



Optimal Bounds on the Positivity of a Matrix from a Few Moments

Gemma De las Cuevas¹ , Tobias Fritz², Tim Netzer³

¹ Institute of Theoretical Physics, Technikerstr. 21a, 6020 Innsbruck, Austria.

E-mail: gemma.delascuevas@uibk.ac.at

² Perimeter Institute for Theoretical Physics, 31 Caroline St N, Waterloo, ON N2L 2Y5, Canada

³ Department of Mathematics, Technikerstr. 13, 6020 Innsbruck, Austria

Received: 10 September 2018 / Accepted: 23 December 2019

Published online: 11 March 2020 – © Springer-Verlag GmbH Germany, part of Springer Nature 2020

Abstract: In many contexts one encounters Hermitian operators M on a Hilbert space whose dimension is so large that it is impossible to write down all matrix entries in an orthonormal basis. How does one determine whether such M is positive semidefinite? Here we approach this problem by deriving asymptotically optimal bounds to the distance to the positive semidefinite cone in Schatten p -norm for all integer $p \in [1, \infty)$, assuming that we know the moments $\mathrm{tr}(M^k)$ up to a certain order $k = 1, \dots, m$. We then provide three methods to compute these bounds and relaxations thereof: the sos polynomial method (a semidefinite program), the Handelman method (a linear program relaxation), and the Chebyshev method (a relaxation not involving any optimization). We investigate the analytical and numerical performance of these methods and present a number of example computations, partly motivated by applications to tensor networks and to the theory of free spectrahedra.

1. Introduction

1.1. The setting. Positive semidefinite matrices and operators come up in a large number of contexts throughout mathematics, physics and computer science, including the theory of operator algebras, computations with quantum states, or semidefinite relaxations of optimization problems. A problem that arises frequently in many of these contexts is determining whether a given matrix M or operator is positive semidefinite (psd). In certain applications—we describe two of them in Sects. 1.2 and 1.3— M is of such a large dimension that it is impossible to express M in an orthonormal basis and store the result in a computer, let alone diagonalize it or compute determinants of principal minors. Instead, we assume that one can compute a few of the normalized moments $\mathrm{tr}(M^k)$, where we define the normalized trace as

$$\mathrm{tr}(M) := s^{-1} \mathrm{tr}(M),$$

where s is the size of M . In this paper, we answer the following questions:

- (i) Given the first m normalized moments $\text{tr}(M^k)$ for $k = 1, \dots, m$ of a Hermitian operator M with $\|M\|_\infty \leq 1$, can one show that M is not psd?
- (ii) Given these moments and a $p \in [1, \infty)$, can one optimally bound the distance of M to the psd cone from above and below in Schatten p -norm?

Since both the moments and the positive semidefiniteness of a Hermitian operator M are characterized by the distribution of eigenvalues—or more generally by the corresponding Borel measure on the spectrum of M —we are secretly concerned with a version of the truncated Hausdorff moment problem. In these terms, the above two questions become:

- (i) Given only the moments

$$\mathbb{E}_\mu[x^k] := \int_{-1}^{+1} x^k d\mu$$

of a compactly supported Borel measure μ on $[-1, 1]$ for $k = 1, \dots, m$, can one show that μ is not supported on $[0, 1]$?

- (ii) Given these moments and a $p \in [1, \infty)$, can one optimally bound the p -Wasserstein distance [27] between μ and the set of probability measures supported on $[0, 1]$ from above and from below?

Given this connection with the moment problem, it should not come as a surprise that our methods also work in certain infinite-dimensional situations. While we focus on the matrix case in the main text, we sketch the extension to testing positivity of a Hermitian element in a von Neumann algebra equipped with a finite faithful trace in “Appendix A”.

Let us now motivate the assumption that we have access to a few moments of M . In the applications we have in mind, the space where M lives has a natural tensor product structure, with respect to which M can be expressed as

$$M = \sum_{j=1}^r A_j^{[1]} \otimes \dots \otimes A_j^{[n]}, \tag{1}$$

where each $A_j^{[i]}$ is Hermitian and of reasonably small dimension.¹ Note that every M on a tensor product space can be written this way, for large enough r . In our example applications below, r is taken to be fixed and typically small, or not scaling with n . Thus, the naturally available operations are those that can be directly performed in terms of the local matrices $A_j^{[i]}$. This includes taking powers of M and taking the trace, which gives us access to the moments of M :

$$\text{tr}(M^k) = \sum_{j_1, \dots, j_k=1}^r \text{tr}(A_{j_1}^{[1]} \dots A_{j_k}^{[1]}) \dots \text{tr}(A_{j_1}^{[n]} \dots A_{j_k}^{[n]}), \tag{2}$$

for $k \in \mathbb{N}$ much smaller than the size of the matrix M , namely $s \times s$.

¹ If $r = 1$, then there is a simple criterion to determine whether M is psd: M is psd if and only if each $A^{[i]}$ is either psd or negative semidefinite, and the number of negative semidefinite matrices is even. But we are not aware of any such simple criterion for $r > 1$.

1.2. Example application: tensor networks. Our first example application concerns quantum states of a many-body system, which are modelled by psd matrices on a Hilbert space $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \dots \otimes \mathcal{H}_n$, where \mathcal{H}_i is the Hilbert space associated to subsystem i . Typically, all \mathcal{H}_i are of the same finite dimension. Since the dimension of \mathcal{H} grows exponentially with n , physicists have attempted to develop a scalable description of quantum many-body systems, that is, one which grows only polynomially in n . This is the objective of the program of *tensor networks* [6, 19, 20, 26]. While this program has been very successful for pure states (i.e. psd matrices of rank 1), its success for mixed states has been more limited. One of the reasons for that is the *positivity problem*, which is the following. In the tensor network paradigm, it is natural to use a few matrices for each local Hilbert space \mathcal{H}_i . For example, the state of the system in one spatial dimension (with periodic boundary conditions) is described by

$$M = \sum_{j_1, \dots, j_n=1}^r A_{j_1, j_2} \otimes A_{j_2, j_3} \otimes \dots \otimes A_{j_n, j_1}, \tag{3}$$

where each $A_{j_l, j_{l+1}}$ is a Hermitian matrix in \mathcal{H}_l [4, 25]. (In general, the local matrices $A_{j_l, j_{l+1}}$ may also depend on the site l , but we do not consider this case for notational simplicity.) Now, while M must be psd to describe a quantum state, each of the local matrices $A_{j_l, j_{l+1}}$ need not be psd. While there is a way of imposing positivity in the local matrices (resulting in the ‘local purification form’), this generally comes at the price of a very large increase in the number of matrices, thus making the representation very inefficient [4].

Indeed, the hardness of deciding whether objects of the kind of (3) are psd has been studied. Specifically:

Problem 1. Given $\{A_{j, j'} \in \text{Her}_s\}_{j, j'=1}^r$ with $s, r \geq 7$, let

$$M(n) := \sum_{j_1, \dots, j_n=1}^r A_{j_1, j_2} \otimes A_{j_2, j_3} \otimes \dots \otimes A_{j_n, j_1}.$$

Decide whether $M(n) \geq 0$ for all n .

Proposition 1. [5] *Problem 1 is undecidable.*

This holds true even if all matrices $A_{j, j'}$ are diagonal and their entries rational. Variants of this problem are also undecidable [14], and deciding whether it is psd for a finite number of system sizes and with open boundary conditions is NP-complete [14]. See also [28] for further perspectives on this problem.

1.3. Example application: free spectrahedra. Our second example application is in the area of convex optimization, where we find the same algebraic structures as the ones we have considered so far, albeit often studied from a different angle. Namely, given a tuple (B_1, \dots, B_r) of Hermitian matrices, its associated *spectrahedron* [3] is defined as

$$S(B_1, \dots, B_r) = \left\{ (y_1, \dots, y_r) \in \mathbb{R}^r \mid \sum_{i=1}^r y_i B_i \geq 0 \right\},$$

where \geq denotes positive semidefiniteness. Note that it is the intersection of an affine space with the set of psd matrices. Spectrahedra are precisely the feasible sets of semidefinite programming (SDP). Characterizing which convex sets are spectrahedra is a main objective in the area of algebraic convexity [3], and has direct implications for the many applications of semidefinite programming.

Recently, a non-commutative generalization of spectrahedra has been proposed, called *free spectrahedra* [10], defined as

$$\text{FS}(B_1, \dots, B_r) = \bigcup_{s=1}^{\infty} \left\{ (A_1, \dots, A_d) \in \text{Her}_s^r \mid \sum_{i=1}^r A_i \otimes B_i \geq 0 \right\}. \quad (4)$$

Thus, asking whether $\sum_{i=1}^r A_i \otimes B_i$ is psd is equivalent to asking whether (A_1, \dots, A_r) is in the free spectrahedron defined by (B_1, \dots, B_r) . This is again a problem of the form (1), with $n = 2$. Surprisingly, many things about standard spectrahedra can be learned by examining their free versions. For example, the inclusion test of spectrahedra proposed in [1] was fully understood in [11] as an inclusion test of free spectrahedra, opening the way to analyzing the exactness of the method (see [7] and the references therein). Also, two minimal matrix tuples defining the same spectrahedron are unitarily equivalent if and only if they define the same free spectrahedron [11]. So the different free spectrahedra over a standard spectrahedron characterize equivalence classes of its defining matrix tuples. Applying these results often means checking whether a free spectrahedron contains a certain matrix tuple, i.e. whether $\sum_i A_i \otimes B_i$ is psd. Since the matrices might be of very large size, this is again a context in which our methods can be applied.

1.4. Related work. The methods we use to solve the problems described above are fairly standard: a combination of techniques used for moment problems combined with results on sums of squares representations of positive polynomials. One might therefore expect there to exist a substantial amount of literature on the problem which we solve in this work, but this does not seem to be the case.

There is only one work that we are aware of Lasserre [17] has investigated the smallest interval $[a_m, b_m]$ on which the support of a measure on \mathbb{R} can be shown to be contained, given only its first m moments. One can think of this as providing a solution to the problems discussed above in the case $p = \infty$. Lasserre found that the lower bound a_m and upper bound b_m are the optimal solutions of two simple semidefinite programs involving the moments.

1.5. Overview. The rest of this paper is structured as follows. In Sect. 2, we characterize the distance of a matrix to the psd cone. In Sect. 3, we provide upper and lower bounds to the distance to the psd cone by using a few moments of the matrix. In Sect. 4, we provide three methods to compute these bounds: the sos polynomial method, the Handelman method and the Chebyshev method. In Sect. 5 we analyse the numerical performance of these methods. Finally, in ‘‘Appendix A’’, we will sketch the extension to von Neumann algebras.

2. The Distance to the psd Cone

We will start with some preliminaries (Sect. 2.1) and then define the negative part function and the distance to the psd cone (Sect. 2.2).

2.1. Preliminaries. Let us first fix some notation. For a matrix M , we denote its Hermitian conjugate by M^* . Her_s denotes the set of Hermitian matrices of size $s \times s$. For $M \in \text{Her}_s$, $M \geq 0$ denotes that M is psd, and $M \leq 0$ that it is negative semidefinite (i.e. $-M \geq 0$).

We now state some basic facts about matrices and their norms. Consider $M \in \text{Her}_s$ and its spectral decomposition $M = U^*DU$, where U is a unitary matrix, $D = \text{diag}(\lambda_1, \dots, \lambda_s)$, and $\{\lambda_1, \dots, \lambda_s\} =: \text{sp}(M)$ is the spectrum of M (considered as a multiset). Any real-valued function f defined on $\text{sp}(M)$ can be defined on M by setting

$$f(M) := U^*f(D)U,$$

where $f(D) = \text{diag}(f(\lambda_1), \dots, f(\lambda_s))$. For example, the absolute value $|M|$ of M is defined this way. For $p \in [1, \infty)$, we define the Schatten p -norm of M as

$$\|M\|_p := (\text{tr}(|M|^p))^{1/p},$$

where taking the normalized trace instead of the usual trace introduces an additional factor of $s^{-1/p}$ with respect to the usual definition. This definition guarantees that $\|I\|_p = 1$, where I is the identity matrix (of any size). The case $p = 2$ corresponds to the normalized Frobenius norm, and the case $p = 1$ to the normalized trace norm. Note also that if p is even, the norm is easier to compute, since in this case the absolute value is superfluous, so that $\|M\|_p = (\text{tr}(M^p))^{1/p}$, which is simply the p th root of the p th moment of M . The operator norm of M induced by the standard Euclidean norm on \mathbb{C}^s is defined as

$$\|M\|_\infty := \max_{\lambda_i \in \text{sp}(M)} |\lambda_i|.$$

We have that $\|M\|_\infty = \lim_{p \rightarrow \infty} \|M\|_p$.

Remark 1. In the following we will often assume that $\|M\|_\infty \leq 1$, i.e. that the spectrum of M lies in $[-1, 1]$. This can clearly be achieved by a suitable scaling of M . Of course, since our main problem is that $\text{sp}(M)$ cannot be computed, we cannot simply scale by $\|M\|_\infty$. But for $1 \leq p \leq q \leq \infty$, we have

$$\|M\|_p \leq \|M\|_q \leq s^{\frac{1}{p} - \frac{1}{q}} \|M\|_p.$$

So we can divide M by $s^{1/p} \|M\|_p$ (for any $p \in \mathbb{N}$) in order to achieve $\|M\|_\infty \leq 1$. Moreover, since $s^{\frac{1}{\log s}} = e$, the norm $\|M\|_{\log s}$ is only a constant factor away from $\|M\|_\infty$.²

2.2. The negative part function and the distance to the psd cone. Given $p \in \mathbb{N}$, let us define the *negative part function* f_p as

$$f_p(x) := \begin{cases} 0 & x \geq 0 \\ |x|^p & x < 0. \end{cases} \quad (5)$$

² This has been pointed out to us by Richard Kueng.

So, for example, $f_1(x)$ is the absolute value of the negative part of x . For $M \in \text{Her}_s$, we set

$$M_- := f_1(M) \quad \text{and} \quad M_+ := M + M_-,$$

so that we obtain the natural definition $M = M_+ - M_-$, where both the positive part M_+ and the negative part M_- are psd. (To define functions of matrices we follow the standard procedure [13]. Namely, if the matrix is diagonalisable, as in our case $M = UDU^*$, then $f(M) = Uf(D)U^*$, where $f(D) = \text{diag}(f(\lambda_1), f(\lambda_2), \dots)$, where $\{\lambda_i\}$ are the eigenvalues.)

Given a Hermitian matrix M of size s , we are interested in its distance to the cone of psd matrices of size s , PSD_s , with respect to the Schatten p -norm, namely

$$d_p(M) := \inf_{N \in \text{PSD}_s} \|M - N\|_p. \quad (6)$$

We now show that this is given by the negative part of M . That is, the best psd approximation to a Hermitian matrix is given by the positive part of this matrix.

Lemma 1. *For $p \in [1, \infty)$ and $M \in \text{Her}_s$, the matrix M_+ is the point in the psd cone that is closest to M , with respect to any Schatten p -norm. That is,*

$$d_p(M) = \|M_-\|_p = \mathbf{tr}(f_p(M))^{1/p}.$$

Proof. Clearly $M_+ \geq 0$, so that the distance can at most be $\|M_-\|_p$. Now, for any matrix $A \in \text{Her}_s$, denote by $\sigma(A)$ the diagonal matrix with the eigenvalues of A on the diagonal, in decreasing order. It follows from [2] (IV.62) that

$$\|P - M\|_p \geq \|\sigma(P) - \sigma(M)\|_p \quad (7)$$

holds for all $P \in \text{Her}_s$ and all p . If $P \geq 0$, then clearly

$$\|P - M\|_p \geq \|\sigma(P) - \sigma(M)\|_p \geq \left(\frac{1}{s} \sum_{\lambda \in \text{sp}(M), \lambda < 0} |\lambda|^p \right)^{1/p} = \|M_-\|_p.$$

This proves the claim. \square

Remark 2. Note that $\|\sigma(P) - \sigma(M)\|_p$ is precisely the p -Wasserstein distance [27] between the spectral measures $s^{-1} \sum_i \delta_{\sigma(P)_{ii}}$ and $s^{-1} \sum_i \delta_{\sigma(M)_{ii}}$ of P and M . So $d_p(M)$ is in fact also the p -Wasserstein distance of the spectral measure of M to the cone of probability measures supported on $[0, \infty)$. \square

3. Bounds on the Distance

We will now bound the negative part function f_p by polynomials (Sect. 3.1) and then show the asymptotic optimality of our bounds (Sect. 3.2).

3.1. Bounding the negative part function by polynomials. Clearly, M is psd if and only if its distance to the psd cone is zero, $d_p(M) = 0$, and $d_p(M)$ is a measure of how far M differs from being psd. Thus, ideally, we would like to compute $d_p(M)$. Since we only have access to a few moments of M , we will use them to estimate $d_p(M)$ as accurately as possible.

We start by showing that $d_p(M)$ can be approximated arbitrarily well by the trace of polynomial expressions in M . For $q \in \mathbb{R}[x]$, we write $f_p \leq q$ if this holds pointwise on the interval $[-1, 1]$.

Lemma 2. *Suppose $M \in \text{Her}_s$ with $\|M\|_\infty \leq 1$ and $p \in \mathbb{N}$. Then*

$$d_p(M)^p = \inf_{f_p \leq q \in \mathbb{R}[x]} \mathbf{tr}(q(M)) = \sup_{f_p \geq q \in \mathbb{R}[x]} \mathbf{tr}(q(M)).$$

Proof. From $f_p \leq q$ we obtain that $f_p(M) \leq q(M)$, thus $\mathbf{tr}(f_p(M)) \leq \mathbf{tr}(q(M))$ and finally

$$d_p(M)^p = \mathbf{tr}(f_p(M)) \leq \mathbf{tr}(q(M)).$$

Conversely, by standard Weierstrass approximation for continuous functions on compact sets, for any $\varepsilon > 0$ there exists a polynomial $q \in \mathbb{R}[x]$ with

$$f_p \leq q \leq f_p + \varepsilon.$$

The same argument as above then shows

$$\mathbf{tr}(q(M)) \leq d_p(M)^p + \varepsilon.$$

This proves the first equation, and the second follows in the same fashion. \square

Thus, any polynomials q_1 and q_2 such that $q_1 \leq f_p \leq q_2$ give the bounds

$$\mathbf{tr}(q_1(M)) \leq d_p(M)^p \leq \mathbf{tr}(q_2(M)).$$

In particular, this means that if $\mathbf{tr}(q_1(M)) > 0$ then M is not psd. The quality of the bounds depends on the quality of the approximation of f_p by q_1, q_2 on $\text{sp}(M)$, or more generally on $[-1, 1]$ —see Fig. 1 for an example.

More generally, given a polynomial q which does not satisfy $q \leq f_p$ or $q \geq f_p$, what we can say is that

$$q - \|(q - f_p)_+\|_\infty \leq f_p \leq q + \|(q - f_p)_-\|_\infty, \tag{8}$$

where we use notation for positive and negative part as in the matrix case, and similarly $\|g\|_\infty := \sup_{x \in [-1, 1]} |g(x)|$. Here, the function on the left hand side corresponds to shifting q by an additive constant until it is below f_p , and similarly for the right hand side—see Fig. 2 for an example. This leads to the following result.

Theorem 1. *Let $M \in \text{Her}_s$ with $\|M\|_\infty \leq 1$, $p \in \mathbb{N}$, as well as $q \in \mathbb{R}[x]$. Then*

$$\mathbf{tr}(q(M)) - \|(q - f_p)_+\|_\infty \leq d_p(M)^p \leq \mathbf{tr}(q(M)) + \|(q - f_p)_-\|_\infty. \tag{9}$$

In particular:

- (i) If $f_p \leq q$, then $d_p(M) \leq \mathbf{tr}(q(M))^{1/p}$.
- (ii) If $\mathbf{tr}(q(M)) > \|(q - f_p)_+\|_\infty$, then M is not psd.
- (iii) If $q \leq f_p$ and $\mathbf{tr}(q(M)) > 0$, then M is not psd.




Fig. 1. Polynomials of degree 7 that approximate f_2 from below and from above, together with the spectrum of M (black dots). The polynomials have been obtained with the Handelman method, to be described in Sect. 4.2




Fig. 2. A polynomial q which approximates f_2 but does not satisfy $q \leq f_2$ or $q \geq f_2$. The polynomial is a Chebyshev polynomial of degree 2, obtained with the Chebyshev method (Sect. 4.3). Then it is shifted up and down by $\|q - f_2\|_\infty$

Remark 3. We can also try to prove that a matrix is psd with the same approach, although this seems to be much harder in practice. Note the following: if for $M \in \text{Her}_s$ we find $d_p(M) \leq \varepsilon$, then $d_\infty(M) \leq s^{1/p}\varepsilon$, and therefore $M + s^{1/p}\varepsilon I_s \geq 0$. So if we first replace M by

$$M_\varepsilon := M - s^{1/p}\varepsilon I_s,$$

and then show that $d_p(M_\varepsilon) \leq \varepsilon$, then we have proven that $M \geq 0$. By Theorem 1, this can be achieved by finding $q \in \mathbb{R}[x]$ with

$$\text{tr}(q(M_\varepsilon)) + \|(q - f_p)_-\|_\infty \leq \varepsilon^p. \tag{10}$$

Note that the k th moment of M_ε can be computed from the moments of M of order $\leq k$. Further, if M is strictly positive definite, then this strategy does indeed work: there is $\varepsilon > 0$ with $M_\varepsilon \geq 0$, i.e. $d_p(M_\varepsilon) = 0$. In view of Lemma 2, we can also find some q to make (10) hold, so that this method can indeed detect that M is psd. However, we might need to test for both very small ε and polynomials q of very high degree before obtaining a positive answer. \square

3.2. *Optimality of the bounds.* We now show that given the first m moments of M , the best bounds to the distance to the psd cone which are independent of the size of M are given by a polynomial approximation. More precisely, given $M \in \text{Her}_s$ and $p, m \in \mathbb{N}$, define

$$\begin{aligned} d_{p,m}^+(M) &:= \inf_{f_p \leq q, \deg(q) \leq m} \text{tr}(q(M))^{1/p} \\ d_{p,m}^-(M) &:= \sup_{q \leq f_p, \deg(q) \leq m} \text{tr}(q(M))^{1/p}. \end{aligned} \quad (11)$$

Clearly, these numbers depend only on the first m moments of M , and they lower and upper bound $d_p(M)$,

$$0 \leq d_{p,m}^-(M) \leq d_p(M) \leq d_{p,m}^+(M). \quad (12)$$

We now show that these are the optimal upper and lower bounds to $d_p(M)$ which can be obtained from the first m normalized moments of M and which are independent of the size of M . We thus call them *asymptotically optimal*.

The following result is a variation on a classical result from the theory of moment problems [15, Theorem 4.1(a)], for which we offer a self-contained proof.

Theorem 2. *For any matrix $M \in \text{Her}_s$ with $\|M\|_\infty \leq 1$ and any $m \in \mathbb{N}$,*

– *For every $\varepsilon > 0$, there are $N_1, N_2 \in \text{Her}_t$ (for suitably large t) such that*

$$|\text{tr}(M^k) - \text{tr}(N_i^k)| \leq \varepsilon \quad \text{for } k = 1, \dots, m,$$

and for which

$$d_p(N_1) \geq d_{p,m}^+(M), \quad d_p(N_2) \leq d_{p,m}^-(M).$$

– *There are operators N_1 and N_2 in a finite von Neumann algebra \mathcal{N} such that*

$$\text{tr}(M^k) = \text{tr}(N_i^k) \quad \text{for } k = 1, \dots, m,$$

and which saturate the bounds,

$$d_p(N_1) = d_{p,m}^+(M), \quad d_p(N_2) = d_{p,m}^-(M).$$

Note that in the first case the moments are approximately reproduced but N_i has finite size — we will see an example thereof in Example 1. In the second case, the moments are exactly reproduced but the size of N_i may need to be infinite.

Proof. We only construct N_1 , since N_2 is obtained similarly.

Consider the linear functional

$$\begin{aligned} \varphi: \mathbb{R}[x]_{\leq m} &\rightarrow \mathbb{R} \\ q &\mapsto \text{tr}(q(M)) \end{aligned}$$

on the space of polynomials of degree at most m , which clearly maps polynomials nonnegative on $[-1, 1]$ to nonnegative numbers. Let us define the real vector space

$$V := \{q + rf_p \mid q \in \mathbb{R}[x]_{\leq m}, r \in \mathbb{R}\}.$$

We extend φ to a linear functional ψ on V by setting

$$\psi(f_p) := \inf_{f_p \leq q \in \mathbb{R}[x]_{\leq m}} \varphi(q).$$

We claim that ψ still maps nonnegative functions in V to nonnegative numbers. So let $h + rf_p \geq 0$ for some $h \in \mathbb{R}[x]_{\leq m}$ and $r \in \mathbb{R}$. The case $r = 0$ is clear, so assume $r < 0$. Then $-\frac{1}{r}h \geq f_p$, and thus

$$\begin{aligned} \psi(h + rf_p) &= \varphi(h) + r \inf_{f_p \leq q} \varphi(q) \\ &\geq \varphi(h) + r\varphi\left(-\frac{1}{r}h\right) \\ &= \varphi(h) - \varphi(h) = 0. \end{aligned}$$

If $r > 0$ instead, then $f_p \leq q$ implies $0 \leq h + rf_p \leq h + rq$, and thus

$$\varphi(h) + r\varphi(q) = \varphi(h + rq) \geq 0.$$

Passing to the infimum over these q proves the statement.

Since $\mathbb{R}[x]_{\leq m}$ already contains an interior point of the convex cone of nonnegative continuous functions on $[-1, 1]$ (such as the constant function 1), we can further extend ψ to a positive linear functional Ψ on the whole of $\mathcal{C}([-1, 1])$, using the Riesz Extension Theorem [21]. By the Riesz Representation Theorem [23], there exists a positive measure μ on $[-1, 1]$ such that

$$\Psi(f) = \int f d\mu$$

for all $f \in \mathcal{C}([-1, 1])$. From $\mu([-1, 1]) = \int 1 d\mu = \Psi(1) = \varphi(1) = 1$, we see that μ is a probability measure. We now take $\mathcal{N} := L^\infty([-1, +1], \mu)$, equipped with integration against μ as a finite normalized trace, and define N_1 to be the multiplication operator by the identity function, $N_1 : f \mapsto xf$.

So for $k = 0, \dots, m$, the k th moment of N_1 is given by

$$\int x^k d\mu = \Psi(x^k) = \varphi(x^k) = \mathbf{tr}(M^k),$$

and we also have

$$d_p(N_1) = \int f_p d\mu = \Psi(f_p) = \psi(f_p) = \inf_{f_p \leq q \in \mathbb{R}[x]_{\leq m}} \varphi(q) = d_{p,m}^+(M)^p,$$

which establishes the first claim.

Concerning the realization by finite-dimensional matrices, we use the well-known fact that each probability measure on $[-1, 1]$ can be approximated arbitrarily well by uniform atomic measures with respect to the weak- $*$ topology. Concretely, we can approximate μ by the uniform atomic measure $\nu_t := \frac{1}{t} \cdot \sum_{i=1}^t \delta_{a_i}$, where the a_i are the right t -quantiles,

$$a_i := \inf \left\{ r \in [-1, 1] \mid \mu([r, 1]) \geq \frac{i}{t} \right\}.$$

This choice guarantees that the cumulative distribution function of ν_t dominates the one of μ_t . Therefore the expectation value of μ is not smaller than that of ν_t on any monotonically nonincreasing function. In particular, we have

$$\int f_p d\nu_t \geq \int f_p d\mu.$$

Furthermore, since ν_t weak-* converges³ to μ as $t \rightarrow \infty$, we can moreover choose t large enough such that

$$\left| \int x^k d\mu - \int x^k d\nu_t \right| \leq \varepsilon \quad \text{for } k = 0, \dots, m.$$

For $N_1 := \text{diag}(a_1, \dots, a_t) \in \text{Her}_t$, we then have $\|N_1\|_\infty \leq 1$ and

$$\text{tr}(N_1^k) = \frac{1}{t} \sum_{i=1}^t a_i^k = \int x^k d\nu_t$$

as well as

$$\int f_p d\nu_t = \frac{1}{t} \sum_{a_i < 0} |a_i|^p = \|N_1\|_p^p = d_p(N_1)^p,$$

which gives the desired bounds. This altogether finishes the proof. \square

Example 1. We show that $\varepsilon = 0$ can in general not be achieved in the second part of Theorem 2. Taking $m = 2$, let us consider the matrix⁴

$$M = \text{diag} \left[c + \sqrt{c(1-c)}, c - \sqrt{c(1-c)} \right],$$

for $c \in (0, 1/2)$. Since the lower right entry is negative, M is not psd. Its first moment is c , the second moment is $\frac{1}{2}(2c^2 + 2c(1-c)) = c$, equal to the first moment. Looking for a probability measure μ on $[0, 1]$ with these moments, we must have $\mathbb{E}_\mu[x(1-x)] = 0$, which implies that μ must be supported on $\{0, 1\}$ only. Since $\mu = (1-c)\delta_0 + c\delta_1$ does indeed have these moments, we conclude that it is the unique measure on $[0, 1]$ with these moments; and the fact that such a μ exists implies $d_{p,2}^-(M) = 0$, irrespectively of the value of p . However, as soon as c is irrational, the measure $(1-c)\delta_0 + c\delta_1$ is not of the form $s^{-1} \sum_{i=1}^s \delta_{\lambda_i}$ for any finite sequence $(\lambda_1, \dots, \lambda_s)$. In particular, there is no psd matrix of finite size with the same moments as M , and $\varepsilon = 0$ cannot be achieved in Theorem 2. By a standard compactness argument, this also implies that one must take $t \rightarrow \infty$ as $\varepsilon \rightarrow 0$ in the theorem. \square

³ One way to see this is by the Portmanteau theorem: the cumulative distribution functions of μ and ν_t differ by at most t^{-1} at every point, and therefore we have (even uniform) convergence as $t \rightarrow \infty$, which implies weak-* convergence $\nu_t \rightarrow \mu$.

⁴ Although M is only of size 2×2 , one can clearly achieve the same moments on larger matrices by simply repeating the eigenvalues.

Remark 4. An interesting question⁵ is how close our bounds $d_{p,m}^+(M)$ and $d_{p,m}^-(M)$ are guaranteed to be to the actual value $d_{p,m}(M)$. In the worst case, the actual value will coincide with one of the two bounds, in which case the other bound differs from the actual value by

$$d_{p,m}^+(M) - d_{p,m}^-(M) \leq 12 \frac{p}{m},$$

which follows from (29) and (26).

4. Algorithms

We now present our numerical methods to compute lower and upper bounds to the distance to the psd cone, in order of decreasing accuracy and complexity: the sos polynomial method (Sect. 4.1), the Handelman method (Sect. 4.2) and the Chebyshev method (Sect. 4.3). The sos polynomial method involves solving a semidefinite program, the Handelman method involves solving a linear program, and the Chebyshev method does not require any optimization. We will compare the numerical performance of three methods in Sect. 5.

Throughout, we fix nonzero $p \in \mathbb{N}$.

4.1. The sos polynomial method. The sos polynomial method solves the optimization problems of Eq. (11) exactly,⁶ and thereby computes $d_{p,m}^+(M)$ and $d_{p,m}^-(M)$. We start by explaining how to compute the upper bound $d_{p,m}^+(M)$ via a semidefinite program.

To be able to talk about polynomials only, we first split the condition $f_p \leq q$ into two parts:

$$0 \leq q(-x) - x^p \quad \text{and} \quad 0 \leq q(x), \quad \text{both for all } x \in [0, 1]. \tag{13}$$

Now note that any polynomial of the form

$$\sigma_0 + \sigma_1 x + \sigma_2(1 - x) + \sigma_3 x(1 - x), \tag{14}$$

where the σ_i are sums of squares of polynomials, is nonnegative on $[0, 1]$. In fact, the converse is true as well:

Theorem 3 [16, 18]. *If $q \in \mathbb{R}[x]_{\leq m}$ is nonnegative on $[0, 1]$, then there exist sums of squares $\sigma_0, \sigma_1, \sigma_2, \sigma_3 \in \mathbb{R}[x]$ such that*

$$q(x) = \sigma_0 + \sigma_1 x + \sigma_2(1 - x) + \sigma_3 x(1 - x)$$

where the degree of each σ_i can be chosen such that each summand has degree at most $\leq m$.

⁵ Communicated to us by Boaz Barak.

⁶ This is not to be confused with the sos polynomial method of [4], which is a semidefinite program that computes $\min_p \|M - p(M)\|_1$, where p is a sos polynomial of given degree m . The goal of the method of [4] is to approximate M as well as possible with a sos polynomial (as this provides a purification), which is possible only if M is psd. Note moreover that $\|M - p(M)\|_1$ cannot be computed from the moments of M .

So assume that we find such representations for both polynomials in (13),

$$\begin{aligned} q(-x) - x^p &= \sigma_0 + \sigma_1 x + \sigma_2(1-x) + \sigma_3 x(1-x), \\ q(x) &= \tau_0 + \tau_1 x + \tau_2(1-x) + \tau_3 x(1-x), \end{aligned}$$

with sums of squares (sos) $\sigma_i, \tau_i \in \mathbb{R}[x]$. Then we have clearly ensured $f_p \leq q$. This can be rewritten as

$$\begin{aligned} q(x) &= (-x)^p + \tilde{\sigma}_0 - \tilde{\sigma}_1 x + \tilde{\sigma}_2(1+x) - \tilde{\sigma}_3 x(1+x) \\ &= \tau_0 + \tau_1 x + \tau_2(1-x) + \tau_3 x(1-x), \end{aligned} \tag{15}$$

where the $\tilde{\sigma}_i(x) := \sigma_i(-x)$ are again sos.

Now assume that every term in (15) has degree $\leq m$. This imposes obvious degree bounds on the sums of squares $\tilde{\sigma}_i, \tau_i$, namely $\deg(\tilde{\sigma}_0) \leq m$, and $\deg(\tilde{\sigma}_1), \deg(\tilde{\sigma}_2) \leq m-1$ as well as $\deg(\tilde{\sigma}_3) \leq m-2$, and analogously for the τ_i , and note that every sos polynomial must have an even degree. It is easy to see that every sum of squares can be written as

$$\begin{aligned} \tilde{\sigma}_i &= (1, x, \dots, x^{l_i}) S_i (1, x, \dots, x^{l_i})^t, \quad S_i \geq 0 \\ \tau_i &= (1, x, \dots, x^{l_i}) T_i (1, x, \dots, x^{l_i})^t, \quad T_i \geq 0 \end{aligned} \tag{16}$$

where $l_i = \deg(\tilde{\sigma}_i)/2$ and similarly for τ_i . Writing each $\sigma_i, \tilde{\tau}_i$ in such a way, using matrices with unknown entries, and then comparing coefficients with respect to x in (15), leads to the problem of finding psd matrices with certain linear constraints on the entries. Any solution to this problem will provide a polynomial $q \in \mathbb{R}[x]_{\leq m}$ with $f_p \leq q$. Among all of them, we want to minimize $\mathbf{tr}(q(M))$, which is a linear function in the entries of the unknown matrices, having the moments of M as coefficients. Optimizing a linear function over an affine section of a cone of psd matrices is known as semidefinite programming, for which there exist efficient algorithms.

The derivation of the lower bound $d_{p,m}^-(M)$ is entirely parallel, except for the fact that in (13) the two inequalities are reversed. This implies that

$$\begin{aligned} -q(-x) + x^p &= \sigma_0 + \sigma_1 x + \sigma_2(1-x) + \sigma_3 x(1-x), \\ -q(x) &= \tau_0 + \tau_1 x + \tau_2(1-x) + \tau_3 x(1-x), \end{aligned}$$

where σ_i, τ_i are sos, and thus

$$\begin{aligned} q(x) &= (-x)^p - \tilde{\sigma}_0 + \tilde{\sigma}_1 x - \tilde{\sigma}_2(1+x) + \tilde{\sigma}_3 x(1+x) \\ &= -\tau_0 - \tau_1 x - \tau_2(1-x) - \tau_3 x(1-x). \end{aligned} \tag{17}$$

In summary, we have obtained:

Proposition 2 (Sos polynomial method). *The sos polynomial method at level m computes the upper bound $d_{p,m}^+(M)^p$ (defined in (11)) by solving the semidefinite program*

$$\begin{aligned} \min \quad & \mathbf{tr}(q(M)) \\ \text{subject to} \quad & \text{Eq. (15)} \\ & \text{and Eq. (16)} \end{aligned} \tag{18}$$




Fig. 3. Lower and upper sos polynomial approximation of degree 7 for $p = 2$ and for the matrix M whose spectrum is shown in black dots

It also computes the lower bound $d_{p,m}^-(M)^p$ (defined in (11)) by solving the semidefinite program

$$\begin{aligned} & \max \operatorname{tr}(q(M)) \\ & \text{subject to Eq. (17)} \\ & \text{and Eq. (16)} \end{aligned} \tag{19}$$

As an example, Fig. 3 shows the sos polynomial approximation of degree 7 obtained for a matrix M with the indicated spectrum. We will discuss numerical results more systematically in Sect. 5.

Remark 5. For computational purposes, it may be better to use the ansatz

$$\sigma_0 + \sigma_1 x(1 - x) \tag{20}$$

for a polynomial nonnegative on $[0, 1]$, instead of (14). The advantage is that this reduces the number of sums of squares from 4 to 2. An analogue of Theorem 3 still holds [16], but with slightly weaker degree bounds: a degree m polynomial nonnegative on $[0, 1]$ has a representation $\sigma_0 + \sigma_1 x(1 - x)$ with $\deg(\sigma_0) \leq m + 1$ and $\deg(\sigma_1) \leq m - 1$ only.⁷ So if we set up our optimization problem as above, but with the simpler representation (20) where we demand $\deg(\sigma_0) \leq m$ and $\deg(\sigma_1) \leq m - 2$ (since moments of M are available only up to order m), then we will obtain a bound on $d_p(M)$ which lies between $d_{p,m}^+(M)^p$ and $d_{p,m-1}^+(M)^p$. \square

4.2. The Handelman method. The Handelman method relaxes the optimization of Eq. (11) by using another ansatz for nonnegative polynomials on $[0, 1]$. This results in a linear optimization problem, which can usually be solved much faster than a semidefinite problem.

We start by splitting the condition $f_p \leq q$ as in (13). This time note that any polynomial of the form

$$\sum_{\alpha \in \mathbb{N}^2} b_\alpha x^{\alpha_1} (1 - x)^{\alpha_2}$$

⁷ This can also be deduced directly from Theorem 3, using $x = x^2 + x(1 - x)$ and $1 - x = (1 - x)^2 + x(1 - x)$.

with $b_\alpha \geq 0$ (and only finitely many different from 0) is nonnegative on $[0, 1]$. So if we find coefficients $b_\alpha, c_\alpha \geq 0$ with

$$\begin{aligned} q(-x) - x^p &= \sum_{\alpha \in \mathbb{N}^2} b_\alpha x^{\alpha_1} (1-x)^{\alpha_2}, \\ q(x) &= \sum_{\alpha \in \mathbb{N}^2} c_\alpha x^{\alpha_1} (1-x)^{\alpha_2}, \end{aligned}$$

then we can be sure to have $f_p \leq q$. Assume that the degree of the polynomials is a priori bounded by m . Then comparing coefficients with respect to x yields a finite system of linear equations:

$$\begin{aligned} q(x) &= (-x)^p + \sum_{|\alpha| \leq m} b_\alpha (-x)^{\alpha_1} (1+x)^{\alpha_2} \\ &= \sum_{|\alpha| \leq m} c_\alpha x^{\alpha_1} (1-x)^{\alpha_2}, \end{aligned} \tag{21}$$

where $|\alpha| = \alpha_1 + \alpha_2$. We look for solutions under the constraint that $(b_\alpha, c_\alpha)_{|\alpha| \leq m}$, and among all these solutions, we look for the one that minimizes the quantity $\mathbf{tr}(q(M))$. This is precisely a linear optimization problem, where information about our matrix M enters through the objective function.

The derivation of the lower bound of $d_{p,m}^-(M)$ is analogous, except that in this case $-q(-x) + x^p$ and $-q(x)$ must be nonnegative polynomials on $[0, 1]$. This leads to

$$\begin{aligned} q(x) &= (-x)^p - \sum_{|\alpha| \leq m} b_\alpha (-x)^{\alpha_1} (1+x)^{\alpha_2} \\ &= \sum_{|\alpha| \leq m} -c_\alpha x^{\alpha_1} (1-x)^{\alpha_2}, \end{aligned} \tag{22}$$

In summary we have obtained:

Proposition 3 (Handelman method). *The Handelman method at level m computes an upper bound of $d_{p,m}^+(M)$ by solving the linear program*

$$\begin{aligned} &\min \mathbf{tr}(q(M)) \\ &\text{subject to Eq. (21)} \\ &\text{and } b_\alpha \geq 0, c_\alpha \geq 0 \end{aligned} \tag{23}$$

It also computes a lower bound of $d_{p,m}^-(M)$ by solving the linear program

$$\begin{aligned} &\max \mathbf{tr}(q(M)) \\ &\text{subject to Eq. (22)} \\ &\text{and } b_\alpha \geq 0, c_\alpha \geq 0 \end{aligned} \tag{24}$$

Note that linear optimization problems are easy to solve algorithmically (for example, interior point methods have polynomial complexity, but the simplex algorithm and its variants often work best in practice). See Fig. 1 above for upper and lower Handelman approximations of f_2 for a given matrix M .

Although this method computes only a relaxation of $d_{p,m}^+(M)$ (and analogously for $d_{p,m}^-(M)$), the following special case of Handelman's Positivstellensatz for polytopes [9] ensures that these relaxations converge to the exact value $d_p(M)$ in the limit $m \rightarrow \infty$:

Theorem 4 (Handelman). *If $q \in \mathbb{R}[x]$ is strictly positive on $[0, 1]$, then*

$$q(x) = \sum_{\alpha \in \mathbb{N}^2} a_\alpha x^{\alpha_1} (1-x)^{\alpha_2}$$

for certain $a_\alpha \geq 0$, only finitely many different from 0.

Note that this result leads directly to the standard solution of the Hausdorff moment problem in terms of complete monotonicity [24, Theorem 1.5]. We now have:

Corollary 1. *Let $\tilde{d}_{p,m}^+(M)$ denote the p th root of the optimal value of the linear program described in Method 3. Then we have $d_p(M) \leq d_{p,m}^+(M) \leq \tilde{d}_{p,m}^+(M)$, and*

$$\tilde{d}_{p,m}^+(M) \xrightarrow{m \rightarrow \infty} d_p(M).$$

Proof. Every feasible point in the program leads to a polynomial $q \in \mathbb{R}[x]_{\leq m}$ with $f_p \leq q$. This proves $d_{p,m}^+(M) \leq \tilde{d}_{p,m}^+(M)$. Now if $q \in \mathbb{R}[x]$ fulfills $f_p \leq q$, then for any $\varepsilon > 0$ there is some $m \in \mathbb{N}$, such that $q + \varepsilon$ corresponds to a feasible point in the respective program, by Handelman’s theorem. This proves the claim. \square

4.3. The Chebyshev method. The Chebyshev method chooses q_m as the Chebyshev polynomial of degree m that best approximates f_p , and uses (9) to derive bounds on $d_p(M)$. This has the advantage of not involving any optimization at all, i.e. one need only compute the Chebyshev polynomials q_m once, and the method can be applied to any matrix M .

Ideally, the best bounds of Eq. (9) are given by the polynomial q of degree m that minimizes $\|q_m - f_p\|_\infty$, which we call q_m^* . By Jackson’s theorem (see e.g. [22, Theorem 1.4]) we have that

$$\|f - q_m^*\|_\infty \leq 6 \omega(f, 1/m). \tag{25}$$

where $\omega(f, \delta)$ denotes the modulus of continuity of a uniformly continuous function f , defined as

$$\omega(f, \delta) = \sup_{\substack{x_1, x_2 \in [-1, 1] \\ |x_1 - x_2| \leq \delta}} |f(x_1) - f(x_2)|.$$

For the negative part function f_p , we thus obtain that

$$\|f_p - q_m^*\|_\infty \leq 6 \left(1 - \left(1 - m^{-1}\right)^p\right) \leq 6 \frac{p}{m}, \tag{26}$$

where the second estimate is by Bernoulli’s inequality, which is tight up to an error of $O(m^{-2})$.

However, it is in general not possible to find q_m^* (called the minimax approximation) analytically. Instead, Chebyshev polynomials provide a good proxy: they are close to the minimax approximation, and are straightforward to compute explicitly [22]. To obtain an analytical upper bound of the additional error, it is known [8] that the Chebyshev approximation q_m of a continuous function f satisfies that

$$\|q_m - f\|_\infty \leq C \omega(f, 1/m) \log m$$




Fig. 4. $\|q_m - f_p\|_\infty$ (obtained numerically) for the Chebyshev polynomial q_m of (28) as a function of m for several p

for some constant C . For our negative part function f_p , it thus holds that

$$\|q_m - f_p\|_\infty \leq Cm^{-1} \log m. \tag{27}$$

Numerically we can see the behaviour of $\|q_m - f_p\|_\infty$ as a function of m and p in Fig. 4.

Let us now construct the Chebyshev interpolating polynomial of degree m , q_m , explicitly. It interpolates f_p at the Chebyshev nodes

$$a_k = \cos(\pi(k + 1/2)/(m + 1)) \quad \text{for } k = 0, 1, \dots, m,$$

so that q_m is given by

$$q_m = \sum_{k=0}^m f_p(a_k) \ell_k,$$

where the ℓ_k are the Lagrange interpolation polynomials

$$\ell_k(x) = \frac{1}{\prod_{i \neq k} (a_k - a_i)} \prod_{i \neq k} (x - a_i).$$

One can also express q_m in terms of the Chebyshev polynomials of the first kind, which are defined as $t_k(x) = \cos(k \arccos(x))$ for $|x| \leq 1$. In this basis, q_m takes the form

$$q_m = \sum_{k=0}^m c_k t_k - c_0/2, \quad \text{with } c_k = \frac{2}{m+1} \sum_{j=0}^m f_p(a_j) t_k(a_j). \tag{28}$$

In summary, we have obtained:

Proposition 4 (Chebyshev method). *The Chebyshev method at level m computes the polynomial q_m of Eq. (28), which provides the following upper and lower bounds to $d_p(M)$,*

$$\begin{aligned} d_p(M)^p &\geq \mathbf{tr}(q_m(M)) - \|q_m - f_p\|_\infty \\ d_p(M)^p &\leq \mathbf{tr}(q_m(M)) + \|q_m - f_p\|_\infty. \end{aligned} \tag{29}$$

Note that by minimizing $\|q_m - f_p\|_\infty$ one minimises $\|(q_m - f_p)_+\|_\infty$ and $\|(q_m - f_p)_-\|_\infty$ simultaneously. As an example, Fig. 5 shows two Chebyshev approximations of f_2 .




Fig. 5. The Chebyshev approximation of degree 2 and 3 of f_2

Table 1. Average running time (in seconds) to compute the first m moments for n tensor factors for a matrix of the form of (1)

$n \backslash m$	8	12	16
8	0.02	0.60	13.82
16	0.09	1.44	28.01
32	0.18	2.60	55.52
64	0.33	4.88	111.44

5. Numerical Results and Examples

We now discuss the numerical performance of the three methods presented in Sect. 4. All computations were done with Mathematica on a standard computer with a 2.9 GHz processor and 16 GB RAM. Our Mathematica code is available with this paper.

Since all three methods take the vector of the first m normalized moments of a matrix M as input, we first study how long it takes to compute these moments for a matrix of the form (1). Specifically, we consider random local matrices $A_j^{[i]}$ of size 2, a sum length of $d = 2$ and several values for the number n of tensor factors. Note that this is a natural scenario in the tensor network paradigm, where one could imagine having a numerical tensor of the form (1) and having to decide, or give necessary or sufficient conditions, on its positive semidefiniteness. Table 1 shows the average running time to compute the first m moments. Note that already for $n = 16$ we were not able to explicitly compute and store the matrix M , let alone decide its semidefiniteness on our computer.

Now we study the running time of the three methods (with $p = 2$) as a function of the number m of moments used (Table 2). Note that this does not include the computation of the moments, since they are part of the input. The moments have been produced from random instances of matrices of the form (1) with local matrices $A_j^{[i]}$ of size 2, $d = 2$ and $n = 8$. Note that the running time includes the full setup of all problems, as provided in the attached Mathematica file. In particular, for the Chebyshev method this includes the computation of the Chebyshev polynomial (28) of degrees $k = 1, \dots, m$ which interpolates f_p . These, however, do not depend on the input, and thus could be computed in advance. The algorithm itself just has to compute the inner product of the coefficient vector and the moment vector, which can be done in almost no time.

To examine the qualitative performance of the three methods, we generated 10,000 random numbers uniformly in $[-\varepsilon, 1]$, for several small values of $\varepsilon > 0$, and took

Table 2. Average running time (in seconds) of the three methods, using the first m moments

m \ method	sos polynomial	Handelman	Chebyshev
8	10.20	0.16	0.37
12	18.39	0.40	0.76
16	49.71	2.22	1.17
32	717.60	23.85	3.84

Table 3. Average number of moments needed to detect non-positivity of a matrix with 10,000 random eigenvalues in the interval $[-\varepsilon, 1]$

ε \ method	sos polynomial	Handelman	Chebyshev
1/2	3	5	4
1/4	4	7	8
1/8	5	8	20
1/16	6	17	?

them as eigenvalues of our matrix M . Note that the corresponding spectral measure is a close approximation of the uniform Lebesgue measure on this interval. We then computed the corresponding normalized moments, and checked how many moments each method needed to produce a positive lower bound on the distance to the psd cone, i.e. to detect non-positivity. Note that the smaller ε , the harder the task. Table 3 shows the average number of moments needed for each method, depending on ε . For $\varepsilon = 1/16$, the Chebyshev algorithm did never provide positive bounds when using a number of moments that we could compute without running into numerical problems (for example when computing $\|f_p - q_m\|_\infty$, which is needed in the algorithm).

Let us summarize and compare the three methods. Concerning the running time, the Chebyshev method is clearly the best. In addition, as mentioned above, its running time can in practice be reduced to almost zero by computing the approximating polynomials beforehand. The Handelman method is also quite fast, in particular when compared to the sos polynomial method. On the other hand, the sos polynomial method needs significantly fewer moments than the other methods in order to produce the same qualitative results. Computing many moments can also require a lot of time, depending on the format of the matrix.

In order to compare both effects (running time versus number of moments needed), we conducted a final experiment. We produced random matrices again of the form (1) with local matrices of size 2, $d = 2$ and different values for the number n of tensor factors. For each method we first checked how many moments were needed to detect non-positivity, and then added the time to compute these moments to the actual running time of the method. Note that in practice one does not know in advance how many moments are needed to detect non-positivity. Interestingly, the Chebyshev method falls behind the other two by far, due to the large number of moments it needs. The comparison of the Handelman and the sos polynomial method is summarized in Table 4. Note that also the Handelman method did not produce any meaningful result in the case of a very large matrix.

In summary, a general and clear decision between the three methods based on their performance cannot be made. They differ greatly in terms of running time and in terms of the moments needed to obtain meaningful results. However, the Chebyshev method seems to fall behind the other two in many relevant instances, at least for single matrices.

Table 4. Total running time (in seconds) to detect non-positivity, involving the time to compute the moments and the running time of the algorithm, for n tensor factors

n \ method	sos polynomial	Handelman
16	0.63	0.01
24	1.33	0.01
32	7.34	1.46
40	??	?

If the matrices are not extremely large and the eigenvalues do not show an extreme behavior, the Handelman methods seems to perform best. The sos polynomial method can however solve some of these extreme cases in which the other two methods fail.

Files attached to this manuscript:

- `PSDBounds.nb`: This Mathematica package contains the three functions `sosApp`, `handelApp` and `chebyApp` that provide the upper and lower bounds, together with the approximating polynomials corresponding to the sos polynomial method, Handelman method, and Chebyshev method, respectively.
- `TensorExample.nb`: This Mathematica notebook provides some examples to illustrate the use of the above package, in particular for matrices of the form (1).

Acknowledgements. GDLC acknowledges funding of the Elise Richter Program of the FWF. TN acknowledges funding through the FWF Project P 29496-N35 (free semialgebraic geometry and convexity). Most of this work was conducted while TF was at the Max Planck Institute for Mathematics in the Sciences. We thank Hilary Carteret and Andreas Thom for helpful comments on the topic.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

A. Extension to Von Neumann Algebras

Let \mathcal{N} be a von Neumann algebra equipped with a faithful normal trace \mathbf{tr} satisfying $\mathbf{tr}(1) = 1$. For example, \mathcal{N} may be the group von Neumann algebra of a discrete group Γ , defined as the weak operator closure of the group algebra $\mathbb{C}[\Gamma]$ as acting on $\ell^2(\Gamma)$; the trace is given by $\mathbf{tr}(x) := \langle e, xe \rangle$ with $e \in \Gamma$ being the unit. In the case where $\mathcal{N} = M_s(\mathbb{C})$ is just the matrix algebra, the discussion presented here specializes to that of the main text.

For a Hermitian element $M \in \mathcal{N}$, the Schatten p -norm is again defined as

$$\|M\|_p := (\mathbf{tr}(|M|^p))^{1/p},$$

where the absolute value and power functions are defined in terms of functional calculus. We can now ask the same question as in the main text: suppose we have M of which we only know the values of the first m moments

$$\mathbf{tr}(M^k) \text{ for } k = 0, \dots, m.$$

Then what can we say about the p -distance from M to the cone of positive elements?

It is straightforward to see that essentially all of the methods of the main text still apply without any change. The only differences are the following:

- Remark 1 no longer applies, since the estimate that we used there involves the matrix size explicitly. Hence there is no bound on $\|M\|_\infty$ that could be computed from the moments. So in order to scale M such that $\|M\|_\infty \leq 1$ is guaranteed, an a priori bound on $\|M\|_\infty$ needs to be known in addition to the moments. Otherwise our methods will not apply in their current form.
- The proof of Proposition 1 is still essentially the same, but σ needs to be generalized to the *spectral scale*, and the corresponding inequality is [12, Corollary 3.3(1)].

Everything else is completely unchanged, including the fact that we are secretly addressing a version of the Hausdorff moment problem.

References

1. Ben-Tal, A., Nemirovski, A.: On tractable approximations of uncertain linear matrix inequalities affected by interval uncertainty. *SIAM J. Optim.* **12**(3), 811–833 (2002)
2. Bhatia, R.: *Matrix Analysis*. Springer, Berlin (1991)
3. Blekherman, G., Parrilo, P.A., Thomas, R.R. (eds.): *Semidefinite optimization and convex algebraic geometry*. MOS-SIAM Series on Optimization (2013)
4. De las Cuevas, G., Schuch, N., Perez-Garcia, D., Cirac, J.I.: Purifications of multipartite states: limitations and constructive methods. *New J. Phys.* **15**, 123021 (2013)
5. De las Cuevas, G., Cubitt, T.S., Cirac, J.I., Wolf, M.M., Perez-Garcia, D.: Fundamental limitations in the purifications of tensor networks. *J. Math. Phys.* **57**, 071902 (2016)
6. Fannes, M., Nachtergaele, B., Werner, R.F.: Finitely correlated states on quantum spin chains. *Commun. Math. Phys.* **144**, 443 (1992)
7. Fritz, T., Netzer, T., Thom, A.: Spectrahedral containment and operator systems with finite-dimensional realization. *SIAM J. Appl. Algebra Geom.* **1**, 556 (2016)
8. Gil, A., Segura, J., Temme, N.M.: *Numerical Methods for Special Functions*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia (2007)
9. Handelman, D.: Representing polynomials by positive linear functions on compact convex polyhedra. *Pac. J. Math.* **132**, 35 (1988)
10. Helton, J.W., Klep, I., McCullough, S.: Free convex algebraic geometry. In: *Semidefinite Optimization and Convex Algebraic Geometry*, volume 13 of MOS-SIAM Ser. Optim., pp. 341–405. SIAM, Philadelphia, PA (2013)
11. Helton, J.W., Klep, I., McCullough, S.: The matricial relaxation of a linear matrix inequality. *Math. Program.* **138**(1–2, Ser. A), 401–445 (2013)
12. Hiai, F., Nakamura, Y.: Majorizations for generalized S-numbers in semifinite von Neumann algebras. *Math. Z.* **195**(1), 17–27 (1987)
13. Higham, N.J.: *Functions of Matrices*. SIAM, Philadelphia (2008)
14. Kliesch, M., Gross, D., Eisert, J.: Matrix product operators and states: NP-hardness and undecidability. *Phys. Rev. Lett.* **113**, 160503 (2014)
15. Kreuin, M.G., Nudelman, A.A.: *The Markov Moment Problem and Extremal Problems*. American Mathematical Society, Providence (1977)
16. Kuhlmann, S., Marshall, M., Schwartz, N.: Positivity, sums of squares and the multi-dimensional moment problem. *Adv. Geom.* **5**(4), 583–606 (2005)
17. Lasserre, J.B.: Bounding the support of a measure from its marginal moments. *Proc. Am. Math. Soc.* **139**(9), 3375–3382 (2011)
18. Marshall, M., Kuhlmann, S.: Positivity, sums of squares and the multi-dimensional moment problem. *Trans. Am. Math. Soc.* **354**(11), 4285–4301 (2002)
19. Orús, R.: A practical introduction to tensor networks: matrix product states and projected entangled pair states. *Ann. Phys.* **349**, 117 (2014)
20. Perez-Garcia, D., Verstraete, F., Wolf, M.M., Cirac, J.I.: Matrix product state representations. *Quantum Inf. Comput.* **7**, 401 (2007)
21. Riesz, M.: Sur le problème des moments. Troisième note, *Arkiv för Matematik, Astronomi och Fysik* **17**(16) (1923)
22. Rivlin, T.J.: *An Introduction to the Approximation of Functions*. Dover, New York (1981)
23. Rudin, W.: *Functional Analysis*. International Series in Pure and Applied Mathematics, 2nd edn. McGraw-Hill Inc, New York (1991)
24. Shohat, J.A., Tamarkin, J.D.: *The Problem of Moments*, 4th edn. American Mathematical Society, New York (1970)

25. Verstraete, F., Garcia-Ripoll, J.J., Cirac, J.I.: Matrix product density operators: simulation of finite-temperature and dissipative systems. *Phys. Rev. Lett.* **93**, 207204 (2004)
26. Vidal, G.: Efficient classical simulation of slightly entangled quantum computation. *Phys. Rev. Lett.* **91**, 147902 (2003)
27. Villani, C.: Optimal transport: old and new. In: *Grundlehren der mathematischen Wissenschaften*, chapter 338. Springer (2009)
28. Werner, A.H., Jaschke, D., Silvi, P., Calarco, T., Eisert, J., Montangero, S.: A positive tensor network approach for simulating open quantum many-body systems. *Phys. Rev. Lett.* **116**, 23720 (2016)

Communicated by M. M. Wolf