

# A Framework for Quadratic Form Maximization over Convex Sets through Nonconvex Relaxations

Vijay Bhattiprolu\*<sup>†</sup>  
 vspvijay@math.ias.edu  
 Institute for Advanced Study  
 Princeton University  
 USA

Euiwoong Lee\*  
 euiwoong@umich.edu  
 University of Michigan  
 USA

Assaf Naor\*<sup>‡</sup>  
 naor@math.princeton.edu  
 Princeton University  
 USA

## ABSTRACT

We investigate the approximability of the following optimization problem. The input is an  $n \times n$  matrix  $A = (A_{ij})$  with real entries and an origin-symmetric convex body  $K \subseteq \mathbb{R}^n$  that is given by a membership oracle. The task is to compute (or approximate) the maximum of the quadratic form  $\sum_{i=1}^n \sum_{j=1}^n A_{ij} x_i x_j = \langle x, Ax \rangle$  as  $x$  ranges over  $K$ . This is a rich and expressive family of optimization problems; for different choices of matrices  $A$  and convex bodies  $K$  it includes a diverse range of optimization problems like max-cut, Grothendieck/non-commutative Grothendieck inequalities, small set expansion and more. While the literature studied these special cases using case-specific reasoning, here we develop a general methodology for treatment of the approximability and inapproximability aspects of these questions.

The underlying geometry of  $K$  plays a critical role; we show under commonly used complexity assumptions that polytime constant-approximability necessitates that  $K$  has type-2 constant that grows slowly with  $n$ . However, we show that even when the type-2 constant is bounded, this problem sometimes exhibits strong hardness of approximation. Thus, even within the realm of type-2 bodies, the approximability landscape is nuanced and subtle.

However, the link that we establish between optimization and geometry of Banach spaces allows us to devise a generic algorithmic approach to the above problem. We associate to each convex body a new (higher dimensional) auxiliary set that is not convex, but is approximately convex when  $K$  has a bounded type-2 constant. If our auxiliary set has an approximate separation oracle, then we design an approximation algorithm for the original quadratic optimization problem, using an approximate version of the ellipsoid method. Even though our hardness result implies that such an oracle does not exist in general, this new question can be solved in specific cases of interest by implementing a range of classical tools from functional analysis, most notably the deep factorization theory of linear operators.

Beyond encompassing the scenarios in the literature for which constant-factor approximation algorithms were found, our generic framework implies that for convex sets with bounded type-2 constant, constant factor approximability is preserved under the following basic operations: (a) Subspaces, (b) Quotients, (c) Minkowski Sums, (d) Complex Interpolation. This yields a rich family of new examples where constant factor approximations are possible, which

were beyond the reach of previous methods. We also show (under commonly used complexity assumptions) that for symmetric norms and unitarily invariant matrix norms the type-2 constant nearly characterizes the approximability of quadratic maximization.

## 1 INTRODUCTION

Suppose that  $n \in \mathbb{N}$  and that  $K \subseteq \mathbb{R}^n$  is a convex body (i.e.,  $K$  is convex, closed, bounded and has nonempty interior) that is origin-symmetric (i.e.,  $x \in K$  if and only if  $-x \in K$ ). We will assume throughout that  $K$  is given by a membership oracle, so that the efficiency of the ensuing algorithms is measured in terms of the dependence on  $n$  and the number of oracle calls.

In this article, we will investigate the approximability of the following optimization problem, special cases of which have been extensively studied in the literature (we will discuss that background after first presenting the problem and our main algorithm). The input is an  $n \times n$  matrix with real entries  $A = (A_{ij}) \in M_n(\mathbb{R})$ , and the task is to evaluate the quantity

$$Q_K^{\max}(A) \stackrel{\text{def}}{=} \max_{x \in K} \sum_{i=1}^n \sum_{j=1}^n A_{ij} x_i x_j = \max_{x \in K} \langle x, Ax \rangle. \quad (1)$$

In (1) and throughout this text,  $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$  is the standard scalar product on  $\mathbb{R}^n$ , namely  $\langle x, y \rangle = \sum_{i=1}^n x_i y_i$  for every two vectors  $x = (x_1, \dots, x_n)$ ,  $y = (y_1, \dots, y_n) \in \mathbb{R}^n$ . Also, we will adhere throughout to the common convention that even though within any in-line discussion the elements of  $\mathbb{R}^n$  are written as row vectors, for the purpose of any linear-algebraic consideration we consider them as column vectors, i.e., members of the  $n \times 1$  matrix space  $M_{n \times 1}(\mathbb{R})$ .

The literature also considers a bilinear variant of (1) in which one is given  $m, n \in \mathbb{N}$ , two convex origin-symmetric bodies  $K \subseteq \mathbb{R}^n$  and  $L \subseteq \mathbb{R}^m$ , and an  $n \times m$  matrix  $B = (B_{ij}) \in M_{n \times m}(\mathbb{R})$ , and the task is to evaluate (or estimate) the quantity

$$\begin{aligned} \text{Op}_{K,L}^{\max}(B) &\stackrel{\text{def}}{=} \max_{\substack{x \in K \\ y \in L}} \sum_{i=1}^n \sum_{j=1}^m B_{ij} x_i y_j = \max_{\substack{x \in K \\ y \in L}} \langle x, By \rangle \\ &= \frac{1}{2} \max_{z \in K \times L} \left\langle z, \begin{pmatrix} 0 & B \\ B^* & 0 \end{pmatrix} z \right\rangle, \end{aligned} \quad (2)$$

where  $B^* = (B_{ji}) \in M_{m \times n}(\mathbb{R})$  is the transpose of  $B$ . The final equality in (2) shows that (2) is a special case of (1), which is why we will mostly focus on (1). But, it is beneficial to consider the bilinear variant separately because sometimes it exhibits better approximation properties than what is possible in the quadratic setting (a notable example is Grothendieck's inequality; see below).

\*Supported by the Simons Collaboration on Algorithms and Geometry.

<sup>†</sup>Supported indirectly by the National Science Foundation Grant No. CCF-1900460.

<sup>‡</sup>Supported in part by the BSF, the Packard Foundation and the Simons Foundation.

Another important special case of (1) which the literature sometimes treats separately is when the input matrix  $A$  is symmetric and positive semidefinite (PSD). In that case

$$\begin{aligned} Q_K^{\max}(A) &= \max_{x \in K} \|A^{\frac{1}{2}}x\|_{\ell_2^n}^2 = \max_{\substack{x \in K \\ y \in \text{Ball}(\ell_2^n)}} \langle A^{\frac{1}{2}}x, y \rangle^2 \\ &= \left( \text{Op}_{K, \text{Ball}(\ell_2^n)}^{\max}(A^{\frac{1}{2}}) \right)^2, \end{aligned}$$

where  $\|\cdot\|_{\ell_2^n}$  is the standard Euclidean norm on  $\mathbb{R}^n$  and  $\text{Ball}(\ell_2^n) = \{x \in \mathbb{R}^n : x_1^2 + \dots + x_n^2 \leq 1\}$  is the corresponding Euclidean ball of radius 1. Thus, the PSD case of (1) is a special case of the aforementioned bilinear variant of (1), which explains why it has better properties (another reason is that in this case  $L$  is a Euclidean ball rather than a more general convex body).

The above framework is a rich and expressive family of optimization problems which contains many discrete and continuous optimization problems as special cases (corresponding to choices of matrices and convex bodies) that occur in several areas, including combinatorial optimization, computational complexity, graph theory, quantum information theory, statistical physics, machine learning, game theory and functional analysis. In fact, we suspect that many readers have already spotted familiar questions as such special cases, but in order to first discuss the contribution of the present work, we will defer specifying a variety of such examples to Section 1.4.

While the literature contains investigations of such special cases using case-specific reasoning, here we develop a general methodology for treatment of the approximability and inapproximability aspects of these questions. We devise an overarching method for obtaining constant factor approximation algorithms that includes the prior cases in the literature for which this was achieved, as well as many more new cases.

The precursor (and inspiration) of the present article is the manuscript [NS09] that has not yet been published but was circulated widely over the years and will be published soon (it is available on request). The goal of [NS09] was to broach the same issue of finding an algorithmic approach to the optimization problem (1) which treats a class of convex bodies  $K$  that is more general than the special cases that have been previously studied, as an extension of the study of the ball of  $\ell_p^n$  that was conducted in [KNS08] (see [GRSW16] for the corresponding hardness result under a weaker hypothesis than that of [KNS08]). The success of [NS09] was partial, as it pertains only to a certain subclass of convex bodies  $K$  that satisfies the following symmetry condition.

$$\forall (x_1, \dots, x_n) \in \mathbb{R}^n, (x_1, \dots, x_n) \in K \iff (|x_1|, \dots, |x_n|) \in K. \quad (3)$$

When (3) holds, there is an obvious vector relaxation of (1) that is given by the maximization

$$\max_{\substack{x_1, \dots, x_n \in \mathbb{R}^n \\ (\|x_1\|_{\ell_2^n}, \dots, \|x_n\|_{\ell_2^n}) \in K}} \sum_{i=1}^n \sum_{j=1}^n A_{ij} \langle x_i, x_j \rangle. \quad (4)$$

The utility of such a relaxation was investigated in [Nes98, KNS08, NS09], where further geometric assumptions on  $K$  were isolated that guarantee that (4) is a convex program that has bounded integrality gap (see below). Note that (3) probes only the intersection

$K \cap [0, \infty)^n$  of  $K$  with the positive orthant, which is why it is natural to study it only when (4) holds; otherwise  $K$  need not be determined by the region of space to which the relaxation (4) is sensitive.

This was the starting point of our work. Namely, for convex bodies that do not satisfy the symmetry assumption (3), there is no longer an obvious vector relaxation. Note that (3) is a stringent assumption that fails for many norms of interest; e.g. for unit balls of matrix norms such as the Schatten–von Neumann trace classes (see below) where the norm of the entry-wise absolute value ( $|A_{ij}|$ ) of a given matrix  $A = (A_{ij}) \in M_n(\mathbb{R})$  can be drastically different from the norm of  $A$ . To overcome this conceptual obstacle, we devise an entirely different algorithmic methodology. Before proceeding, we set some notation and record some basic definitions.

## 1.1 Notation and Preliminaries

It is most natural to present our approach in the (equivalent) setting of normed spaces rather than origin-symmetric convex bodies. Specifically, let  $\|\cdot\|_X : \mathbb{R}^n \rightarrow [0, \infty)$  be a norm on  $\mathbb{R}^n$  and denote the corresponding normed space  $(\mathbb{R}^n, \|\cdot\|_X)$  by  $X$ . The (closed) unit ball of  $X$  will be denoted throughout what follows by

$$\text{Ball}(X) \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n : \|x\|_X \leq 1\}.$$

The standard correspondence is that  $\text{Ball}(X)$  is an origin-symmetric convex body, and conversely any  $K \subseteq \mathbb{R}^n$  as above is equal to  $\text{Ball}(X)$  for some  $X = (\mathbb{R}^n, \|\cdot\|_X)$ , where the norm  $\|x\|_X$  of each  $x \in \mathbb{R}^n \setminus \{0\}$  is the unique scaling factor  $s > 0$  for which  $\frac{1}{s}x$  belongs to the boundary of  $K$ .

In accordance with the above convention for convex bodies, we will tacitly assume throughout the ensuing discussion that all normed spaces  $X = (\mathbb{R}^n, \|\cdot\|_X)$  are given by a membership oracle for  $\text{Ball}(X)$ . By binary search for the smallest  $r \geq 0$  such that  $x \in \mathbb{R}^n$  belongs to  $r \cdot \text{Ball}(X)$ , such an oracle directly yields also a norm-evaluation oracle.

So, given a normed space  $X = (\mathbb{R}^n, \|\cdot\|_X)$  and a matrix  $A \in M_n(\mathbb{R})$ , denote

$$Q_X^{\max}(A) \stackrel{\text{def}}{=} Q_{\text{Ball}(X)}^{\max}(A).$$

Observe in passing that the bilinear variant (2) when  $K = \text{Ball}(X)$  and  $L = \text{Ball}(Y)$  for normed spaces  $X = (\mathbb{R}^n, \|\cdot\|_X)$  and  $Y = (\mathbb{R}^m, \|\cdot\|_Y)$ , respectively, is nothing more than the operator norm of the matrix  $B \in M_{n \times m}(\mathbb{R})$  when it is viewed as an operator from  $Y$  to the dual  $X^*$  of  $X$ . Namely,

$$\text{Op}_{K, L}^{\max}(B) = \|B\|_{Y \rightarrow X^*} = \|B^*\|_{X \rightarrow Y^*}, \quad (5)$$

where the first equality in (5) can be taken to be the definition of the corresponding operator norm and it is equal to the more common definition  $\|B\|_{Y \rightarrow X^*} = \max_{y \in \text{Ball}(Y)} \|By\|_{X^*}$  by duality (Hahn–Banach). The second equality in (5) is the fact that the norm of an operator between Banach spaces is equal to the norm of its adjoint. See e.g. the textbook [Rud73] for this standard material.

## Type and Cotype

It is beneficial to introduce the following convention regarding random variables that will be used extensively in what follows. We will work throughout with the families of random variables  $\{\varepsilon_i : i \in \mathbb{N}\}$ ,  $\{g_i : i \in \mathbb{N}\}$  and  $\{g_{ij} : i, j \in \mathbb{N}\}$ , where it will always be tacitly understood that they are independent,  $\{\varepsilon_i : i \in \mathbb{N}\}$  are  $\pm 1$

Bernoulli random variables, i.e., distributed uniformly over  $\{-1, 1\}$ , and  $\{g_i : i \in \mathbb{N}\}$  and  $\{g_{ij} : i, j \in \mathbb{N}\}$  are standard Gaussian random variables. All the expectations that appear below are with respect to the joint distribution of these random variables. We will always denote the standard Gaussian random vector in  $\mathbb{R}^n$  by  $g = (g_1, \dots, g_n)$ .

The (Rademacher) type 2 constant [DPR72] of a normed space  $X = (\mathbb{R}^n, \|\cdot\|_X)$ , denoted  $T_2(X)$ , is the smallest  $T > 0$  such that for every  $m \in \mathbb{N}$ , every  $x_1, \dots, x_m \in \mathbb{R}^n$  satisfy

$$\mathbb{E} \left[ \left\| \sum_{i=1}^m \varepsilon_i x_i \right\|_X^2 \right] \leq T^2 \cdot \sum_{i=1}^m \|x_i\|_X^2, \quad (6)$$

Correspondingly, the (Rademacher) cotype 2 constant of  $X$ , denoted  $C_2(X)$ , is the smallest  $C > 0$  such that for every  $m \in \mathbb{N}$ , every choice of vectors  $x_1, \dots, x_m \in \mathbb{R}^n$  satisfies

$$\sum_{i=1}^m \|x_i\|_X^2 \leq C^2 \cdot \mathbb{E} \left[ \left\| \sum_{i=1}^m \varepsilon_i x_i \right\|_X^2 \right]. \quad (7)$$

These invariants of normed spaces are of immense importance to various areas; see the survey [Mau03] for an indication of (part of) this body of work, as well as its history. Here we show that they are closely related to the computational complexity of the quadratic optimization problem (1), and, in fact, under common complexity assumptions, they govern it in a sense that will be made precise later. For concreteness, we record the following asymptotic evaluations<sup>1</sup> of these constants when  $X = \ell_p^n$  for some integer  $n \geq 2$  and  $p \in [1, \infty]$ , all of which can be found in [MS86].

$$\begin{aligned} T_2(\ell_p^n) &\asymp \begin{cases} n^{\frac{1}{p}-\frac{1}{2}} & \text{if } 1 \leq p \leq 2, \\ \sqrt{\min\{p, \log n\}} & \text{if } 2 \leq p \leq \infty, \end{cases} \quad \text{and} \\ C_2(\ell_p^n) &\asymp \begin{cases} 1 & \text{if } 1 \leq p \leq 2, \\ n^{\frac{1}{2}-\frac{1}{p}} & \text{if } 2 \leq p \leq \infty. \end{cases} \end{aligned} \quad (8)$$

We also record the following duality relations that hold for any normed space  $X$ .

$$C_2(X^*) \leq T_2(X) \leq C_2(X^*) \log(\dim(X) + 1). \quad (9)$$

The first inequality in (9) is straightforward [MP76] and the second inequality in (9) is from [Pis80].

## 1.2 A Generic Framework

We are now ready to describe our algorithmic approach, starting with a simpler “warm-up” algorithm which covers many new instances of (1). Fix an integer  $n \geq 2$ . The set of symmetric positive definite matrices with real entries will be denoted by  $\text{PSD}^n \subseteq M_n(\mathbb{R})$ . For  $\mathbb{W} \in M_n(\mathbb{R})$  we will use the notation  $\mathbb{W} \geq 0$  to indicate that  $\mathbb{W} \in \text{PSD}^n$ . We associate to a normed space  $X = (\mathbb{R}^n, \|\cdot\|_X)$  the following subset  $\mathcal{U}(X)$  of  $\text{PSD}^n$  that we call the *upper covariance*

<sup>1</sup>In addition to the usual  $o(\cdot), O(\cdot), \Omega(\cdot), \Theta(\cdot)$  notation for asymptotic relations, we will also use throughout the following (standard) asymptotic notation. For  $P, Q > 0$ , the notations  $P \leq Q$  and  $Q \geq P$  mean that  $P \leq KQ$  for a universal constant  $K > 0$ . The notation  $P \asymp Q$  stands for  $(P \leq Q) \wedge (Q \leq P)$ .

body of  $X$ .

$$\begin{aligned} \mathcal{U}(X) &\stackrel{\text{def}}{=} \\ &\bigcup_{m=1}^{\infty} \left\{ \sum_{i=1}^m w_i w_i^* : w_1, \dots, w_m \in \mathbb{R}^n \text{ and } \mathbb{E} \left[ \left\| \sum_{i=1}^m g_i w_i \right\|_X^2 \right] \leq 1 \right\}. \end{aligned} \quad (10)$$

For every  $w_1, \dots, w_m \in \mathbb{R}^n$ , the random vectors  $\sum_{i=1}^m g_i w_i$  and  $\mathbb{W}^{\frac{1}{2}} g$ , where  $\mathbb{W} = \sum_{i=1}^m w_i w_i^* \geq 0$ , are equi-distributed, since they are both Gaussian vectors whose covariance matrix is  $\mathbb{W}$ . Thus,

$$\mathcal{U}(X) = \left\{ \mathbb{W} \in \text{PSD}^n : \mathbb{E} \left[ \left\| \mathbb{W}^{\frac{1}{2}} g \right\|_X^2 \right] \leq 1 \right\}. \quad (11)$$

This observation explains our choice of nomenclature, namely  $\mathcal{U}(X)$  consists of those covariance matrices of Gaussian vectors in  $\mathbb{R}^n$  whose expected squared  $X$ -norm is bounded from above by 1. An important property of  $\mathcal{U}(X)$  is that one can relate quadratic optimization over  $\text{Ball}(X)$  to linear optimization over  $\mathcal{U}(X)$ :

## From Quadratic Optimization to Linear Optimization

Observe that for any  $A = (A_{ij}) \in M_n(\mathbb{R})$ , any normed space  $X = (\mathbb{R}^n, \|\cdot\|_X)$  satisfies

$$Q_X^{\max}(A) = \max_{\mathbb{W}=(w_{ij}) \in \mathcal{U}(X)} \sum_{i=1}^n \sum_{j=1}^n A_{ij} w_{ij} = \max_{\mathbb{W} \in \mathcal{U}(X)} \langle A, \mathbb{W} \rangle. \quad (12)$$

Indeed, if  $\mathbb{W} \in \text{PSD}^n$  satisfies  $\mathbb{E} \left[ \left\| \mathbb{W}^{\frac{1}{2}} g \right\|_X^2 \right] \leq 1$ , then

$$\begin{aligned} \langle A, \mathbb{W} \rangle &= \text{Tr}(A\mathbb{W}) = \text{Tr}(\mathbb{W}^{\frac{1}{2}} A \mathbb{W}^{\frac{1}{2}}) = \mathbb{E} \left[ \langle \mathbb{W}^{\frac{1}{2}} g, A \mathbb{W}^{\frac{1}{2}} g \rangle \right] \\ &\leq \mathbb{E} \left[ Q_X^{\max}(A) \cdot \left\| \mathbb{W}^{\frac{1}{2}} g \right\|_X^2 \right] \leq Q_X^{\max}(A). \end{aligned}$$

This shows that right hand side of (12) is at most the left hand side of (12). The reverse inequality follows by noting that if  $w \in \text{Ball}(X)$ , then  $ww^* \in \mathcal{U}(X)$  and  $\langle A, ww^* \rangle = \langle w, Aw \rangle$ .

## Approximate Convexity of $\mathcal{U}(X)$

The body  $\mathcal{U}(X)$  need not be convex, but it is  $T_2(X)^2$ -*approximately convex* in the sense that

$$\mathcal{U}(X) \subseteq \text{conv}(\mathcal{U}(X)) \subseteq T_2(X)^2 \cdot \mathcal{U}(X), \quad (13)$$

where, given a subset  $S$  of some  $\mathbb{R}^d$ , we denote the convex hull of  $S$  by  $\text{conv}(S)$ . To justify (13), fix  $k \in \mathbb{N}$  and suppose that  $\mathbb{W}_1, \dots, \mathbb{W}_k \in \mathcal{U}(X)$  and  $s_1, \dots, s_k \in [0, 1]$  satisfy  $\sum_{j=1}^k s_j = 1$ . The goal is to demonstrate that  $T_2(X)^{-2} \sum_{j=1}^k s_j \mathbb{W}_j \in \mathcal{U}(X)$ . For each  $j \in \{1, \dots, k\}$ , the assumption  $\mathbb{W}_j \in \mathcal{U}(X)$  means that for some  $m(j) \in \mathbb{N}$  there are vectors  $w_{1,j}, \dots, w_{m(j),j} \in \mathbb{R}^n$  such that

$$\mathbb{W}_j = \sum_{i=1}^{m(j)} w_{ij} w_{ij}^* \quad \text{and} \quad \mathbb{E} \left[ \left\| \sum_{i=1}^{m(j)} g_{ij} w_{ij} \right\|_X^2 \right] \leq 1.$$

Hence,

$$\begin{aligned} \mathbb{E} \left[ \left\| \sum_{j=1}^k \sum_{i=1}^{m(j)} g_{ij} \sqrt{s_j} w_{ij} \right\|_X^2 \right] &= \mathbb{E} \left[ \left\| \sum_{j=1}^k \varepsilon_j \sqrt{s_j} \sum_{i=1}^{m(j)} g_{ij} w_{ij} \right\|_X^2 \right] \\ &\leq T_2(X)^2 \cdot \sum_{j=1}^k s_j \mathbb{E} \left[ \left\| \sum_{i=1}^{m(j)} g_{ij} w_{ij} \right\|_X^2 \right] \leq T_2(X)^2. \end{aligned}$$

Therefore

$$T_2(X)^{-2} \cdot \sum_{j=1}^k \sum_{i=1}^{m(j)} (\sqrt{s_j} w_{ij}) (\sqrt{s_j} w_{ij})^* = T_2(X)^{-2} \cdot \sum_{j=1}^k s_j \mathbb{W}_j$$

indeed belongs to  $\mathcal{U}(X)$ .

Motivated by (13), we set the following terminology.

*Definition 1.1.* Suppose that  $S \subseteq \mathbb{R}^n$  is star-shaped with respect to the origin, i.e.,  $tx \in S$  for every  $x \in S$  and  $t \in [0, 1]$ . Given  $\alpha \in [1, \infty)$ , we say that  $S$  is  $\alpha$ -approximately convex if  $\text{conv}(S) \subseteq \alpha S$ .

The two observations (12) and (13) highlight the following important facts. Firstly, the relaxation of  $\text{Ball}(X) \subseteq \mathbb{R}^n$  to the upper covariance body  $\mathcal{U}(X) \subseteq M_n(\mathbb{R})$  is lossless, i.e., it reduces the maximization over  $\text{Ball}(X)$  of a quadratic form to a maximization over  $\mathcal{U}(X)$  of a linear function. Secondly, the geometry of  $X$ , through the extent to which it has type 2, plays a role by ensuring that the potentially complicated set  $\mathcal{U}(X)$  is at the very least approximately convex. It is thus natural to investigate the efficient optimization of linear functions over approximately convex sets. However, the following theorem (see Section 9 in [BLN21]) shows that this is a subtle matter, because even when the type-2 constant of  $X$  is small, the computational complexity of approximating  $Q_X^{\max}(A)$  could be poor.

**THEOREM 1.2 (IMPOSSIBILITY OF QUADRATIC MAXIMIZATION ASSUMING ONLY BOUNDED TYPE-2).**

For every  $n \in \mathbb{N}$  and  $0 < \varepsilon < 1$  there exists a distribution  $\mathbb{P} = \mathbb{P}_{n,\varepsilon}$  over random normed spaces  $X = (\mathbb{R}^n, \|\cdot\|_X)$  and  $p_n \in (0, 1)$  with  $\lim_{n \rightarrow \infty} p_n = 1$ , such that the following properties hold.

- (1)  $\mathbb{P}_{n,\varepsilon}[T_2(X) \leq 1] = 1$ .
- (2)  $\mathbb{P}_{n,\varepsilon}[|S \cap \text{Ball}(X) = S \cap \text{Ball}(\ell_2^n)| \geq p_n \text{ for every } S \subseteq \mathbb{R}^n \text{ with } |S| \leq \exp(n^\varepsilon)] \geq p_n$ .
- (3)  $\mathbb{P}_{n,\varepsilon}[Q_X^{\max}(I_n) \geq n^{1-\varepsilon}] \geq p_n$ , where  $I_n \in M_n(\mathbb{R})$  is the identity matrix.

Theorem 1.2 demonstrates that if there were an algorithm that takes as input a normed space  $X$  whose type-2 constant is  $O(1)$  and outputs a number that is guaranteed to be within a factor that is  $o(n^{1-\varepsilon})$  of  $Q_X^{\max}(I_n)$ , then that algorithm must necessarily make more than  $\exp(n^\varepsilon)$  membership queries to  $\text{Ball}(X)$ . Indeed,  $Q_X^{\max}(I_n) = 1$  when  $X = \ell_2^n$ , while if  $X$  is the random normed space of Theorem 1.2, then  $T_2(X) \leq 1$  and with high probability  $Q_X^{\max}(I_n) \geq n^{1-\varepsilon}$ . However, if  $S$  is the set of points that the algorithm queried, then with high probability the algorithm did not obtain any information that distinguishes  $X$  from  $\ell_2^n$ .

Thus, even if  $X$  has a small type-2 constant, this does not suffice for the existence of an efficient algorithm for approximating  $Q_X^{\max}(\cdot)$ , but, as we have seen, requiring this property is a good place to start because it ensures that the upper covariance

body is approximately convex. The following theorem establishes a further connection between type 2 and the computational complexity of approximating  $Q_X^{\max}(\cdot)$  by providing evidence (under a commonly used complexity assumption, namely the Small Set Expansion Hypothesis) that if the type 2 constant of  $X$  is very large, then there is no polynomial time algorithm that obtains a  $O(1)$ -approximation to  $Q_X^{\max}(\cdot)$ . Further hardness results (with and without non-uniform complexity assumptions and with weaker assumptions on the growth of the type-2 constant assuming (necessarily) the Exponential Time Hypothesis), are derived in the full version [BLN21].

**THEOREM 1.3 (IMPOSSIBILITY OF QUADRATIC MAXIMIZATION WHENEVER TYPE-2 IS GROWING POLYNOMIALLY).**

Fix a sequence of normed spaces  $\{X^n = (\mathbb{R}^n, \|\cdot\|_{X^n})\}_{n=1}^\infty$  satisfying  $T_2(X^n) = n^{\Omega(1)}$ . We assume that they are given to us algorithmically in the sense that there is a polynomial time algorithm that takes as input  $x \in \mathbb{R}^n$  and determines whether or not  $x \in \text{Ball}(X^n)$ . Then, assuming the Small Set Expansion Hypothesis and that  $\text{NP} \not\subseteq \text{P/poly}$ , there is no polynomial time algorithm that takes as input a matrix  $A \in M_n(\mathbb{R})$  and approximates  $Q_{X^n}^{\max}(A)$  up to a universal constant factor.

**REMARK 1.** The Small Set Expansion Hypothesis (SSEH) is a commonly used hardness assumption that was formulated in [RS10] and is recalled in the full version [BLN21]. Of course, the SSEH is less standard than, say,  $\text{NP} \not\subseteq \text{P/poly}$ , so one should take Theorem 1.3 as evidence that if the underlying norm has large type-2 constant, then it is unlikely that there is an efficient constant-factor algorithm for (1), namely by designing such an algorithm one would refute the SSEH, thus making a major breakthrough in complexity theory.

**REMARK 2.** Recalling (8), Theorem 1.3 applies in particular to  $X^n = \ell_p^n$  when  $1 \leq p < 2$ , thus demonstrating the computational difficulty of the  $\ell_p$  Grothendieck problem, which was left open in [KNS08], where it was shown that this problem does have a  $O_p(1)$  approximation algorithm when  $2 \leq p < \infty$ . In the unpublished manuscript [Alo06] it was proved that a  $O(1)$  approximation algorithm exists when  $p = 1$  provided that all of the diagonal entries of the input matrix  $A$  vanish; see the exposition in [KN12]. In the full version [BLN21] we show that if  $X^n = \ell_p^n$  and  $1 < p < 2$ , then the hardness statement of Theorem 1.3 holds even when the diagonal of  $A$  vanishes, so in this setting we obtain rigorous evidence for an interesting complexity theoretic terrain: The  $\ell_p$  Grothendieck problem is approximable when  $p = 1$  or  $2 \leq p < \infty$ , but likely hard to approximate when  $1 < p < 2$  or  $p = \infty$  (see [ABH<sup>+</sup>05] for hardness when  $p = \infty$ ).

## Approximation Algorithms From Upper Covariance Separation Oracle

Recall that Theorem 1.2 implies that even though the (random) upper covariance body  $\mathcal{U}(X)$  is  $O(1)$ -approximately convex (as  $X$  has bounded type 2 constant), with high probability one cannot optimize linear functionals over  $\mathcal{U}(X)$  efficiently. It turns out that the issue at hand is that even if one permits the algorithm to make oracle norm-evaluation queries for  $X$ , the auxiliary body  $\mathcal{U}(X)$  need not even have an efficient ‘‘approximate separation oracle,’’ which we define as follows.

*Definition 1.4.* Fix  $\alpha \geq 1$ . Let  $S \subseteq \mathbb{R}^n$  be star shaped with respect to the origin and  $\alpha$ -approximately convex. An  $\alpha$ -approximate separation oracle for  $S$  is a function  $\mathcal{O}$  defined on  $\mathbb{R}^n$  that outputs to each input  $x \in \mathbb{R}^n$  either “Inside” or an affine hyperplane of  $\mathbb{R}^n$ . The requirements for  $\mathcal{O}$  are as follows.

- If the output  $\mathcal{O}(x)$  is “Inside,” then necessarily  $x \in \alpha S$ .
- If the output  $\mathcal{O}(x)$  is a hyperplane  $H \subseteq \mathbb{R}^n$ , then  $H$  must separate  $x$  from  $S$ , i.e.,  $x$  and  $S$  are contained in different sides of  $H$ . Note that this implies in particular that  $x \notin \mathbf{conv}(S)$ .

Observe that these requirements are not dichotomic, i.e., they are ambiguous when  $x \in (\alpha S) \setminus \mathbf{conv}(S)$  (recall that  $\mathbf{conv}(S) \subseteq \alpha S$  since  $S$  is  $\alpha$ -approximately convex). Namely, if  $x \in (\alpha S) \setminus \mathbf{conv}(S)$ , then the oracle is allowed to either output a hyperplane or output “Inside.”

Using a natural approximate version of the ellipsoid method, we prove the following theorem (see the full version [BLN21]).

**THEOREM 1.5 (APPROXIMATE ELLIPSOID METHOD).**

Fix  $\alpha \geq 1$  and  $R \geq r > 0$ . Suppose that  $S \subseteq \mathbf{PSD}^n$  is star shaped with respect to the origin,  $\alpha$ -approximately convex, and has an  $\alpha$ -approximate separation oracle. Suppose also that

$$r \cdot \text{Ball}(\ell_2^{n^2}) \subseteq S \subseteq R \cdot \text{Ball}(\ell_2^{n^2}),$$

where we use the natural identification of  $M_n(\mathbb{R})$  with  $\mathbb{R}^{n^2}$ . Then, there exists an algorithm that takes as input a matrix  $A \in M_n(\mathbb{R})$ , makes a number of oracle calls that grows polynomially in  $n$ ,  $\log R$ ,  $\log(1/r)$  and the length of the bit description of  $A$ , and outputs a matrix  $\mathbb{W} \in S$  that satisfies

$$\langle \mathbb{W}, A \rangle \geq \frac{1 - o(1)}{\alpha} \sup_{V \in S} \langle V, A \rangle.$$

For the sake of the discussion within this extended abstract, it will be convenient to always assume tacitly that  $X = (\mathbb{R}^n, \|\cdot\|_X)$  is a normed space whose upper covariance body satisfies

$$e^{-n^{O(1)}} \cdot \text{Ball}(\ell_2^{n^2}) \subseteq \mathcal{U}(X) \subseteq e^{n^{O(1)}} \cdot \text{Ball}(\ell_2^{n^2}). \quad (14)$$

Such a normalization, which is mechanical to verify in all the cases that we examined, removes the need to state running times in terms of  $r, R$  as done in Theorem 1.5. Another simplifying assumption that we will make throughout this extended abstract is that the length of the bit description of all inputs (namely matrices) to algorithms is  $n^{O(1)}$ .

Using Theorem 1.5 and applying it to  $\mathcal{U}(X)$ , we readily deduce the following approximation algorithm for quadratic maximization (see the full version [BLN21])

**PROPOSITION 1.6 (QUADRATIC MAXIMIZATION GIVEN SEPARATION ORACLE FOR UPPER COVARIANCE BODY).**

Given access to an  $\alpha$ -approximate separation oracle for  $\mathcal{U}(X)$ , there is an algorithm that on any input  $A \in M_n(\mathbb{R})$  runs in polynomial time and returns a  $(1 + o(1))\alpha$ -approximation to  $\mathbf{Q}_X^{\max}(A)$ .

The upshot of the above result is that it refocuses our attention to the task of designing an approximate separation oracle for the upper covariance body. Using this approach, we are already able to conclude new results for quadratic maximization by applying tools from classical analysis to design an approximate separation oracle for  $\mathcal{U}(X)$ . In some cases, however, it is quite difficult to design such

an oracle directly for  $\mathcal{U}(X)$ . Inspired by deep tools from functional analysis, specifically the factorization theory of linear operators (see the monograph [Pis86]), we will prove that under the assumption of having a bounded type-2 constant it suffices to design a separation oracle for the *lower covariance region* of  $X$  which we define in (15) below.

To give a couple of examples, it is easy to design a lower covariance separation oracle for the Minkowski sum  $\ell_4^n + \ell_5^n$  (see Section 1.3) or for the quotient  $\ell_4^n / \ell_5^m$ , while on the other hand it is unclear how to directly describe an upper covariance separation oracle in these cases (see the full version [BLN21] for more details). Another advantage which will become apparent soon is that lower covariance separation oracles allow for provably better approximation factors than the upper covariance separation oracles in the special cases of PSD quadratic maximization and bilinear maximization (the difference can be as big as  $\log n$ , as can be seen in the familiar example of  $X = \ell_\infty^n$ ). Below we give a proof sketch for main “framework” theorem, namely an approximation algorithm for quadratic maximization (resp. bilinear maximization) when type-2 (resp. dual cotype-2) is bounded assuming access to only a separation oracle for the lower covariance region.

**REMARK 3.** In the interest of simplicity, the proof sketch below assumes we only desire to approximate the optimal value (and not produce solution vectors). For this simpler goal it suffices to use certain factorization theorems (see the full version [BLN21] for a detailed introduction to factorization and the relevant theorems we use) as a black box. For the full proof in [BLN21], we give rounding algorithms as well. For technical reasons, it was necessary to “open the factorization black box” and make some parts of the argument constructive, in addition to dualizing the entire argument. We therefore caution the reader that the full proof in [BLN21] is syntactically different from the ensuing overview.

## Lower Covariance Region

We define the lower covariance region as follows:

$$\begin{aligned} \mathcal{L}(X) &\stackrel{\text{def}}{=} \\ &\bigcup_{m=1}^{\infty} \left\{ \sum_{i=1}^m w_i w_i^* : w_1, \dots, w_m \in \mathbb{R}^n \text{ and } \mathbb{E} \left[ \left\| \sum_{i=1}^m g_i w_i \right\|_{X^*}^2 \right] \geq 1 \right\} \\ &= \left\{ \mathbb{W} \in \mathbf{PSD}^n : \mathbb{E} \left[ \left\| \mathbb{W}^{\frac{1}{2}} g \right\|_{X^*}^2 \right] \geq 1 \right\}, \end{aligned} \quad (15)$$

where the second inequality in (15) is justified the same way as (11). Note that because  $\mathcal{L}(X)$  is equal to

$$\mathbf{PSD}^n \setminus \left\{ \mathbb{W} \in \mathbf{PSD}^n : \mathbb{E} \left[ \left\| \mathbb{W}^{\frac{1}{2}} g \right\|_{X^*}^2 \right] < 1 \right\},$$

the lower covariance region of  $X$  is the complement in  $\mathbf{PSD}^n$  of the interior of the upper covariance body of  $X^*$ . As such, it is a complement of a set that is star shaped with respect to the origin, and therefore  $s \cdot \mathcal{L}(X) \supseteq \mathcal{L}(X)$  for every  $0 < s \leq 1$ .

## Approximate Convexity of Lower Covariance Region.

By reasoning analogously to the proof of (13), we see that

$$\mathcal{L}(X) \subseteq \mathbf{conv}(\mathcal{L}(X)) \subseteq \frac{1}{C_2(X^*)} \mathcal{L}(X). \quad (16)$$

Thus, the lower covariance region of  $X$  is  $C_2(X^*)^2$ -approximately convex in the following sense, which is the natural adaptation of Definition 1.1 to regions that are complements of star shape sets. Recall that by (9) if  $X$  has bounded type 2 constant, then  $X^*$  has bounded cotype 2 constant.

*Definition 1.7.* Let  $T \subseteq \mathbb{R}^n$  satisfy  $[1, \infty)T \subseteq T$  (equivalently,  $\mathbb{R}^n \setminus T$  is star shaped with respect to the origin). Given  $\alpha \geq 1$ , we say that  $T$  is  $\alpha$ -approximately convex if  $\mathbf{conv}(T) \subseteq \frac{1}{\alpha}T$ .

With this definition at hand, the natural adaptation of Definition 1.4 is as follows.

*Definition 1.8.* Fix  $\alpha \geq 1$ . Suppose that  $T \subseteq \mathbb{R}^n$  satisfies  $[1, \infty)T \subseteq T$  and that  $T$  is  $\alpha$ -approximately convex. An  $\alpha$ -approximate separation oracle for  $T$  is a function  $\mathcal{O}$  defined on  $\mathbb{R}^n$  that outputs to each input  $x \in \mathbb{R}^n$  either “Inside” or an affine hyperplane of  $\mathbb{R}^n$ . The requirements for  $\mathcal{O}$  are as follows.

- If the output  $\mathcal{O}(x)$  is “Inside,” then necessarily  $x \in \frac{1}{\alpha}T$ .
- If the output  $\mathcal{O}(x)$  is a hyperplane  $H \subseteq \mathbb{R}^n$ , then  $H$  must separate  $x$  from  $T$ .

If  $x \in (\frac{1}{\alpha}T) \setminus \mathbf{conv}(T)$ , then  $\mathcal{O}$  is allowed to either output a hyperplane or output “Inside”.

## Approximation Algorithms from Lower Covariance Separation Oracle

With these notions at hand, if the lower covariance region of  $X$  has an  $\alpha$ -approximate separation oracle for some  $\alpha \geq C_2(X^*)$ , then by analysing a natural approximate version of the ellipsoid method we obtain an (oracle-time) efficient algorithm for approximating certain convex programs up to factor  $(1 + o(1))\alpha$ , in the spirit of Theorem 1.5. For the sake of simplicity, rather than explaining this methodology in the introduction in its full generality, we state the following two consequences of it and refer to the full version [BLN21] for a complete treatment.

**THEOREM 1.9 (QUADRATIC/BILINEAR MAXIMIZATION GIVEN SEPARATION ORACLE FOR LOWER COVARIANCE REGION).**

*Suppose that  $X = (\mathbb{R}^n, \|\cdot\|_X)$  is a normed space such that  $\mathcal{L}(X)$  has an  $\alpha$ -approximate separation oracle for some  $\alpha \geq C_2(X^*)$ . Then, there is an algorithm that given an input matrix  $A \in M_n(\mathbb{R})$  makes polynomially many oracle calls and runs in time  $n^{O(1)}$ , and outputs a matrix  $\mathbb{W} \in \mathbf{PSD}^n$  with  $\mathbb{W} \succeq A$  that satisfies*

$$\inf \{Q_X^{\max}(\mathbb{M}) : \mathbb{M} \in \mathbf{PSD}^n \text{ and } \mathbb{M} \succeq A\} \geq \frac{Q_X^{\max}(\mathbb{W})}{\alpha}. \quad (17)$$

*Also, if  $X = (\mathbb{R}^n, \|\cdot\|_X)$ ,  $Y = (\mathbb{R}^m, \|\cdot\|_Y)$  are normed spaces such that  $\mathcal{L}(X)$ ,  $\mathcal{L}(Y)$  have  $\alpha$ -approximate separation oracles for  $\alpha \geq \max\{C_2(X^*), C_2(Y^*)\}$ , then there is an algorithm that given an input matrix  $B \in M_{n \times m}(\mathbb{R})$  makes polynomially many oracle calls and runs in time that is polynomial in  $n, m$ , and outputs a matrices  $\mathbb{W} \in$*

$\mathbf{PSD}^n$ ,  $\mathbb{V} \in \mathbf{PSD}^m$  with  $\begin{pmatrix} \mathbb{W} & 0 \\ 0 & \mathbb{V} \end{pmatrix} \succeq \begin{pmatrix} 0 & B \\ B^* & 0 \end{pmatrix}$  and

$$\inf \left\{ Q_X^{\max}(\mathbb{M}_1) + Q_Y^{\max}(\mathbb{M}_2) : (\mathbb{M}_1, \mathbb{M}_2) \in \mathbf{PSD}^n \times \mathbf{PSD}^m \right. \\ \left. \text{and } \begin{pmatrix} \mathbb{M}_1 & 0 \\ 0 & \mathbb{M}_2 \end{pmatrix} \succeq \begin{pmatrix} 0 & B \\ B^* & 0 \end{pmatrix} \right\} \geq \frac{Q_X^{\max}(\mathbb{W}) + Q_Y^{\max}(\mathbb{V})}{\alpha}. \quad (18)$$

We will next explain the ingredients that go into (17); the justification of (18) is similar and will be carried out in the full version [BLN21]. The reason why we include (18) here is that it is important for the bilinear variant (2), namely for the question of approximating the operator norm  $\|A\|_{Y \rightarrow X^*}$ .

The goal of (17) is to  $O(\alpha)$ -approximately minimize the convex function  $\mathbb{M} \mapsto Q_X^{\max}(\mathbb{M})$  over the convex set  $\{\mathbb{M} \in \mathbf{PSD}^n : \mathbb{M} \succeq A\}$ . In the full version [BLN21], we will show that in order to efficiently find a  $(1 + o(1))\alpha$ -approximate minimizer, it suffices to show that each of the corresponding sub-level sets  $\{\mathbb{M} \in \mathbf{PSD}^n : Q_X^{\max}(\mathbb{M}) \leq t\} : t \in \mathbb{R}$  has a  $(1 + o(1))\alpha$ -approximate separation oracle. By homogeneity, we therefore need to show that under the assumptions of Theorem 1.9, the convex set  $\{\mathbb{M} \in \mathbf{PSD}^n : Q_X^{\max}(\mathbb{M}) \leq 1\}$  has a  $(1 + o(1))\alpha$ -approximate separation oracle.

To this end, fix  $\mathbb{M} \in \mathbf{PSD}^n$  and consider the following optimization problem.

$$\max \left\{ \mathbb{E} \left[ \|\mathbb{M}^{\frac{1}{2}} \mathbb{V}^{\frac{1}{2}} \mathbf{g}\|_{X^*}^2 \right] : \mathbb{V} \in \mathbf{PSD}^n \text{ and } \text{Tr}(\mathbb{V}) \leq 1 \right\}. \quad (19)$$

We claim that one can find in polynomial time and with polynomially many oracle calls a matrix  $\mathbb{V} \in \mathbf{PSD}^n$  that attains this maximum up to a factor of  $(1 + o(1))\alpha$ . Indeed, in the full version [BLN21] we will show that for this it suffices to check that each of the corresponding super-level sets

$$\left\{ \mathbb{V} \in \mathbf{PSD}^n : \mathbb{E} \left[ \|\mathbb{M}^{\frac{1}{2}} \mathbb{V}^{\frac{1}{2}} \mathbf{g}\|_{X^*}^2 \right] \geq t \right\} : t \in \mathbb{R} \quad (20)$$

has an  $\alpha$ -approximate separation oracle. Since each of the sets appearing in (20) is (by definition) a linear transformation of the lower covariance body of  $X$ , the assumption of Theorem 1.9 ensures that the desired oracle exists. Therefore, we can find  $\mathbb{V} \in \mathbf{PSD}^n$  with  $\text{Tr}(\mathbb{V}) \leq 1$  at which the maximum in (19) is attained up to a factor of  $(1 + o(1))\alpha$ .

Finally, we can describe what the desired oracle for  $\{\mathbb{M}' \in \mathbf{PSD}^n : Q_X^{\max}(\mathbb{M}') \leq 1\}$  will output for the input matrix  $\mathbb{M}$ . For each realization of the Gaussian vector  $\mathbf{g} \in \mathbb{R}^n$ , let  $x_{\mathbf{g}} \in \text{Ball}(X)$  be the random vector that is given by

$$x_{\mathbf{g}} \stackrel{\text{def}}{=} \operatorname{argmax}_{x' \in \text{Ball}(X)} \langle x', \mathbb{M}^{\frac{1}{2}} \mathbb{V}^{\frac{1}{2}} \mathbf{g} \rangle.$$

Note that  $x_{\mathbf{g}}$  can be found efficiently using polynomially many membership queries to  $\text{Ball}(X)$ , using the classical theory of convex programming [GLS93]. If

$$\frac{\|\mathbb{M}^{\frac{1}{2}} \mathbb{V}^{\frac{1}{2}} \mathbf{g}\|_{X^*}}{\|\mathbb{V}^{\frac{1}{2}} \mathbf{g}\|_{\ell_2^n}} \leq 1,$$

then the oracle outputs “Inside.” Otherwise, the oracle outputs the hyperplane

$$\{\mathbb{M}' \in M_n(\mathbb{R}) : \langle \mathbb{M}' x_{\mathbf{g}}, x_{\mathbf{g}} \rangle = 1\}.$$

By tracking the above definitions, one checks that this oracle satisfies the desired properties with positive probability. One gets this to hold with sufficiently high probability (to account for the polynomially many oracle calls) by repeating the above procedure with  $n^{O(1)}$  independent samples from  $\mathfrak{g}$  rather than only one such sample; the details appear in the full version [BLN21].

With the algorithmic groundwork of Theorem 1.9 complete, our final algorithm relies on the analytic inequalities that are contained in the following theorem (see the full version [BLN21] for proofs).

**THEOREM 1.10 (FACTORIZATION INEQUALITIES).**

For every normed space  $X = (\mathbb{R}^n, \|\cdot\|_X)$  and  $A \in M_n(\mathbb{R})$  we have

$$\begin{aligned} Q_X^{\max}(A) &\leq \inf \{Q_X^{\max}(\mathbb{W}) : \mathbb{W} \in \text{PSD}^n \text{ and } \mathbb{W} \geq A\} \\ &\leq T_2(X)^2 \cdot Q_X^{\max}(A). \end{aligned} \quad (21)$$

Also, for every two normed spaces  $X = (\mathbb{R}^n, \|\cdot\|_X)$ ,  $Y = (\mathbb{R}^m, \|\cdot\|_Y)$ , and every  $B \in M_{n \times m}(\mathbb{R})$ , denote

$$\begin{aligned} \gamma_2^{Y \rightarrow X^*}(B) &\stackrel{\text{def}}{=} \\ &\inf \left\{ \frac{Q_X^{\max}(\mathbb{W}) + Q_Y^{\max}(\mathbb{V})}{2} : (\mathbb{W}, \mathbb{V}) \in \text{PSD}^n \times \text{PSD}^m \right. \\ &\quad \left. \text{and } \begin{pmatrix} \mathbb{W} & 0 \\ 0 & \mathbb{V} \end{pmatrix} \geq \begin{pmatrix} 0 & B \\ B^* & 0 \end{pmatrix} \right\}. \end{aligned} \quad (22)$$

Then,

$$\begin{aligned} \|B\|_{Y \rightarrow X^*} &\leq \gamma_2^{Y \rightarrow X^*}(B) \\ &\leq C_2(X^*)C_2(Y^*) \log(C_2(X^*)C_2(Y^*)) \cdot \|B\|_{Y \rightarrow X^*}. \end{aligned} \quad (23)$$

We chose the notation  $\gamma_2^{Y \rightarrow X^*}(B)$  in (22) purposefully to coincide with the classical functional analytic notation for factorization norms [Pis86], namely it is the  $\gamma_2$  norm of  $B$  when it is viewed as an operator from  $Y$  to  $X^*$ . The equality (22) is therefore a variational characterization of the classical quantity in the left hand side in terms of the infimum on the right hand side; we prove this identity in the full version [BLN21]. With this identity at hand, the inequality (23) is an application of a deep factorization theorem of Pisier [Pis80]. The inequality (21) is inspired by the aforementioned factorization theory, but it seems to be new; it could be viewed as a factorization theorem for quadratic forms (see the full version [BLN21] for details) and it would be interesting to study its ramifications within functional analysis.

By combining Theorem 1.9 with Theorem 1.10, we get the following algorithmic result.

**THEOREM 1.11 (GENERIC FRAMEWORK).**

Suppose that  $X = (\mathbb{R}^n, \|\cdot\|_X)$  is a normed space such that  $\mathcal{L}(X)$  has an  $\alpha$ -approximate separation oracle for some  $\alpha \geq C_2(X^*)$ . Then, there is an algorithm that given an input matrix  $A \in M_n(\mathbb{R})$  makes polynomially many oracle calls and runs in time  $n^{O(1)}$ , and outputs a number  $\text{Alg}_1$  that is guaranteed to satisfy

$$Q_X^{\max}(A) \leq \text{Alg}_1 \leq \alpha T_2(X)^2 \cdot Q_X^{\max}(A).$$

For the bilinear case, if  $X = (\mathbb{R}^n, \|\cdot\|_X)$ ,  $Y = (\mathbb{R}^m, \|\cdot\|_Y)$  are normed spaces such that  $\mathcal{L}(X)$ ,  $\mathcal{L}(Y)$  have  $\alpha$ -approximate separation oracles for some  $\alpha \geq \max\{C_2(X^*), C_2(Y^*)\}$ , then there is an algorithm that given an input matrix  $B \in M_{n \times m}(\mathbb{R})$  makes polynomially many

oracle calls and runs in time that is polynomial in  $n, m$ , and outputs a number  $\text{Alg}_2$  that is guaranteed to satisfy

$$\begin{aligned} \|B\|_{Y \rightarrow X^*} &\leq \text{Alg}_2 \\ &\leq \alpha C_2(X^*)C_2(Y^*) \log(C_2(X^*)C_2(Y^*)) \cdot \|B\|_{Y \rightarrow X^*}. \end{aligned}$$

**REMARK 4.** One often wishes to approximate efficiently not only the values of the quantities  $Q_X^{\max}(A)$  and  $\|B\|_{Y \rightarrow X^*}$ , but also to find efficiently the vector  $x \in \mathbb{R}^n$  at which  $Q_X^{\max}(A)$  is approximately attained, and correspondingly the vectors  $x \in \mathbb{R}^n$  and  $y \in \mathbb{R}^m$  at which  $\|B\|_{Y \rightarrow X^*}$  is approximately attained. For the latter, we need a constructive version of Pisier's factorization theorem that entails several adjustments of its classical proof; the details appear in the full version [BLN21]. For this variant (namely, finding almost maximizing vectors rather than only estimating the quantity  $\|B\|_{Y \rightarrow X^*}$ ), we get the slightly worse approximation factor

$$O(\alpha C_2(X^*)C_2(Y^*) \log(\alpha C_2(X^*)C_2(Y^*)))$$

in the second part of Theorem 1.11.

### 1.3 Examples of Applications

Theorem 1.11 focuses our attention to designing approximate separation oracles for lower covariance bodies. In the specific cases that we examined, it turns out that this task is tractable because it reduces to probabilistic (Khinchine-type) inequalities that are available in the literature. We will examine such applications next. The advantage of the above approach is that it shifts our focus to a new algorithmic task. This task most likely cannot always be achieved due to the aforementioned hardness results, but in specific cases it becomes a concrete new question that lends itself to classical tools that may have not seemed relevant in the initial formulation of the problem. This reframing also allows us to prove various closure properties for the class of convex bodies for which efficient quadratic or bilinear maximization is possible.

### Closure Properties

Given normed spaces  $X = (\mathbb{R}^n, \|\cdot\|_X)$  and  $Y = (\mathbb{R}^m, \|\cdot\|_Y)$ , one can obtain various other normed spaces. The most basic examples are passing to a subspace or a quotient of  $X$ . One can also consider the normed spaces  $X + Y = (\mathbb{R}^n, \|\cdot\|_{X+Y})$  and  $X \cap Y = (\mathbb{R}^n, \|\cdot\|_{X \cap Y})$  whose unit balls are  $\text{Ball}(X) + \text{Ball}(Y) = \{x + y : (x, y) \in \text{Ball}(X) \times \text{Ball}(Y)\}$  and  $\text{Ball}(X) \cap \text{Ball}(Y)$ , respectively; we call the former the Minkowski sum of  $X$  and  $Y$  and we call the latter the intersection of  $X$  and  $Y$ . A further operation of great importance is the 1-parameter family of complex<sup>2</sup> interpolation spaces  $\{[X, Y]_\theta\}_{\theta \in [0,1]}$  whose definition is recalled in the full version [BLN21] (see the monograph [BL76] for a thorough account). There are of course more such operations (a notable example is duality), but the above list of constructions is singled out because it always results in a normed space whose type 2 constant does not exceed  $O(\max\{T_2(X), T_2(Y)\})$ , which is crucial for us due to Theorem 1.3.

In the full version [BLN21], we use the above framework to show that the class of normed spaces  $X$  with  $T_2(X) = O(1)$  for which

<sup>2</sup>The real interpolation method (see [BL76]) furnishes another such 1-parameter family of intermediate norms, but in the present work we will investigate only the complex interpolation method and we expect that it would be mechanical to obtain the analogous results for real interpolation using the same ideas.

there exists a polynomial time  $O(1)$ -approximation algorithm for  $\mathcal{Q}_X^{\max}(A)$  is preserved under subspaces, quotients, Minkowski sums, intersection and complex interpolation. Among these operations, passing to subspaces is quite straightforward, but the rest rely on the methodology that is developed here. Beyond the intrinsic interest of such closure properties, we remark that if one starts with the many examples of spaces that belong to the aforementioned class (see below), then these operations produce a rich variety of new examples that were beyond the reach of previous methods. Also, observe that these closure properties do not assume any information whatsoever on the initial algorithms: These algorithms are used as a “black box” to design an approximate separation oracle for the lower covariance body of the resulting normed space, after which one applies the first part of Theorem 1.11. An analogous treatment of the bilinear case is carried out using the second part of Theorem 1.11, where closure under quotients and Minkowski sums is derived under the assumption that cotype 2 constants of the duals of the initial spaces are  $O(1)$ ; we do not treat the rest of the above-listed operations because they do not necessarily preserve this bounded cotype 2 assumption on the dual.

## Symmetric Norms

A norm  $\|\cdot\|$  on  $\mathbb{R}^n$  is said to be a symmetric norm if  $\|x\| \asymp \|(\varepsilon_1 x_{\pi(1)}, \dots, \varepsilon_n x_{\pi(n)})\|$  for any  $x \in \mathbb{R}^n$ , any permutation  $\pi$  of  $\{1, \dots, n\}$ , and any choice of signs  $\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}$ .<sup>3</sup> This is a well studied class of norms occurring frequently in the computer science, learning and optimization literature. Several papers have attempted to characterize the symmetric norms that are appropriate for various algorithmic tasks; see e.g. [LNRW19, ANN<sup>+</sup>17, BBC<sup>+</sup>17, ALS<sup>+</sup>18, SWZ19, SWY<sup>+</sup>19].

In the full version [BLN21], we use Theorem 1.11 to give a constant-factor approximation algorithm for quadratic (respectively bilinear) maximization over unit balls of symmetric norms whose type-2 constant (respectively the cotype-2 constant of their dual) is  $O(1)$ . Combined with Theorem 1.3, we obtain a near characterization of those symmetric norms for which quadratic maximization admits a constant factor approximation algorithm.

The class of those symmetric norms that have a bounded (or slowly growing) type-2 constant contains many examples that are not covered by the available literature. Below we will list some explicit examples of symmetric norms appearing in the optimization literature for various algorithmic tasks and for which we can conclude either a new quadratic maximization approximation algorithm or a new inapproximability result.

- (1) An Orlicz norm  $\ell_\phi^n$  is defined by setting for every  $x \in \mathbb{R}^n$ ,

$$\|x\|_{\ell_\phi^n} \stackrel{\text{def}}{=} \inf \left\{ \lambda > 0 : \sum_{i=1}^n \phi\left(\frac{|x_i|}{\lambda}\right) \leq 1 \right\},$$

where  $\phi : [0, \infty) \rightarrow [0, \infty)$  is a convex function satisfying  $\phi(0) = 0$  and  $\phi(t) > 0$  for all  $t > 0$ . Thus, in the special case  $\phi(t) = t^p$  for some  $p \geq 1$  we have  $\ell_\phi^n = \ell_p^n$ . Among the many applications of Orlicz norms, we note that they are important for the study of tail behaviour of random variables and are

<sup>3</sup>One could replace the exact invariance under permutations and signs by the analogous approximate requirement  $\|x\| \asymp \|(\varepsilon_1 x_{\pi(1)}, \dots, \varepsilon_n x_{\pi(n)})\|$ . We will not do so here, though our results work under that assumption as well.

studied in statistics/machine learning [CW14] as examples of M-estimators with (convex loss functions).

The class of Orlicz norms with bounded type-2 constant has a complete description [Kat98] as the set of norms  $\ell_\phi^n$  where  $\phi : [0, \infty) \rightarrow [0, \infty)$  satisfies the following two conditions.

- (a) There are constants  $K, \delta, c > 0$  such that for all  $t > 0$ , if  $\phi(t) \leq \delta$ , then  $\phi(2t) \leq K\phi(t) + c$ .
  - (b) There is  $\psi : [0, \infty) \rightarrow [0, \infty)$  such that  $t \mapsto \psi(\sqrt{t})$  is convex and  $\phi$  is equivalent to  $\psi$  in the sense that there are constants  $K_1, K_2, \delta_1, \delta_2, c_1, c_2 > 0$  such that  $\psi(t) \leq \delta_1$  implies  $\phi(K_1 t) \leq \psi(t) + c_1$  and  $\phi(t) \leq \delta_2$  implies  $\psi(K_2 t) \leq \phi(t) + c_2$  for all  $t > 0$ .
- (2) Norms whose unit balls are of the form  $\text{Ball}(\ell_p^n) \cap (\alpha \text{Ball}(\ell_q^n))$  have a  $O(1)$  type-2 constant (i.e., independent of  $n, \alpha$ ) whenever  $2 \leq p, q < \infty$ . Quadratic maximization over such norms is considered in order to capture optimization problems with a sparsity restriction. For instance, the densest  $k$ -subgraph and  $k$ -sparse principal component analysis, which are extensively studied optimization problems, can be cast as quadratic maximization by taking the underlying norm to be  $\text{Ball}(\ell_\infty^n) \cap (k \text{Ball}(\ell_1^n))$  and  $\text{Ball}(\ell_2^n) \cap (\sqrt{k} \text{Ball}(\ell_1^n))$ , respectively; note that these norms have polynomially large type-2 constant due to the  $\ell_1$  component, which is consistent with the widespread belief that densest  $k$ -subgraph and  $k$ -sparse principal component analysis are hard to approximate. The above examples with  $2 \leq p, q < \infty$  can be viewed as smoothed out versions of these classical algorithmic questions which do admit a polynomial time constant factor approximation algorithm.
- (3) Motivated by applications to kernel pattern matching, [NS09] gave an approximation algorithm for the following symmetric norm that has slowly growing type-2 constant.

$$\|(x_1, \dots, x_n)\|_{p, \infty} \stackrel{\text{def}}{=} \max_{i \in \{1, \dots, n\}} i^{\frac{1}{p}} x_i^*,$$

where  $p \geq 2$  and  $x_i^*$  denotes the entry of  $(|x_1|, \dots, |x_n|)$  with the  $i$ -th largest magnitude.

- (4) Order statistics norms are defined as the inner product of a non-increasing vector  $a$  with the sorted vector  $x^*$ . This class is well studied in the clustering literature [BSS18, CS19a, CS19b] and includes e.g. the top- $k$  norm (sum of top  $k$  magnitudes of  $x$ ). The type-2 constant of such norms is bounded whenever  $a$  has bounded support.

## Unitarily Invariant Matrix Norms

A norm  $\|\cdot\| : M_n(\mathbb{C}) \rightarrow [0, \infty)$  on the space  $M_n(\mathbb{C})$  of  $n \times n$  matrices with complex entries is said to be unitarily invariant if  $\|UAV\| = \|A\|$  for any matrix  $A \in M_n(\mathbb{C})$  and any two unitary matrices  $U, V \in UM_n(\mathbb{C})$ ; this can be defined analogously for matrices with real entries (using orthogonal matrices), as well as for rectangular matrices, and all of our results hold in these settings. Key examples include the Schatten–von Neumann trace class  $S_p$  for  $p \in [1, \infty)$ ,

which is defined by

$$\forall A \in M_n(\mathbb{C}), \quad \|A\|_{S_p} \stackrel{\text{def}}{=} \left( \text{Tr}((AA^*)^{\frac{p}{2}}) \right)^{\frac{1}{p}} = \left( \text{Tr}((A^*A)^{\frac{p}{2}}) \right)^{\frac{1}{p}} \\ = \left( \sum_{j=1}^n \sigma_j(A)^p \right)^{\frac{1}{p}},$$

where  $\sigma_1(A) \geq \dots \geq \sigma_n(A) \geq 0$  are the singular values of  $A$ . Thus,  $\|A\|_{S_\infty} = \|A\|_{\ell_2^n(\mathbb{C}) \rightarrow \ell_2^n(\mathbb{C})}$  is the usual operator norm of  $A$ . Another example is the Ky-Fan  $k$ -norm  $\|\cdot\|_{(k)}$  for each  $k \in \{1, \dots, n\}$ , which is the sum of the top  $k$  singular values, i.e.,

$$\forall A \in M_n(\mathbb{C}), \quad \|A\|_{(k)} \stackrel{\text{def}}{=} \sum_{j=n-k+1}^n \sigma_j(A).$$

More generally, if  $E = (\mathbb{R}^n, \|\cdot\|_E)$  is a symmetric normed space, then the following norm is unitarily invariant and any unitarily invariant norm is obtained in this way (the fact that this defines a norm is not immediate; see e.g. [Bha97] for a proof).

$$\forall A \in M_n(\mathbb{C}), \quad \|A\|_{S_E} \stackrel{\text{def}}{=} \|(\sigma_1(A), \dots, \sigma_n(A))\|_E.$$

A (substantial) theorem of [GTJ83] asserts that  $\|\cdot\|_E$  has  $O(1)$  type 2 or cotype 2 constant if and only if  $\|\cdot\|_E$  does.

In the full version [BLN21] we use Theorem 1.11 to obtain a constant-factor approximation algorithm for quadratic (respectively bilinear) maximization over unitarily invariant norms with bounded type-2 constant (respectively whose dual has cotype-2 constant). In particular, this provides a different rounding algorithm for the non-commutative Grothendieck problem [NRV13] (namely, bilinear maximization over the operator norm), albeit with a worse universal constant than in [NRV13]. As another concrete example, this gives a constant factor approximation algorithm for bilinear maximization over Ky-Fan  $k$ -norms when  $k = O(1)$ . Combined with Theorem 1.3, we thus obtain a near characterization of unitarily invariant matrix norms over which quadratic maximization admits a constant factor approximation algorithm.

## Robust Principle Component Analysis

In [NRV13], efficient bilinear maximization over the operator norm (Schatten- $\infty$ ) was used to give a constant factor approximation algorithm for the following subspace approximation problem, called  $R_1$ -PCA, which was introduced in [DZHZ06]. Given a set of vectors  $v_1, \dots, v_m \in \mathbb{R}^n$  find a  $k$ -dimensional subspace  $S \subseteq \mathbb{R}^n$  maximizing the sum of the Euclidean lengths of the orthogonal projections  $\Pi_S v_1, \dots, \Pi_S v_m$  of  $v_1, \dots, v_m$  onto  $S$ . Thus, the goal of  $R_1$ -PCA is to find a  $k$ -dimensional subspace  $S \subseteq \mathbb{R}^n$  for which the quantity  $\sum_{i=1}^m \|\Pi_S v_i\|_{\ell_2}^2$  is (approximately) minimized.

Our framework implies that a more general class of robust PCA variants admits constant factor approximation algorithms. Given a normed space  $X = (\mathbb{R}^m, \|\cdot\|_X)$ , one can use it to aggregate the length of the projections, thus leading to the following subspace approximation problem.

$$\text{OPT} \stackrel{\text{def}}{=} \max_{\dim(S)=k} \|(\|\Pi_S v_1\|_2, \dots, \|\Pi_S v_m\|_2)\|_X.$$

Let  $T$  denote the linear operator taking an  $m \times k$  matrix  $U$  with column vectors  $u_1, \dots, u_k \in \mathbb{R}^n$  as input and outputting the vector

$$\langle u_1, v_1 \rangle, \dots, \langle u_k, v_1 \rangle \oplus \dots \oplus \langle u_1, v_m \rangle, \dots, \langle u_k, v_m \rangle,$$

where  $\oplus$  denotes vector-concatenation. Let  $\|\cdot\|_{X(\ell_2^k)}$  be a norm defined over the set of sequences  $(a_i)_{i=1}^m \in (\mathbb{R}^k)^m$  of  $k$ -dimensional vectors and given by

$$\|(a_i)_{i=1}^m\|_{X(\ell_2^k)} \stackrel{\text{def}}{=} \|(\|a_1\|_2, \dots, \|a_m\|_2)\|_X.$$

Then, one can cast OPT as a bilinear maximization problem in the following way.

$$\text{OPT} = \max_{U \in O_n} \|T(U)\|_{X(\ell_2^k)} = \max_{\|U\|_{S_\infty} \leq 1} \|T(U)\|_{X(\ell_2^k)} \\ = \|T\|_{S_\infty \rightarrow X(\ell_2^k)},$$

where  $O_n \subseteq M_n(\mathbb{R})$  is the set of orthogonal matrices. The second equality above follows since the extreme points of  $\text{Ball}(S_\infty)$  are precisely  $O_n$ , and the maximum of a convex function over a convex set occurs at an extreme point.

Thanks to this bilinear maximization formulation, Theorem 1.11 may be combined with the lower covariance separation oracles constructed in the full version [BLN21], to provide good approximation algorithms for a variety of norms  $\|\cdot\|_X$ , like constant approximations for sign-invariant norms with 2-concavity constant 1 or symmetric norms with bounded cotype-2 constant. We illustrate the versatility of our framework by providing a more intricate example; by combining Theorem 1.11 with the separation oracles constructed in the full version [BLN21] and using algorithmic closure properties for complex interpolation (see the full version [BLN21]), we obtain a  $(\log n)^{O(1)}$ -factor approximation algorithm for the following refinement of robust-PCA: Find a  $k$ -dimensional subspace  $S \subseteq \mathbb{R}^n$  (approximately) maximizing

$$\|(\Pi_S v_i)_{i=1}^m\|_{[X_0, X_1]_\theta},$$

where  $[\cdot, \cdot]_\theta$  denotes complex interpolation,  $\alpha \geq 0$  is a parameter, and

$$\|(\Pi_S v_i)_{i=1}^m\|_{X_0} \stackrel{\text{def}}{=} \sum_{i=1}^m \|\Pi_S v_i\|_2 \quad \text{and} \\ \|(\Pi_S v_i)_{i=1}^m\|_{X_1} \stackrel{\text{def}}{=} \alpha \cdot \sum_{i=1}^m \sum_{j=1}^m \|\Pi_S v_i - \Pi_S v_j\|_2.$$

As defined above,  $X_1$  is a semi-norm but can be made into a norm by adding a sufficiently small multiple of  $\ell_2^n$  which would cause negligible change to the objective value. By tuning the parameters  $\alpha \geq 0$  and  $\theta \in [0, 1]$ , the above optimization problem intuitively asks for a subspace maximizing its correlation with the given vectors  $\{v_i\}_{i=1}^m$ , while also requiring that the orthogonal projections onto  $S$  of these vectors are not clustered together much on average.

## 1.4 Brief Summary of the Literature and Problems Captured by Quadratic Maximization

Here we will mention some of what is known about the quadratic and bilinear optimization problem over convex bodies. Quadratic/bilinear maximization over  $\text{Ball}(\ell_2^n)$  correspond to the familiar linear-algebraic quantities maximum eigenvalue/maximum singular value. The (non-origin-symmetric) case of (1) when  $K$  is a simplex has been investigated in [HH06, dKLP06], partly in connection to problems in computational biology. The case when  $K$  is

a polytope with polynomially many facets is classical. It is among the most important non-linear optimization problems, with a wide range of applications in operations research, computational biology and economics. See [FL92, BR95, Bri02] for more information on the computational complexity of such problems.

Perhaps the first nontrivial and most influential case of bilinear maximization is Grothendieck’s classical inequality [Gro53] and its more common formulation in [LP68], which corresponds to the case  $K = \text{Ball}(\ell_\infty^n)$ . This leads to a constant factor polynomial time algorithm, as shown in [AN06] (see [BMMN13] for the best known approximation factor), with a variety of applications to combinatorial optimization. The quadratic maximization problem over  $\text{Ball}(\ell_\infty)$  was studied in [CW04] with application to correlation clustering, and the matching integrality-gap lower bound in this case was obtained in [AMMN06]. Hardness results in these settings (under various complexity assumptions) were obtained in [AN06, ABH<sup>+</sup>05, KO09, RS09]. The survey [KN12] is devoted to the use of Grothendieck-type inequalities in combinatorial optimization.

Krivine [Kri73] (see also [Pis12]) observed that Grothendieck’s inequality generalizes (with the same constant) to the class of norms of the form

$$\| (x_1, \dots, x_n) \| \stackrel{\text{def}}{=} \| (x_1^2, \dots, x_n^2) \|_Y^{\frac{1}{2}},$$

where  $\| \cdot \|_Y$  is a norm on  $\mathbb{R}^n$  that satisfies the symmetry condition (3). Such norms are clearly invariant to flipping signs of the entries and are precisely those norms having a 2-convexity constant of 1 (see the full version [BLN21] for definitions). Hereafter, we shall refer to them as exactly 2-convex norms. Note in particular that the above class includes the norm  $\ell_p^n$  whenever  $p \geq 2$ . Underlying Krivine’s observation is a constant factor bound on the integrality gap of the bilinear analogue of the convex programming relaxation (4) over exactly 2-convex norms; in [Nes98], a different proof of this was obtained. The problem of quadratic maximization over exactly 2-convex norms was investigated in [NS09], where a constant factor approximation algorithm was obtained under the additional (necessary) assumption of bounded  $q$ -concavity for some finite  $q$  (see the full version [BLN21] for the definition); this was used in [NS09] to obtain a  $(\log \log n)^{O(1)}$ -approximation algorithm for a special case of the quadratic assignment problem. It can also be shown that the  $(\log n)$ -approximation algorithm for vertex expansion of a graph due to [LRV13] is a consequence of the algorithm of [NS09].

Implicit in the non-commutative Khintchine inequality [LPP91] is a constant factor convex programming algorithm for Quadratic Maximization over Schatten- $p$  when  $2 \leq p < \infty$  (and a  $\log n$ -approximation when  $p = \infty$ ). In the bilinear Schatten- $\infty$  case, Grothendieck [Gro53] conjectured a noncommutative version of his inequality which was proven in [Pis78] (the sharp constant was obtained in [Haa85]). In [NRV13], algorithmic proofs of the non-commutative Grothendieck inequality were derived, thereby obtaining efficient constant factor rounding algorithms for bilinear maximization over Schatten- $\infty$ . This was used in [NRV13] to give approximation algorithms for robust principal component analysis and a generalization of the orthogonal Procrustes problem. In [RV15], it was shown how this can be used to bound the power of entanglement in quantum XOR games. A corresponding (sharp)

hardness result was obtained in [BRS15] (see also [HV16] for a different proof).

## Other Problems in the Literature Captured by Quadratic Maximization

The bilinear  $\ell_p$  case captures the problem of certifying hypercontractivity which in turn has connections to small set expansion and quantum separability ([BBH<sup>+</sup>12]). Vertex expansion and a related analytic proxy ([LRV13]) can be cast as quadratic maximization, and so can densest- $k$ -subgraph, sparse-PCA, the spread constant of a metric [ABS98], and the Poincaré constant (in discrete domains). Approximability/inapproximability aspects of these expansion-type problems have been the subject of a large body of work. Expansion-type problems are of interest in part due to their connection to the unique games conjecture, and also due to their relevance to hardness results for optimization over pseudo-random instances.

For appropriate choices of linear maps and convex sets, quadratic maximization also captures (upto constants) the maximization of the absolute value of homogeneous polynomials of any constant degree. Homogeneous polynomial maximization is a very expressive class of problems in its own right, and has connections to quantum information theory [BBH<sup>+</sup>12], refuting random constraint satisfaction problems [RRS16], statistical physics, tensor principal component analysis and tensor decomposition [BKS15, GM15, MR14, HSS15], game theory, control theory and population dynamics [DK08].

Quadratic maximization also captures problems of interest in compressed sensing and coding theory, like subspace distortion, or the sparsest vector in a subspace.

## ACKNOWLEDGMENTS

The first and second authors would like to thank Venkatesan Guruswami and Madhur Tulsiani for useful discussions on these topics over the last few years.

The authors would like to thank the anonymous STOC’21 referees for useful comments on the writing of the manuscript.

## REFERENCES

- [ABH<sup>+</sup>05] Sanjeev Arora, Eli Berger, Elad Hazan, Guy Kindler, and Muli Safra. On non-approximability for quadratic programs. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 206–215. IEEE, 2005.
- [ABS98] Noga Alon, Ravi Boppana, and Joel Spencer. An asymptotic isoperimetric inequality. *Geometric & Functional Analysis GAFA*, 8(3):411–436, 1998.
- [Ala06] Noga Alon. Maximizing a quadratic form on the  $\ell_1^n$  unit ball. Unpublished manuscript, 2006.
- [ALS<sup>+</sup>18] Alexandr Andoni, Chengyu Lin, Ying Sheng, Peilin Zhong, and Ruiqi Zhong. Subspace embedding and linear regression with orlicz norm. *arXiv preprint arXiv:1806.06430*, 2018.
- [AMMN06] Noga Alon, Konstantin Makarychev, Yury Makarychev, and Assaf Naor. Quadratic forms on graphs. *Inventiones mathematicae*, 163(3):499–522, 2006. Conference version in STOC ’05.
- [AN06] Noga Alon and Assaf Naor. Approximating the cut-norm via Grothendieck’s inequality. *SIAM J. Comput.*, 35(4):787–803, 2006.
- [ANN<sup>+</sup>17] Alexandr Andoni, Huy L Nguyen, Aleksandar Nikolov, Ilya Razenshteyn, and Erik Waingarten. Approximate near neighbors for general symmetric norms. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 902–913, 2017.
- [BBC<sup>+</sup>17] Jarosław Błasiok, Vladimir Braverman, Stephen R Chestnut, Robert Krauthgamer, and Lin F Yang. Streaming symmetric norms via measure concentration. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 716–729, 2017.

- [BBH<sup>+</sup>12] Boaz Barak, Fernando GSL Brandao, Aram W Harrow, Jonathan Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 307–326. ACM, 2012.
- [Bha97] Rajendra Bhatia. *Matrix analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.
- [BKS15] Boaz Barak, Jonathan A Kelner, and David Steurer. Dictionary learning and tensor decomposition via the sum-of-squares method. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 143–151. ACM, 2015.
- [BL76] Jöran Bergh and Jörgen Löfström. *Interpolation spaces. An introduction*. Springer-Verlag, Berlin-New York, 1976. Grundlehren der Mathematischen Wissenschaften, No. 223.
- [BLN21] Vijay Bhattiprolu, Euiwoong Lee, and Assaf Naor. A framework for quadratic form maximization over convex sets through nonconvex relaxations. *Soon to appear on arxiv*, 2021.
- [BMMN13] Mark Braverman, Konstantin Makarychev, Yuri Makarychev, and Assaf Naor. The Grothendieck constant is strictly smaller than Krivine’s bound. In *Forum of Mathematics, Pi*, volume 1. Cambridge University Press, 2013. Conference version in FOCS ’11.
- [BR95] Mihir Bellare and Phillip Rogaway. The complexity of approximating a nonlinear program. *Math. Programming*, 69(3, Ser. A):429–441, 1995.
- [Bri02] Andreas Brieden. Geometric optimization problems likely not contained in  $\text{APX}$ . *Discrete Comput. Geom.*, 28(2):201–209, 2002.
- [BRS15] Jop Briët, Oded Regev, and Rishi Saket. Tight hardness of the non-commutative Grothendieck problem. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 1108–1122. IEEE, 2015.
- [BSS18] Jaroslaw Byrka, Krzysztof Sornat, and Joachim Spoerhase. Constant-factor approximation for ordered  $k$ -median. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 620–631, 2018.
- [CS19a] Deeparnab Chakrabarty and Chaitanya Swamy. Approximation algorithms for minimum norm and ordered optimization problems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 126–137, 2019.
- [CS19b] Deeparnab Chakrabarty and Chaitanya Swamy. Simpler and better algorithms for minimum-norm load balancing. *arXiv preprint arXiv:1905.00044*, 2019.
- [CW04] Moses Charikar and Anthony Wirth. Maximizing quadratic programs: Extending grothendieck’s inequality. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 54–60. IEEE Computer Society, 2004.
- [CW14] Kenneth L Clarkson and David P Woodruff. Sketching for  $m$ -estimators: A unified approach to robust regression. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 921–939. SIAM, 2014.
- [DK08] Etienne De Klerk. The complexity of optimizing over a simplex, hypercube or sphere: a short survey. *Central European Journal of Operations Research*, 16(2):111–125, 2008.
- [dKLP06] Etienne de Klerk, Monique Laurent, and Pablo A Parrilo. A PTAS for the minimization of polynomials of fixed degree over the simplex. *Theoret. Comput. Sci.*, 361(2-3):210–225, 2006.
- [DPR72] Ed Dubinsky, A. Pelczyński, and H. P. Rosenthal. On Banach spaces  $X$  for which  $\Pi_2(L_\infty, X) = B(L_\infty, X)$ . *Studia Math.*, 44:617–648, 1972.
- [DZHZ06] Chris Ding, Ding Zhou, Xiaofeng He, and Hongyuan Zha.  $r_1$ -pca: Rotational invariant  $L_1$ -norm principal component analysis for robust subspace factorization. In *Proceedings of the 23rd International Conference on Machine Learning (ICML’06)*, pages 281–288, 2006.
- [FL92] U. Feige and L. Lovász. Two-prover one-round proof systems, their power and their problems. In *STOC’92: Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pages 733–744, New York, 1992. ACM.
- [GLS93] M Gotschel, L Lovász, and A Schrijver. Geometric algorithms and combinatorial optimization, 1993.
- [GM15] Rong Ge and Tengyu Ma. Decomposing overcomplete 3rd order tensors using sum-of-squares algorithms. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, page 829, 2015.
- [Gro53] Alexander Grothendieck. Résumé de la théorie métrique des produits tensoriels topologiques. *Bol. Soc. Mat. São Paulo* 8, 1953.
- [GRSW16] Venkatesan Guruswami, Prasad Raghavendra, Rishi Saket, and Yi Wu. Bypassing UGC from some optimal geometric inapproximability results. *ACM Transactions on Algorithms (TALG)*, 12(1):6, 2016. Conference version in SODA ’12.
- [GTJ83] D. J. H. Garling and N. Tomczak-Jaegermann. The cotype and uniform convexity of unitary ideals. *Israel J. Math.*, 45(2-3):175–197, 1983.
- [Haa85] Uffe Haagerup. The grothendieck inequality for bilinear forms on  $C^*$ -algebras. *Advances in Mathematics*, 56(2):93–116, 1985.
- [HH06] Eran Halperin and Elad Hazan. HAPLOFREQ—estimating haplotype frequencies efficiently. *J. Comput. Biol.*, 13(2):481–500 (electronic), 2006.
- [HSS15] Samuel B Hopkins, Jonathan Shi, and David Steurer. Tensor principal component analysis via sum-of-square proofs. In *Proceedings of The 28th Conference on Learning Theory*, pages 956–1006, 2015.
- [HV16] Steven Heilman and Thomas Vidick. A moment majorization principle for random matrix ensembles with applications to hardness of the noncommutative grothendieck problem. Manuscript available at <https://arxiv.org/abs/1603.05620>, 2016.
- [Kat98] Eleni Katirtzoglou. Type and cotype of musielak–orlicz sequence spaces. *Journal of mathematical analysis and applications*, 226(2):431–455, 1998.
- [KN12] Subhash Khot and Assaf Naor. Grothendieck-type inequalities in combinatorial optimization. *Communications on Pure and Applied Mathematics*, 65(7):992–1035, 2012.
- [KNS08] Guy Kindler, Assaf Naor, and Gideon Schechtman. The UGC hardness threshold of the  $l_p$  Grothendieck problem. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 64–73. ACM, New York, 2008.
- [KO09] Subhash Khot and Ryan O’Donnell. SDP gaps and UGC-hardness for Max-Cut-Gain. *Theory OF Computing*, 5:83–117, 2009.
- [Kri73] JL Krivine. Théorèmes de factorisation dans les espaces réticulés. *Séminaire Analyse fonctionnelle (dit “Maurey-Schwartz”)*, pages 1–22, 1973.
- [LNRW19] Jerry Li, Aleksandar Nikolov, Ilya Razenshteyn, and Erik Waingarten. On mean estimation for general norms with statistical queries. *arXiv preprint arXiv:1902.02459*, 2019.
- [LP68] Joram Lindenstrauss and Aleksander Pelczyński. Absolutely summing operators in  $l_p$ -spaces and their applications. *Studia Mathematica*, 29(3):275–326, 1968.
- [LPP91] Françoise Lust-Piquard and Gilles Pisier. Noncommutative Khintchine and Paley inequalities. *Ark. Mat.*, 29(2):241–260, 1991.
- [LRV13] Anand Louis, Prasad Raghavendra, and Santosh Vempala. The complexity of approximating vertex expansion. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 360–369. IEEE, 2013.
- [Mau03] Bernard Maurey. Type, cotype and  $K$ -convexity. In *Handbook of the geometry of Banach spaces, Vol. 2*, pages 1299–1332. North-Holland, Amsterdam, 2003.
- [MP76] Bernard Maurey and Gilles Pisier. Séries de variables aléatoires vectorielles indépendantes et propriétés géométriques des espaces de Banach. *Studia Math.*, 58(1):45–90, 1976.
- [MR14] Andrea Montanari and Emile Richard. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems*, pages 2897–2905, 2014.
- [MS86] Vitali D. Milman and Gideon Schechtman. *Asymptotic theory of finite-dimensional normed spaces*, volume 1200 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1986. With an appendix by M. Gromov.
- [Nes98] Yurii Nesterov. Semidefinite relaxation and nonconvex quadratic optimization. *Optimization methods and software*, 9(1-3):141–160, 1998.
- [NRV13] Assaf Naor, Oded Regev, and Thomas Vidick. Efficient rounding for the noncommutative grothendieck inequality. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 71–80. ACM, 2013.
- [NS09] Assaf Naor and Gideon Schechtman. An approximation scheme for quadratic form maximization on convex bodies. 2009.
- [Pis78] Gilles Pisier. Grothendieck’s theorem for noncommutative  $C^*$ -algebras, with an appendix on grothendieck’s constants. *Journal of Functional Analysis*, 29(3):397–415, 1978.
- [Pis80] Gilles Pisier. Un théorème sur les opérateurs linéaires entre espaces de Banach qui se factorisent par un espace de Hilbert. *Ann. Sci. École Norm. Sup. (4)*, 13(1):23–43, 1980.
- [Pis86] Gilles Pisier. *Factorization of linear operators and geometry of Banach spaces*. American Mathematical Soc., 1986.
- [Pis12] Gilles Pisier. Grothendieck’s theorem, past and present. *Bulletin of the American Mathematical Society*, 49(2):237–323, 2012.
- [RRS16] Prasad Raghavendra, Satish Rao, and Tselil Schramm. Strongly refuting random csps below the spectral threshold. *arXiv preprint arXiv:1605.00058*, 2016.
- [RS09] Prasad Raghavendra and David Steurer. Towards computing the Grothendieck constant. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 525–534. Society for Industrial and Applied Mathematics, 2009.
- [RS10] Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *Proceedings of the 42nd annual ACM Symposium on Theory of Computing*, pages 755–764. ACM, 2010.
- [Rud73] Walter Rudin. *Functional analysis*. McGraw-Hill Book Co., New York-Düsseldorf-Johannesburg, 1973. McGraw-Hill Series in Higher Mathematics.
- [RV15] Oded Regev and Thomas Vidick. Quantum xor games. *ACM Transactions on Computation Theory (ToCT)*, 7(4):15, 2015.

[SWY<sup>+</sup>19] Zhao Song, Ruosong Wang, Lin Yang, Hongyang Zhang, and Peilin Zhong. Efficient symmetric norm regression via linear sketching. In *Advances in Neural Information Processing Systems*, pages 830–840, 2019.

[SWZ19] Zhao Song, David Woodruff, and Peilin Zhong. Towards a zero-one law for column subset selection. In *Advances in Neural Information Processing Systems*, pages 6123–6134, 2019.