ARTICLE IN PRESS

Discrete Applied Mathematics xxx (xxxx) xxx

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

A geometric lower bound on the extension complexity of polytopes based on the f-vector

Julien Dewez [a,*], Nicolas Gillis [b], François Glineur [a]

[a] *Université catholique de Louvain, ICTEAM/INMA and CORE, Avenue Georges Lemaître, 4-6/L4.05.01, B-1348 Louvain-la-Neuve, Belgium*

[b] *Université de Mons, Department of Mathematics and Operational Research, Rue de Houdain, 9, B-7000 Mons, Belgium*

ARTICLE INFO

Article history:
Received 15 January 2020
Received in revised form 6 September 2020
Accepted 30 September 2020
Available online xxxx

Keywords:
Combinatorial optimization
Convex polytopes
Linear extension
Extended formulation
Extension complexity
Lower bound
f-vector
Slack matrices
Nonnegative matrix factorization
Nonnegative rank
Restricted nonnegative rank

ABSTRACT

A linear extension of a polytope Q is any polytope which can be mapped onto Q via an affine transformation. The extension complexity of a polytope is the minimum number of facets of any linear extension of this polytope. In general, computing the extension complexity of a given polytope is difficult. The extension complexity is also equal to the nonnegative rank of any slack matrix of the polytope.

In this paper, we introduce a new *geometric* lower bound on the extension complexity of a polytope, i.e., which relies only on the knowledge of some of its geometric characteristics. It is based on the monotone behaviour of the f-vector of polytopes under affine maps. We present numerical results showing that this bound can improve upon existing geometric lower bounds, and provide a generalization of this lower bound for the nonnegative rank of any matrix.

1. Introduction

A (linear) extension of a polytope can be defined as follows [23].

Definition 1.1. Let Q be a polytope, a *linear extension* for Q is a polytope P together with an affine map π such that $Q = \pi(P)$.

Many combinatorial optimization problems can be formulated as the optimization of a linear objective over a polytope. However, polytopes encountered in many applications frequently have very large numbers of facets (possibly growing exponentially with the size of the problem). Because of the high number of facets, computing the optimal solution of this linear optimization problem can be time consuming, and even impractical already for small to medium size problems. Note that the number of vertices also has an influence on the computation time; if it grows moderately with the size of the problem, one can simply solve by enumerating all vertices. However, if one can find an extension for this polytope with a moderate number of facets (growing for example polynomially with the size of the problem), one can solve the optimization problem over the extension efficiently, and project its optimal solution back to the original polytope of feasible solutions. Therefore, we would like to know the minimum number of facets of any extension of these combinatorial polytopes. This number is called the extension complexity of the polytope.

* Corresponding author.
E-mail addresses: julien.dewez@uclouvain.be (J. Dewez), nicolas.gillis@umons.ac.be (N. Gillis), francois.glineur@uclouvain.be (F. Glineur).

https://doi.org/10.1016/j.dam.2020.09.028

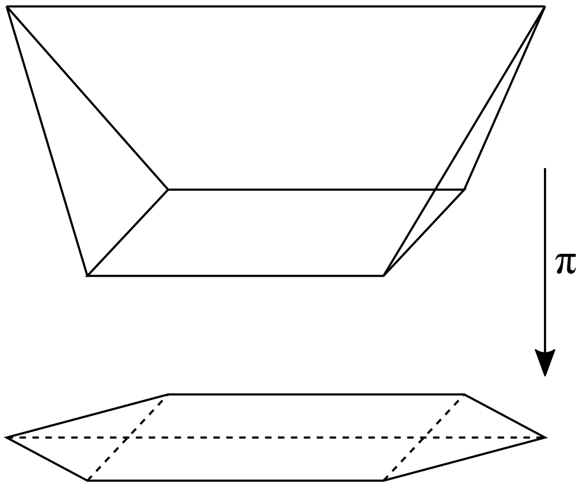

Fig. 1. Extension for the hexagon with 5 facets.

Definition 1.2. Let P be a polytope, the *extension complexity* of P, denoted $\mathrm{xc}(P)$, is the minimum number of facets of any extension of P.

An example of extension is given on Fig. 1 for the hexagon.

In particular, if the extension complexity grows exponentially with the size of the problem, i.e., the dimension of the polytope, then the polytope does not admit any extension with a polynomial number of facets and we will not be able to solve the problem in polynomial time with an LP-based approach that requires a compact formulation. For example, Rothvoss proved that the matching polytope has exponential extension complexity [22].

The problem of finding an extension of a polytope is closely related to the problem of finding a nonnegative factorization of a nonnegative matrix. Crucially, any polytope can be represented by a nonnegative slack matrix.

Definition 1.3. Let P be a polytope with v vertices denoted $x^{(j)}$, $1 \leq j \leq v$ and f facets expressed as linear inequalities $a_i^\top x \leq b_i$, $1 \leq i \leq f$. A *slack matrix* for P is an $f \times v$ matrix S whose entry at position (i, j), $1 \leq i \leq f$, $1 \leq j \leq v$ is the slack of vertex j with respect to facet i, i.e., $S_{i,j} = b_i - a_i^\top x^{(j)} \geq 0$. Thus, S is an element-wise nonnegative matrix.

Note that the slack matrix of a polytope is not unique. The slack matrix of a polytope P whose affine hull is n-dimensional has rank $n + 1$ [17].

An important notion for element-wise nonnegative matrices is the nonnegative rank which is defined as follows [8].

Definition 1.4. Let $M \in \mathbb{R}_+^{m\times n}$ be an element-wise nonnegative matrix. The *nonnegative rank* of M, denoted $\mathrm{rank}_+(M)$, is the minimum number of terms required to represent M as a sum of nonnegative rank-one matrices.

The nonnegative rank of $M \in \mathbb{R}_+^{m\times n}$ is also the smallest integer k such that there exists nonnegative factorization of dimension k, i.e., a matrix $U \in \mathbb{R}_+^{m\times k}$ and a matrix $V \in \mathbb{R}_+^{k\times n}$ such that $M = UV$. Indeed,

$$M = UV = \sum_{i=1}^{k} U_{:,i} V_{i,:},$$

where $U_{:,i}$ denotes the ith column of U and $V_{i,:}$ denotes the ith row of V. The products $U_{:,i}V_{i,:}$ are nonnegative rank-one matrices.

Yannakakis [27] proved that the extension complexity of a polytope equals the nonnegative rank of any slack matrix for P (see also [13]), hence all slack matrices of a certain polytope have the same nonnegative rank.

Theorem 1.5 (*[27]*)**.** *Let P be a polytope and let S be a slack matrix for P. Then,*

$$\mathrm{xc}(P) = \mathrm{rank}_+(S).$$

Vavasis [26] proved that, in general, computing this rank and a corresponding factorization is NP-hard. This hints that, in general, computing the extension complexity of a polytope could also be NP-hard. However, if the dimension of factorization is fixed, computing a factorization (when it exists) can be done in polynomial time [2]. Therefore, we would like to have bounds (upper and lower bounds), as strong as possible, for the nonnegative rank and/or for the extension complexity. A consequence of Theorem 1.5 is that, for slack matrices of polytopes, any bound on the nonnegative rank of a matrix can be used to bound the extension complexity of a polytope.

1.1. Previous work

Upper bounds on the nonnegative rank can be obtained from any factorization of the given nonnegative matrix. For example, an interesting upper bound for regular polygons has been deduced from heuristically-computed factorization [25].

Determining nontrivial lower bounds on the nonnegative rank is in general more challenging. Thus, an important research area in nonnegative matrix factorization is to compute strong lower bounds for the nonnegative rank. Several approaches have already been explored, some of them are described briefly here after. Most of these bounds were initially developed for the nonnegative rank of any nonnegative matrix, but can be used to bound the extension complexity, as explained previously. We distinguish between two kinds of bounds: those which come from a geometric representation of the matrix, i.e, representing the matrix as a polytope or a pair of polytopes and bounding the nonnegative rank based on characteristics of this representation, and those which do not. Note that this is a rather arbitrary classification of lower bounds, because most bounds can be given a geometric interpretation. However, this classification is motivated by the way these bounds were derived, using a geometric representation or not.

Yannakakis [27] observed that the extension complexity of a polytope is lower bounded by the logarithm of its number of faces, and hence the nonnegative rank of its slack matrix by the logarithm of its number of rows. Another bound which comes from a geometric representation was introduced by Goemans [16], first as a bound on the extension complexity then generalized to the nonnegative rank of any nonnegative matrix. This bound relies on a relation between the number of vertices of the polytope and the number of facets of any extension. Another geometric bound was introduced by Gillis and Glineur [14,15], this bound is based on a representation of the matrix as a pair of nested polytopes. This last bound will be described in more detail in Section 4.

Some other bounds do not come from a geometric representation. These include combinatorial bounds such as the fooling sets bound [1] and the rectangle covering number [13] which are based solely on the position of the non-zeros elements of the matrix. Fawzi and Parrilo also introduced bounds based on semidefinite programming [11,12]. More precisely, they introduce some quantities which are solutions of difficult optimization problems involving the nonnegative matrix. These quantities provide lower bounds on the nonnegative rank and can be nicely under-approximated via tractable semidefinite relaxations. These bounds do not rely solely on the position of the non-zeros elements, the values of the elements influence the bound. This means that this bound is more efficient for matrices with few or no zeros.

1.2. Contribution and outline of the paper

Our main contributions are a new geometric lower bound on the extension complexity based on the monotonic behaviour of the f-vector under projections and a comparison of our bound with other lower bounds on a set of benchmark polytopes.

In Section 2, we recall some useful notions and results about polytopes, projections and the polar transformation of a polytope. In Section 3, we introduce our main contribution, a new geometric lower bound on the extension complexity of a polytope based on the monotone behaviour of its f-vector under projections. We first give a proof that the f-vector decreases monotonically under projections, and use this result to define our new lower bound. In Section 4, we show that our bound is a stronger version of the bound from [15]. In Section 5, we give some numerical results on low-dimensional regular polytopes comparing our bound with existing bounds; we also compare their associated computational times. In Section 6, we use the notion of flag numbers of a polytope to strengthen our bound. Finally, in Section 7, we generalize our bound for nonnegative matrices which are not necessarily slack matrices.

2. Preliminaries

In this section, we recall useful notions about polytopes, the definition of projections and the polar transformation for polytopes.

2.1. About polytopes

First, let us recall the definition of the affine hull of a polytope.

Definition 2.1. Let $P \subset \mathbb{R}^n$ be a polytope, the *affine hull* of P denoted $\text{aff}(P)$ is the set of all affine combinations of elements of P, i.e.,

$$\text{aff}(P) = \left\{ \sum_{i=1}^{k} \alpha_i x_i \mid k > 0, \underbrace{x_i \in P, \alpha_i \in \mathbb{R}}_{i \in 1,\dots,k}, \sum_{i=1}^{k} \alpha_i = 1 \right\}.$$

The affine hull of P is the smallest affine subset of $\mathbb{R}^n$ which contains P.

For the rest of the paper, we will consider the following definition for the dimension of a polytope.

Definition 2.2. Let $P \subset \mathbb{R}^n$ be a polytope, the *dimension* of P is the dimension of its affine hull, i.e., $\dim(P) = \dim(\text{aff}(P))$.

This allows us to work in an n-dimensional space with polytopes that can have a dimension strictly smaller than n. Such polytopes have empty interiors, therefore we will need the notion of relative interior.

Definition 2.3. Let $P \subset \mathbb{R}^n$ be a polytope, the *relative interior* of P denoted $\text{ri}(P)$ is the interior of P within its affine hull, i.e.,

$$\text{ri}(P) = \{x \in P \mid \exists \epsilon > 0, B_\epsilon(x) \cap \text{aff}(P) \subseteq P\},$$

where $B_\epsilon(x)$ denotes the ball of radius ϵ centred at x.

Let us now define the k-faces of the polytope.

Definition 2.4. Let $P \subset \mathbb{R}^n$ be a polytope, a linear inequality $ax \le b$ is *valid* for P if all points of P satisfy it. A *face* of P is any intersection of P with a set of the form $\{x \in \mathbb{R}^n : ax = b\}$, where $ax \le b$ is a valid inequality for P. A *k-face* of P is a face of P which has an affine hull of dimension k.

Let us also recall the definition of supporting hyperplane.

Definition 2.5. Let $v \in \mathbb{R}^n$ such that $v \neq 0$ and $h \in \mathbb{R}$, the set

$$H = \left\{x \in \mathbb{R}^n \mid v^\top x = h\right\}$$

is a *hyperplane* in $\mathbb{R}^n$. Let $P \subset \mathbb{R}^n$ be a polytope, H is a *supporting hyperplane* of P if $H \cap P \neq \emptyset$ and P lies entirely in one of the two closed half-spaces bounded by the hyperplane.

Note that the non-trivial faces of a polytope, i.e., excluding the polytope itself and the empty set, can be expressed as the intersection of the polytope with one of its supporting hyperplanes.

Let us recall the following proposition which states that faces of a polytope are also polytopes. In this proposition, $\text{vertices}(P)$ denotes the set of all vertices of P.

Proposition 2.6 (*[28, p. 53]*). *Let $P \subseteq \mathbb{R}^n$ be a polytope, and $V := \text{vertices}(P)$. Let F be a face of P.*

- *The face F is a polytope, with $\text{vertices}(F) = F \cap V$.*
- *Every intersection of faces of P is a face of P.*
- *The faces of F are exactly the faces of P that are contained in F.*
- *$F = P \cap \text{aff}(F)$.*

Let us define the f-vector, which consists of the collection of the numbers of k-faces of a polytope.

Definition 2.7. Given a polytope $P \subset \mathbb{R}^n$, the *f-vector* of P is the vector collecting the numbers of k-faces of P for each value of k between 0 and $n-1$,

$$f = (f_0, f_1, f_2, \ldots, f_{n-1}),$$

where f_k is the number of k-faces of P.

2.2. Projections

The projection of a polytope is the result we get when we apply a linear map to the polytope.

Definition 2.8. Let $A \in \mathbb{R}^{n\times n}$ and $T_A : x \mapsto Ax$, a linear mapping. Let $P \subseteq \mathbb{R}^n$ be a polytope, the *projection* of P with respect to T_A is the set

$$Q = T_A(P) = \{T_A(x) | x \in P\}.$$

Remark 2.9. An affine mapping $x \to Ax + b$, can be considered as a succession of transformations; a linear map $x \to Ax$ followed by a translation $x \to x + b$. Since a translation does not modify the geometry (dimension, numbers of faces, angles, …) of a polytope, we consider only linear mappings.

Note that each point has a single image by a linear map. However, several points can have the same image by a linear map. This implies that the dimension of a set can only decrease after a linear map is applied.

Property 2.10. *Let $A \in \mathbb{R}^{n\times n}$ and $T_A : x \mapsto Ax$, a linear mapping. Let $P \subseteq \mathbb{R}^n$ be a polytope and let $Q = T_A(P)$, then $\dim(Q) \le \dim(P)$.*

Projections have very interesting properties, one that will be useful to us is that convexity is preserved under a linear transformation. More precisely, linear transformations map a convex combination of some points onto a convex combination of the image of those points with the same coefficients of combination. This also means that if points are co-linear, their images after a linear map are also co-linear. Another interesting consequence is that if some points $x^{(i)}$ have the same image by a linear map, any point in the convex hull of these $x^{(i)}$ also has the same image. Moreover, we also know that the projection of a polytope is also a polytope [28, p. 29].

Definition 2.11. Let $T_A : x \mapsto Ax$ be a linear map. The *preimage* of a point $y \in \mathbb{R}^n$ with respect to T_A, denoted $T_A^{-1}(y)$, is the set

$$T_A^{-1}(y) = \{x \in \mathbb{R}^n | T_A(x) = y\}.$$

Similarly, the preimage of a set $S \subset \mathbb{R}^n$ with respect to T_A, denoted $T_A^{-1}(S)$, is the set

$$T_A^{-1}(S) = \{x \in \mathbb{R}^n | T_A(x) \in S\}.$$

Let $P \subseteq \mathbb{R}^n$ be a polytope, the preimage of $y \in \mathbb{R}^n$ restricted to P denoted $T_A^{-1}(y)|_P$ is the set

$$T_A^{-1}(y)|_P = \{x \in P | T_A(x) = y\}.$$

Note that the preimage of a point can be the empty set, a single point or a set of points. However, the preimage is always an affine subspace.

Corollary 2.12. *Let $T_A : x \mapsto Ax$ be a linear map. For any $y \in \mathbb{R}^n$, $T_A^{-1}(y)$ is an affine subspace of $\mathbb{R}^n$.*

Proof. Let x_i, $1 \leq i \leq k$ be some points in the preimage of y, i.e., $T_A(x_i) = y$, $1 \leq i \leq k$. Let $\tilde{x} = \sum_{i=1}^{k} \alpha_i x_i$ such that $\alpha_i \in \mathbb{R}$, $\sum_{i=1}^{k} \alpha_i = 1$ be an affine combination of these points and let us prove that it is also in the preimage of y.

$$T_A(\tilde{x}) = T_A\left(\sum_{i=1}^{k} \alpha_i x_i\right) = \sum_{i=1}^{k} T_A(\alpha_i x_i) = \sum_{i=1}^{k} \alpha_i T_A(x_i) = \sum_{i=1}^{k} \alpha_i y = y.$$

Thus, $\tilde{x}$ is in the preimage of y. Since, this reasoning is valid for any affine combination of the points in the preimage of y, it must be an affine subspace of $\mathbb{R}^n$. □

Now that the preimage is defined, we can define slices as preimages of polytopes with respect to linear maps.

Definition 2.13. Let $A \in \mathbb{R}^{n\times n}$ and $T_A : x \mapsto Ax$, a linear mapping. Let $Q \subseteq \mathbb{R}^n$ be a polytope, the *slice* of Q with respect to T_A is the set $T_A^{-1}(Q)$.

Note that slices are not necessarily polytopes, as they can be unbounded polyhedra.

Remark 2.14. If $Q = T_A(P)$, then $T_A^{-1}(Q)|_P = P$.

2.3. Polar transformation

Let us now define the polar transformation of a polytope.

Definition 2.15. Let $P \subset \mathbb{R}^n$ be a polytope, the *polar* of P denoted P^0 is the set defined as

$$P^0 = \{y \in \mathbb{R}^n \mid y^\top x \leq 1\, \forall x \in P\}.$$

Property 2.16. *If $P \subset \mathbb{R}^n$ is a polytope and $0 \in \text{int}(P)$, then P^0 is a polytope.*

Note that we can assume, without loss of generality, that every full-dimensional polytope contains the origin in its interior. If it is not the case, we translate the polytope such that the origin is in its interior which does not modify the geometry of the polytope, i.e., the translated polytope has the same properties as the original polytope such as the same dimension, the same numbers of faces, the same distances between vertices, etc. Most importantly, its extension complexity is preserved.

Remark 2.17. The polar of a polytope which does not contain the origin in its interior is an unbounded polyhedron.

This remark is important because if we have a polytope which is not full-dimensional, it cannot contain the origin in its interior (even after a translation). However, we can still compute the polar set $\{y \in \mathbb{R}^n | y^\top x \leq 1\, \forall x \in P\}$ which is an unbounded polyhedron.

Fig. 2 provides an example of two polytopes, one being the translation of the other, and their polars. We observe that the shape of the polar depends on where the origin is in the interior of the polytope. By translating the polytope (from

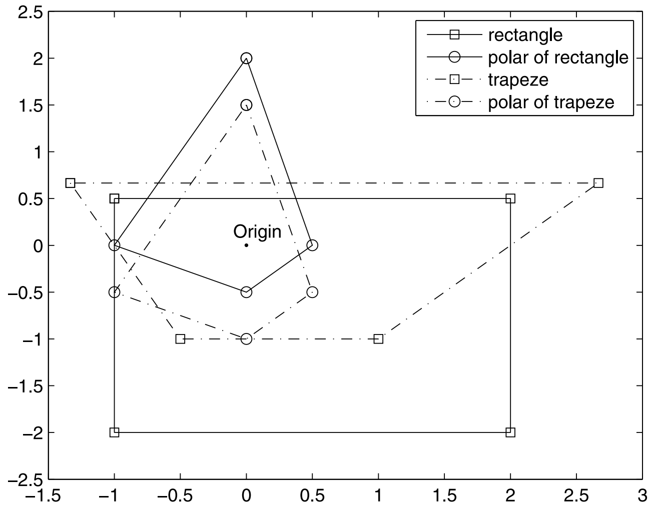

Fig. 2. Example of two polytopes and their polar. The shape of the polar depends on the position with respect to the origin.

the polytope in plain lines to the one in dashed lines), the polar is no longer a rectangle, it is a trapezoid. However, the number of faces does not change (this observation will be useful later). Let us recall the following well-known result.

Lemma 2.18 ([28]). *Let P be a polytope containing the origin in its interior. Then P^0 is a polytope containing the origin in its interior. Moreover, $(P^0)^0 = P$ and P has m facets if and only if P^0 has m vertices.*

This lemma states that vertices are converted into facets via polar transformation and vice versa, and that if we apply the polar transformation twice, we get the original polytope. A stronger result [28, p. 64] states that the f-vector is reversed by polarity, i.e., if the polytope P has the following f-vector $(f_0, f_1, \dots, f_{n-1})$, its polar P^0 will have the following f-vector $(f_{n-1}, \dots, f_1, f_0)$. For example, the cube, which has 8 vertices, 12 edges and 6 facets, has the following f-vector: $f = (8, 12, 6)$ if we consider a 3-dimensional space. Its polar has the following f-vector: $f^0 = (6, 12, 8)$ which is the f-vector of a (not necessarily regular) octahedron.

Let us show the following result which states that the polar of a projection of a polytope can be computed from the original polar: it is equal to its slice with respect to the adjoint linear map.

Property 2.19. *Let $P \subset \mathbb{R}^n$ be a polytope such that $0 \in \operatorname{int} P$ and let $T_A : x \to Ax$ be a linear map, with $A \in \mathbb{R}^{n\times n}$. Let $T_{A^\top} : y \to A^\top y$ be the adjoint linear map. Then,*

$$[T_A(P)]^0 = T_{A^\top}^{-1}(P^0).$$

Proof. The transformation of P with respect to T_A is given by

$$T_A(P) = \{Ax \mid x \in P\}.$$

The polar of $T_A(P)$ is given by

$$[T_A(P)]^0 = \left\{y \mid y^\top x \le 1 \,\forall x \in T_A(P)\right\}.$$

All x in $T_A(P)$ can be written as $x = Ax'$ for some $x' \in P$. Therefore, the polar of $T_A(P)$ can be written as

$$[T_A(P)]^0 = \left\{y \mid y^\top Ax' \le 1 \,\forall x' \in P\right\} = \left\{y \mid (A^\top y)^\top x' \le 1 \,\forall x' \in P\right\},$$
$$= \left\{y \mid A^\top y \in P^0\right\} = T_{A^\top}^{-1}(P^0). \quad \square$$

The duality between projection and slice is illustrated on Fig. 3. This notion of duality between projections and slices has a very important consequence: every property of the projection (resp. slice) of a polytope implies a related property for the slice (resp. projection) of a polytope. This means that any result that applies to projections also applies to slices.

Let us take a look at a small example in two dimensions to illustrate this duality. Let P be the square with vertices $(1, 1)$, $(1, -1)$, $(-1, 1)$ and $(-1, -1)$ (square with circle vertices on Fig. 4). Its polar, P^0, is the square with vertices $(0, 1)$, $(1, 0)$, $(-1, 0)$ and $(0, -1)$ (square with square vertices on Fig. 4). Let $A = \left(\begin{smallmatrix}1 & 0\\ 1 & 0\end{smallmatrix}\right)$. The projection T_1 of P with respect to A is

$$T_1(P) = \left\{A\binom{x}{y} \mid \binom{x}{y} \in P\right\} = \left\{\binom{x}{x} \mid \binom{x}{y} \in P\right\},$$

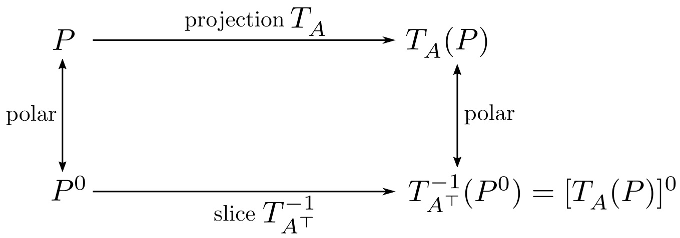

Fig. 3. Duality between projection and slice.

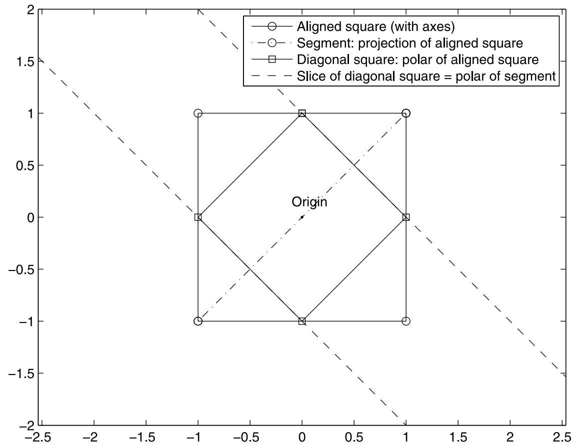

Fig. 4. Example of duality between projection and slice.

which is the line segment going from $(-1,-1)$ to $(1,1)$ (dot-dashed line segment on Fig. 4). Furthermore, the slice of P^0 with respect to $A^\top$ is

$$T_2^{-1}(P^0) = \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \mid A^\top \begin{pmatrix} x \\ y \end{pmatrix} \in P^0 \right\} = \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \mid \begin{pmatrix} x+y \\ 0 \end{pmatrix} \in P^0 \right\},$$

which implies that all the points $\begin{pmatrix} x \\ y \end{pmatrix}$ such that $-1 \leq x+y \leq 1$ belong to the preimage of P^0 (which is the region between the two dashed lines on Fig. 4). One can check (via the definition of polar transformation) that, in the two-dimensional space, the line segment obtained by projecting P and the region obtained by computing the preimage of P^0 are polar to each other.

3. Lower bound on the extension complexity

In this section, we introduce a new geometric lower bound for the extension complexity based on the behaviour of the f-vector of a polytope when projected on a lower dimensional subspace.

3.1. Monotonicity of the f-vector under linear maps

Our new lower bound on the extension complexity is based on the fact that the f-vector of any polytope decreases monotonically after a linear map, i.e, for a polytope P and any projection $T(P)$, we have $f_k(T(P)) \leq f_k(P)\,\forall k$. The reason we are interested in the behaviour of the f-vector under a linear map is that any extension for a given polytope is a higher dimensional polytope which projects on the original polytope. Therefore, for a polytope Q and any of its extensions P, we must have $f_k(Q) \leq f_k(P)\,\forall k$. The main result on which our bound is based is the following.

Lemma 3.1. *Let P be a polytope and let $T : x \mapsto Ax$ be a linear map. Then, $f_k(T(P)) \leq f_k(P)\,\forall k$.*

We will prove this property by identifying, for each k-face F of $T(P)$, a k-face F' of P such that $T(F') \subseteq F$ and $T(F') \nsubseteq \hat{F}$, for any other k-face $\hat{F} \neq F$ of $T(P)$. A k-face of P satisfies this condition if and only if one of its point is mapped in the relative interior of F.

Property 3.2. *Let P be a polytope, $T : x \mapsto Ax$ be a linear map, F be a k-face of $T(P)$ and F' be a k-face of P. If $T(F') \subseteq F$ and $\exists x \in F'$ such that $T(x) \in \mathrm{ri}(F)$, then $T(F') \nsubseteq \hat{F}$, for any other k-face $\hat{F} \neq F$ of $T(P)$.*

Proof. Let F be a k-face of $T(P)$, let F' be a k-face of P such that $T(F') \subseteq F$ and $\exists x \in F'$ such that $T(x) \in \text{ri}(F)$. From the second assertion in Proposition 2.6, we know that the intersection of faces of a polytope is also a face of this polytope. This implies that the intersection of two different k-faces is either empty of one of their faces which means that no two different k-faces have their relative interior in common. Thus, for any k-face $\hat{F} \neq F$ of $T(P)$, we have

$$T(x) \in \text{ri}(F) \Rightarrow T(x) \notin F \cap \hat{F} \Rightarrow T(x) \notin \hat{F} \Rightarrow T(F') \nsubseteq \hat{F}. \quad \square$$

To prove the main result, we need the following Lemma which states that the preimage restricted to the polytope of a face of its projection is a face of the polytope.

Lemma 3.3 (*[28, p. 196]*)**.** *Let $P \subseteq \mathbb{R}^n$ be a polytope and let $T : x \mapsto Ax$ be a linear map. Then for every face F of $T(P)$, the preimage $T^{-1}(F)|_P$ is a face of P.*

Furthermore, if F, G are faces of $T(P)$, then $F \subseteq G$ holds if and only if $T^{-1}(F) \subseteq T^{-1}(G)$.

Remark 3.4. From Property 2.10, if G is the face of P which is the preimage of a face F of $T(P)$, then $\dim G \geq \dim F$.

Let us prove a last intermediate result that will help us prove Lemma 3.1.

Property 3.5. *Let $P \subseteq \mathbb{R}^n$ be a polytope. Let A be a d-dimensional affine subspace of $\mathbb{R}^n$ such that $P \cap A \neq \emptyset$. Then, $I := P \cap A$ contains at least one point that belongs to a $(n-d)$-face of P.*

Proof. Let us prove this result by induction on d. We know that any intersection of a polytope with an affine subspace is a polytope [28, p. 29]. Thus, I is a polytope of dimension $\leq d$.

1. Base case: $d = 1$. In this case, I is either a single point on the boundary of P or I is a line segment in P with its extremities on the boundary of P. Any point on the boundary of P belongs to at least one facet of P, which is a $(n-1)$-faces of P.
2. Induction case: let us assume that if A is a d-dimensional affine subspace, I contains at least one point that belongs to a $(n-d)$-face of P. And, let us prove that, if A is a $(d+1)$-dimensional affine subspace, then I contains at least one point that belongs to a $(n-(d+1))$-face of P.
Let F be a facet of P such that $F \cap A \neq \emptyset$ and let H be the supporting hyperplane such that $F = P \cap H$. Such a facet always exists, otherwise $P \cap A = \emptyset$. By Proposition 2.6, we know that F is a $(n-1)$-dimensional polytope. Since A does not support P, we have $A \nsubseteq H$. Then, $\tilde{A} = A \cap H$ is an affine subspace of dimension d of H. Since we assumed that for a d-dimensional affine subspace, the intersection contains at least one point in a $(n-d)$-face of the polytope, we have that $F \cap \tilde{A}$ contains at least a point that belongs to a $((n-1)-d)$-face of F, i.e., a $(n-(d+1))$-face of P. Let x be this point, thus $x \in F \cap \tilde{A}$. Since $x \in \tilde{A}$ and $\tilde{A} \subset A$, $x \in A$. Thus, $x \in I = A \cap P$, i.e., I contains at least a point that belongs to a $(n-(d+1))$-face of P. $\square$

We now prove the main result of monotonic decrease of the f-vector under projections.

Proof of Lemma 3.1. Let F be a k-face of $T(P)$ and let $G = T^{-1}(F)|_P$ be the preimage of F, G is a $(k+\lambda)$-face of P, where λ is a nonnegative integer. Let $y \in \text{ri}(F)$ be a point in the relative interior of F. The preimage of y, $A_y := T^{-1}(y)$ is an affine subspace of dimension λ. Since $y \in F$, by definition of the preimage, $\exists x \in G$ such that $T(x) = y$. Thus, $G \cap A_y \neq \emptyset$ and, by Property 3.5, we know that $G \cap A_y$ contains at least one point that belongs to a k-face of G, let us denote by F' this k-face of G. Thus, by Proposition 2.6, F' is a k-face of P and F' has one of its point mapped in the relative interior of F. By Property 3.2, F' is such that $T(F') \subseteq F$ and $T(F') \nsubseteq \hat{F}$, for any other k-face $\hat{F} \neq F$ of $T(P)$ and we can conclude that $f_k(T(P)) \leq f_k(P)$. $\square$

Remark 3.6. This result is stated as an exercise without solution in [18, p. 37], however, we found no proof of it in the literature.

Using the duality between projections and slices defined before, we have the following corollary for slices.

Corollary 3.7. *Let P be an n-dimensional polytope and let $T : x \to A^\top x$ be a linear map such that $T^{-1}(P)$ is n'-dimensional. Then, $f_k(T^{-1}(P)) \leq f_{n-n'+k}(P)\, \forall k$.*

Proof. Let $(f_0, f_1, \ldots, f_{n-1})$ be the f-vector of P. By polarity, the f-vector of P^0 is

$$(g_0, g_1, \ldots, g_{n-1}) \text{ such that } g_k = f_{n-1-k}\, \forall k.$$

According to Lemma 3.1, let $T_2 : x \to Ax$ be a linear map, the f-vector of $T_2(P^0)$ is

$$(g'_0, \ldots, g'_{n'-1}) \text{ such that } g'_k \leq g_k\, \forall k.$$

According to Property 2.19, $T^{-1}(P) = [T_2(P^0)]^0$, thus the f-vector of $T^{-1}(P)$ is

$$(f'_0, \ldots, f'_{n'-1}) \text{ such that } f'_k = g'_{n'-1-k}\, \forall k.$$

Table 1
McMullen's upper bound for 6-dimensional polytopes.

$m = 7$	$f_0 \le 7$	$f_1 \le 21$	$f_2 \le 35$	$f_3 \le 35$	$f_4 \le 21$	$f_5 = 7$
$m = 8$	$f_0 \le 16$	$f_1 \le 48$	$f_2 \le 68$	$f_3 \le 56$	$f_4 \le 28$	$f_5 = 8$
$m = 9$	$f_0 \le 30$	$f_1 \le 90$	$f_2 \le 117$	$f_3 \le 84$	$f_4 \le 36$	$f_5 = 9$
$m = 10$	$f_0 \le 50$	$f_1 \le 150$	$f_2 \le 185$	$f_3 \le 120$	$f_4 \le 45$	$f_5 = 10$

This implies that, $\forall k$:

$$f'_k = g'_{n'-1-k} \le g_{n'-1-k} = f_{n-1-(n'-1-k)} = f_{n-n'+k},$$

hence that $f_k(T^{-1}(P)) \le f_{n-n'+k}(P)\,\forall k$. □

We have now proved that, for P, the original polytope and $X(P)$, an extension for P, we must have $f_k(X(P)) \ge f_k(P)\,\forall k$, according to Lemma 3.1 meaning that any extension of a polytope must always have more k-faces than the polytope itself.

Note that however, any polytope Q which satisfies $f_k(Q) \ge f_k(P)\,\forall k$ is not necessarily an extension for P. For example, if Q is a pentagon with f-vector (5,5) and P is the square with f-vector (4,4), the inequality is satisfied but it is not possible to project a pentagon onto a square.

3.2. A new geometric lower bound

If P is n-dimensional and has f_{n-1} facets, its extension complexity is then bounded below by the solution of

$$\min_{n<i<f_{n-1}(P)} \min_{Q,\dim(Q)=i} f_{i-1}(Q) \quad \text{s.t.} \quad f_k(Q) \ge f_k(P)\,\forall k = 0, \dots, n-1. \tag{1}$$

This optimization problem can be interpreted as follows. In dimension i, we look for all the polytopes Q which satisfy the projection theorem and we minimize the number of facets $f_{i-1}(Q)$ of such polytopes. Then, we take the minimum number of facets among all dimensions $n < i < f_{n-1}(P)$ between n, the dimension of P and $f_{n-1}(P)$, its number of facets. However, solving this optimization problem is very hard because there are many polytopes that satisfy the projection property and we cannot check them all. To compute a solution, we relax the problem using McMullen's upper bound theorem [20].

Theorem 3.8 (*Upper Bound Theorem, McMullen [1970]*)**.** *Let P be a n-dimensional convex polytope with m vertices. Then,*

$$f_k(P) \le \mathcal{F}(m, n, k)\,\forall k = 1, \dots, n-1,$$

where

$$\mathcal{F}(m, n, k-1) = \sum_{i=0}^{\lfloor \frac{n}{2} \rfloor} \frac{2 - \delta(2i, n)}{2} \left(\binom{n-i}{k-i} + \binom{i}{k-n+i} \right) \binom{m-n-1+i}{i},$$

and δ denotes the Kronecker delta function, i.e., $\delta(x, y) = 1$ if $x = y$ and $\delta(x, y) = 0$ if $x \ne y$.

This upper bound on the number of k-faces of any polytope is tight since it is attained by cyclic polytopes.

Remark 3.9. Note that, in the literature, McMullen's Upper Bound Theorem has several formulations [6,24]; for example,

$$f_k(P) \le \sum_{i=0}^{n} \binom{i}{k} \binom{m-1-\max(i, n-i)}{\min(i, n-i)},$$

or

$$f_k(P) \le \sum_{i=0}^{\lfloor \frac{n}{2} \rfloor} \left[\binom{i}{k} + (1 - \delta(2i, n)) \binom{n-i}{k} \right] \binom{m-n+i-1}{i}.$$

These formulas give the same upper bound on the number of k-faces of the polytope. An example is given in Table 1 for 6-dimensional polytopes ($n = 6$) and for several numbers of facets (m).

Furthermore, we know that vertices of any polytope give facets of its polar and conversely. Therefore, we can also compute the upper bounds knowing the number of facets rather than the number of vertices. More precisely, we have the following corollary.

Corollary 3.10. *Let P be a n-dimensional convex polytope with m facets. Then,*

$$f_k(P) \le \mathcal{F}(m, n, n-k-1)\,\forall k = 1, \dots, n-1.$$

Thus, all i-dimensional polytopes with f_{i-1} facets will have their numbers of k-faces bounded above by $f_k \leq \mathcal{F}(f_{i-1}, i, i-k-1)\,\forall k$. We can then replace the constraints

$$f_k(Q) \geq f_k(P)\,\forall k = 0, \ldots, n-1$$

by

$$\mathcal{F}(f_{i-1}, i, i-k-1) \geq f_k(P)\,\forall k = 0, \ldots, n-1$$

in the minimization problem (1). We have then a relaxation of the problem which implies the following Lemma.

Lemma 3.11. *Let P be a n-dimensional polytope and let $f_k(k = -1, \ldots, n)$ be the number of k-faces of P. The extension complexity of P can be bounded below as follows:*

$$\mathrm{xc}(P) \geq b(P) := \min_{n<i<f_{n-1}} \min_{f_{i-1}>i} f_{i-1}$$
$$\text{s.t.} \quad \mathcal{F}(f_{i-1}, i, i-k-1) \geq f_k(P)\,\forall k = 0, \ldots, n-1.$$

This optimization problem can be solved easily as follows. For every dimension i ($n < i < f_{n-1}$), we fix the number of facets $f_{i-1} = i+1$. We check the constraints $\mathcal{F}(f_{i-1}, i, i-k-1) \geq f_k(P)\,\forall k = -1, \ldots, n$ and if they are not satisfied, we increment f_{i-1} by 1. We stop incrementing as soon as the constraints are all satisfied. This gives us a minimum value of f_{i-1} for each i. Then, we take the minimum value of f_{i-1} among all dimensions i.

As explained before, this lower bound on the extension complexity of the polytope is also a lower bound on the nonnegative rank of its slack matrix. An important result for this bound is that our bound is always greater than the rank of this slack matrix.

Theorem 3.12. *Let P be a n-dimensional polytope and let $b(P)$ be the lower bound on the extension complexity of P described in Lemma 3.11. Then,*

$$b(P) \geq n+1.$$

Proof. Since P is n-dimensional, the bound from Lemma 3.11 will look for polytopes in dimensions greater or equal to n such that they satisfy the projection property. Moreover, all polytopes in dimensions greater or equal to n must have at least $n+1$ facets, as explained before. Therefore, the minimum number of facets among all these polytopes will always be greater or equal to $n+1$. This means that $b(P) \geq n+1$. □

As explained previously, the extension complexity of a polytope P is equal to the nonnegative rank of its slack matrix, thus $b(P)$ is a lower bound on the nonnegative rank of any slack matrix of P. Moreover, the rank of the slack matrix of a n-dimensional polytope is always $n+1$.

This lower bound can be improved if we take into account the fact that a polytope and its polar have the same extension complexity, meaning that the size of the optimal solution found must be the same for the polytope and its polar. Thus we can solve a second minimization problem with constraints

$$\mathcal{F}(f_{i-1}, i, i-k-1) \geq f_{n-k-1}(P)\,\forall k = 0, \ldots, n-1$$

and we have the following Lemma.

Lemma 3.13. *Let P be a n-dimensional polytope and let $f_k(k = 0, \ldots, n-1)$ be the number of k-faces of P. The extension complexity of P is bounded by*

$$\mathrm{xc}(P) \geq \max\{b(P), b(P^0)\},$$

with

$$b(P) = \min_{n<i<f_{n-1}} \min_{f_{i-1}>i} f_{i-1}$$
$$\text{s.t.} \quad \mathcal{F}(f_{i-1}, i, i-k-1) \geq f_k(P)\,\forall k = 0, \ldots, n-1,$$
$$b(P^0) = \min_{n<i<f_{n-1}} \min_{f_{i-1}>i} f_{i-1}$$
$$\text{s.t.} \quad \mathcal{F}(f_{i-1}, i, i-k-1) \geq f_{n-k-1}(P)\,\forall k = 0, \ldots, n-1.$$

Once again, if we use polarity, we can rewrite Lemma 3.13 as follows.

Lemma 3.14. *Let P be a n-dimensional polytope and let $f_k(k = -1, \ldots, n)$ be the number of k-faces of P. The extension complexity of P is bounded by*

$$\mathrm{xc}(P) \geq \max\{\bar{b}(P), \bar{b}(P^0)\},$$

with

$$\bar{b}(P) = \min_{n<i<f_{n-1}} \min_{f_0>i} f_0$$
$$\text{s.t.} \quad \mathcal{F}(f_0, i, i-n+k) \geq f_{n-k-1}(P) \, \forall k = 0, \ldots, n-1,$$
$$\bar{b}(P^0) = \min_{n<i<f_{n-1}} \min_{f_0>i} f_0$$
$$\text{s.t.} \quad \mathcal{F}(f_0, i, i-n+k) \geq f_k(P) \, \forall k = 0, \ldots, n-1.$$

This last Lemma can be interpreted as minimizing the number of vertices of a polytope such that the original polytope can be obtained as a slice of this polytope. This is equivalent to minimizing the number of facets such that the polar of the original polytope can be obtained as a projection by duality.

Note that this bound is monotone in the elements of the f-vector. Actually, if element f_i of the f-vector increases, the corresponding constraint will be strengthened in the optimization problems of Lemmas 3.13 and 3.14, therefore the optimal value can only increase.

3.3. Potential strengthening our bound

Since our bound can be expressed either as finding a polytope which projects onto the original polytope or as finding a polytope for which the original polytope is a slice, any property of the projection or the slice of polytopes could be used to improve the bound. For example, if we assumed that

Conjecture 3.15. *A two-dimensional hyperplane can cut at most two thirds of the edges of a three-dimensional convex polytope.*

is true, we could replace the constraint

$$\mathcal{F}(f_0, 3, 1) \geq f_0(P)$$

by

$$\mathcal{F}(f_0, 3, 1) \geq \frac{3}{2} f_0(P),$$

when trying to bound the extension complexity of a two-dimensional polytope. However, for two-dimensional polytope, $f_0(P) = f_1(P)$ and one can also check that for any value of f_0, we have $\mathcal{F}(f_0, 3, 1) = \frac{3}{2}\mathcal{F}(f_0, 3, 2)$. Therefore,

$$\mathcal{F}(f_0, 3, 2) \geq f_1(P) \Leftrightarrow \frac{2}{3}\mathcal{F}(f_0, 3, 1) \geq f_0(P),$$
$$\Leftrightarrow \mathcal{F}(f_0, 3, 1) \geq \frac{3}{2} f_0(P),$$

which means that the strengthened constraint was already satisfied and the lower bound is not improved. However, if similar conjectures can be generalized and proved for higher-dimensional polytopes and for all the components of the f-vector, i.e., the k-faces, we could strengthen the lower bound. More precisely, if we could find coefficients $c(n, i, k)$ which indicates the maximal proportion of the k-faces of an n-dimensional polytope that can be sliced by an i-dimensional affine subspace, we could replace the constraints

$$\mathcal{F}(f_0, i, i-n+k) \geq f_k(P)$$

by

$$\mathcal{F}(f_0, i, i-n+k) \geq \frac{1}{c(i, n, i-n+k)} f_k(P).$$

Given a general expression for the coefficients $c(n, i, k)$, we would be able to strengthen our lower bound on the extension complexity. Note that, if the conjecture is verified, $c(3, 2, 1) = \frac{2}{3}$.

4. Comparison with the bound from [15]

In this section, we recall the geometric lower bound on the nonnegative rank introduced by Gillis and Glineur [14,15] and we make a link between this bound and ours.

The authors introduced a new notion of rank for nonnegative matrices which leads to some bounds for the nonnegative rank. This rank itself is an upper bound for the nonnegative rank and its geometric interpretation leads to the lower bound on the nonnegative rank described below.

Definition 4.1. Let M be a $m \times n$ nonnegative matrix, the *restricted nonnegative rank* of M is the smallest positive integer r such that there exists an $m \times r$ nonnegative matrix U and an $r \times n$ nonnegative matrix V such that $M = UV$ and $\text{rank}(U) = \text{rank}(M)$. This rank is denoted $\text{rank}^*_+(M)$.

One can see that this new rank is an upper bound for the nonnegative rank ($\text{rank}_+(M) \leq \text{rank}^*_+(M)$) since any pair (U, V) such that $M = UV$ and $\text{rank}(U) = \text{rank}(M)$ is a nonnegative factorization.

The problem of finding this restricted nonnegative rank and a corresponding factorization can be formulated as follows:

(RNR) Given a nonnegative matrix $M \in \mathbb{R}^{m\times n}_+$, find $k = \text{rank}^*_+(M)$ and compute $U \in \mathbb{R}^{m\times k}_+$ and $V \in \mathbb{R}^{k\times n}_+$ such that $M = UV$ and $\text{rank}(U) = \text{rank}(M)$.

They claimed and partially proved that there is a polynomial time reduction from this problem to the nested polytopes problem. A complete proof was later given in [7].

(NPP) Given a bounded polyhedron

$$P = \{x \in \mathbb{R}^{r-1} \mid 0 \leq Cx + d\},$$

with $(C, d) \in \mathbb{R}^{m\times r}$ of rank r, and a set S of n points in P not contained in any hyperplane (i.e., the convex hull of S is full-dimensional), find the minimum number k of points in P whose convex hull T contains S, i.e., $S \subseteq T \subseteq P$.

This equivalence is a generalization of the result of Vavasis who showed the equivalence between exact NMF and intermediate simplex [26]. This means that any nonnegative matrix can be represented as two nested polytopes ($S \subseteq P$) and that the restricted nonnegative rank of the matrix would then be equal to the minimum number of vertices of any polytope T such that $S \subseteq T \subseteq P$. More precisely, we have another definition for the restricted nonnegative rank based on this equivalence.

Definition 4.2. Let M be a nonnegative matrix and let $S \subseteq P$ be the two polytopes of an NPP instance corresponding to M. Then, the *restricted nonnegative rank* of M is the minimum number of vertices of any polytope T such that $S \subseteq T \subseteq P$.

Then, Gillis and Glineur proved the following theorem

Theorem 4.3 (*[15]*). *Let M be a nonnegative matrix with $r = \text{rank}(M)$ and $r_+ = \text{rank}_+(M)$. Then,*

$$1 \leq \max_{r\leq r_u\leq r_+} \min\left(\frac{\mathcal{F}(r_+, r_u - 1, r_u - r)}{\text{rank}^*_+(M)}, \frac{\mathcal{F}(r_+, r_u - 1, r_u - 2)}{\text{rank}^*_+(M^\top)}\right),$$

where $\mathcal{F}$ denotes McMullen's upper bound theorem for polytopes.

The nonnegative rank of M is bounded from below by the smallest value of r_+ such that this inequality is satisfied.

If we know, or if we can compute, the restricted nonnegative rank of a matrix and its transpose, we can then compute a lower bound for the nonnegative rank. They also proved that if M is a slack matrix for a polytope P with v vertices and f facets then

$$\text{rank}^*_+(M) = v \quad \text{and} \quad \text{rank}^*_+(M^\top) = f.$$

Furthermore, if P is n-dimensional, then its slack matrix has rank $n + 1$.

Thus, for a slack matrix, the bound from [14,15] is the smallest value of r_+ such that the following inequality holds

$$1 \leq \max_{n\leq i\leq r_+} \min\left(\frac{\mathcal{F}(r_+, i, i - n)}{\#vertices(P)}, \frac{\mathcal{F}(r_+, i, i - 1)}{\#facets(P)}\right),$$

This lower bound can be rewritten as

$$\min_{n\leq i\leq r_+} r_+ \quad \text{s.t.} \quad \mathcal{F}(r_+, i, i - n) \geq f_0(P),$$
$$\mathcal{F}(r_+, i, i - 1) \geq f_{n-1}(P).$$

Remember that $\mathcal{F}(m, n, k)$ is the maximum number of k-faces of an n-dimensional polytope with m vertices. Thus, this optimization problem consists in the minimization of the number of vertices (r_+) of an i-dimensional polytope such that this polytope has more facets than the original polytope ($\mathcal{F}(r_+, i, i-1) \geq f_{n-1}(P)$) and more $(i-n)$-faces than the original polytope has vertices ($\mathcal{F}(r_+, i, i - n) \geq f_0(P)$). Actually, these constraints must be satisfied by any higher-dimensional polytope for which the original polytope can be obtained as a slice as stated in Corollary 3.7. Thus, the geometric interpretation of this bound (for slack matrices) is that we look for a higher-dimensional polytope with minimum number of vertices such that the polytope represented by the slack matrix can be obtained as the slice of this polytope. Therefore, our bound from Lemma 3.14, which has exactly the same geometric interpretation, can be seen as a strengthening of this bound since it uses the whole f-vector of the polytope and not just two components of this f-vector.

Table 2
Lower bounds on the extension complexity of low-dimensional regular polytopes.

	Dimension	# vertices	# facets	Bound from [16]	Bound from [15]	Bound from [14]	Our bound	Best known extension
				Geometric lower bounds				Upper bound
Square	2	4	4	2	4	4	**4**	4
Pentagon	2	5	5	3	4	5	**5**	5
Hexagon	2	6	6	3	4	5	**5**	5
Octagon	2	8	8	3	5	6	**6**	6
16-gon	2	16	16	4	6	8	**8**	8
32-gon	2	32	32	5	8	10	**10**	10
64-gon	2	64	64	6	9	11	**11**	12
2^{10}-gon	2	2^{10}	2^{10}	10	14	17	**17**	20
Cube	3	8	6	3	5	6	**6**	6
Dodecahedron	3	20	12	5	6	8	**8**	9
Cuboctahedron	3	12	14	4	6	7	**7**	8
Icosidodecahedron	3	30	32	5	7	10	**10**	?
24-cell	4	24	24	5	7	9	**10**	12
600-cell	4	120	600	7	9	16	**16**	47
6-cube	6	64	12	6	9	11	**11**	12

Table 3
Lower bounds on the extension complexity of uniform k_{21} polytopes.

	Dimension	# vertices	# facets	Bound from [16]	Bound from [15]	Bound from [14]	Our bound	Best known extension
				Lower bounds (Higher is better)				Upper bound
-1_{21}	3	6	5	3	5	5	**5**	?
0_{21}	4	10	10	4	6	7	**7**	?
1_{21}	5	16	26	4	7	9	**10**	?
2_{21}	6	27	99	5	8	12	**13**	?
3_{21}	7	56	702	6	10	17	**18**	?
4_{21}	8	240	19 440	8	12	24	**25**	?

5. Numerical results

We compute the lower bound from Lemma 3.13 for some low-dimensional regular polytopes and compare it with other geometric lower bounds (in Table 2) and non-geometric lower bounds (in Table 4). The polytopes were chosen because they possess a lot of symmetry, therefore we expect the extension complexity to be the smallest among the polytopes with same f-vector. The bound will then be closer to the actual extension complexity. Moreover, for these polytopes, we also know some upper bounds on the extension complexity coming from the number of facets of in known extensions of the polytope.

The last column indicates the size (number of facets) of the best known extension[1] of these polytopes. The four geometric lower bounds in Table 2 rely on some components of the f-vector of the polytope to bound the extension complexity. The bound from [16] uses only the number of vertices of the polytope, the bound from [15] uses the number of vertices and the dimension and the bound from [14] uses the numbers of vertices and facets and the dimension.

We observe that our bound, which uses all the components of the f-vector, is always equal or better than the other three. As expected, for two-dimensional polytopes (n-gons), our bound does not improve the bound from [15] since we do not add any information. However, the bound can be better for higher-dimensional polytopes; for example, for the 24-cell.

In Table 3, we present the results of our bound for a family of convex uniform polytopes called the k_{21} polytopes. Once again, we observe that our bound improves the bound from [15] for the higher-dimensional polytopes. Note that the three lower-dimensional polytopes in this family are better known under other names; the -1_{21} polytope is the triangular prism, the 0_{21} polytope is the rectified 5-cell and the 1_{21} polytope is the demipenteract.

This family of polytopes can also be used to show that the bound from Lemma 3.11 can be strictly lower than the bound from Lemma 3.13. Indeed, for the 3_{21} polytope, the bound from Lemma 3.11 is 16 and the bound from Lemma 3.13 is 18. Similarly, for the 4_{21} polytope, the bound from Lemma 3.11 is 22 and the bound from Lemma 3.13 is 25. The explanation is that these polytopes have much fewer vertices (56 and 240 respectively) than facets (702 and 19 440 respectively). This asymmetry causes the bound from Lemma 3.11 to be weaker for the polytope than for its polar.

1 Extensions which can be computed numerically or analytically.

Table 4
Lower bounds on the extension complexity of low-dimensional regular polytopes.

	Dimension	# vertices	# facets	Lower bounds: Rectangle covering number	Bound from [12]	Our bound	Upper bound: Best known extension
Square	2	4	4	4	4	**4**	4
Pentagon	2	5	5	5	5	**5**	5
Hexagon	2	6	6	5	5	**5**	5
Octagon	2	8	8	6	6	**6**	6
16-gon	2	16	16			**8**	8
32-gon	2	32	32			**10**	10
64-gon	2	64	64			**11**	12
2^{10}-gon	2	2^{10}	2^{10}			**17**	20
Cube	3	8	6	6	6	**6**	6
Dodecahedron	3	20	12	9		**8**	9
Cuboctahedron	3	12	14	8	8	**7**	8
Icosidodecahedron	3	30	32			**10**	?
24-cell	4	24	24			**10**	12
600-cell	4	120	600			**16**	47
6-cube	6	64	12			**11**	12

Table 5
Lower bounds on the extension complexity of 4-dimensional 0/1-polytopes.

	Gap between lower bound and xc: # 0-gap	Mean	Std	Max
Our bound	81	0.6931	0.6347	2
Bound from [16]	0	4.3515	0.7792	6
Bound from [15]	12	2.0446	0.9687	4
Bound from [14]	77	0.8762	0.8402	3
Bound from [11]	37	1.5149	1.0661	4
Bound from [12]	189	0.0644	0.2460	1
Fooling set bound	139	0.3762	0.6043	2
Rectangle covering number	195	0.0347	0.1834	1
	Computational time for lower bound: Mean	**Std**	**Max**	**Min**
Our bound	0.0012	6e−4	0.0041	<0.01
Bound from [16]	8e−6	1.5e−5	2e−4	<0.01
Bound from [15]	2.5e−4	3e−4	0.0021	<0.01
Bound from [14]	5.5e−4	2.7e−4	0.0022	<0.01
Bound from [11]	0.6417	0.1678	0.9980	0.3070
Bound from [12]	27.8910	33.8154	131.1148	0.1897
Fooling set bound	0.1729	0.0244	0.4566	0.1541
Rectangle covering number	0.2816	0.1451	1.3233	0.1500

In this next set of results, the two lower bounds reported in Table 4 (rectangle covering number and conic programming bound) are not based on a geometric interpretation. They compute a lower bound on the nonnegative rank of the slack matrix, i.e., the extension complexity of the polytope from the slack matrix.

Some entries of Table 4 are empty because the computational time is too large. For example, computing the bound from [12] for the cuboctahedron already takes around 1 h and 45 min on a laptop with 2 cores (4 threads) and 4 GB of RAM.

We also computed the value of our bound on the family of affinely independent 4-dimensional 0/1-polytopes. These are the polytopes which have all their vertices in $\{0, 1\}^4$ (see, for example, [21] for a description of these polytopes). For all these polytopes, the extension complexity is known exactly and satisfies $5 \leq \text{xc} \leq 10$. Table 5 displays the results for the different lower bounds presented before on this family of polytopes. For each polytope (there are 202 polytopes), we compute the lower bound then check the gap between its value and the actual extension complexity of the polytope. In the table, we show the number of polytopes for which the bound is tight, the average gap, the standard deviation and the largest gap. Table 5 also displays the average computational time, the standard deviation, the largest and smallest computational time (in seconds).

We observe that our bound is stronger than geometric bounds, while staying very fast to compute. It is sometimes slightly lower than some non-geometric bounds but is computationally much more efficient, hence can be applied to polytopes with significantly more vertices and facets.

6. Strengthening of our bound using flag numbers

In this section, we use a generalization of the f-vector to strengthen our bound. Let us first give the definition of flag numbers.

Definition 6.1. Let $P \subset \mathbb{R}^n$ be a polytope, a chain of faces of P, $F_1 \subset F_2 \subset \ldots \subset F_k \subset P$, is called an S-*flag*, where $S = \{\dim F_i : 1 \leq i \leq k\}$. The number of S-flags of P is denoted $f_S(P)$, and together these *flag numbers* form the *flag vector*, $(f_S(P))_{S\subseteq\{0,1,\ldots,n-1\}}$.

The flag vector is a generalization of the f-vector in which not only the faces are counted but also the number of inclusions of lower-dimensional faces into higher-dimensional faces. For example, $f_{\{0,1\}}$ is the number of inclusions of vertices into edges.

If $S = \{k\}$ for $0 \leq k \leq n-1$, then $f_S(P) = f_k(P)$, the number of k-faces of P.

We can strengthen our bound by adding constraints on the flag numbers of an extension of the polytope. Indeed, Lemma 3.1 can be generalized to flag numbers.

Property 6.2. *Let $P \subset \mathbb{R}^n$ be a polytope and let $T : x \mapsto Ax$ be a linear map. Then, $f_S(T(P)) \leq f_S(P) \, \forall S \subset \{0, 1, \ldots, n\}$.*

Proof. By Lemma 3.1, for any k-face of $T(P)$, we can identify a k-face of P that is mapped onto it, and, we know that inclusions are preserved after projections. Thus, for each flag of $T(P)$, we can also identify a flag of P. □

To add constraints on the flag numbers in our bound, we need an upper bound on the flag numbers of a polytope which depends on the dimension and the number of vertices. The following Lemma [3] is a generalization of McMullen's upper bound theorem.

Lemma 6.3. *Let $P \subset \mathbb{R}^n$ be a polytope with v vertices. Then,*

$$f_S(P) \leq f_S(C(n, v)) \, \forall S \subseteq \{0, 1, \ldots, n-1\},$$

where $C(n, v)$, denotes the n-dimensional cyclic polytope with v vertices.

Once again, using polarity, we have the following corollary.

Corollary 6.4. *Let $P \subset \mathbb{R}^n$ be a polytope with m facets. Then,*

$$f_S(P) \leq f_S(C^*(n, m)) \, \forall S \subseteq \{0, 1, \ldots, n-1\},$$

where $C^*(n, m)$, denotes the n-dimensional cyclic polytope with m facets.

And, since cyclic polytopes are simplicial, we can compute their flag numbers easily using the following property.

Property 6.5 (*[3]*). *Let $P \subset \mathbb{R}^n$ be a simplicial polytope and let $S = \{i_1, \ldots, i_s\}$ with $i_1 < i_2 < \cdots < i_s < n$. Then,*

$$f_S(P) = f_{i_s}(P) f_{S\setminus\{i_s\}}(T^{i_s}),$$

where T^{i_s} is the i_s-dimensional simplex and $f_{i_s}(P)$ is the number of i_s-faces of P.

We can then strengthen our bound by adding constraints on the flag numbers. The formulation is similar as before.

Lemma 6.6. *Let P be a n-dimensional polytope and let $f_S(S \subseteq \{0, 1, \ldots, n-1\})$ be the number of S-flags of P. The extension complexity of P is bounded by*

$$\mathrm{xc}(P) \geq \max\{\beta(P), \beta(P^0)\},$$

with

$$\beta(P) = \min_{n<i<f_{n-1}} \min_{f_{i-1}>i} f_{i-1} \quad \text{s.t.} \quad f_S(C^*(i, f_{i-1})) \geq f_S(P) \, \forall S \subseteq \{0, 1, \ldots, n-1\},$$

$$\beta(P^0) = \min_{n<i<f_{n-1}} \min_{f_{i-1}>i} f_{i-1} \quad \text{s.t.} \quad f_S(C^*(i, f_{i-1})) \geq f_S(P^0) \, \forall S \subseteq \{0, 1, \ldots, n-1\}.$$

Let us call $b_1(P)$ the bound from Lemma 3.13 and let us call $b_2(P)$ the bound from Lemma 6.6. Since there are more constraints in the formulation, we know that $b_2(P) \geq b_1(P)$. However, we may wonder if there are cases where $b_2(P) > b_1(P)$ for some polytope P.

For any polytope, the flag numbers must satisfy a set of equalities and inequalities. Firstly, for any n-dimensional polytope, the face numbers (elements of the f-vector) must satisfy the Euler relation, i.e.,

$$\sum_{i=0}^{n-1}(-1)^i f_i = 1-(-1)^n.$$

Secondly, the flag numbers must satisfy the so-called generalized Dehn–Sommerville equations (see [4] for details).

Theorem 6.7. *Let P be a n-dimensional polytope, and $S \subseteq \{0, 1, \ldots, n-1\}$. If $\{i, k\} \subseteq S \cup \{-1, n\}$, $i < k-1$, and S contains no j such that $i < j < k$, then*

$$\sum_{j=i+1}^{k-1}(-1)^{j-i-1} f_{S\cup j}(P) = f_S(P)(1-(-1)^{k-i-1}).$$

Lastly, for any dimension, a set of inequalities can be defined. For example, in dimension 4, two inequalities that the flag numbers of any polytope must satisfy are $f_0 \geq 5$ and $f_3 \geq 5$. Then, we can generate all vectors that satisfy this set of equalities and inequalities and check if one of these vectors is such that $b_2 > b_1$. In dimension $n \leq 3$, we always have $b_2 = b_1$ because all flag numbers that are not elements of the f-vector can be expressed as linear combinations of the elements of the f-vector. In dimension $n = 4$, there is one flag number that cannot be expressed as linear combinations of the elements of the f-vector. However, it can be bounded above by linear combinations of the elements of the f-vector. Thus, in dimension 4, we also have $b_2 = b_1$ for any vector. A set of inequalities that the flag numbers of any 4-dimensional polytope must satisfy is given in [5]. In dimension $n = 5$, we could find several vectors that satisfy the set of equalities and inequalities and such that $b_2 > b_1$. For example, the following vector

$$\begin{aligned}&(f_0, f_1, f_2, f_3, f_4, f_{\{0,1\}}, f_{\{0,2\}}, f_{\{0,3\}}, f_{\{0,4\}}, f_{\{1,2\}}, f_{\{1,3\}}, f_{\{1,4\}}, f_{\{2,3\}}, f_{\{2,4\}},\\ &f_{\{3,4\}}, f_{\{0,1,2\}}, f_{\{0,1,3\}}, f_{\{0,1,4\}}, f_{\{0,2,3\}}, f_{\{0,2,4\}}, f_{\{0,3,4\}}, f_{\{1,2,3\}}, f_{\{1,2,4\}}, f_{\{1,3,4\}},\\ &f_{\{2,3,4\}}, f_{\{0,1,2,3\}}, f_{\{0,1,2,4\}}, f_{\{0,1,3,4\}}, f_{\{0,2,3,4\}}, f_{\{1,2,3,4\}}, f_{\{0,1,2,3,4\}}) =\\ &(10, 25, 38, 32, 11, 50, 172, 244, 122, 172, 366, 244, 546, 546, 64, 344, 732,\\ &488, 732, 732, 488, 732, 732, 732, 1092, 1464, 1464, 1464, 1464, 1464, 2928)\end{aligned}$$

satisfies the set of equalities and inequalities and is such that $b_1 = 8$ and $b_2 = 10$. A set of inequalities that the flag numbers of any 5-dimensional polytope must satisfy is given in [19].

However, we do not know if this vector is the flag vector of an actual polytope but this hints that using flag numbers could improve the value of our bound for polytopes of dimension higher or equal to 5.

7. Generalization of the bound for any nonnegative matrix

We presented a stronger lower bound for slack matrices of polytopes that relies on the f-vector of the polytope. Now we would like to generalize this bound for nonnegative matrices which are not slack matrices. This means that we should define a new vector for nonnegative matrices with the constraint that this new vector must be equal to the f-vector in the particular case of a slack matrix.

In [15], the authors defined a new rank which is equal to the number of vertices of the polytope in the case of a slack matrix. We would like to generalize this rank for all the k-faces. From Definition 4.2, the generalization is rather straightforward.

Definition 7.1. Let M be a nonnegative matrix and let $S \subseteq P$, be the two polytopes of an NPP instance corresponding to M, the *k-restricted nonnegative rank* of M is the minimum number of k-faces of any polytope T such that $S \subseteq T \subseteq P$. We denote it $k\text{-rank}^*_+(M)$.

The restricted nonnegative rank defined in [15] is then equal to the 0-restricted nonnegative rank, $\text{rank}^*_+(M) = 0\text{-rank}^*_+(M)$. The restricted nonnegative rank of the transpose matrix is then equal to the $(r-2)$-restricted nonnegative rank, $\text{rank}^*_+(M^\top) = (r-2)\text{-rank}^*_+(M)$, where r is the classical rank of M.

Note that, for slack matrices, the k-restricted nonnegative rank is equal to the number of k-faces of the polytope it represents. This is a consequence of the fact that, the NPP instance corresponding to a slack matrix is such that $S = P$ which is the polytope represented by the matrix. Thus, the only polytope T such that $S \subseteq T \subseteq P$ is the polytope itself.

Given these new ranks, we can generalize our bound for any nonnegative matrix.

Lemma 7.2. *Let M be a nonnegative matrix, the nonnegative rank of M is bounded by*

$$\begin{aligned}\text{rank}_+(M) \geq \min_{i\geq n} & f_{i-1}\\ \text{s.t.}\quad & \mathcal{F}(f_{i-1}, i, i-k-1) \geq k\text{-}\,\text{rank}^*_+(M)\ \forall k = -1, \ldots, n.\end{aligned}$$

However, computing the restricted nonnegative rank is NP-hard when the classical rank is larger than 4 [9,10,15]. So there is no reason to believe that computing the k-restricted nonnegative ranks could be done in polynomial time, unless P = NP. Nonetheless, if we could compute good estimates of the value of these k-restricted nonnegative ranks (more precisely, lower bounds), we would be able to compute a lower bound for the nonnegative rank of any nonnegative matrix. To use our bound in the general case, it is then required to define lower bounds on the k-restricted nonnegative ranks of a nonnegative matrix.

8. Conclusion

In this paper, we have introduced a new lower bound for the extension complexity of a polytope based on its f-vector. We have also shown that this bound could be improved by using flag numbers, and can be applied to any nonnegative matrix by generalizing the notion of restricted nonnegative rank.

This bound improves upon existing geometric bounds for slack matrices and can be computed very efficiently, which we illustrated on many different polytopes.

Acknowledgements

This paper presents research results of the Concerted Research Action (ARC) programme supported by the Federation Wallonia-Brussels, Belgium (contract ARC 14/19-060). This work was also supported by the Fonds de la Recherche Scientifique — FNRS, Belgium under Grant EOS 30468160 (SeLMA — Structured low-rank matrix/tensor approximation: numerical optimization-based algorithms and applications). The authors acknowledge the support by the Fonds de la Recherche Scientifique — FNRS, Belgium and the Fonds Wetenschappelijk Onderzoek — Vlanderen (FWO), Belgium under EOS Project no O005318F-RG47. N. Gillis acknowledges the support by the European Research Council (ERC starting grant n° 679515).

References

[1] A.V. Aho, J.D. Ullman, M. Yannakakis, On notions of information transfer in vlsi circuits, in: Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing, ACM, 1983, pp. 133–139.
[2] S. Arora, R. Ge, R. Kannan, A. Moitra, Computing a nonnegative matrix factorization–provably, in: Proceedings of the 44th Symposium on Theory of Computing, ACM, 2012, pp. 145–162.
[3] M. Bayer, L. Billera, Counting Faces and Chains in Polytopes and Posets, Tech. rep., Cornell University Operations Research and Industrial Engineering, 1983.
[4] M.M. Bayer, L.J. Billera, Generalized dehn-sommerville relations for polytopes, spheres and eulerian partially ordered sets, Invent. Math. 79 (1) (1985) 143–157.
[5] M.M. Bayer, C.W. Lee, Combinatorial aspects of convex polytopes, in: P. Gruber, J. Wills (Eds.), Handbook of Convex Geometry, North-Holland, Amsterdam, 1993, pp. 485–534.
[6] A. Bondesen, A. Bronsted, A dual proof of the upper bound conjecture for convex polytopes, Math. Scand. 46 (1980) 95–102.
[7] D. Chistikov, S. Kiefer, I. Marusic, M. Shirmohammadi, J. Worrell, On restricted nonnegative matrix factorization, in: I. Chatzigiannakis, M. Mitzenmacher, Y. Rabani, D. Sangiorgi (Eds.), 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016), in: Leibniz International Proceedings in Informatics (LIPIcs), vol. 55, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2016, pp. 103:1–103:14.
[8] J.E. Cohen, U.G. Rothblum, Nonnegative ranks, decompositions, and factorizations of nonnegative matrices, Linear Algebra Appl. 190 (1993) 149–168.
[9] G. Das, M.T. Goodrich, On the complexity of optimization problems for 3-dimensional convex polyhedra and decision trees, Comput. Geom. 8 (3) (1997) 123–137.
[10] G. Das, D. Joseph, The complexity of minimum convex nested polyhedra, in: Proc. 2nd Canad. Conf. Comput. Geom, 1990, pp. 296–301.
[11] H. Fawzi, P.A. Parrilo, New Lower Bounds on Nonnegative Rank using Conic Programming, Tech. rep. Citeseer, 2012.
[12] H. Fawzi, P.A. Parrilo, Self-scaled bounds for atomic cone ranks: applications to nonnegative rank and cp-rank, Math. Program. 158 (1–2) (2016) 417–465.
[13] S. Fiorini, V. Kaibel, K. Pashkovich, D.O. Theis, Combinatorial bounds on nonnegative rank and extended formulations, Discrete Math. 313 (1) (2013) 67–83.
[14] N. Gillis, Nonnegative Matrix Factorization: Complexity, Algorithms and Applications (Ph.D. thesis), Univeristé catholique de Louvain, 2011.
[15] N. Gillis, F. Glineur, On the geometric interpretation of the nonnegative rank, Linear Algebra Appl. 437 (11) (2012) 2685–2712.
[16] M.X. Goemans, Smallest compact formulation for the permutahedron, Math. Program. 153 (1) (2015) 5–11.
[17] J. Gouveia, R. Grappe, V. Kaibel, K. Pashkovich, R.Z. Robinson, R.R. Thomas, Which nonnegative matrices are slack matrices? Linear Algebra Appl. 439 (10) (2013) 2921–2933.
[18] B. Grünbaum, Convex Polytopes, second ed., Springer, 2003.
[19] G. Kalai, On low-dimensional faces that high-dimensional polytopes must have, Combinatorica 10 (3) (1990) 271–280.
[20] P. McMullen, The maximum numbers of faces of a convex polytope, Mathematika 17 (02) (1970) 179–184.
[21] M. Oelze, A. Vandaele, S. Weltge, Computing the extension complexities of all 4-dimensional 0/1-polytopes, arXiv preprint arXiv:1406.4895.
[22] T. Rothvoß, The matching polytope has exponential extension complexity, in: Proceedings of the 46th Annual ACM Symposium on Theory of Computing, ACM, 2014, pp. 263–272.
[23] A. Schrijver, Combinatorial Optimization: Polyhedra and Efficiency, Vol. 24, Springer Science & Business Media, 2003.
[24] R. Seidel, The upper bound theorem for polytopes: an easy proof of its asymptotic version, Comput. Geom. 5 (2) (1995) 115–116.
[25] A. Vandaele, N. Gillis, F. Glineur, D. Tuyttens, Heuristics for exact nonnegative matrix factorization, J. Global Optim. (2015) 1–32.
[26] S.A. Vavasis, On the complexity of nonnegative matrix factorization, SIAM J. Optim. 20 (3) (2009) 1364–1377.
[27] M. Yannakakis, Expressing combinatorial optimization problems by linear programs, J. Comput. System Sci. 43 (1991) 441–466.
[28] G.M. Ziegler, Lectures on Polytopes, Vol. 152, Springer, 1995.