uplus \dots \uplus R_\ell = \{1, \dots, n\}$ (in the same way as the column partition) so that property 4 is satisfied. Finally, note that the parts in the partition are in correspondence with faces of P (as was the case with the column partition), giving $\ell \leq p^d$. This completes the construction. \square

Let X be a set of points in \mathbb{R}^d . We say a set of polyhedra P_1, \dots, P_k is a *polyhedral covering* of X if $(P_1 \cap X) \cup \dots \cup (P_k \cap X) = X$. We say a partition $X_1 \uplus \dots \uplus X_k = X$ is *induced by a polyhedral covering* if there exists a polyhedral covering P_1, \dots, P_k of X and $X_1 = P_1 \cap X$ such that $X_i = (P_i \cap X_i) \setminus (X_1 \cup \dots \cup X_{i-1})$ for $i = 2, \dots, k$. A (k_1, k_2) -*polyhedral partition* of X is a partition induced by a polyhedral covering of X with at most k_1 polyhedra and each polyhedron is described by at most k_2 inequalities.

Lemma 2. *Let k_1, k_2 be fixed natural numbers and let X be a set of points in \mathbb{R}^d . The number of (k_1, k_2) -polyhedral partitions is at most $O((2^d m^d)^{k_1 k_2})$ and one can enumerate these partitions in time $O((2^d m^d)^{k_1 k_2})$, where $m = |X|$.*

Proof. Let us first count the number of subsets of X of the form $P \cap X$ where P is a polyhedron with at most k_2 inequalities. As observed in Arora et al [1], this can be reduced to counting the number of subsets of the form $H \cap X$ where H is a halfspace. The number of such subsets is $O(2^d m^d)$ and can be enumerated in the same amount of time (as was shown in Arora et al [1] by a simple iterative procedure). To choose a subset of the form $P \cap X$ where P is a polyhedron with at most k_2 inequalities, one simply needs to iteratively choose k_2 subsets given by halfspace intersections. Thus, there are $O((2^d m^d)^{k_2})$ such subsets and these can be enumerated in this iterative fashion.

To finally get partitions induced by polyhedral coverings, one needs to iteratively choose k_1 subsets of the form $P \cap X$ where P is a polyhedron with at most k_2 inequalities. The result follows. \square

Using these tools, we can now prove Theorem 4.

Proof (of Theorem 4). By Lemma 1, there exists a partition $C_1 \uplus C_2 \uplus \dots \uplus C_k = \{1, \dots, m\}$, a partition $R_1 \uplus R_2 \uplus \dots \uplus R_\ell = \{1, \dots, n\}$, and matrices $\bar{U} \in \mathbb{R}^{n \times d}$, $\bar{V} \in \mathbb{R}^{d \times m}$ such that conditions 1, 2, 3, and 4 in Lemma 1 hold.

Condition 3 from Lemma 1 and Proposition 3 imply that for each $j \in \{1, \dots, k\}$, there exist $\dim(\text{aff}(M^{C_j})) + 1 \leq d + 1$ columns of \bar{V}^{C_j} , such that every other column in \bar{V}^{C_j} can be expressed as linear combinations of these columns. Moreover, the coefficients in these linear combinations can be computed from the columns of M^{C_j} . Similarly, Condition 4 from Lemma 1 and Proposition 3 imply that for every $i = 1, \dots, \ell$, the rows of \bar{U}_{R_i} can be expressed as linear combinations of $\dim(\text{aff}(M_{R_i})) + 1 \leq d + 1$ rows of \bar{U}_{R_i} .

Moreover, for any fixed $j \in \{1, \dots, k\}$, since $M^{\{s\}} = \bar{U} \bar{V}^{\{s\}}$ for every $s \in C_j$, we have $\{M^{\{s\}} : s \in C_j\} \subseteq \bar{U}(P)$ since every column of \bar{V} is in P by Condition 2 from Lemma 1. Invoking Proposition 2, we obtain that $\bar{U}(P)$ is described using at most p^{2^d} inequalities. By Lemma 1, k is bounded by p^d . Therefore, C_1, \dots, C_k is a (p^d, p^{2^d}) -polyhedral partition of $\{1, \dots, m\}$. By Lemma 2, we can enumerate such partitions in time $O((2^d m^d)^{p^{d+2^d}})$. Similarly, one can enumerate all possible partitions R_1, \dots, R_ℓ satisfying the conditions of Lemma 1 in time $O((2^d n^d)^{p^{d+2^d}})$.

Our algorithm will find \bar{U} and \bar{V} from Lemma 1. By conditions 1 and 2 of Lemma 1, these matrices form the desired factorization of M . By the discussion above, it suffices to find the partition $C_1 \uplus C_2 \uplus \dots \uplus C_k = \{1, \dots, m\}$, the partition $R_1 \uplus R_2 \uplus \dots \uplus R_\ell = \{1, \dots, n\}$ such that conditions 1, 2, 3, and 4 in Lemma 1 hold, the $\dim(\text{aff}(M^{C_j})) + 1$ columns of \bar{V}^{C_j} that form a basis for the other columns of \bar{V}^{C_j} for each $j = 1, \dots, k$, and the $\dim(\text{aff}(M_{R_i})) + 1$ rows of \bar{U}_{R_i} that form a basis for the other rows of \bar{U}_{R_i} for each $i = 1, \dots, \ell$. Once we have all of these, we can reconstruct the full \bar{U} and \bar{V} .

Finding these partitions and bases can be done by enumerating all possible (p^d, p^{2^d}) -polyhedral partitions C_1, \dots, C_k of the columns of M , all possible (p^d, p^{2^d}) -polyhedral partitions R_1, \dots, R_ℓ of the rows of M , and for each choice of such partitions, introducing variables for the entries of the $\dim(\text{aff}(M^{C_j})) + 1$ special columns of \bar{V}^{C_j} , $j = 1, \dots, k$, and $\dim(\text{aff}(M_{R_i})) + 1$ special rows of \bar{U}_{R_i} , $i = 1, \dots, \ell$. Finally, set up a system of polynomial equalities in these variables that represent $M = \bar{U}\bar{V}$. Notice that this system has only $O((k + \ell)d^2)$ variables which is a constant since p, d are constants, and $k, \ell \leq p^d$. We finally invoke Theorem 3 to test the feasibility of such a system. Note that the requirement that $M = \bar{U}\bar{V}$ can be expressed using nm polynomial equalities, where each polynomial is a quadratic. \square

4 Proof of Theorem 1

For any square matrix $X \in \mathbb{R}^{r \times r}$, we use $\|X\|_{sp} := \max_{y \in \mathbb{R}^r \setminus \{0\}} \frac{\|Xy\|_2}{\|y\|_2}$ to denote the spectral norm of X . The algorithm depends on this key result (paraphrased here) from [7].

Theorem 5. *Let M be an $n \times m$ matrix with nonnegative entries. If M has a PSD factorization $M = UV$ such that the rows of U and columns of V are in \mathcal{S}_+^r , then there exists a PSD factorization $M = \bar{U}\bar{V}$ such that the rows of \bar{U} and the columns of \bar{V} have spectral norm bounded by $\sqrt{r}\|M\|_\infty$.*

We outline the steps of the algorithm in Theorem 1. Let $f(r)$ be such that for every matrix $X \in \mathcal{S}^r$, $\|X\|_\infty \leq f(r)\|X\|_{sp}$. Such an $f(r)$ must exist because all norms are equivalent on a Euclidean space, i.e., their values are the same upto a factor depending only on the dimension of the space.

1. Given M , let $\Delta = \|M\|_\infty$. Construct a polyhedral ϵ -approximation of \mathcal{S}_+^r with respect to the $\|\cdot\|_\infty$ norm on \mathcal{S}^r – see Theorem 2 and Remark 1. Let P be the polyhedron formed by the intersection of this polyhedral approximation with the cube $\{x \in \mathcal{S}^r : \|x\|_\infty \leq f(r)\sqrt{r}\Delta\}$.
2. By Theorem 5 and the assumption that M has PSD rank r , we know there exists a factorization $M = \bar{U}\bar{V}$ such that the rows of \bar{U} and columns of \bar{V} are in the PSD cone, and their spectral norm is at most $\sqrt{r}\Delta$. Therefore, for every row u of \bar{U} , we have $\|u\|_\infty \leq f(r)\sqrt{r}\Delta$ and similarly for the columns of \bar{V} . This implies that the rows of \bar{U} and columns of \bar{V} are in P . Since \bar{U}, \bar{V} exist, we can employ Theorem 4 to construct a factorization $M = U'V'$ such that the rows of U' and the columns of V' are in P . Note that the algorithm of Theorem 4 may not produce a PSD factorization. To obtain an approximate PSD factorization, we construct matrices U and V by projecting each row of \bar{U} to the nearest point in the PSD cone (according to the $\|\cdot\|_\infty$ norm), and similarly for the columns of \bar{V} . This can be done in polynomial time by solving a semidefinite program – see [5, Section 8.1].

This concludes the description of the algorithm. It remains to prove that

$$\|M - UV\|_\infty \leq \epsilon \|M\|_\infty.$$

Our first step will be to use the fact that we projected from an ϵ -approximation to the PSD cone. In particular, we know that $\|V'^j - V^j\|_\infty \leq \epsilon \|V'^j\|_\infty$ for each $j \in \{1, \dots, m\}$, and similarly $\|U'_i - U_i\|_\infty \leq \epsilon \|U'^i\|_\infty$ for each $i \in \{1, \dots, n\}$. This clearly implies that

$$\|\bar{V} - V\|_\infty \leq \epsilon \|\bar{V}\|_\infty \text{ and } \|\bar{U} - U\|_\infty \leq \epsilon \|\bar{U}\|_\infty. \quad (1)$$

Now we can analyze the approximation of our factorization:

$$\begin{aligned} \|M - UV\|_\infty &= \|\bar{U}\bar{V} - UV\|_\infty \\ &\leq \|\bar{U}\bar{V} - \bar{U}V\|_\infty + \|\bar{U}V - UV\|_\infty \\ &\leq r\|\bar{U}\|_\infty\|\bar{V} - V\|_\infty + r\|\bar{U} - U\|_\infty\|V\|_\infty \\ &\leq r\|\bar{U}\|_\infty(\epsilon\|\bar{V}\|_\infty) + r\epsilon\|\bar{U}\|_\infty\|V\|_\infty \end{aligned}$$

where the first equality is from the fact that $M = \bar{U}\bar{V}$, the first inequality is from the triangle inequality, and the third is from (1). The second inequality follows from the observation that for any matrices $A \in \mathbb{R}^{n \times r}$, $B \in \mathbb{R}^{r \times m}$, $\|AB\|_\infty \leq r\|A\|_\infty\|B\|_\infty$.

Since $\|V\|_\infty \leq (1 + \epsilon)\|\bar{V}\|_\infty$ because of (1), we obtain $\|M - UV\|_\infty \leq 3\epsilon r\|\bar{U}\|_\infty\|\bar{V}\|_\infty$. Since each row u of U' is in P , we have $\|u\|_\infty \leq f(r)\sqrt{r\Delta}$. Therefore, $\|U'\|_\infty \leq f(r)\sqrt{r\Delta}$. Similarly, $\|V'\|_\infty \leq f(r)\sqrt{r\Delta}$. Hence, $\|M - UV\|_\infty \leq 3\epsilon r\|\bar{U}\|_\infty\|\bar{V}\|_\infty \leq 3f(r)r^2\epsilon\Delta$. By redefining ϵ appropriately (in particular, letting ϵ' be the previous ϵ and letting $\epsilon = 3f(r)f\epsilon'$), we get that

$$\|M - UV\|_\infty \leq \epsilon \|M\|_\infty$$

as desired.

5 Open Questions

Question 1 remains the outstanding open question in the line of research on factorization algorithms with polynomial time guarantees. Another interesting direction would be generalize Theorem 1 to approximation guarantees with other norms. For example, the induced norms $\|M\|_{1,2} := \max_{x \in \mathbb{R}^m} \frac{\|Mx\|_2}{\|x\|_1}$ and $\|M\|_{\infty,2} := \max_{x \in \mathbb{R}^m} \frac{\|Mx\|_2}{\|x\|_\infty}$ were used in [13]. The authors show that approximate factorization with respect to these norms give rise to small SDP reformulations whose projections approximate a given polytope, where the geometric approximation is tightly determined by the approximation factor in the matrix factorization.

It would also be interesting to resolve the following question:

Let $r \in \mathbb{N}$ and $\epsilon > 0$ be fixed constants. Let \mathcal{M} be the family of non-negative matrices such that for every $M \in \mathcal{M}$, there exists another

nonnegative matrix \overline{M} such that $\|M - \overline{M}\|_\infty \leq \epsilon \|M\|_\infty$ and \overline{M} admits a rank r PSD factorization.

Does there exist an algorithm which, given any nonnegative matrix $M \in \mathcal{M}$, can find matrices U and V such that each row of U and each column of V are in \mathcal{S}_+^r such that

$$\|M - UV\|_\infty \leq 2\epsilon \|M\|_\infty,$$

and has runtime polynomial in the dimensions of M ? In other words: if the input matrix M is *close* to a matrix with small PSD rank, can we find a low-rank factorization that is a good approximation to M ?

Approximate low-rank nonnegative factorizations of matrices with high nonnegative rank have been extensively studied – see [3] for a survey of the diverse applications, and [1] for a recent algorithm with provable guarantees on the complexity. The corresponding question for PSD factorizations is of similar interest.

References

1. Sanjeev Arora, Rong Ge, Ravindran Kannan, and Ankur Moitra. Computing a nonnegative matrix factorization—provably. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 145–162. ACM, 2012.
2. Saugata Basu, Richard Pollack, and Marie-Françoise Roy. On the combinatorial and algebraic complexity of quantifier elimination. *Journal of the ACM (JACM)*, 43(6):1002–1045, 1996.
3. Michael W Berry, Murray Browne, Amy N Langville, V Paul Pauca, and Robert J Plemmons. Algorithms and applications for approximate nonnegative matrix factorization. *Computational statistics & data analysis*, 52(1):155–173, 2007.
4. Lenore Blum, Mike Shub, and Steve Smale. On a theory of computation and complexity over the real numbers: W-completeness, recursive functions and universal machines. *Bull. Amer. Math. Soc.*, 21(1):1–46, 1989.
5. S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, 2004.
6. Gábor Braun, Jonah Brown-Cohen, Arefin Huq, Sebastian Pokutta, Prasad Raghavendra, Aurko Roy, Benjamin Weitz, and Daniel Zink. The matching problem has no small symmetric sdp. *arXiv preprint arXiv:1504.00703*, 2015.
7. Jop Briët, Daniel Dadush, and Sebastian Pokutta. On the existence of 0/1 polytopes with high semidefinite extension complexity. *Mathematical Programming*, pages 1–21, 2014.
8. Michele Conforti, Gérard Cornuéjols, and Giacomo Zambelli. Extended formulations in combinatorial optimization. *4OR*, 8(1):1–48, 2010.
9. Hamza Fawzi, João Gouveia, Pablo A. Parrilo, Richard Z. Robinson, and Rekha R. Thomas. Positive semidefinite rank. <http://arxiv.org/abs/1407.4095>, 2015.
10. Samuel Fiorini, Volker Kaibel, Kanstantsin Pashkovich, and Dirk Oliver Theis. Combinatorial bounds on nonnegative rank and extended formulations. *Discrete mathematics*, 313(1):67–83, 2013.
11. Samuel Fiorini, Serge Massar, Sebastian Pokutta, Hans Raj Tiwary, and Ronald 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*, pages 95–106. ACM, 2012.

12. Joao Gouveia, Pablo A Parrilo, and Rekha R Thomas. Lifts of convex sets and cone factorizations. *Mathematics of Operations Research*, 38(2):248–264, 2013.
13. João Gouveia, Pablo A Parrilo, and Rekha R Thomas. Approximate cone factorizations and lifts of polytopes. *Mathematical Programming*, 151(2):613–637, 2015.
14. Volker Kaibel. Extended formulations in combinatorial optimization. *arXiv preprint arXiv:1104.1023*, 2011.
15. J Lee, Prasad Raghavendra, David Steurer, and Ning Tan. On the power of symmetric lp and sdp relaxations. In *Computational Complexity (CCC), 2014 IEEE 29th Conference on*, pages 13–21. IEEE, 2014.
16. James R Lee, Prasad Raghavendra, and David Steurer. Lower bounds on the size of semidefinite programming relaxations. *arXiv preprint arXiv:1411.6317*, 2014.
17. Ankur Moitra. An almost optimal algorithm for computing nonnegative rank. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1454–1464. SIAM, 2013.
18. Yuval Rabani and Amir Shpilka. Explicit construction of a small ϵ -net for linear threshold functions. *SIAM Journal on Computing*, 39(8):3501–3520, 2010.
19. Thomas Rothvoß. Some 0/1 polytopes need exponential size extended formulations. *Mathematical Programming*, 142(1-2):255–268, 2013.
20. Thomas Rothvoß. The matching polytope has exponential extension complexity. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 263–272. ACM, 2014.
21. Stephen A Vavasis. On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization*, 20(3):1364–1377, 2009.
22. Mihalis Yannakakis. Expressing combinatorial optimization problems by linear programs. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 223–228. ACM, 1988.
23. G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.