

A Framework for Quadratic Form Maximization over Convex Sets through Non-Convex Relaxations

Vijay Bhattiprolu

Euiwoong Lee

Assaf Naor

vspvijay@math.ias.edu
IAS/Princeton

euiwoong@umich.edu
University of Michigan

naor@math.princeton.edu
Princeton University

September 25, 2021

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 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.

Contents

1	Introduction	1
1.1	Notation and Preliminaries	2
1.2	A Generic Framework	4
1.3	Examples of Applications	11
1.3.1	Closure Properties	11
1.3.2	Symmetric Norms	12
1.3.3	Unitarily Invariant Matrix Norms	13
1.3.4	Robust Principle Component Analysis	14
1.4	Brief Summary of the Literature and Problems Captured by Quadratic Maximization	15
2	Detailed Preliminaries	16
2.1	Vectors and Matrices	16
2.2	Norms	17
2.3	Polar Operations	18
2.4	Quadratic Maximization and Related Optimization Problems	19
2.5	Projective Tensor Norm and Related Measures	21
2.5.1	Polars of Various Sets of Bounded Forms	21
2.5.2	Covariance Regions and their Connection to Projective Norm and Related Measures	23
2.5.3	Verification of Balance	25
2.6	Oracle Algorithms	27
3	Approximate Convex Optimization	27
3.1	Convex Optimization with an Approximate Separation Oracle	27
3.2	Duality of Approximation Algorithms for Linear Function Optimization	33
4	A Generic Framework: Algorithms from Covariance Separation Oracles, and Reductions across Quadratic/Bilinear/PSD Maximization	35
4.1	Approximation Algorithms from Covariance Separation Oracles	36
4.1.1	Quadratic Maximization via Upper Covariance Separation Oracles	36
4.1.2	PSD Quadratic Maximization via Lower Covariance Separation Oracles	37
4.2	Reductions Across Quadratic/Bilinear/PSD Maximization	39
4.2.1	Reducing Quadratic to PSD under Bounded Type-2	40
4.2.2	Reducing Bilinear to PSD under Bounded Dual Cotype-2	41
4.2.3	A Generic Framework for Quadratic/Bilinear Maximization	44
4.2.4	Type-2 Equivalence of Quadratic Maximization/Upper Covariance Separation	45
4.2.5	Cotype-2 Equivalence of PSD Maximization/Lower Covariance Separation	45
4.2.6	Reducing PSD to Bilinear under Bounded Dual Cotype-2	46
5	Factorization through ℓ_2 via Gaussian Rounding + Convex Programming Duality	47

5.1	Warmup: Sign-Invariant Norms	48
5.1.1	Grothendieck's Inequality/Factorization as a Motivating Example	48
5.1.2	Generalizations to 2-Convex Norms.	49
5.2	Factorization Through ℓ_2 without Lattice Assumptions	51
5.2.1	Convex Programming Formulation of $\gamma_2(\cdot)$	52
5.2.2	An Alternate Proof of a Dual Characterization of $\gamma_2(A)$	53
5.3	Factorization Theorem for Quadratic Maximization under Bounded Type-2	55
6	Algorithmic Closure Properties of Quadratic/Bilinear Maximization	56
6.1	Minkowski Sums	57
6.2	Intersection	58
6.3	Quotients	59
6.4	Complex Interpolation	60
6.4.1	Interpolation Preliminaries	60
6.4.2	Approximation Algorithms for Interpolants of Type-2 Norms	61
6.4.3	Proof of Proposition 6.10	61
6.5	Sections (Subspaces)	63
7	Unconditional Algorithms for Special Families	63
7.1	Approximation Algorithms for Sign Invariant Norms	64
7.1.1	p -convexity and q -concavity Preliminaries	64
7.1.2	Approximation Algorithms for Maximization over Exactly 2-Convex Norms	65
7.2	Approximation Algorithms for Symmetric Norms	66
7.2.1	Preliminaries: Linear Optimization over Symmetric Upward/Downward-closed Sets	66
7.2.2	Symmetric Type-2 Norms	69
7.3	Approximation Algorithms for Unitarily Invariant Norms	70
7.3.1	Preliminaries: Non-Commutative Khintchine Inequality	70
7.3.2	Unitarily Invariant Type-2 Matrix Norms	72
8	Hardness in the Absence of Type-2	74
8.1	SSE-Hardness of ℓ_p -Quadratic Maximization when $p < 2$	74
8.2	SSE-Hardness of Approximation when Type-2 Fails	76
8.2.1	Preliminaries: Embedding Copies of ℓ_p^k in X	76
8.2.2	The Final Reduction	78
9	Oracle Lower Bound for General Type-2 Norms	80
9.1	The Construction	80
9.2	Oracle Lower Bound	80

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

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

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

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

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).

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

$$\mathbf{Q}_K^{\max}(A) = \max_{x \in K} \|A^{\frac{1}{2}} x\|_{\ell_2^n}^2 = \max_{\substack{x \in K \\ y \in \text{Ball}(\ell_2^n)}} \langle A^{\frac{1}{2}} x, y \rangle^2 = \left(\mathbf{Op}_{K, \text{Ball}(\ell_2^n)}^{\max}(A^{\frac{1}{2}}) \right)^2,$$

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 ℓ_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.

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

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, it will be convenient to 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 $r\text{Ball}(X)$, 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 ± 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

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

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

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

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 evaluations¹ 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].

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

¹In 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 $P \lesssim Q$ and $Q \gtrsim P$ mean that $P \leq KQ$ for a universal constant $K > 0$. The notation $P \asymp Q$ stands for $(P \lesssim Q) \wedge (Q \lesssim P)$.

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

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

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 $\text{PSID}^n \subseteq M_n(\mathbb{R})$. For $\mathbb{W} \in M_n(\mathbb{R})$ we will use the notation $\mathbb{W} \succcurlyeq 0$ to indicate that $\mathbb{W} \in \text{PSID}^n$. We associate to a normed space $X = (\mathbb{R}^n, \|\cdot\|_X)$ the following subset $\mathcal{U}(X)$ of PSID^n that we call the *upper covariance body* of X .

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

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

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

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

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

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

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

where, given a subset S of some \mathbb{R}^d , we denote the convex hull of S by $\mathbf{conv}(S)$. 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

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

Hence,

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

Therefore $T_2(X)^{-2} \sum_{j=1}^k \sum_{i=1}^{m(j)} (\sqrt{s_j} w_{ij}) (\sqrt{s_j} w_{ij})^* = T_2(X)^{-2} \sum_{j=1}^k s_j \mathbb{W}_j$ 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 α -approximately convex if $\mathbf{conv}(S) \subseteq \alpha S$.

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 (from Section 9) 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. $\mathbb{P}_{n,\varepsilon}[T_2(X) \lesssim 1] = 1$.
2. $\mathbb{P}_{n,\varepsilon}[S \cap \text{Ball}(X) = S \cap \text{Ball}(\ell_2^n)] \geq p_n$ for every $S \subseteq \mathbb{R}^n$ with $|S| \leq \exp(n^\varepsilon)$.
3. $\mathbb{P}_{n,\varepsilon}[Q_X^{\max}(I_n) \gtrsim n^{1-\varepsilon}] \geq p_n$, 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 $T_2(X) \lesssim 1$ and with high probability $Q_X^{\max}(I_n) \gtrsim n^{1-\varepsilon}$. 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 ℓ_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 [Section 8](#).

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 $\text{NP} \not\subseteq \text{P}_{/\text{poly}}$, 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 1.4. *The Small Set Expansion Hypothesis (SSEH) is a commonly used hardness assumption that was formulated in [RS10] and is recalled in [Section 8](#). Of course, the SSEH is less standard than, say, $\text{NP} \not\subseteq \text{P}_{/\text{poly}}$, 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 1.5. *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 ℓ_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 [Section 8](#) 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 ℓ_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⁺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 1.6. *Fix $\alpha \geq 1$. Let $S \subseteq \mathbb{R}^n$ be star shaped with respect to the origin and α -approximately convex. An α -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 $x \notin \text{conv}(S)$.

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 α -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 Section 3).

Theorem 1.7 (Approximate Ellipsoid Method).

Fix $\alpha \geq 1$ and $R \geq r > 0$. Suppose that $S \subseteq \mathbb{PSD}^n$ is star shaped with respect to the origin, α -approximately convex, and has an α -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 $W \in S$ that satisfies

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

For the sake of the discussion within the introduction, 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 1.7. Another simplifying assumption that we will make throughout this introduction is that the length of the bit description of all inputs (namely matrices) to algorithms is $n^{O(1)}$.

Using Theorem 1.7 and applying it to $\mathcal{U}(X)$, we readily deduce the following approximation algorithm for quadratic maximization (see Section 4.1.1)

Proposition 1.8 (Quadratic Maximization Given Separation Oracle for Upper Covariance Body).

Given access to an α -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 $Q_X^{\max}(A)$.

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 norm ℓ_4^n / ℓ_5^m , while on the other hand it is unclear how to directly describe an upper covariance separation oracle in these cases (see Section 6 for more details). Another advantage which will become apparent in 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 1.9. *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 Section 5 for a detailed introduction to factorization and the relevant theorems we use) as a black box. For the full proof in Section 4, 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 thus caution the reader that the full proof in Section 4 is syntactically different from the ensuing overview. In Section 5 we discuss how the results in Section 4 may be viewed as “dual transpositions” of algorithmic factorization theorems.*

Lower Covariance Region. We define the lower covariance region as follows:

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

where the second inequality in (15) is justified the same way as (11). Note that because $\mathcal{L}(X)$ is equal to $\text{PSID}^n \setminus \{W \in \text{PSID}^n : \mathbb{E}[\|W^{\frac{1}{2}}g\|_{X^*}^2] < 1\}$, the lower covariance region of X is the complement in PSID^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 $s\mathcal{L}(X) \supseteq \mathcal{L}(X)$ 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 1.10. *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 α -approximately convex if $\mathbf{conv}(T) \subseteq \frac{1}{\alpha}T$.*

With this definition at hand, the natural adaptation of Definition 1.6 is as follows.

Definition 1.11. *Fix $\alpha \geq 1$. Suppose that $T \subseteq \mathbb{R}^n$ satisfies $[1, \infty)T \subseteq T$ and that T is α -approximately convex. An α -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 α -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 1.7](#). 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 [Section 4](#) for a complete treatment.

Theorem 1.12 (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 α -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 $W \in \text{IPSD}^n$ with $W \succcurlyeq A$ that satisfies

$$\inf \{Q_X^{\max}(\mathbb{M}) : \mathbb{M} \in \text{IPSD}^n \text{ and } \mathbb{M} \succcurlyeq A\} \gtrsim \frac{Q_X^{\max}(W)}{\alpha}. \quad (17)$$

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 α -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 $W \in \text{IPSD}^n, V \in \text{IPSD}^m$ with $\begin{pmatrix} W & 0 \\ 0 & V \end{pmatrix} \succcurlyeq \begin{pmatrix} 0 & B \\ B^* & 0 \end{pmatrix}$ and

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

We will next explain the ingredients that go into (17); the justification of (18) is similar and will be carried out separately in [Section 4](#). 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 $\{\mathbb{M} \in \text{IPSD}^n : \mathbb{M} \succcurlyeq A\}$. In [Section 4](#) 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 $\{\{\mathbb{M} \in \text{IPSD}^n : Q_X^{\max}(\mathbb{M}) \leq t\} : t \in \mathbb{R}\}$ has a $(1 + o(1))\alpha$ -approximate separation oracle. By homogeneity, we therefore need to show that under the assumptions of [Theorem 1.12](#), the convex set $\{\mathbb{M} \in \text{IPSD}^n : Q_X^{\max}(\mathbb{M}) \leq 1\}$ has a $(1 + o(1))\alpha$ -approximate separation oracle.

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

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

We claim that one can find in polynomial time and with polynomially many oracle calls a matrix $\mathbb{V} \in \text{IPSD}^n$ that attains this maximum up to a factor of $(1 + o(1))\alpha$. Indeed, in [Section 4](#) we will show that for this it suffices to check that each of the corresponding super-level sets

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

has an α -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 1.12](#) ensures that the desired oracle exists. Therefore, we can find $\mathbb{V} \in \text{IPSD}^n$ 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 $\{\mathbb{M}' \in \text{PSID}^n : Q_X^{\max}(\mathbb{M}') \leq 1\}$ 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 \mathbf{g} rather than only one such sample; the details appear in Section 4.

With the algorithmic groundwork of [Theorem 1.12](#) complete, our final algorithm relies on the analytic inequalities that are contained in the following theorem (see [Section 5](#) for proofs).

Theorem 1.13 (Factorization Inequalities).

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

$$Q_X^{\max}(A) \leq \inf \{Q_X^{\max}(W) : W \in \text{PSID}^n \text{ and } W \succcurlyeq A\} \leq T_2(X)^2 \cdot Q_X^{\max}(A). \quad (21)$$

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

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

Then,

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

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 γ_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 Section 5.2.1. 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 [Section 5.3](#)) and it would be interesting to study its ramifications within functional analysis.

By combining [Theorem 1.12](#) with [Theorem 1.13](#), we get the following algorithmic result.

Theorem 1.14 (Generic Framework).

Suppose that $X = (\mathbb{R}^n, \|\cdot\|_X)$ is a normed space such that $\mathcal{L}(X)$ has an α -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 that is polynomial in n , and outputs a number Alg_1 that is guaranteed to satisfy

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

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 α -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 Alg_2 that is guaranteed to satisfy

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

Remark 1.15. One often wishes not only to approximate efficiently 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 Section 4.2.2. 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 1.14.

1.3 Examples of Applications

Theorem 1.14 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.

1.3.1 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 complex² interpolation spaces $\{[X, Y]_\theta\}_{\theta \in [0,1]}$ whose definition is recalled in Section 6.4 (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.

²The 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.

In [Section 6](#), we use the above framework to show that the class of normed spaces X with $T_2(X) = O(1)$ for which there exists a polynomial time $O(1)$ -approximation algorithm for $Q_X^{\max}(A)$ 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 1.14](#). An analogous treatment of the bilinear case is carried out in [Section 6](#) using the second part of [Theorem 1.14](#), 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.

1.3.2 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 π of $\{1, \dots, n\}$, and any choice of signs $\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}$.³ 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⁺17](#), [BBC⁺17](#), [ALS⁺18](#), [SWZ19](#), [SWY⁺19](#)].

In [Section 7.2](#), we use [Theorem 1.14](#) 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 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 ℓ_φ^n is defined by setting for every $x \in \mathbb{R}^n$,

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

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

³One 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.

- (a) There are constants $K, \delta, c > 0$ such that for all $t > 0$, if $\varphi(t) \leq \delta$, then $\varphi(2t) \leq K\varphi(t) + c$.
- (b) There is $\psi : [0, \infty) \rightarrow [0, \infty)$ such that $t \mapsto \psi(\sqrt{t})$ is convex and φ is equivalent to ψ 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 $\varphi(K_1 t) \leq \psi(t) + c_1$ and $\varphi(t) \leq \delta_2$ implies $\psi(K_2 t) \leq \varphi(t) + c_2$ 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, α) 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 ℓ_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.

1.3.3 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, $\|A\|_{S_\infty} = \|A\|_{\ell_2(\mathbb{C}) \rightarrow \ell_2(\mathbb{C})}$ 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 in 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 $\|\cdot\|_{S_E}$ has $O(1)$ type 2 or cotype 2 constant if and only if $\|\cdot\|_E$ does.

In Section 7.3 we use Theorem 1.14 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.

1.3.4 Robust Principle Component Analysis

In [NRV13], efficient bilinear maximization over the operator norm (Schatten- ∞) 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 set of extreme points of $\text{Ball}(S_\infty)$ is 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 1.14](#) may be combined with the lower covariance separation oracles constructed in [Section 7](#) 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 1.14](#) with the separation oracles constructed in [Section 7](#) and using algorithmic closure properties for complex interpolation ([Proposition 6.10](#)), 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

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

As defined above, X_1 is a semi-norm but can be made into a norm by adding a sufficiently small multiple of ℓ_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⁺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 Section 7.1.1 for definitions). Hereafter, we shall refer to them as exactly 2-convex norms. Note in particular that the above class includes the norm ℓ_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 Section 7.1.1 for 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- ∞ 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- ∞ . 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 ℓ_p case captures the problem of certifying hypercontractivity which in turn has connections to small set expansion and quantum separability ([BBH⁺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 poincare 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.

The versatility of quadratic maximization is evident from being able to For appropriate choices of linear maps and convex sets, quadratic maximization also captures (upto constants) the maximization (in 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⁺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.

2 Detailed Preliminaries

2.1 Vectors and Matrices

All vectors we consider are finite dimensional and real valued (with the exception of [Section 6.4](#) where we consider complex entries). We denote the elementary basis of \mathbb{R}^n by $e_1, e_2 \dots e_n$. Vectors are denoted by lower-case letters (x, y, z, v, \dots) and we often refer to sequences of vectors $x_1, x_2, \dots \in \mathbb{R}^n$ indexed by subscripts. We also use subscripts to denote entries of a vector (e.g. for $x \in \mathbb{R}^n$, x_i denotes its i -th entry), but the distinction between entries of a single vector and a sequence of vectors will be clear in context. For example, $(x_i)_j$ denotes the j -th entry of vector x_i .

Matrices are always finite dimensional and are denoted by upper-case letters (A, B, C, M, W, \dots). For a field $F \in \{\mathbb{R}, \mathbb{C}\}$, let $M_n(F)$ (resp. $M_{n,m}(F)$) be the set of all $n \times n$ (resp. $n \times m$) matrices whose entries are from F . We use PSID^n to denote the set of symmetric positive semidefinite (henceforth PSD) matrices in $M_n(\mathbb{R})$. We (mostly) use upper case blackboard-bold letters ($\mathbb{X}, \mathbb{Y}, \mathbb{W}, \dots$) to denote matrices that are the indeterminates in a convex program. Given a vector $v \in \mathbb{R}^n$, let $\text{Diag}(v) \in M_n(\mathbb{R})$ be such that $(\text{Diag}(v))_{i,j} = v_i$ if $i = j$ and 0 otherwise. Similarly, given a matrix $A \in M_n(\mathbb{R})$, $\text{diag}(A)$ is an n -dimensional vector defined as $(\text{diag}(A))_i \stackrel{\text{def}}{=} A_{i,i}$. For a vector or matrix, let $\text{bit}(\cdot)$ denote its bit complexity, which is the number of bits used to represent it.

2.2 Norms

In this paper we consider exclusively finite dimensional norms denoted as $\|\cdot\|_X$. Throughout, the underlying vector space is finite dimensional euclidean space \mathbb{R}^n , with the exception of [Section 6.4](#) where we discuss complex interpolation and hence the underlying space is \mathbb{C}^n .

We denote the unit ball of a norm $\|\cdot\|_X$ over \mathbb{R}^n by $\text{Ball}(X) \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid \|x\|_X \leq 1\}$. $\text{Ball}(X)$ is an origin-symmetric (i.e., $\text{Ball}(X) = -\text{Ball}(X)$) convex body. Recall that every origin symmetric convex body $C \subseteq \mathbb{R}^n$ (i.e., $C = -C$) can be realized as the unit ball of a norm defined as $\|x\| \stackrel{\text{def}}{=} \inf_{\lambda > 0} \{x/\lambda \in C\}$ ($\|\cdot\|$ is called the Minkowski functional of C). Thus there is a one-to-one correspondence between norms and origin symmetric convex bodies.

In this work we assume norms are given as input to algorithms via membership oracles, i.e., an oracle that takes $x \in \mathbb{R}^n$ as input and returns “Inside” if $x \in \text{Ball}(X)$ and “Outside” otherwise. Equivalently (using binary search) one can assume there is an oracle returning $\|x\|_X$ given input $x \in \mathbb{R}^n$.

Recall the dual norm $\|\cdot\|_{X^*}$ over \mathbb{R}^n is given by

$$\|\zeta\|_{X^*} \stackrel{\text{def}}{=} \max_{x \in \text{Ball}(X)} \langle \zeta, x \rangle$$

and that $\text{Ball}(X^*) = \text{Ball}(X)^\circ$ where the polar of a set B is defined as $\{\zeta \mid \langle x, \zeta \rangle \leq 1 \forall x \in B\}$. By the bipolar theorem ([Theorem 2.4](#)) we also conclude $\text{Ball}(X) = \text{Ball}(X^*)^\circ$ which implies $\|\cdot\|_{(X^*)^*} = \|\cdot\|_X$.

We will repeatedly use the following inequality due to Kahane [[Kah64](#)].

Theorem 2.1 (Kahane-Khintchine Inequality). *Consider any $1 \leq p < \infty$. Then there are constants*

$C_p, C'_p \lesssim \sqrt{p}$, such that for every finite sequence of vectors $\{x_i\}$ in a normed space X ,

$$\frac{1}{C'_p} \leq \frac{\mathbb{E} [\|\sum_i g_i x_i\|_X^p]^{1/p}}{\mathbb{E} [\|\sum_i g_i x_i\|_X]} \leq C_p.$$

Type-2/cotype-2 constants of a norm are powerful classification tools from Banach space theory and informally speaking, provide upper/lower bounds on the expected deviation (measured according to the ambient norm) of a random walk from the origin. Formally,

Definition 2.2 (Type-2/Cotype-2). *The (Gaussian) type-2 constant of a normed space X , denoted by $\tilde{T}_2(X)$, is the smallest constant C such that for every finite sequence of vectors $\{x_i\}$ in X ,*

$$\mathbb{E} [\|\sum_i g_i \cdot x_i\|_X^2] \leq C^2 \cdot \sum_i \|x_i\|_X^2$$

where each g_i is an independent standard Gaussian.

The (Gaussian) cotype-2 constant of a normed space X , denoted by $\tilde{C}_2(X)$, is the smallest constant C such that for every finite sequence of vectors $\{x_i\}$ in X ,

$$\mathbb{E} [\|\sum_i g_i \cdot x_i\|_X^2] \geq \frac{1}{C^2} \cdot \sum_i \|x_i\|_X^2$$

$\tilde{C}_2(X^*)$ and $\tilde{T}_2(X)$ are closely related; in particular $\tilde{C}_2(X^*) \leq \tilde{T}_2(X)$ and if X is n -dimensional then $\tilde{T}_2(X) \lesssim \tilde{C}_2(X^*) \cdot \log n$.

The Gaussian and rademacher type-2 (resp. cotype-2) constants are equivalent within universal constants. More specifically one has $\tilde{T}_2(X) \leq T_2(X) \lesssim \tilde{T}_2(X)$ and $\tilde{C}_2(X) \leq C_2(X) \lesssim \tilde{C}_2(X)$ (see [MP76]). In the sequel we work mostly with the Gaussian type-2/cotype-2 constants since in certain places this allows us to provide slightly more precise estimates.

2.3 Polar Operations

In this paper, we use three different notions of polars, and use them to derive equivalences between approximate optimization and approximate separation over various sets. Here we introduce the three notions, beginning with the usual notion of polar.

Definition 2.3 (Standard Polar in \mathbb{R}^n). *Let $B \subseteq \mathbb{R}^n$. The polar of B is $B^\circ \stackrel{\text{def}}{=} \{\xi \in \mathbb{R}^n \mid \langle \xi, x \rangle \leq 1 \ \forall x \in B\}$. Note that B° is always convex and moreover $\mathbf{conv}(B)^\circ = B^\circ$.*

This notion has the following nice duality statement.

Theorem 2.4 (Standard Bipolar Theorem). *Let $B \subseteq \mathbb{R}^n$ be a origin-symmetric convex body. Then we have $(B^\circ)^\circ = B$.*

The other two notions of polars are defined with respect to a self-dual cone $K \subseteq \mathbb{R}^n$. In this paper we will consider either the non-negative orthant or the positive semidefinite cone. We first define the notions of upward-closedness and downward-closedness. Given $B \subseteq K$, let $\uparrow B \stackrel{\text{def}}{=} (B + K) \cap K$ and $\downarrow B \stackrel{\text{def}}{=} (B - K) \cap K$ denote upward and downward closures respectively, where $+$ and $-$ denote Minkowski addition and subtraction respectively. B is said to be *upward-closed* when $\uparrow B = B$ and *downward-closed* when $\downarrow B = B$.

We now define a conic version of the polar operation for a subset $B \subseteq K$.

Definition 2.5 (Conic Polars). Let $K \subseteq \mathbb{R}^n$ be a self-dual cone. We define the polar of a set $B \subseteq K$ (with respect to K) as $B^{\circ K} \stackrel{\text{def}}{=} \{\xi \in K \mid \langle \xi, x \rangle \leq 1 \ \forall x \in B\}$. Note that $B^{\circ K}$ is always convex and moreover $\text{conv}(B)^{\circ K} = B^{\circ K}$.

For any $A \subseteq K$, $A^{\circ K}$ is downward-closed; indeed if $y \in A^{\circ K}$ and $z = y - w$ for some $w \in K$, then $\langle z, x \rangle \leq \langle y, x \rangle \leq 1$ for all $x \in A$. We require a conic version of the bipolar theorem. As we could not locate a reference for this, we include a proof here.

Fact 2.6 (Conic Bipolar Theorem). Let $K \subseteq \mathbb{R}^n$ be a self-dual cone. Let $B \subseteq K$ be a bounded, closed and convex set containing the origin. Then $(B^{\circ K})^{\circ K} = \downarrow B$.

Proof. By definition, $B \subseteq (B^{\circ K})^{\circ K}$. Since $A^{\circ K}$ is downward-closed for any A , $\downarrow B \subseteq (B^{\circ K})^{\circ K}$. For the other direction, assume towards contradiction that $x \in (B^{\circ K})^{\circ K} \setminus \downarrow B$. Since $x \in K$, it implies $x \notin B - K$. Since $B - K$ is also convex and closed (this follows since B is closed and bounded and K is closed), by Hahn-Banach separation theorem there exists $y \in \mathbb{R}^n$ such that $\langle x, y \rangle > c$ and $\langle z, y \rangle < c$ for all $z \in B - K$. We claim that $y \in K$; otherwise, there exists $w \in K$ with $\langle y, w \rangle < 0$, and $\langle y, -\alpha w \rangle = -\alpha \langle y, w \rangle > c$ for large enough $\alpha > 0$ with $-\alpha w \in B - K$, leading to contradiction. Since B contains 0 , $c > 0$. Then $y/c \in B^{\circ K}$, so $\langle x, y/c \rangle > 1$ contradicts that $x \in (B^{\circ K})^{\circ K}$. ■

We define a third notion of polar, namely the *inverse polar*, again with respect to a self-dual cone K , but this time it is the set of elements whose inner product with every element of the original set is at least 1.

Definition 2.7 (Inverse Polar). Let $K \subseteq \mathbb{R}^n$ be a self-dual cone. For a set $B \subseteq K$ and $c \in \mathbb{R}$, we define $B_c^{\circ} \stackrel{\text{def}}{=} \{\xi \mid \langle \xi, x \rangle \geq c \ \forall x \in B\}$, and let $B^{\circ} := B_1^{\circ}$. Note that B° is always convex and moreover $\text{conv}(B)^{\circ} = B^{\circ}$.

We will consider B contained in a self-dual cone $K \subseteq \mathbb{R}^n$. By definition, $B \subseteq (B^{\circ})^{\circ}$ for any B . Under the above condition on B , they are indeed equal. We include a proof below.

Fact 2.8 (Inverse Bipolar Theorem). Let B be a closed, convex, and upward-closed set contained in some self-dual cone $K \subseteq \mathbb{R}^n$. Then (1) B° is a closed, convex and upward-closed set contained in K , and (2) $(B^{\circ})^{\circ} = B$.

Proof. Closure follows straightforwardly from the definition.

We next show that B° is also upward-closed. For any $x \in B^{\circ}$, $y \in K$, and $z \in B$, the definition of B° implies $\langle x, z \rangle \geq 1$, and the self-duality of K implies $\langle y, z \rangle \geq 0$. Therefore, $\langle x + y, z \rangle \geq 1$, implying $x + y \in B^{\circ}$.

We also show that $B^{\circ} \subseteq K$. More strongly, we show that the set B_c° for any $c \in \mathbb{R}$ is contained in K . Assume towards contradiction that there exists $y \in B_c^{\circ} \setminus K$. By the self-duality of K , there exists $z \in K$ with $\langle z, y \rangle < 0$. Let $x \in B$. By the upward-closedness of B , $x + \alpha z \in B$ for any $\alpha \geq 0$, but $\langle y, x + \alpha z \rangle = \langle y, x \rangle + \alpha \langle y, z \rangle$ will be strictly less than c for large enough α , achieving the desired contradiction.

Therefore, B° is a convex, and upward-closed set contained in K . The same argument implies convexity and upward-closedness of $(B^{\circ})^{\circ}$.

Finally, we prove $(B^{\circ})^{\circ} = B$. Assume towards contradiction that there exists $x \in (B^{\circ})^{\circ} \setminus B$. Since B is closed and convex, there exist $y \in \mathbb{R}^n$ and $c \in \mathbb{R}$ such that $\langle x, y \rangle < c$ and $\langle z, y \rangle > c$ for all $z \in B$. In particular, $y \in B_c^{\circ}$, which implies that $y \in K$. Since $x \in (B^{\circ})^{\circ} \subseteq K$, $\langle x, y \rangle \geq 0$, so $c > 0$. Then y/c satisfies $\langle z, y/c \rangle > 1$, so $y/c \in B^{\circ}$, but $\langle x, y/c \rangle < 1$, so it contradicts that $x \in (B^{\circ})^{\circ}$. ■

2.4 Quadratic Maximization and Related Optimization Problems

We are primarily interested in approximation algorithms for maximizing a quadratic form over an origin-symmetric convex body. Formally we define,

Definition 2.9 (Quadratic Maximization).

For an $n \times n$ matrix A and a norm $\|\cdot\|_X$ over \mathbb{R}^n , we define quadratic maximization of A over X as

$$Q_X^{\max}(A) \stackrel{\text{def}}{=} \max_{x \in \text{Ball}(X)} \langle x, Ax \rangle. \quad (24)$$

We will also be interested in two special cases of quadratic maximization; namely, quadratic maximization of PSD instances and bilinear form maximization over $C_1 \times C_2$ for origin-symmetric convex bodies C_1, C_2 . We define some notation for the bilinear case:

Definition 2.10 (Bilinear Maximization/Operator Norm).

For an $n \times m$ matrix A and norms $\|\cdot\|_X, \|\cdot\|_Y$ over \mathbb{R}^n and \mathbb{R}^m respectively, we define bilinear maximization of A over (X, Y) , as

$$\text{Op}_{X,Y}^{\max}(A) \stackrel{\text{def}}{=} \max_{\substack{x \in \text{Ball}(X) \\ y \in \text{Ball}(Y)}} \langle x, Ay \rangle. \quad (25)$$

By definition of the dual norm, bilinear maximization can be cast as the operator norm/distortion of A when A is thought of as a linear map from Y to X^* . Formally we have,

$$\begin{aligned} \|A\|_{Y \rightarrow X^*} &\stackrel{\text{def}}{=} \max_{y \in \text{Ball}(Y)} \|Ay\|_{X^*} = \max_{\substack{x \in \text{Ball}(X) \\ y \in \text{Ball}(Y)}} \langle x, Ay \rangle = \max_{\substack{x \in \text{Ball}(X) \\ y \in \text{Ball}(Y)}} \langle A^*x, y \rangle = \|A\|_{X \rightarrow Y^*} \\ \text{thus } \|A\|_{Y \rightarrow X^*} &= \|A\|_{X \rightarrow Y^*} = \text{Op}_{X,Y}^{\max}(A). \end{aligned} \quad (26)$$

Throughout this document, we will switch notation between operator norms and bilinear maximization based on convenience and context.

Bilinear maximization can be reduced to quadratic maximization using the following identity

$$\text{Op}_{X,Y}^{\max}(A) = Q_{X \oplus_{\infty} Y}^{\max}(B) \quad \text{where } B \stackrel{\text{def}}{=} \frac{1}{2} \cdot \begin{bmatrix} 0 & A \\ A^* & 0 \end{bmatrix}$$

where $\|\cdot\|_{X \oplus_{\infty} Y}$ is a norm over \mathbb{R}^{n+m} defined as $\|(x, y)\|_{X \oplus_{\infty} Y} \stackrel{\text{def}}{=} \max\{\|x\|_X, \|y\|_Y\}$ and whose unit ball is given simply by $\text{Ball}(X) \times \text{Ball}(Y)$.

We will repeatedly use the following simple property that relates quadratic maximization of a PSD matrix over X to bilinear maximization over (ℓ_2^n, X) :

$$Q_X^{\max}(B^*B) = \max_{x \in \text{Ball}(X)} \|Bx\|_2^2 = \|B\|_{X \rightarrow \ell_2^n}^2 = \|B^*\|_{\ell_2^n \rightarrow X^*}^2. \quad (27)$$

Intuitively, this is useful because it transforms a search problem over $\text{Ball}(X)$ to a search problem over $\text{Ball}(\ell_2^n)$. Quadratic maximization of a PSD matrix $W \succeq 0$ over X is in fact equal to bilinear maximization of W over (X, X) . Indeed by Cauchy-Schwarz we have

$$\text{Op}_{X,X}^{\max}(B^*B) = \max_{x,y \in \text{Ball}(X)} \langle x, B^*By \rangle = \max_{x,y \in \text{Ball}(X)} \langle Bx, By \rangle \leq \max_{x,y \in \text{Ball}(X)} \|Bx\|_2 \|By\|_2 = \|B\|_{X \rightarrow \ell_2}^2.$$

The inequality in the reverse direction follows by considering the substitution $y = x = x'$ where x' is the vector maximizing $\|B\|_{X \rightarrow 2}$.

We will also consider quadratic maximization over sections of convex sets, i.e., maximization over $V \cap \text{Ball}(X)$ where V is a subspace. Upto a factor of $1 + o(1)$, maximization over $V \cap \text{Ball}(X)$ can be captured by the maximization over $\text{Ball}(X)$ by adding a large negative multiple of the projector to the orthogonal complement V^\perp . This is formalized in [Observation 6.13](#).

The final special case of quadratic maximization we will consider is the minimum factor by which a linear map shrinks a unit vector. (it can be seen as a special case of quadratic maximization by reducing first to subspace quadratic maximization).

Definition 2.11 (Contractivity of a Linear Map).

For an $n \times m$ matrix A and norms $\|\cdot\|_X, \|\cdot\|_Y$ over \mathbb{R}^n and \mathbb{R}^m respectively, we define the $Y \rightarrow X$ contractivity of A as

$$\|A\|_{Y \rightarrow X}^{\min} \stackrel{\text{def}}{=} \inf_{\|y\|_Y=1} \|Ay\|_X = \inf_{y \neq 0} \|Ay\|_X / \|y\|_Y \quad (28)$$

We will be concerned with the case of invertible $n \times n$ matrices A , where one has the following useful connection with operator norm (bilinear maximization):

$$\begin{aligned} \|A\|_{Y \rightarrow X}^{\min} &= \inf_{\|y\|_Y=1} \|Ay\|_X = \inf_{\|Ay\|_X=1} \frac{1}{\|y\|_Y} \\ &= \inf_{\|x\|_X=1} \frac{1}{\|A^{-1}x\|_Y} = \inf_{x \neq 0} \frac{\|x\|_X}{\|A^{-1}x\|_Y} = \frac{1}{\|A^{-1}\|_{X \rightarrow Y}}. \end{aligned} \quad (29)$$

2.5 Projective Tensor Norm and Related Measures

Let $\|\cdot\|_X, \|\cdot\|_Y$ be norms over \mathbb{R}^n and \mathbb{R}^m respectively. As we make extensive use of duality (and algorithmic versions thereof) in this work, it will be useful for us to investigate the polars of the following level sets:

$$\{A \in M_n(\mathbb{R}) \mid Q_X^{\max}(A) \leq 1\}, \{A \in M_{n,m}(\mathbb{R}) \mid \text{Op}_{X,Y}^{\max}(A) \leq 1\}.$$

2.5.1 Polars of Various Sets of Bounded Forms.

Polar of Bounded Bilinear Forms. We begin by discussing the bilinear case where the polar has been extensively studied in the context of the metric theory of tensor products. $\text{Op}_{X,Y}^{\max}(\cdot)$ is easily checked to be a norm on the space $M_n(\mathbb{R}) \equiv \mathbb{R}^n \otimes \mathbb{R}^n$ and is known as the injective tensor norm. While the injective tensor norm is typically denoted as $\|\cdot\|_{X^* \otimes Y^*}$, we choose to use the unorthodox notation of $\text{Op}_{X,Y}^{\max}(\cdot)$ for convenience in later sections.

The dual norm of $\|\cdot\|_{X^* \otimes Y^*}$ is called the projective tensor norm $\|\cdot\|_{X \otimes Y}$ and has the following representation

$$\|\Xi\|_{X \otimes Y} \stackrel{\text{def}}{=} \inf \sum_i \|x_i\|_X \cdot \|y_i\|_Y = \inf (\sum_i \|x_i\|_X^2)^{1/2} \cdot (\sum_i \|y_i\|_Y^2)^{1/2} \quad (30)$$

where the infimum runs over all finitary decompositions $\Xi = \sum_i x_i y_i^*$. The unit ball of the projective norm is given simply by $\text{Ball}(X \otimes Y) \stackrel{\text{def}}{=} \mathbf{conv}(\{xy^* \mid x \in \text{Ball}(X), y \in \text{Ball}(Y)\})$. We include a derivation of the injective/projective duality which arises quite naturally. Observe that

$$\text{Op}_{X,Y}^{\max}(A) = \max_{x \in \text{Ball}(X), y \in \text{Ball}(Y)} \langle x, Ay \rangle = \max_{x \in \text{Ball}(X), y \in \text{Ball}(Y)} \langle A, xy^* \rangle = \max_{\Xi \in \text{Ball}(X \otimes Y)} \langle A, \Xi \rangle.$$

From the final equality we see that the set of bounded bilinear forms $\{A \in M_{n,m}(\mathbb{R}) \mid \text{Op}_{X,Y}^{\max}(A) \leq 1\} = \{A \in M_{n,m}(\mathbb{R}) \mid \max_{\Xi \in \text{Ball}(X \hat{\otimes} Y)} \langle A, \Xi \rangle \leq 1\}$ is precisely $\text{Ball}(X \hat{\otimes} Y)^\circ$. It is then easily checked that $\|\cdot\|_{X \hat{\otimes} Y}$ as defined above is the Minkowski functional of $\text{conv}(\{xy^* \mid x \in \text{Ball}(X), y \in \text{Ball}(Y)\})$. Thus $\text{Op}_{X,Y}^{\max}(\cdot)$ and $\|\cdot\|_{X \hat{\otimes} Y}$ are dual norms and therefore also satisfy

$$\|\Xi\|_{X \hat{\otimes} Y} = \max_{\text{Op}_{X,Y}^{\max}(A) \leq 1} \langle A, \Xi \rangle.$$

Polar of Bounded Quadratic Forms. Let $B_\wedge^{\text{Sym}}(X) \stackrel{\text{def}}{=} \text{conv}(\{xx^* \mid x \in \text{Ball}(X)\})$. Similar to the bilinear case, we observe that $Q_X^{\max}(A) = \max_{\mathbb{X} \in B_\wedge^{\text{Sym}}(X)} \langle A, \mathbb{X} \rangle$ and therefore the set of bounded quadratic forms $\{A \in M_n(\mathbb{R}) \mid Q_X^{\max}(A) \leq 1\}$ is the polar of $B_\wedge^{\text{Sym}}(X)$. Restricted to the PSD cone PSID^n , we may consider the Minkowski functional of $B_\wedge^{\text{Sym}}(X)$ which is given by

$$\wedge_X^{\text{Sym}}(\mathbb{X}) \stackrel{\text{def}}{=} \inf \sum_i \|x_i\|_X^2$$

where the infimum runs over all finitary decompositions $\mathbb{X} = \sum_i x_i x_i^*$. Note also that by definition of $\wedge_X^{\text{Sym}}(\cdot)$, $B_\wedge^{\text{Sym}}(X)$ can be alternatively described as $\{\mathbb{X} \in \text{PSID}^X \mid \wedge_X^{\text{Sym}}(\mathbb{X}) \leq 1\}$ (note that this is not the unit ball of any norm). We caution the reader that even though $\wedge_X^{\text{Sym}}(\cdot)$ is homogeneous and satisfies triangle inequality, it is not a norm as it is defined only within the PSD cone. It may be useful to think of $\wedge_X^{\text{Sym}}(\mathbb{X})$ as a symmetric version of the projective tensor norm $\|\cdot\|_{X \hat{\otimes} X}$.

Conic Polar of Bounded PSD Quadratic Forms. Since $B_\wedge^{\text{Sym}}(X)$ is a subset of the PSD cone which is self-dual, it is possible and beneficial to study its conic polar $B_\wedge^{\text{Sym}}(X) \circ \text{PSID}^n$ which is in fact the set of bounded PSD quadratic forms, i.e.,

$$B_\wedge^{\text{Sym}}(X) \circ \text{PSID}^n = \{W \in \text{PSID}^n \mid Q_X^{\max}(W) \leq 1\}.$$

Thus by the conic version of the bipolar theorem (Fact 2.6) we have $(B_\wedge^{\text{Sym}}(X) \circ \text{PSID}^n) \circ \text{PSID}^n = \downarrow B_\wedge^{\text{Sym}}(X)$. From this we conclude $\downarrow B_\wedge^{\text{Sym}}(X)$ and the set of bounded PSD quadratic forms are (conically) polar to one another:

$$\begin{aligned} (\downarrow B_\wedge^{\text{Sym}}(X)) \circ \text{PSID}^n &= \{W \in \text{PSID}^n \mid Q_X^{\max}(W) \leq 1\} \\ \downarrow B_\wedge^{\text{Sym}}(X) &= \{W \in \text{PSID}^n \mid Q_X^{\max}(W) \leq 1\} \circ \text{PSID}^n. \end{aligned} \quad (31)$$

Restricted to the PSD cone PSID^n , we may describe the Minkowski functional of $\downarrow B_\wedge^{\text{Sym}}(X)$

$$\wedge_X^{\downarrow \text{Sym}}(\mathbb{X}) \stackrel{\text{def}}{=} \inf_{\mathbb{Y} \succeq \mathbb{X}} \wedge_X^{\text{Sym}}(\mathbb{Y})$$

which can be thought of as a symmetric and monotone (in the Loewner ordering) variant of the projective tensor norm $\|\cdot\|_{X \hat{\otimes} X}$.

Inverse Polar of Non-Contractive Linear Maps. Recall that $B^\circ \stackrel{\text{def}}{=} \{y \mid \langle y, x \rangle \geq 1 \ \forall x \in B\}$ and that $\|U\|_{X^* \rightarrow 2}^{\min} \stackrel{\text{def}}{=} \inf_{\|\xi\|_{X^*}=1} \|U\xi\|_2$. In this subsection we wish to describe $\{U^*U \mid \|U\|_{X^* \rightarrow 2}^{\min} \geq 1\}^\circ$. Let $B_\wedge^\uparrow(X) \stackrel{\text{def}}{=} \text{conv}(\{\xi\xi^* \mid \|\xi\|_{X^*} \geq 1\})$. Observe that

$$(\|U\|_{X^* \rightarrow 2}^{\min})^2 = \inf_{\|\xi\|_{X^*} \geq 1} \|U\xi\|_2^2 = \inf_{\|\xi\|_{X^*} \geq 1} \langle U^*U, \xi\xi^* \rangle = \inf_{W \in B_\wedge^\uparrow(X)} \langle U^*U, W \rangle.$$

Thus we have $\{U^*U \mid \|U\|_{X^* \rightarrow 2}^{\min} \geq 1\} = B_\wedge^\uparrow(X)^\circ$. Furthermore since $\{U^*U \mid \|U\|_{X^* \rightarrow 2}^{\min} \geq 1\}$ is a closed, convex and upward-closed (in the Loewner ordering) subset of $\mathbb{P}\text{SD}^n$, we may apply a version of the bipolar theorem for upward-closed sets (Fact 2.8) to obtain

$$B_\wedge^\uparrow(X)^\circ = \{U^*U \in M_n(\mathbb{R}) \mid \|U\|_{X^* \rightarrow 2}^{\min} \geq 1\} \quad (32)$$

and moreover that $B_\wedge^\uparrow(X)$ is closed, convex and upward-closed.

We give a second description of $B_\wedge^\uparrow(X)$. To this end we define for any $W \in \mathbb{P}\text{SD}^n$,

$$\wedge_X^\uparrow(W) \stackrel{\text{def}}{=} \sup \sum_i \|w_i\|_{X^*}^2. \quad (33)$$

where the supremum runs over all finitary decompositions $W = \sum_i w_i w_i^*$. It is then easily checked that $B_\wedge^\uparrow(X) = \{W \succeq 0 \mid \wedge_X^\uparrow(W) \geq 1\}$.

2.5.2 Covariance Regions and their Connection to Projective Norm and Related Measures

We define the upper (resp. lower) covariance region which, as discussed in the introduction allow us to formulate a generic approximate convex optimization approach to quadratic (resp. PSD quadratic) maximization in the presence of type-2 (resp. dual cotype-2).

Definition 2.12 (Gaussian Rounding Function).

For a norm X over \mathbb{R}^n and i.i.d. standard Gaussians $\mathbf{g} = (g_1, \dots, g_n)$ we define the Gaussian Rounding Function $\mathcal{N}_X(\cdot) : \mathbb{P}\text{SD}^n \rightarrow \mathbb{R}_{\geq 0}$ as $\mathcal{N}_X(\mathbb{X}) \stackrel{\text{def}}{=} \mathbb{E} [\|\mathbb{X}^{1/2} \mathbf{g}\|_X^2]$.

Remark 2.13. Since a Gaussian distribution is uniquely determined by its first two moments, we may alternatively define $\mathcal{N}_X(\mathbb{X})$ as $\mathbb{E} [\|\sum_i g_i \cdot x_i\|_X^2]$ where (x_i) is any finite sequence satisfying $\mathbb{X} = \sum_i x_i x_i^*$. In fact $\| (x_i) \| \stackrel{\text{def}}{=} \mathbb{E} [\|\sum_i g_i \cdot x_i\|_X^2]^{1/2}$ defines a norm on the space of sequences of a fixed length.

Definition 2.14 (Upper Covariance Body).

We define the Upper Covariance Body denoted by $\mathcal{U}(X)$ as $\{\mathbb{X} \succeq 0 \mid \mathcal{N}_X(\mathbb{X}) \leq 1\}$.

Definition 2.15 (Lower Covariance Region).

We define the Lower Covariance Region denoted by $\mathcal{L}(X)$ as $\{W \succeq 0 \mid \mathcal{N}_{X^*}(W) \geq 1\}$.

The lower covariance region is the complement of (the interior of) the upper covariance region of the dual, i.e., $\mathcal{L}(X) = \mathbb{P}\text{SD}^n \setminus \text{Int}(\mathcal{U}(X^*))$.

$\mathcal{N}_X(\cdot)$ is known to be non-decreasing (in the Loewner ordering). We include a proof below:

Fact 2.16. $\mathcal{N}_X(\cdot)$ is non-decreasing in the Loewner ordering.

Proof. It suffices to show that for any $\mathbb{X}, \mathbb{Y} \in \mathbb{P}\text{SD}^n$, we have $\mathcal{N}_X(\mathbb{X}) \leq \mathcal{N}_X(\mathbb{X} + \mathbb{Y})$. Let $g_1, \dots, g_n, h_1, \dots, h_n, z_1, \dots, z_n$ be i.i.d. standard Gaussians. Observe that the vector $\mathbb{X}^{1/2} \mathbf{g}$ has covariance \mathbb{X} and that the vectors $\mathbb{X}^{1/2} \mathbf{g} \pm \mathbb{Y}^{1/2} \mathbf{h}$, both have covariance $\mathbb{X} + \mathbb{Y}$. Thus $(\mathbb{X} + \mathbb{Y})^{1/2} \mathbf{z}$ has the same distribution as each of the vectors $\mathbb{X}^{1/2} \mathbf{g} \pm \mathbb{Y}^{1/2} \mathbf{h}$. Therefore we have,

$$\begin{aligned} \mathcal{N}_X(\mathbb{X})^{1/2} &= \mathbb{E}_{\mathbf{g}} [\|\mathbb{X}^{1/2} \mathbf{g}\|_X^2]^{1/2} \\ &\leq \frac{1}{2} \cdot \mathbb{E}_{\mathbf{g}, \mathbf{h}} [\|\mathbb{X}^{1/2} \mathbf{g} + \mathbb{Y}^{1/2} \mathbf{h}\|_X^2]^{1/2} + \frac{1}{2} \cdot \mathbb{E}_{\mathbf{g}, \mathbf{h}} [\|\mathbb{X}^{1/2} \mathbf{g} - \mathbb{Y}^{1/2} \mathbf{h}\|_X^2]^{1/2} \\ &= \mathbb{E}_{\mathbf{z}} [\|(\mathbb{X} + \mathbb{Y})^{1/2} \mathbf{z}\|_X^2]^{1/2} \end{aligned}$$

$$= \mathcal{N}_X(\mathbb{X} + \mathbb{Y})^{1/2}$$

where the first inequality is an application of the triangle inequality in the space $L_2(\mathbb{R}^n, \gamma)$ where γ denotes the n -variate Gaussian measure. ■

The next observation relates the bodies $\mathcal{U}(X)$, $B_\wedge^{\text{Sym}}(X)$, $\downarrow B_\wedge^{\text{Sym}}(X)$, $\text{Ball}(X \widehat{\otimes} X) \cap \text{PSID}^n$.

Observation 2.17 (Relating Projective Measures to Gaussian Rounding Function).

Let X be a finite dimensional normed space. For any $\mathbb{X} \in \text{PSID}^n$, we have,

$$\wedge_X^{\downarrow \text{Sym}}(\mathbb{X}) \leq \|\mathbb{X}\|_{X \widehat{\otimes} X} \leq \wedge_X^{\text{Sym}}(\mathbb{X}) \leq \mathcal{N}_X(\mathbb{X}) \leq \widetilde{T}_2(X)^2 \cdot \wedge_X^{\downarrow \text{Sym}}(\mathbb{X}).$$

This implies the inclusions

$$\begin{aligned} & \widetilde{T}_2(X)^{-2} \cdot \downarrow B_\wedge^{\text{Sym}}(X) \\ & \subseteq \mathcal{U}(X) \\ & \subseteq B_\wedge^{\text{Sym}}(X) = \mathbf{conv}(\{xx^* \mid x \in \text{Ball}(X)\}) \\ & \subseteq \text{Ball}(X \widehat{\otimes} X) \cap \text{PSID}^n = \mathbf{conv}(\{xy^* \mid x, y \in \text{Ball}(X)\}) \cap \text{PSID}^n \\ & \subseteq \downarrow B_\wedge^{\text{Sym}}(X) = \downarrow \mathbf{conv}(\{xx^* \mid x \in \text{Ball}(X)\}). \end{aligned}$$

Proof. Let $\mathbb{X} \succeq 0$. For the first inequality, assume $\|\mathbb{X}\|_{X \widehat{\otimes} X} \leq 1$. Then by definition there exist finite sequences $(u_i), (v_i)$ such that $\mathbb{X} = \sum_i u_i v_i^*$ and $\sum_i \|u_i\|_X^2, \sum_i \|v_i\|_X^2 \leq 1$. Also we have $\sum_i (u_i u_i^* + v_i v_i^*)/2 \succeq \sum_i u_i v_i^*/2 + \sum_i v_i u_i^*/2 = \sum_i u_i v_i^* = \mathbb{X}$, where the second last equality follows by symmetry of \mathbb{X} . By assumption on $(u_i), (v_i)$, we conclude $\wedge_X^{\text{Sym}}(\sum_i (u_i u_i^* + v_i v_i^*)/2) \leq 1$. Thus $\wedge_X^{\downarrow \text{Sym}}(\mathbb{X}) = \inf_{\mathbb{Y} \succeq \mathbb{X}} \wedge_X^{\text{Sym}}(\mathbb{Y}) \leq 1$.

The second inequality is immediate.

For the third inequality we approximate the Gaussian integral $\mathbb{E}[\|\mathbb{X}^{1/2} \mathbf{g}\|_X^2]$ by a sequence of finite sums to obtain a sequence of upper bounds on $\wedge_X^{\text{Sym}}(\mathbb{X})$ that converge to $\mathcal{N}_X(\mathbb{X})$.

For the fourth inequality, consider any $\mathbb{Y} \succeq \mathbb{X}$ and any fixed decomposition $\mathbb{Y} = \sum_i y_i y_i^*$. Now since $\mathcal{N}_X(\cdot)$ is non-decreasing in the Loewner ordering (Fact 2.16), we have

$$\mathcal{N}_X(\mathbb{X}) \leq \mathcal{N}_X(\mathbb{Y}) \leq \widetilde{T}_2(X)^2 \cdot \sum_i \|y_i\|_X^2$$

where the final inequality in the preceding equation follows from the definition of the Gaussian type-2 constant. Taking infimum over all $\mathbb{Y} \succeq \mathbb{X}$ and all decompositions of \mathbb{Y} completes the proof of the claim. ■

Remark 2.18. It is clear from the proof that $\widetilde{T}_2(X)^2$ is the best possible constant in the fourth inequality.

Observation 2.17 allows us to prove approximate convexity of the upper covariance body.

Observation 2.19 (Convex Hull of $\mathcal{U}(X)$ and Approximate Convexity).

$\mathbf{conv}(\mathcal{U}(X)) = B_\wedge^{\text{Sym}}(X) = \mathbf{conv}(\{xx^* \mid x \in \text{Ball}(X)\})$ and moreover $\mathcal{U}(X)$ is $\widetilde{T}_2(X)^2$ -approximately convex.

Proof. By Observation 2.17 we have $\mathcal{U}(X) \subseteq B_\wedge^{\text{Sym}}(X)$ and so $\mathbf{conv}(\mathcal{U}(X)) \subseteq B_\wedge^{\text{Sym}}(X)$. On the other hand $\mathcal{U}(X) \supseteq \{xx^* \mid x \in \text{Ball}(X)\}$ and so $\mathbf{conv}(\mathcal{U}(X)) \supseteq \mathbf{conv}(\{xx^* \mid x \in \text{Ball}(X)\}) = B_\wedge^{\text{Sym}}(X)$. We conclude that $\mathbf{conv}(\mathcal{U}(X)) = B_\wedge^{\text{Sym}}(X)$. This yields the first claim.

By Observation 2.17 again, we have $\mathcal{U}(X) \subseteq \mathbf{conv}(\mathcal{U}(X)) \subseteq \widetilde{T}_2(X)^2 \cdot \mathcal{U}(X)$. This yields the second claim. ■

The next observation relates the regions $\mathcal{L}(X)$, $B_\lambda^\dagger(X)$.

Observation 2.20. For any $\mathbb{W} \in \text{PSID}^n$, we have $\tilde{C}_2(X^*)^{-2} \cdot \wedge_X^\dagger(\mathbb{W}) \leq \mathcal{N}_{X^*}(\mathbb{W}) \leq \wedge_X^\dagger(\mathbb{W})$. Equivalently, $B_\lambda^\dagger(X) \supseteq \mathcal{L}(X) \supseteq \tilde{C}_2(X^*)^2 \cdot B_\lambda^\dagger(X)$.

Proof. For the first inequality we approximate the Gaussian integral $\mathbb{E}[\|\mathbb{W}^{1/2}\mathbf{g}\|_{X^*}^2]$ by a sequence of finite sums to obtain a sequence of lower bounds on $\wedge_X^\dagger(\mathbb{W})$ that converge to $\mathcal{N}_{X^*}(\mathbb{W})$.

For the second inequality, we observe a stronger fact: namely it is true for any fixed decomposition $\mathbb{W} = \sum_i w_i w_i^*$. Indeed $\mathcal{N}_{X^*}(\mathbb{W}) \geq \tilde{C}_2(X^*)^2 \cdot \sum_i \|w_i\|_{X^*}^2$ simply by the definition of the Gaussian cotype-2 constant. The claim follows. ■

Remark 2.21. It is clear from this proof that $\tilde{C}_2(X^*)^2$ is the best possible constant for which the first inequality holds.

Observation 2.20 allows us to prove approximate convexity of the lower covariance region.

Observation 2.22 (Convex Hull of $\mathcal{U}(X)$ and Approximate Convexity).

$\mathbf{conv}(\mathcal{L}(X)) = B_\lambda^\dagger(X) = \mathbf{conv}(\{xx^* \mid \|x\|_{X^*} \geq 1\})$ and moreover $\mathcal{L}(X)$ is $\tilde{C}_2(X^*)^2$ -(inverse) approximately convex.

Proof. By Observation 2.20 we have $\mathcal{L}(X) \subseteq B_\lambda^\dagger(X)$ and so $\mathbf{conv}(\mathcal{U}(X)) \subseteq B_\lambda^\dagger(X)$. On the other hand $\mathcal{U}(X) \supseteq \{xx^* \mid \|x\|_{X^*} \geq 1\}$ and so $\mathbf{conv}(\mathcal{L}(X)) \supseteq \mathbf{conv}(\{xx^* \mid \|x\|_{X^*} \geq 1\}) = B_\lambda^\dagger(X)$. We conclude that $\mathbf{conv}(\mathcal{U}(X)) = B_\lambda^\dagger(X)$. This yields the first claim.

By Observation 2.20 again, we have $\mathcal{L}(X) \subseteq \mathbf{conv}(\mathcal{L}(X)) \subseteq \tilde{C}_2(X^*)^{-2} \cdot \mathcal{L}(X)$. This yields the second claim. ■

Remark 2.23. While we work predominantly over \mathbb{R}^n , in Section 6.4 we consider finite dimensional complex normed spaces. Every definition/claim in this Section 2.5 generalizes verbatim to the complex case where we replace $M_{n,m}(\mathbb{R})$ by $M_{n,m}(\mathbb{C})$ and PSID^n by $\text{PSID}^{\mathbb{C}^n} \stackrel{\text{def}}{=} \mathbf{conv}(\{xx^* \mid x \in \mathbb{C}^n\})$. For $A, B \in \text{PSID}^{\mathbb{C}^n}$, we write $A \succeq B$ if $A - B \in \text{PSID}^{\mathbb{C}^n}$.

2.5.3 Verification of Balance

The purpose of this section is to verify that all of the sets/functions to which we apply the approximate ellipsoid method satisfy certain sanity-conditions that are necessary for multiplicative approximation guarantees. We refer to these conditions as “balance”, as defined below.

We say a norm $\|\cdot\|_X$ is (R, r) -balanced if $\text{Ball}(X)$ contains a euclidean ball of radius r and is contained in a euclidean ball of radius R . We say a set $B \subseteq \text{PSID}^n$ is (R, r, PSID^n) -balanced if $r(\text{PSID}^n \cap \text{Ball}(\ell_2^{n \times n})) \subseteq B \subseteq R(\text{PSID}^n \cap \text{Ball}(\ell_2^{n \times n}))$. We also say that a set $B \subseteq \text{PSID}^n$ is inverse- (R, r, PSID^n) -balanced if $r \cdot \text{Ball}(\ell_2^{n \times n}) \cap B = \emptyset$ and $\{M \in \text{PSID}^n : \|M\|_{\ell_2^{n \times n}} = R\} \subseteq B$.

For a cone K , we say a function $f : K \rightarrow \mathbb{R}$ is (R, r, K) -balanced if (1) satisfies $r \cdot f(x) \leq \|x\|_2 \leq R \cdot f(x)$ for all $x \in K$, and (2) satisfies $\|f(x) - f(y)\| \leq R\|x - y\|_2$ for all $x, y \in K$. We say $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is (R, r) -balanced if it is (R, r, \mathbb{R}^n) -balanced.

Throughout this text we make the assumption in our algorithmic results that $\|\cdot\|_X$ is (R, r) -balanced. In this section we verify that such an assumption on $\|\cdot\|_X$ implies that the associated sets/functions in Section 2.5 also satisfy appropriate versions of balance, which is required in order to optimize over them.

Lemma 2.24. Let $X = (\mathbb{R}^n, \|\cdot\|_X), Y = (\mathbb{R}^m, \|\cdot\|_Y)$ be (R, r) -balanced. Then there exist $R', 1/r' = \text{poly}(n, R, 1/r)$ and $R'', 1/r'' = \text{poly}(n, m, R, 1/r)$ such that

- (1) $\text{Ball}(X \widehat{\otimes} Y), \text{Op}_{X,Y}^{\max}(\cdot)$ and $\text{Ball}(\text{Op}_{X,Y}^{\max}(\cdot))$ are (R'', r'') -balanced.
- (2) $\mathcal{U}(X), B_{\wedge}^{\text{Sym}}(X), \downarrow B_{\wedge}^{\text{Sym}}(X)$ and $\text{Q}_{X,X}^{\max}(\cdot)$ are (R', r', PSID^n) -balanced.
- (3) $\mathcal{L}(X), B_{\wedge}^{\uparrow}(X)$ are inverse- (R', r', PSID^n) -balanced.

Proof. We begin by showing (1). Recall the $\ell_2^m \rightarrow \ell_2^n$ -operator norm can be written as

$$\|M\|_{2 \rightarrow 2} = \max_{\|x\|_2, \|y\|_2 \leq 1} \langle x, My \rangle$$

and combining with the balance condition on X, Y , we obtain

$$r^2 \cdot \|M\|_{2 \rightarrow 2} \leq \text{Op}_{X,Y}^{\max}(M) \leq R^2 \cdot \|M\|_{2 \rightarrow 2}.$$

Now since $\|M\|_{2 \rightarrow 2} \leq \|M\|_{\ell_2^{n \times m}} \leq \max\{\sqrt{n}, \sqrt{m}\} \cdot \|M\|_{2 \rightarrow 2}$, it follows that $\{A \in M_{n,m}(\mathbb{R}) \mid \text{Op}_{X,Y}^{\max}(A) \leq 1\} = \text{Ball}(\text{Op}_{X,Y}^{\max}(\cdot))$ is $(\max\{\sqrt{n}/r^2, \sqrt{m}/r^2\}, 1/R^2)$ -balanced. By polarity, we conclude $\text{Ball}(X \widehat{\otimes} Y)$ is $(R^2, \min\{r^2/\sqrt{n}, r^2/\sqrt{m}\})$ -balanced. To show balance of $\text{Op}_{X,Y}^{\max}(\cdot)$ is balanced (as a function), we are left with checking Lipschitzness. By triangle inequality we have $|\text{Op}_{X,Y}^{\max}(A) - \text{Op}_{X,Y}^{\max}(B)| \leq \text{Op}_{X,Y}^{\max}(A - B) \leq R^2 \cdot \|A - B\|_{\ell_2^{n \times m}}$. Thus $\text{Op}_{X,Y}^{\max}(\cdot)$ is $(\text{poly}(n, R, 1/r), 1/\text{poly}(n, R, 1/r))$ -balanced.

We now show (2). From (1) we conclude that $\text{Ball}(X \widehat{\otimes} X) \cap \text{PSID}^n$ is $(R^2, \min\{r^2/\sqrt{n}, r^2/\sqrt{m}\}, \text{PSID}^n)$ -balanced. Combining with the equivalences in [Observation 2.17](#) and the fact that $\widetilde{T}_2(X) \leq \sqrt{n}$ implies that $\mathcal{U}(X), B_{\wedge}^{\text{Sym}}(X), \downarrow B_{\wedge}^{\text{Sym}}(X)$ are $(\text{poly}(n, R, 1/r), 1/\text{poly}(n, R, 1/r), \text{PSID}^n)$ -balanced. It remains to show Lipschitzness of $\text{Q}_{X,X}^{\max}(\cdot)$ restricted to PSID^n , but this follows from Lipschitzness of $\text{Op}_{X,X}^{\max}(\cdot)$ since for any $W \succeq 0$, we have $\text{Q}_{X,X}^{\max}(W) = \text{Op}_{X,X}^{\max}(W)$. Thus $\text{Q}_{X,X}^{\max}(\cdot)$ is $(\text{poly}(n, R, 1/r), 1/\text{poly}(n, R, 1/r), \text{PSID}^n)$ -balanced.

Finally we establish (3). We first show $B_{\wedge}^{\uparrow}(X)$ is inverse- $(\text{poly}(n, R, 1/r), 1/\text{poly}(n, R, 1/r), \text{PSID}^n)$ -balanced. Then by the equivalence in [Observation 2.20](#) and the fact that $\widetilde{C}_2(X^*) \leq \sqrt{n}$, we also obtain that $\mathcal{L}(X)$ is inverse- $(\text{poly}(n, R, 1/r), 1/\text{poly}(n, R, 1/r), \text{PSID}^n)$ -balanced. To this end, observe that for any $W \succeq 0$ and any finite decomposition $W = \sum_i w_i w_i^*$, we have

$$\text{Tr}(W) = \sum_i \|w_i\|_2^2.$$

From this as well as the balance of X^* (which follows from balance of X) we conclude that $\wedge_X^{\uparrow}(W)$ and $\text{Tr}(W)$ are equivalent within $\text{poly}(R, 1/r)$. Now since $\text{Tr}(W)$ and $\|W\|_{\ell_2^{n \times n}}$ are equivalent within \sqrt{n} (assuming $W \succeq 0$), we conclude that $\wedge_X^{\uparrow}(W)$ and $\|W\|_{\ell_2^{n \times n}}$ are equivalent within $\text{poly}(n, R, 1/r)$ which immediately implies the desired inverse-balance of $B_{\wedge}^{\uparrow}(X)$. This completes the proof. \blacksquare

Remark 2.25. In [Section 3](#), we use an approximate version of the ellipsoid method to give multiplicative approximation algorithms for various tasks ([Proposition 3.11, Proposition 3.12, Theorem 3.14, Theorem 3.16](#)), provided the appropriate balance conditions are satisfied. In all the settings in this text wherein we apply the above theorems, the associated balance conditions are satisfied as a consequence of [Lemma 2.24](#).

2.6 Oracle Algorithms

In this work, an oracle algorithm denoted $\text{ALG}(\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{O}_1, \mathcal{O}_2, \dots)$ running in time $T = T(\text{size}(\mathcal{I}_1) + \text{size}(\mathcal{I}_2) + \dots)$ is an algorithm that takes in a finite set of inputs $\mathcal{I}_1, \mathcal{I}_2, \dots$ and a finite set of oracles $\mathcal{O}_1, \mathcal{O}_2, \dots$, and terminates with an output after T steps, where at any given step the algorithm is allowed to make black-box queries to any of the oracles and this counts as a single step in the runtime. For a fixed oracle algorithm, the set of queries made to the oracles is determined completely by the input $\mathcal{I}_1, \mathcal{I}_2, \dots$ (and in the case of randomized algorithms, the distribution of the set of queries is determined entirely by the input).

We define below some oracles making a repeated appearance in the sequel.

Definition 2.26 (Approximate Optimization Oracles.).

1. For a norm $\|\cdot\|_X$ over \mathbb{R}^n an α -approximate search oracle for quadratic maximization over X is an oracle that on any input $A \in M_n(\mathbb{R})$ outputs a vector $x \in \text{Ball}(X)$ satisfying $\langle x, Ax \rangle \geq Q_X^{\max}(A)/\alpha$.
2. For a norm $\|\cdot\|_X$ over \mathbb{R}^n an α -approximate search oracle for PSD quadratic maximization over X is an oracle that on any PSD input $A \succeq 0$ outputs a vector $x \in \text{Ball}(X)$ satisfying $\langle x, Ax \rangle \geq Q_X^{\max}(A)/\alpha$.
3. For norms $(\mathbb{R}^n, \|\cdot\|_X), (\mathbb{R}^m, \|\cdot\|_Y)$, an α -approximate search oracle for bilinear maximization over X, Y is an oracle that on any input $A \in M_{n,m}(\mathbb{R})$ outputs vectors $x \in \text{Ball}(X), y \in \text{Ball}(Y)$ satisfying $\langle x, Ay \rangle \geq \text{Op}_{X,Y}^{\max}(A)/\alpha$.

3 Approximate Convex Optimization

In this section, we provide optimization tools used in the paper. Intuitively, we show that if a set is “approximately convex” and there is an “approximate separation oracle,” the ellipsoid algorithm can be used to approximately minimize a convex function over the set.

3.1 Convex Optimization with an Approximate Separation Oracle

We start with our definition of approximate convexity. Recall that a set $B \subseteq \mathbb{R}^n$ is called *star-shaped with respect to the origin* when $[0, 1]B \subseteq B$. We say that $B \subseteq \mathbb{R}^n$ is *inverse star-shaped with respect to the origin* when $[1, \infty)B \subseteq B$, which is equivalent to that $\mathbb{R}^n \setminus B$ is star-shaped with respect to the origin.

Definition 3.1 (Approximately Convex Body).

1. Consider a body $B \subseteq \mathbb{R}^n$ star-shaped with respect to the origin. For $\alpha \geq 1$, we shall say B is α -approximately convex if every convex combination of points in B is contained in $\alpha \cdot B$. Equivalently we have the inclusions $B \subseteq \mathbf{conv}(B) \subseteq \alpha \cdot B$.
2. Consider a set $B \subseteq \mathbb{R}^n$ inverse star-shaped with respect to the origin. For $\alpha \geq 1$, we shall say B is α -inverse approximately convex if every convex combination of points in B is contained in B/α . Equivalently we have the inclusions $B \subseteq \mathbf{conv}(B) \subseteq B/\alpha$.

Since $\mathbf{conv}(B)$ is contained in $\alpha \cdot B$ (case 1) or B/α (case 2) in the above definition, the existence of a separating hyperplane for a convex set implies the following.

Definition 3.2 (Approximately Separable Body).

1. Let $B \subseteq \mathbb{R}^n$ be a body star-shaped with respect to the origin. For $\alpha \geq 1$, we shall say B is α -(forward) separable if whenever x is outside $\alpha \cdot B$, there exists a hyperplane separating x from B .
2. Let $B \subseteq \mathbb{R}^n$ be a set inverse star-shaped with respect to the origin. For $\alpha \geq 1$, we shall say B is α -(inversely) separable if whenever x is outside B/α , there exists a hyperplane separating x from B .

Henceforth we drop the term “forward” or “inversely” whenever it is clear which definition is being referred to.

Remark: α -approximate convexity implies α -separability.

Next, we define *approximate separation oracles* that return approximately separating hyperplanes.

Definition 3.3 (Approximate Separation Oracle).

1. For $\alpha \geq 1$, an α -approximate separation oracle for an α -separable body $B \subseteq \mathbb{R}^n$, is an oracle that on any input point $x \in \mathbb{R}^n$, either correctly outputs “Inside” when $x \in \alpha \cdot B$ or outputs a hyperplane separating x from B .

Remark: In any ambiguous case, i.e., when $x \notin B$ and $x \in \alpha \cdot B$, the oracle is allowed to either output a hyperplane separating x from B or output “Inside”.

2. For $\alpha \geq 1$, an α -approximate (inverse) separation oracle for an α -(inversely) separable set $B \subseteq \mathbb{R}^n$, is an oracle that on any input point $x \in \mathbb{R}^n$, either correctly outputs “Inside” when $x \in B/\alpha$ or outputs a hyperplane separating x from B .

Remark: In any ambiguous case, i.e., when $x \notin B$ and $x \in B/\alpha$, the oracle is allowed to either output a hyperplane separating x from B or output “Inside”.

Henceforth we drop the term “inverse” whenever it is clear which definition is being referred to.

The following simple observation will be useful for establishing approximate separation oracles in various contexts.

Observation 3.4 (Oracle from Equivalence). Consider bodies $B_1, B_2 \subseteq \mathbb{R}^n$ satisfying $B_2 \subseteq B_1 \subseteq \alpha \cdot B_2$ and let \mathcal{SO} be a β -approximate separation oracle for B_1 . Then \mathcal{SO} is also an $\alpha\beta$ -approximate separation oracle for B_2 .

Proof. Indeed if $\mathcal{SO}(x)$ returns “Inside”, then $x \in \alpha \cdot B_1 \Rightarrow x \in \alpha\beta \cdot B_2$. On the other hand if $\mathcal{SO}(x)$ returns a hyperplane separating B_1 from x , it of course separates B_2 from x as desired.

The claim for approximate inverse separation oracles follows analogously. ■

A standard simplification step in convex optimization is to use binary search to reduce the problem to the following decision task: given (exact) separation oracles for two convex bodies, decide whether their intersection is empty. The ellipsoid method is a popular method for solving this subroutine. We use a variant of the ellipsoid method to show that given approximate separation oracles (forward and inverse) for a set, we can efficiently decide whether the intersection of two bodies is close to empty.

Proposition 3.5 (Folklore: Approximate Ellipsoid Method).

Let C_1 have an α_1 -approximate separation oracle \mathcal{SO}_1 , C_2 have an α_2 -approximate inverse separation oracle \mathcal{SO}_2 . Additionally, assume that C_1 is contained in a (Euclidean) ball of radius R . There is an algorithm $\text{ALG}(R, \varepsilon, \mathcal{SO}_1, \mathcal{SO}_2)$ running in time $\text{poly}(n, \log R, \log(1/\varepsilon))$ returning one of the following:

1. A point x inside $(\alpha_1 C_1) \cap (C_2/\alpha_2)$.
2. $C_1 \cap C_2$ does not contain a (Euclidean) ball of radius ε .

Proof. We follow the ellipsoid algorithm. Start from an ellipsoid containing the ball of radius R containing $C_1 \cap C_2$. If both oracles return “Inside” when queried on the center x , then we know x is in $(\alpha_1 C_1) \cap (C_2/\alpha_2)$. Otherwise, by definition of the separation oracles, we have for some $i \in [2]$, a hyperplane separating x and C_i . Lemma (3.2.10) of [GLS93] shows that one can compute a new ellipsoid of volume $e^{-1/5n}$ times that of the original ellipsoid, while still maintaining containment of $C_1 \cap C_2$. If we continue this process T times without terminating and T satisfies

$$(\varepsilon/R)^n > e^{-T/5n} \Leftrightarrow T > 5n^2 \log(R/\varepsilon),$$

then the volume of the T -th ellipsoid (and therefore the volume of $C_1 \cap C_2$) is smaller than that of the Euclidean ball of radius ε .

We further note that Lemma 3.2.10 of [GLS93] ensures that each ellipsoid $\{y : \langle y - x, A^{-1}(y - x) \rangle \leq 1\}$ involved in the algorithm is described by a center x and a PSD matrix A of bit complexities at most $p = \text{poly}(n, \log R, \log(1/\varepsilon))$, where the result of each computation (possibly involving irrational numbers) is rounded down to the closest integer multiple of 2^{-p} . ■

We require some additional notation before proceeding. For a set $K \subseteq \mathbb{R}^n$ and $\varepsilon > 0$, let $K_\varepsilon^+ \stackrel{\text{def}}{=} K + \varepsilon \cdot \text{Ball}(\ell_2^n)$ and $K_\varepsilon^- \stackrel{\text{def}}{=} \{x : x + \varepsilon \cdot \text{Ball}(\ell_2^n) \subseteq K\}$, where $+$ denotes the minkowski sum. We may now combine binary search with [Proposition 3.5](#) to approximately maximize a concave function over a body with a (forward) approximate separation oracle.

Proposition 3.6 (Concave Maximization with an Approximate Separation Oracle).

Let C_1 be a closed set satisfying $x \in C_1 \Rightarrow \alpha x \in C_1$ for all $\alpha \in [0, 1]$, with an α_1 -approximate separation oracle \mathcal{SO}_1 . Let f be a homogeneous concave function such that $f_{\geq \lambda} \stackrel{\text{def}}{=} \{x : f(x) \geq \lambda\}$ has an α_2 -approximate inverse separation oracle \mathcal{SO}_2 and $f_{\geq \lambda} \cap C_1 \subseteq R \cdot \text{Ball}(\ell_2^n)$, for any $\lambda > 0$. Let $\text{OPT} = \sup_{x \in C_1} f(x)$, and assume that $\Lambda > \text{OPT}$ is given. Let $\varepsilon > 0$ and $\text{OPT}_\varepsilon = \sup_{x \in (C_1)_\varepsilon^-} f(x)$. There is an algorithm $\text{ALG}(R, \Lambda, \varepsilon, \mathcal{SO}_1, \mathcal{SO}_2)$ that, in time $\text{poly}(n, \log R, \log \Lambda, \log(1/\varepsilon))$, returns $y \in C_1$ that satisfies

$$f(y) \geq (\text{OPT}_\varepsilon - \varepsilon)/(\alpha_1 \alpha_2).$$

Proof. Starting from $\lambda = \Lambda$, we apply [Proposition 3.5](#) ($C_1 \leftarrow C_1$ and $C_2 \leftarrow f_{\geq \lambda}$), and perform binary search. In $\text{poly}(\log \Lambda, \log(1/\varepsilon))$ steps, we can find $\lambda \geq 0$ such that

1. There is a point $x \in (\alpha_1 \cdot C_1) \cap (f_{\geq \lambda}/\alpha_2)$.
2. $C_1 \cap f_{\geq \lambda+\varepsilon}$ does not contain a ball of radius ε .

(The first item is satisfied by $\lambda = 0$, and the second item is satisfied by $\lambda = \Lambda$.) The second item implies $\text{OPT}_\varepsilon \leq \lambda + \varepsilon$. Therefore, $y = x/\alpha_1$ satisfies $y \in C_1$ and

$$f(y) = f(x)/\alpha_1 \geq \lambda/(\alpha_1 \alpha_2) \geq (\text{OPT}_\varepsilon - \varepsilon)/(\alpha_1 \alpha_2),$$

which proves the claim. ■

Similarly we can approximately minimize a concave function over a set with an (inverse) approximate separation oracle.

Proposition 3.7 (Convex Minimization with an Approximate Separation Oracle).

Let C_1 be a closed set with an α_1 -approximate inverse separation oracle \mathcal{SO}_1 . Let f be a nonnegative, homogeneous, and convex function such that $f_{\leq \lambda} \stackrel{\text{def}}{=} \{x : f(x) \leq \lambda\}$ has an α_2 -approximate separation oracle \mathcal{SO}_2 for any $\lambda > 0$. Assume $\Lambda > \text{OPT}$ and $R > 0$ are given such that $\text{OPT} = \inf_{x \in C_1} f(x) = \inf_{x \in C_1 \cap R \cdot \text{Ball}(\ell_2^n)} f(x)$. Let $\varepsilon > 0$ and $\text{OPT}_\varepsilon = \inf_{x \in (C_1 \cap R \cdot \text{Ball}(\ell_2^n))_\varepsilon^-} f(x)$. There is an algorithm $\text{ALG}(R, \Lambda, \varepsilon, \mathcal{SO}_1, \mathcal{SO}_2)$ that, in time $\text{poly}(n, \log R, \log \Lambda, \log(1/\varepsilon))$, returns $y \in C_1$ that satisfies

$$f(y) \leq (\alpha_1 \alpha_2) \cdot (\text{OPT}_\varepsilon + \varepsilon).$$

Proof. Starting from $\lambda = \Lambda$, we apply [Proposition 3.5](#) ($C_1 \cap R \cdot \text{Ball}(\ell_2^n) \leftarrow f_{\leq \lambda}$ and $C_2 \leftarrow C_1$), and perform binary search. In $\text{poly}(\log \Lambda, \log(1/\varepsilon))$ steps, we can find $\lambda \geq 0$ such that

- There is a point inside $x \in (\alpha_2 \cdot f_{\leq \lambda}) \cap (C_1/\alpha_1)$.
- $(f_{\leq \lambda - \varepsilon} \cap R \cdot \text{Ball}(\ell_2^n) \cap C_1)$ does not contain a ball of radius ε .

(The first item is satisfied by $\lambda = \Lambda$, and the second item is satisfied by $\lambda = 0$.) The second item implies $\text{OPT}_\varepsilon \geq \lambda - \varepsilon$. Therefore, $y = \alpha_1 x$ satisfies $y \in C_1$ and

$$f(y) = \alpha_1 f(x) \leq \alpha_1 \alpha_2 \lambda \leq (\alpha_1 \alpha_2) \cdot (\text{OPT}_\varepsilon + \varepsilon),$$

which proves the claim. ■

A subtle issue is that we require versions of [Proposition 3.6](#) and [Proposition 3.7](#) with multiplicative approximation guarantees. For this it is necessary to assume C_1 and f have additional properties, which is in contrast to the exact setting where such assumptions are not required. In our work we only consider sets that are contained in some self-dual cone with the following additional properties. For a vector b , let $\text{bit}(b)$ be its bit complexity.

Definition 3.8 (Tractable Cone). We will say a closed, self-dual cone $K \subseteq \mathbb{R}^n$ is tractable if the following properties hold:

- (T1) Given $y \notin K$, there is a polynomial time algorithm to find a hyperplane separating y and K .
- (T2) Given $y \in K$ that admits $x \in K$ with $\langle x, y \rangle = 0$, there is a polynomial time algorithm to find such a vector x .
- (T3) If $y \in K$ does not admit such a x , $\min_{x \in K, \|x\|_2=1} \langle x, y \rangle > 2^{-\text{poly}(\text{bit}(y))}$.
- (T4) For any $y \in \mathbb{R}^n$, if $\max_{x \in K, \|x\|_2 \leq 1} \langle x, y \rangle \neq 0$, then it is at least $2^{-\text{poly}(\text{bit}(y))}$.
- (T5) For any $\varepsilon > 0$, there exists $y \in K$ with $\|y\|_2 \leq \text{poly}(n) \cdot \varepsilon$ such that $y + \varepsilon \cdot \text{Ball}(\ell_2^n) \subseteq K$.

We will only consider cases wherein K is the positive orthant or the positive semidefinite cone.

Definition 3.9 (Balanced Set/Function).

Let K be a tractable cone or \mathbb{R}^n . We say a set $B \subseteq \mathbb{R}^n$ is (R, r, K) -balanced if $(K \cap r \cdot \text{Ball}(\ell_2^n)) \subseteq B \subseteq (K \cap R \cdot \text{Ball}(\ell_2^n))$. When K is a tractable cone, we also say that a set $B \subseteq K$ is inverse (R, r, K) -balanced if $r \cdot \text{Ball}(\ell_2^n) \cap B = \emptyset$ and $\{x \in K : \|x\|_2 = R\} \subseteq B$. We say a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is (R, r, K) -balanced if (1) satisfies $r \cdot f(x) \leq \|x\|_2 \leq R \cdot f(x)$ for all $x \in K$, and (2) satisfies $\|f(x) - f(y)\| \leq R \|x - y\|_2$ for all $x, y \in K$.

We can deduce the following property which states that if C is balanced, C_ε^- is not too much smaller than C .

Lemma 3.10. *Let $K \subseteq \mathbb{R}^n$ be a tractable cone or \mathbb{R}^n itself. Let B be a closed, convex, and (R, r, K) -balanced set. For any $x \in B$ and $\varepsilon \in (0, r/\text{poly}(n))$, there exists $y \in B$ such that $\|x - y\|_2 \leq \varepsilon$ and $y + \frac{r\varepsilon}{R \cdot \text{poly}(n)} \cdot \text{Ball}(\ell_2^n) \subseteq B$.*

Proof. By property (T5) of K and the fact that $(K \cap r \cdot \text{Ball}(\ell_2^n)) \subseteq B$, for some $\alpha = r/\text{poly}(n)$, there exists $z \in K$ such that $z + \alpha \cdot \text{Ball}(\ell_2^n) \subseteq B$ and $\|z\|_2 \leq r/2$. (If $K = \mathbb{R}^n$, we can take $z = 0$.)

Fix $x \in B$ and let $T := \{t \in z + \alpha \cdot \text{Ball}(\ell_2^n) : \|z - t\|_2 = \alpha, \langle z - x, z - t \rangle = 0\}$, and consider $\text{conv}(x \cup T) \subseteq B$. Let θ be the angle $\angle zxt$ for some $t \in T$, which does not depend on the choice of t . Since $\|t - z\|_2 = \alpha$ and $\|x - t\|_2 \leq \|x - z\|_2 + \|z - t\|_2 \leq R + r/2 + \alpha \leq 2R$, $\sin \theta \geq \frac{\alpha}{2R}$. Let y be the point on \overline{xz} which is at distance ε from x . Then the distance from y to the boundary of $\text{conv}(x \cup T)$ is at least $\varepsilon \cdot \sin \theta \geq \frac{\varepsilon \alpha}{2R}$. Therefore, $y + \frac{\varepsilon \alpha}{2R} \cdot \text{Ball}(\ell_2^n) = y + \frac{r\varepsilon}{R \cdot \text{poly}(n)} \cdot \text{Ball}(\ell_2^n)$ is contained in $\text{conv}(x \cup T) \subseteq B$. ■

We now prove a multiplicative approximation version of [Proposition 3.6](#).

Proposition 3.11 (Concave Maximization with an Approximate Separation Oracle).

Let K be a tractable cone or \mathbb{R}^n and let C_1 be an (R, r, K) -balanced set with an α_1 -approximate separation oracle \mathcal{SO}_1 . Let f be either a linear function $\langle \xi, x \rangle$ for some $\xi \in \mathbb{R}^n$ (in which case $\alpha_2 = 1$) or an (R, r, K) -balanced, nonnegative, homogeneous, and concave function such that $f_{\geq \lambda} \stackrel{\text{def}}{=} \{x : f(x) \geq \lambda\}$ has an α_2 -approximate inverse separation oracle \mathcal{SO}_2 for any $\lambda > 0$. Let $\text{OPT} = \sup_{x \in C_1} f(x)$. There is an algorithm $\text{ALG}(R, r, \mathcal{SO}_1, \mathcal{SO}_2)$ that, in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(\xi))$, returns $y \in C_1$ that satisfies

$$f(y) \geq \frac{(1 - 1/n)}{\alpha_1 \alpha_2} \cdot \text{OPT}.$$

Proof. If f is a linear function, since we try to get a multiplicative approximation, assume without loss of generality that $f(x) = \langle \xi, x \rangle$ for some ξ with $\|\xi\|_2 = 1$. We first find an upper bound on OPT . By (R, r, K) -balancedness of C_1 , whether f is a linear or (R, r, K) -balanced,

$$\text{OPT} \leq \sup_{x \in C_1} f(x) \leq \max(\|x\|_2/r, \|x\|_2) \leq R \max(1, 1/r) =: \text{OPT}_{\max}$$

We also find a lower bound on OPT . If f is (R, r, K) -balanced,

$$\text{OPT} = \sup_{x \in C_1} f(x) \geq \sup_{x \in K \cap r \cdot \text{Ball}(\ell_2^n)} f(x) \geq \frac{r}{R} =: \text{OPT}_{\min}.$$

If f is linear, by property (T4) of K , unless $\text{OPT} = 0$ (in which case we can just return 0 as the optimal solution),

$$\text{OPT} \geq \sup_{x \in K \cap r \cdot \text{Ball}(\ell_2^n)} \langle \xi, x \rangle = r \cdot \sup_{x \in K \cap \text{Ball}(\ell_2^n)} \langle \xi, x \rangle \geq r \cdot 2^{-\text{poly}(\text{bit}(\xi))} =: \text{OPT}_{\min}$$

Finally, we bound the difference between OPT_ε and OPT for small enough $\varepsilon > 0$. Let $x^* \in \arg\max_{x \in C_1} \langle \xi, x \rangle$. By [Lemma 3.10](#), For any $\varepsilon \in (0, r/\text{poly}(n))$, there exists y^* such that $y \in (C_1)_\varepsilon^-$

and $\|x^* - y^*\|_2 \leq \varepsilon \cdot p(n)$ for a fixed polynomial $p(n)$. Since f is Lipschitz with the constant $\max(R, 1)$,

$$\text{OPT}_\varepsilon := \sup_{y \in (C_1)^\varepsilon} f(y) \geq f(y^*) \geq f(x^*) - \varepsilon \cdot p(n) \cdot \max(R, 1) = \text{OPT} - \varepsilon \cdot p(n) \cdot \max(R, 1).$$

We apply [Proposition 3.6](#) with $C_1 \leftarrow C_1$, $f \leftarrow f$, $\Lambda \leftarrow \text{OPT}_{\max}$, $R \leftarrow R$, and $\varepsilon \leftarrow \text{OPT}_{\min} / (2np(n) \max(R, 1))$. The running time is $\text{poly}(n, \log R, \log \Lambda, \log(1/\varepsilon)) = \text{poly}(n, \log R, \log(1/r), \text{bit}(\xi))$ and we find $x \in C_1$ that satisfies

$$f(x) \geq \alpha_1 \alpha_2 (\text{OPT}_\varepsilon - \varepsilon) \geq \alpha_1 \alpha_2 \left(\text{OPT} - \varepsilon \cdot 2p(n) \max(R, 1) \right) \geq \alpha_1 \alpha_2 \left(1 - \frac{1}{n} \right) \text{OPT},$$

which proves the claim. \blacksquare

We also give a multiplicative approximation version of [Proposition 3.7](#).

Proposition 3.12 (Convex Minimization with an Approximate Separation Oracle).

Let K be a tractable cone. Let $C_1 \subseteq K$ is a closed, upward-closed, and (R, r, K) -inversed balanced set with an α_1 -approximate inverse separation oracle \mathcal{SO}_1 . Let $f(x) = \langle \xi, x \rangle$ be a linear function for some $\xi \in K$ (in which case $\alpha_2 = 1$) or an (R, r, K) -balanced, nonnegative, homogeneous, and convex function such that $f_{\leq \lambda} \stackrel{\text{def}}{=} \{x : f(x) \leq \lambda\}$ has an α_2 -approximate separation oracle \mathcal{SO}_2 for any $\lambda > 0$. Let $\text{OPT} = \inf_{x \in C_1} f(x)$. There is an algorithm $\text{ALG}(R, r, \mathcal{SO}_1, \mathcal{SO}_2)$ that, in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(\xi))$, returns $y \in C_1$ that satisfies

$$f(y) \leq \text{OPT} \cdot \alpha_1 \alpha_2 (1 + 1/n).$$

Proof. We first prove upper and lower bounds on OPT . When f is (R, r, K) -balanced, since C_1 is also (R, r, K) -inversed balanced, $\text{OPT} \in [r^2, R^2]$.

Now we consider linear $f(x) = \langle \xi, x \rangle$. Assume without loss of generality that $\|\xi\|_2 = 1$. The inverse (R, r, K) -balancedness of C_1 implies

$$\text{OPT} := \inf_{x \in C_1} \langle \xi, x \rangle = \min_{x \in C_1 : \|x\| \in [r, R]} \langle \xi, x \rangle.$$

If the optimum is 0, there exists $x \in K$ such that $\langle x, \xi \rangle = 0$. Such x can be found by property (T2) of K , and $Rx / \|x\|_2 \in C_1$ is also an optimal solution. Otherwise, the property (T3) of K implies that

$$\text{OPT} \geq \min_{x \in K : \|x\|_2 = r} \langle \xi, x \rangle > \delta = 2^{-\text{poly}(\text{bit}(\xi))} \cdot r.$$

For an upper bound, since C_1 contains a vector of length R , the optimum is also upper bounded by R .

Let $\text{OPT}_{\min} = \min(2^{-\text{poly}(\text{bit}(\xi))} \cdot r, r^2)$, $\text{OPT}_{\max} := \max(R, R^2)$ be the lower and upper bound on OPT whether OPT is linear or (R, r, k) -balanced.

Finally, we bound the difference between OPT_ε and OPT for small enough $\varepsilon > 0$. For any $x^* \in \text{argmin}_{x \in C_1} f(x)$ and $\varepsilon \in (0, R/4)$, the property (T5) of K and the fact that C_1 is upward-closed imply that there exists y^* such that $y^* + \varepsilon \cdot \text{Ball}(\ell_2^n) \in C_1$ and $\|x^* - y^*\|_2 \leq \varepsilon \cdot p(n)$ for a fixed polynomial $p(n)$. Furthermore, both x^* and y^* have ℓ_2 norm at most $R + 2\varepsilon < 1.5R$. Since f is Lipschitz with the constant $\max(R, 1)$,

$$\text{OPT}_\varepsilon := \min_{y \in (C_1 \cap 2R \cdot \text{Ball}(\ell_2^n))^\varepsilon} f(y) \leq f(y^*) \leq f(x^*) + \varepsilon \cdot p(n) \cdot \max(R, 1).$$

We apply [Proposition 3.7](#) with $C_1 \leftarrow C_1$, $f(x) \leftarrow f(x)$, $\Lambda \leftarrow \text{OPT}_{\max}$, $R \leftarrow 2R$, and $\varepsilon \leftarrow \frac{\text{OPT}_{\min}}{2n \cdot p(n) \cdot \max(R, 1)}$. The running time is $\text{poly}(n, \log R, \log \Lambda, \log(1/\varepsilon)) = \text{poly}(n, \text{bit}(\zeta), \log R, \log(1/r))$ and we find x that satisfies

$$f(x) \leq \alpha_1 \alpha_2 (\text{OPT}_\varepsilon + \varepsilon) \leq \alpha_1 \alpha_2 (\text{OPT} + \varepsilon \cdot p(n) + \varepsilon) \leq \alpha_1 \alpha_2 (1 + 1/n) \text{OPT},$$

which proves the claim. \blacksquare

3.2 Duality of Approximation Algorithms for Linear Function Optimization

Systematic use of duality is fundamental in the metric theory of tensor products. Since we are interested in approximation algorithms for the injective tensor norm (read. bilinear maximization), we end up making extensive use of algorithmic versions of duality. Such results however, can fail drastically in the approximate optimization setting [[SV15](#)]. Nevertheless we show that in the special case of optimization over downward/upward-closed subsets of a self-dual cone, certain approximate algorithmic duality statements hold true.

We establish in this subsection a partial converse of [Proposition 3.11](#) and [Proposition 3.12](#), proving that an approximate linear function maximization algorithm for some downward/upward-closed convex subset B of a self-dual cone K , leads to an approximate separation oracle for B . In particular, we are interested in when K is \mathbb{R}^n , the nonnegative orthant, or the positive semidefinite cone.

We first observe that maximization over $B \subseteq K$ leads to separation for the polar of B . Here, two kinds of polars are covered — regular polars [2.3](#) when the underlying set K is \mathbb{R}^n , or conic polars [2.5](#) when the underlying set is a self-dual cone K .

Observation 3.13 (Maximization Oracle implies Polar Separation Oracle).

Let $K \subseteq \mathbb{R}^n$ be a tractable cone or \mathbb{R}^n itself. Let $B \subseteq K$ be a closed convex body and let \mathcal{O} be an oracle that takes a vector $y \in K$ as input and returns a β -approximately optimal solution to $\sup_{x \in B} \langle y, x \rangle$. Then there is an oracle polytime β -approximate separation oracle \mathcal{SO} for $B^{\circ K}$.

Proof. For an input vector $\zeta \in \mathbb{R}^n$, if $\zeta \notin K$, we use property (T1) of tractability to output a hyperplane (exactly) separating ζ and $B^{\circ K} \subseteq K$. For $\zeta \in K$, the description of \mathcal{SO} is as follows: if $\langle \mathcal{O}(\zeta), \zeta \rangle \leq 1$ return “Inside” else return $\{y \mid \langle \mathcal{O}(\zeta), y \rangle = 1\}$ as the hyperplane separating ζ from $B^{\circ K}$.

Indeed if \mathcal{O} returns x such that $\langle x, \zeta \rangle \leq 1$, it is valid to return “Inside” since $\sup_{x \in B} \langle x, \zeta \rangle \leq \beta$ and therefore $\zeta \in \beta \cdot B^{\circ K}$. On the other hand if \mathcal{O} returns $x \in B$ so that $\langle x, \zeta \rangle > 1$, we know $\sup_{\bar{\zeta} \in B^{\circ K}} \langle x, \bar{\zeta} \rangle \leq 1$ by definition. Thus $\{y \mid \langle x, y \rangle = 1\}$ is a hyperplane separating ζ from $B^{\circ K}$. \blacksquare

We now prove a partial converse of [Proposition 3.11](#) for linear maximization over a set with a (forward) approximate separation oracle.

Theorem 3.14 (Duality of Approximate Linear Maximization). Let $K \subseteq \mathbb{R}^n$ be a tractable cone or \mathbb{R}^n itself. Let $B \subseteq K$ be a closed, convex, and (R, r, K) -balanced set. If $K = \mathbb{R}^n$, additionally assume B is origin-symmetric. Consider any $\beta \geq 1$, and let \mathcal{O} be an oracle that $y \in K$ as input and returns a β -approximately optimal solution to $\sup_{x \in B} \langle y, x \rangle$. Then

- (1) There is an algorithm $\text{ALG}(R, r, \mathcal{O})$ that, given $\zeta \in K$, returns a $(1 + o(1))\beta$ -approximately optimal solution to $\sup_{x \in B^{\circ K}} \langle \zeta, x \rangle$ in time $\text{poly}(n, \text{bit}(\zeta), \log R, \log(1/r))$.

(2) There is a $(1 + o(1))\beta$ -separation oracle $\mathcal{SO}(R, r, \mathcal{O})$ for B (if $K = \mathbb{R}^n$) or $\downarrow B$ (if K is a tractable cone) that, on input ξ , runs in time $\text{poly}(n, \text{bit}(\xi), \log R, \log(1/r))$.

Proof. By [Observation 3.13](#) we have a β -approximate separation oracle for $B^{\circ K}$. We use it to get a $(\beta + 1/n)$ -approximation algorithm for computing $\sup_{x \in B^{\circ K}} \langle \xi, x \rangle$ for any $\xi \in K$. Note that (R, r, K) -balancedness of B implies $(K, 1/r, 1/R)$ -balancedness of $B^{\circ K}$. We apply [Proposition 3.11](#) (with $C_1 \leftarrow B^{\circ K}$ and $f(x) \leftarrow \langle \xi, x \rangle$) to have a $(\beta + 1/n)$ -approximation algorithm for computing $\sup_{x \in B^{\circ K}} \langle \xi, x \rangle$ for any $\xi \in K$. so the running time becomes $\text{poly}(n, \log \text{bit}(\xi), \log R, \log(1/r))$ instead of polynomially depending on $\text{bit}(\xi)$.

Finally we apply [Observation 3.13](#) to $B^{\circ K}$ (and use [Fact 2.6](#) for $(B^{\circ K})^{\circ K} = \downarrow B$ if K is a tractable cone or $(B^{\circ K})^{\circ K} = B$ if $K = \mathbb{R}^n$) to obtain the separation oracle claimed in (2). \blacksquare

Now we study a similar phenomenon for minimization using inverse polars ([Definition 2.7](#)). Our goal is to show that when K is a self-dual cone and $B \subseteq K$ is closed, convex, and upward-closed, an algorithm to approximately minimize a linear function over B implies other oracles. We first check that multiplicative approximation is well defined. If B satisfies the above condition, $B_c^\circ \subseteq K$ for any $c \in \mathbb{R}$, implying that the optimal solution $\inf_{x \in B} \langle y, x \rangle = -\infty$ if $y \notin K$ and nonnegative if $y \in K$. Therefore, a β -approximate algorithm is well defined for the problem of computing $\inf_{x \in B} \langle y, x \rangle$; output $-\infty$ if $y \notin K$ and a β -multiplicative approximate solution if $y \in K$.

Observation 3.15 (Minimization Oracle implies Polar Separation Oracle).

Let B be a closed, convex, and upward-closed set contained in a tractable cone $K \subseteq \mathbb{R}^n$. Assume further that $0 \notin B$. Let \mathcal{O} be an oracle that takes a vector $\xi \in K$ as input and returns a β -approximately optimal solution to $\inf_{x \in B} \langle \xi, x \rangle$ for some $\beta \geq 1$. Then there is an oracle polytime β -approximate inverse separation oracle \mathcal{SO} for B° .

Proof. For an input vector $\xi \in \mathbb{R}^n$, the description of \mathcal{SO} is as follows. Since $B^\circ \subseteq K$, given $\xi \notin K$, it outputs a hyperplane (exactly) separating ξ and K . For $\xi \in K$, the β -approximation algorithm for $\inf_{x \in B} \langle \xi, x \rangle$ returns $\mathcal{O}(\xi)$ such that $\langle \mathcal{O}(\xi), \xi \rangle \geq \text{OPT}/\beta$. If $\langle \mathcal{O}(\xi), \xi \rangle \geq 1$ return “Inside” else return $\{y \mid \langle \mathcal{O}(\xi), y \rangle = 1\}$ as the hyperplane separating ξ from B° .

Indeed if \mathcal{O} returns x such that $\langle x, \xi \rangle \geq 1$, it is valid to return “Inside” since $\inf_{x \in B} \langle x, \xi \rangle \geq 1/\beta$ and therefore $\xi \in B^\circ/\beta$.

On the other hand if \mathcal{O} returns $x \in B$ so that $\langle x, \xi \rangle < 1$, we know $\inf_{\bar{\xi} \in B^\circ} \langle x, \bar{\xi} \rangle \geq 1$ since by the Bipolar theorem, $(B^\circ)^\circ = B$ and hence $x \in (B^\circ)^\circ$. Thus $\{y \mid \langle x, y \rangle = 1\}$ is a hyperplane separating ξ from B° . \blacksquare

Finally, we prove a partial converse of [Proposition 3.12](#) for linear minimization over a set with an (inverse) approximate separation oracle.

Theorem 3.16 (Duality of Approximate Linear Minimization). Let B be a closed, convex, and upward-closed set contained in a tractable cone $K \subseteq \mathbb{R}^n$. Assume that either B or B° is (R, r, K) -inverse balanced. Consider any $\beta \geq 1$, and let \mathcal{O} be an oracle that takes $\xi \in K$ as input and returns a β -approximately optimal solution to $\inf_{x \in B} \langle \xi, x \rangle$. Then

- (1) There is an algorithm $\text{ALG}(R, r, \mathcal{O})$ that, given $\xi \in K$, returns a $(1 + o(1))\beta$ -approximately optimal solution to $\inf_{x \in B^\circ} \langle \xi, x \rangle$ in time $\text{poly}(n, \text{bit}(\xi), \log R, \log(1/r))$.
- (2) There is a $(1 + o(1))\beta$ -separation oracle $\mathcal{SO}(R, r, \mathcal{O})$ for B that, on input ξ , runs in time $\text{poly}(n, \text{bit}(\xi), \log R, \log(1/r))$.

Proof. By [Observation 3.15](#) we have a β -approximate separation oracle for B° . We apply [Proposition 3.12](#) ($B \leftarrow B^\circ, f \leftarrow \langle \xi, \cdot \rangle$). Note that the (R, r, K) -balancedness of B implies that (1) $(1/R)\text{Ball}(\ell_2^n)$ is disjoint from B° and (2) there exists $x \in B^\circ$ with $\|x\|_2 \leq \text{poly}(n)/r$, and the above two conditions can replace $(\text{poly}(n)/r, 1/R, K)$ -balancedness of B° in the proof of [Proposition 3.12](#); lower bounding OPT works verbatim as $(1/R)\text{Ball}(\ell_2^n) \cap B^\circ = \emptyset$ and upper bounding OPT just requires one point of $x \in B^\circ$ with $\|x\|_2 \leq \text{poly}(n)/r$. Therefore, we have a $(1 + 1/n)\beta$ -approximation algorithm for computing $\inf_{x \in B^\circ} \langle \xi, x \rangle$ in time $\text{poly}(n, \text{bit}(\xi), \log R, \log(1/r))$. Finally we apply [Observation 3.15](#) to B° (and use the fact that $(B^\circ)^\circ = B$) to obtain the separation oracle claimed in (2). ■

4 A Generic Framework: Algorithms from Covariance Separation Oracles, and Reductions across Quadratic/Bilinear/PSD Maximization

The main result of this section is a proof of [Theorem 1.14](#) which provides generic framework for quadratic/bilinear maximization under bounded type-2/dual-cotype-2. We do this by developing a theory of polynomial time reductions (with multiplicative loss depending only on $\tilde{T}_2(X)$ or $\tilde{C}_2(X^*)$) between the following five oracles

- (O1) Approximate Search Oracle for Quadratic Maximization.
- (O2) Approximate Search Oracle for Bilinear Maximization.
- (O3) Approximate Search Oracle for Quadratic Maximization of PSD instances.
- (SO1) Approximate Separation Oracle for Upper Covariance Body.
- (SO2) Approximate Separation Oracle for Lower Covariance Region.

See [Fig. 1](#) for a pictorial depiction of our reductions. While all of our applications in [Section 6](#) and [Section 7](#) may be proved via our framework theorem some proofs do not require its full power. We choose to cast all proofs as an appropriate combination of reductions between the above oracles, which leads to a more distilled and cohesive presentation.

Our framework theorem can be deduced as a consequence of the reductions $(O1) \rightarrow (O3), (O2) \rightarrow (O3), (O3) \rightarrow (SO2)$. The reductions $(O1) \rightarrow (O3), (O1) \rightarrow (SO1), (SO1) \rightarrow (O1)$ will be used in [Section 6](#) to derive closure properties for quadratic maximization under bounded Type-2. The reductions $(O2) \rightarrow (O3), (O3) \rightarrow (O2)$ will be used in [Section 6](#) to derive closure properties for bilinear maximization under bounded dual Cotype-2. The reduction $(O1) \rightarrow (SO1)$ will be used in [Section 7](#) to derive approximation algorithms for quadratic maximization over special families of norms with bounded Type-2. The reduction $(O2) \rightarrow (SO2)$ will be used in [Section 7](#) to derive approximation algorithms for bilinear maximization over special families of norms with bounded dual Cotype-2.

All our reductions in this section follow from judicious application of one or more of the following three ingredients:

1. Approximate convex optimization under approximate separation oracles ([Proposition 3.6](#), [Proposition 3.7](#)).
2. Algorithmic-duality/bipolarity for an appropriate downward/upward-closed subset of a self-dual cone ([Theorem 3.14](#), [Theorem 3.16](#)).
3. An appropriate Rounding algorithm based on sampling Gaussians.




Figure 1: The figure illustrates algorithmic reductions across Quadratic/Bilinear/PSD maximization and Upper/Lower Covariance Separation oracles. For brevity we only depict bilinear maximization over (X, X) , however all proofs address bilinear maximization in full generality (i.e., over (X, Y)). The notation $P \stackrel{\beta(\alpha)}{\leftarrow \alpha} Q$ denotes that a polytime α -approximation algorithm for task Q implies a polytime $\beta(\alpha)$ -approximation algorithm for task P . The subscript of \leq is omitted when the reduction is lossless.

4.1 Approximation Algorithms from Covariance Separation Oracles

In this section we give approximation algorithms for quadratic and PSD-quadratic maximization over X assuming approximate separation oracles for certain bodies associated to X .

An obstacle to convex programming approaches for quadratic maximization is that for norms that are not exactly 2-convex, it is unclear how to choose a computable convex relaxation of the objective. Here we show that simply by abandoning convexity and appealing to the approximate ellipsoid method one can approximate quadratic maximization over X (conditioned on a separation oracle for a body associated to X) without any structural assumptions on X like sign-invariance or 2-convexity. This serves only as the starting point of our approach as designing such a separation oracle can be a non-trivial task.

4.1.1 Quadratic Maximization via Upper Covariance Separation Oracles

Motivated by Gaussian rounding schemes, we consider the following (not convex but approximately convex) relaxation of $Q_X^{\max}(A)$:

$$\begin{aligned}
 & \text{maximize} && \langle A, \mathbb{X} \rangle \\
 & \text{subject to} && \mathbb{E}[\|\mathbb{X}^{1/2} \mathbf{g}\|_{\mathbb{X}}^2] \leq 1 \\
 & && \mathbb{X} \succeq 0
 \end{aligned} \tag{34}$$

In fact the relaxation is lossless, i.e.,

Observation 4.1. (34) is equal to $Q_X^{\max}(A)$.

Proof. To show (34) $\leq Q_X^{\max}(A)$, let X be an optimal solution of (34). We have

$$(34) = \langle A, X \rangle = \mathbb{E} \left[\langle X^{1/2}g, AX^{1/2}g \rangle \right] \leq \mathbb{E} [Q_X^{\max}(A) \cdot \|X^{1/2}g\|_X^2] \leq Q_X^{\max}(A).$$

The direction (34) $\geq Q_X^{\max}(A)$ follows from considering the substitution $X \stackrel{\text{def}}{=} xx^*$ where $x \in \text{Ball}(X)$ is an optimal solution to $Q_X^{\max}(A)$. ■

We show next how applying the approximate ellipsoid method to (34) can provide a conditional approximation algorithm for quadratic maximization.

Proposition 4.2 (Quadratic Maximization Given Separation Oracle for Upper Covariance Body). *There is an algorithm $\text{ALG}(A, R, r, \mathcal{SO})$ such that if $\|\cdot\|_X$ is an (R, r) -balanced norm over \mathbb{R}^n and \mathcal{SO} is an α -approximate separation oracle for $\mathcal{U}(X)$, then on any input $A \in M_n(\mathbb{R})$, ALG runs in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A))$ and returns a $(1 + o(1))\alpha$ -approximate solution to $Q_X^{\max}(A)$ with probability $1 - 2^{-\Omega(n)}$.*

Proof. We use the approximate ellipsoid method (Proposition 3.11 with $C_1 \leftarrow \mathcal{U}(X)$) to compute a $(1 + o(1))\alpha$ -approximate solution X to (34). By Observation 4.1, we know $\mathbb{E} [\langle X^{1/2}g, AX^{1/2}g \rangle] \geq (1 - o(1))Q_X^{\max}(A)/\alpha$. We define a random variable $G \stackrel{\text{def}}{=} \langle X^{1/2}g, AX^{1/2}g \rangle - Q_X^{\max}(A) \cdot \|X^{1/2}g\|_X^2/\tau$ where $g \in \mathbb{R}^n$ is a vector of independent standard Gaussians and $\tau = \alpha(1 + \varepsilon)$ (where $\varepsilon = o(1)$ is chosen to decay slower than the hidden $o(1)$ factor above). By a routine application of Chebyshev's inequality (and independent resampling) we conclude that polynomially many independent samples of g will guarantee that with probability $1 - 2^{-\Omega(n)}$, $G \geq 0$ for some sample and therefore $X^{1/2}g/\|X^{1/2}g\|_X$ is the desired approximate solution (the argument with all details is identical to the concentration argument made in the proof of Proposition 4.4, we choose not to duplicate it here for brevity). ■

A necessary condition for an α -separation oracle is α -separability. Interestingly, $\mathcal{U}(X)$ is precisely $\tilde{T}_2(X)^2$ -separable (see Section 4.2.4) and from this the connection to type-2 is evident. Unfortunately, (as we show in Section 9) bounded type-2 is not sufficient in general. However, in Section 6 and Section 7 we exhibit several families of norms wherein one can construct approximate separation oracles for the Gaussian body.

In its present form Proposition 4.2 is not easy to use and in most of our examples it takes significant work to design C -separation oracles for $\mathcal{U}(X)$ with C being a function of $\tilde{T}_2(X)$. We show in the next section how for the special class of PSD instances one can obtain an analogue of Proposition 4.2 assuming only a separation oracle for the lower covariance region. Designing such separation oracles turns out to be an easier task in many cases.

4.1.2 PSD Quadratic Maximization via Lower Covariance Separation Oracles

Let $A = BB^*$ where B is an $n \times n$ matrix and let X be a norm over \mathbb{R}^n . We crucially use the simple fact that $Q_X^{\max}(A) = \text{Op}_{X, X}^{\max}(A) = \|BB^*\|_{X \rightarrow X^*} = \|B\|_{2 \rightarrow X^*}^2$, which informally speaking, allows us to convert a search problem over X into a search problem over ℓ_2^n . We consider the following (not convex but approximately convex) relaxation of $\|B\|_{2 \rightarrow X^*}^2 = Q_X^{\max}(A)$:

$$\text{maximize} \quad \mathbb{E} [\|BW^{1/2}g\|_{X^*}^2]$$

$$\begin{aligned} \text{subject to } \quad & \text{Tr}(\mathbb{W}) \leq 1 \\ & \mathbb{W} \in \text{PSD}^n \end{aligned} \tag{35}$$

In fact the relaxation is lossless, i.e.,

Observation 4.3. (35) is equal to $\|B\|_{2 \rightarrow X^*}^2 = Q_X^{\max}(A)$.

Proof. Let \mathbb{W} be the optimal solution of (35).

$$\mathbb{E} [\|B\mathbb{W}^{1/2}\mathbf{g}\|_{X^*}^2] \leq \mathbb{E} [\|B\|_{2 \rightarrow X^*}^2 \cdot \|\mathbb{W}^{1/2}\mathbf{g}\|_2^2] = \|B\|_{2 \rightarrow X^*}^2 \cdot \text{Tr}(\mathbb{W}) \leq \|B\|_{2 \rightarrow X^*}^2$$

which implies one direction of the claim.

The other direction follows from considering the substitution $\mathbb{W} \stackrel{\text{def}}{=} ww^*$ where $w \in \text{Ball}(\ell_2^n)$ is an optimizer of $\|Bw\|_{X^*}$. \blacksquare

We show next how applying the approximate ellipsoid method to (35) can provide an approximation algorithm for PSD quadratic maximization with a weaker assumption than [Proposition 4.2](#).

Proposition 4.4 (PSD Maximization Given Lower Covariance Separation Oracle).

There is an algorithm $\text{ALG}(A, R, r, \mathcal{SO})$ such that if $\|\cdot\|_X$ is an (R, r) -balanced norm over \mathbb{R}^n and \mathcal{SO} is an α -approximate separation oracle for $\mathcal{L}(X)$, then on any input $A \in \text{PSD}^n$, ALG runs in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A))$ and returns a $(1 + o(1))\alpha$ -approximate solution to $Q_X^{\max}(A)$ with probability $1 - 2^{-\Omega(n)}$.

Proof. Without loss of generality we may assume A is invertible otherwise we may add a small enough multiple of the identity without changing the objective value much (and therefore any $n \times n$ matrix B is invertible if it satisfies $A = BB^*$). By [Observation 4.3](#), to estimate $Q_X^{\max}(A)$, it suffices to approximately compute (35). By [Proposition 3.11](#) with $C_1 \leftarrow \text{Ball}(\ell_2^n)$ and $f(\mathbb{W}) \leftarrow \mathbb{E}[\|B\mathbb{W}^{1/2}\mathbf{g}\|_{X^*}^2]$, it suffices to give for every $\lambda > 0$, an α -approximate separation oracle for the set

$$S_\lambda \stackrel{\text{def}}{=} \left\{ \mathbb{W} \succeq 0 \mid \mathbb{E}[\|B\mathbb{W}^{1/2}\mathbf{g}\|_{X^*}^2] \geq \lambda \right\} = \left\{ \mathbb{W} \mid \lambda^{-1} \cdot B\mathbb{W}B^* \in \mathcal{L}(X) \right\}.$$

Since S_λ is a linear transformation of $\mathcal{L}(X)$, it is straightforward to check that \mathcal{SO} can be adapted to provide an α -approximate separation oracle for S_λ . Indeed consider any $\overline{W} \in M_n(\mathbb{R})$ and run $\mathcal{SO}(\lambda^{-1} \cdot B\overline{W}B^*)$. If “Inside” is returned, we know $\mathbb{W} \in \alpha \cdot S_\lambda$ and therefore can return “Inside”. On the other hand if \mathcal{SO} returns M separating $\mathcal{L}(X)$ from $\lambda^{-1} \cdot B\overline{W}B^*$ then $\lambda^{-1} \cdot B^*MB$ separates S_λ from \overline{W} and thus we obtain an α -approximate separation oracle for S_λ .

Our rounding algorithm proceeds by first obtaining w.h.p. a vector $\overline{\mathbf{g}} \in \text{Ball}(\ell_2^n)$ such that $\|B\overline{\mathbf{g}}\|_{X^*}^2 \geq (1 - o(1)) \cdot \|B\|_{2 \rightarrow X^*}^2 / \alpha$. Then the desired solution is any vector $\xi \in \text{Ball}(X)$ satisfying $\langle \xi, B\overline{\mathbf{g}} \rangle = \|B\overline{\mathbf{g}}\|_{X^*}$ (which can be found using standard convex optimization). This is because

$$\langle A, \xi\xi^* \rangle = \|B^*\xi\|_2^2 \geq \langle \xi, B\overline{\mathbf{g}} \rangle^2 = \|B\overline{\mathbf{g}}\|_{X^*}^2 \geq (1 - o(1)) \frac{\|B\|_{2 \rightarrow X^*}^2}{\alpha} = (1 - o(1)) \frac{Q_X^{\max}(A)}{\alpha}.$$

Let \mathbb{W} be a $\alpha/(1 - 1/n)$ -approximately optimal solution to (35). We define a random variable $\mathbf{G} \stackrel{\text{def}}{=} \|B\mathbb{W}^{1/2}\mathbf{g}\|_{X^*}^2 - \|B\|_{2 \rightarrow X^*}^2 \cdot \|\mathbb{W}^{1/2}\mathbf{g}\|_2^2 / \tau$ where \mathbf{g} is an i.i.d. standard Gaussian vector and $\tau > \alpha$ is a parameter to be chosen later. Note that by assumption on \mathbb{W} ,

$$\frac{1 - 1/n}{\alpha} \cdot \|B\|_{2 \rightarrow X^*}^2 \leq \mathbb{E} [\|B\mathbb{W}^{1/2}\mathbf{g}\|_{X^*}^2] \leq \mathbb{E} [\|B\|_{2 \rightarrow X^*}^2 \cdot \|\mathbb{W}^{1/2}\mathbf{g}\|_2^2] \leq \|B\|_{2 \rightarrow X^*}^2$$

$$\Rightarrow \mathbb{E} [\mathbf{G}] / \|B\|_{2 \rightarrow X^*}^2 \in [(1 - 1/n)/\alpha - 1/\tau, 1] \quad (36)$$

For $i \in [t]$ let $\mathbf{G}_i \stackrel{\text{def}}{=} \|B\mathbf{W}^{1/2}\mathbf{g}^i\|_{X^*}^2 - \|B\|_{2 \rightarrow X^*}^2 \cdot \|\mathbf{W}^{1/2}\mathbf{g}^i\|_2^2 / \tau$ be i.i.d. copies of \mathbf{G} where $t = n^{3C}$ for some constant $C > 1$ and each $\mathbf{g}^i \in \mathbb{R}^n$ is a vector of independent standard Gaussians. We set $\bar{\mathbf{g}} \stackrel{\text{def}}{=} \mathbf{W}^{1/2}\mathbf{g}^j / \|\mathbf{W}^{1/2}\mathbf{g}^j\|_2$ for any maximizer $\mathbf{G}_j = \max_{i \in [t]} \mathbf{G}_i$. We argue that $\bar{\mathbf{g}}$ is the desired vector w.h.p. through a routine application of Chebyshev's inequality. We start with a bound on the variance:

$$\begin{aligned} & \text{Var} \left[\sum_{i \in [t]} \mathbf{G}_i \right] \\ &= t \cdot \text{Var} [\mathbf{G}] \\ &\leq t \cdot \mathbb{E} [\mathbf{G}^2] \\ &\leq t \cdot \left(\mathbb{E} [\|B\mathbf{W}^{1/2}\mathbf{g}\|_{X^*}^4] + \mathbb{E} [\|B\|_{2 \rightarrow X^*}^4 \|\mathbf{W}^{1/2}\mathbf{g}\|_2^4 / \tau^2] \right) \\ &\leq C' \cdot t \cdot \left(\mathbb{E} [\|B\mathbf{W}^{1/2}\mathbf{g}\|_{X^*}^2]^2 + \mathbb{E} [\|B\|_{2 \rightarrow X^*}^2 \|\mathbf{W}^{1/2}\mathbf{g}\|_2^2 / \tau]^2 \right) \quad (\text{Khinchine-Kahane Theorem 2.1}) \\ & \hspace{15em} (\text{where } C' \text{ is an absolute const.}) \\ &\leq C' \cdot ((1 - 1/n)/\alpha - 1/\tau) \cdot t \cdot \mathbb{E} [\mathbf{G}]^2 \quad (\text{by (36)}) \end{aligned}$$

By Chebyshev's inequality, with probability at least $1 - 1/k^2$ it holds that

$$\sum_{i \in [t]} \mathbf{G}_i \geq t \cdot \mathbb{E} [\mathbf{G}] - k \cdot \left(C' \cdot t \cdot \left(\frac{1 - 1/n}{\alpha} - \frac{1}{\tau} \right) \right)^{1/2} \cdot \mathbb{E} [\mathbf{G}].$$

By averaging, there must exist $i \in [t]$ such that

$$\frac{\|B\mathbf{W}^{1/2}\mathbf{g}^i\|_{X^*}^2}{\|\mathbf{W}^{1/2}\mathbf{g}^i\|_2^2} \geq \mathbb{E} [\mathbf{G}] \cdot \left(1 - O \left(\frac{k}{\sqrt{t}} \left(\frac{1 - 1/n}{\alpha} - \frac{1}{\tau} \right) \right) \right).$$

Setting $k \stackrel{\text{def}}{=} t^{0.49}$ and $\tau \stackrel{\text{def}}{=} (1 + t^{-0.001})\alpha$, we obtain that with probability at least $1 - n^{-C}$ there exists $i \in [t]$ such that $\mathbf{G}_i \geq \mathbb{E} [\mathbf{G}] (1 - O(t^{-0.009})) \geq 0$. Thus $\|B\mathbf{W}^{1/2}\mathbf{g}^i\|_{X^*}^2 / \|\mathbf{W}^{1/2}\mathbf{g}^i\|_2^2 \geq \|B\|_{2 \rightarrow X^*}^2 / \tau$ as desired. We obtain success probability $1 - 2^{-\Omega(n)}$ by repeating the above process independently for polynomially many rounds. This completes the proof. \blacksquare

Recall a necessary condition for an α -separation oracle is α -separability. Interestingly, $\mathcal{L}(X)$ is precisely $\tilde{C}_2(X^*)^2$ -separable (see Section 4.2.5) and from this the connection to dual cotype-2 is evident. It turns out that in many cases separation oracles are easier to design for $\mathcal{L}(X)$ than $\mathcal{U}(X)$. In Section 6 and Section 7 we explore several situations wherein one can construct such approximate separation oracles.

One of our main technical contributions is using factorization theory to show that quadratic maximization under bounded type-2 (and bilinear maximization under bounded dual cotype-2) reduces (algorithmically) to solving a series of PSD maximization instances. This leads us to several new results for quadratic and bilinear maximization. We now proceed to exhibit the claimed reductions.

4.2 Reductions Across Quadratic/Bilinear/PSD Maximization

In this section we explore in depth polynomial time reductions (with multiplicative loss depending only on $\tilde{T}_2(X)$ or $\tilde{C}_2(X^*)$) between the following three oracles

- (O1) Approximate Search Oracle for Quadratic Maximization.
- (O2) Approximate Search Oracle for Bilinear Maximization.
- (O3) Approximate Search Oracle for Quadratic Maximization of PSD instances.

The reductions from quadratic/bilinear maximization to PSD maximization, as we will see in [Section 5](#), are closely related to the theory of factorization of linear operators through ℓ_2 .

4.2.1 Reducing Quadratic to PSD under Bounded Type-2

In this section we show how a C -approximate search algorithm for quadratic maximization of PSD instances over X can be used to give a $(1 + o(1)) \cdot C \cdot \tilde{T}_2(X)^2$ -approximation algorithm for quadratic maximization of general instances. Our approach is as follows:

1. We will use the PSD oracle and algorithmic-duality for downward-closed sets ([Theorem 3.14](#)) to construct a separation oracle for $\downarrow B_{\wedge}^{\text{Sym}}(X)$.
2. This allows us to approximately solve the following (not exactly computable) vector relaxation of $Q_X^{\max}(A)$:

$$\max\{\langle A, \mathbb{X} \rangle \mid \mathbb{X} \in \downarrow B_{\wedge}^{\text{Sym}}(X)\}. \quad (37)$$

3. Finally we use Gaussian rounding to show that (37) $\leq \tilde{T}_2(X)^2 \cdot Q_X^{\max}(A)$.

We proceed with the proof.

Theorem 4.5 (Type-2 Quadratic Maximization Reduces to PSD Quadratic Maximization).

There is an algorithm $\text{ALG}(A, R, r, \mathcal{O})$ such that if $\|\cdot\|_X$ is an (R, r) -balanced norm over \mathbb{R}^n and \mathcal{O} is an α -approximate search oracle for PSD quadratic maximization over X , then on any input $A \in M_n(\mathbb{R})$, ALG runs in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A))$ and returns a $(1 + o(1)) \cdot \alpha \cdot \tilde{T}_2(X)^2$ -approximate solution to $Q_X^{\max}(A)$ with probability at least $1 - 2^{-\Omega(n)}$.

Proof. By [Proposition 4.2](#) it suffices to exhibit an approximate separation oracle for $\mathcal{U}(X)$.

To this end, note that \mathcal{O} is an oracle returning an α -approximate solution to $Q_X^{\max}(W) = \sup_{\mathbb{X} \in \downarrow B_{\wedge}^{\text{Sym}}(X)} \langle W, \mathbb{X} \rangle$ for any $W \in \text{PSID}^n$. Therefore we may apply [Theorem 3.14](#) to the set $\downarrow B_{\wedge}^{\text{Sym}}(X)$ with $K \stackrel{\text{def}}{=} \text{PSID}^n$ to obtain a $(1 + o(1))\alpha$ -approximate separation oracle for $\downarrow B_{\wedge}^{\text{Sym}}(X)$. (The balancedness of $\downarrow B_{\wedge}^{\text{Sym}}(X)$ from [Lemma 2.24](#).)

Lastly [Observation 3.4](#) yields the desired $(1 + o(1)) \cdot \alpha \cdot \tilde{T}_2(X)^2$ -approximate separation oracle for $\mathcal{U}(X)$ since by [Observation 2.17](#) we have the equivalence $\tilde{T}_2(X)^{-2} \cdot \downarrow B_{\wedge}^{\text{Sym}}(X) \subseteq \mathcal{U}(X) \subseteq \downarrow B_{\wedge}^{\text{Sym}}(X)$. \blacksquare

Remark 4.6. *We make two technical remarks that are required for application of this theorem to algorithmic closure properties for complex interpolation. Note that the proof above can be adapted verbatim for the following additional features:*

1. We may take $\|\cdot\|_X$ to be a norm over \mathbb{C}^n and input $A \in M_n(\mathbb{C})$.
2. We may replace the PSD maximization search oracle \mathcal{O} with a slightly weaker oracle \mathcal{O}' that on input $W \succeq 0$, returns a witness $Z \in \text{Ball}(X \otimes X)$ satisfying $\langle W, Z \rangle \geq Q_X^{\max}(W)/C$. Such an oracle is weaker as Z can be a convex combination of rank-1 matrices of the form xy^* (where $x, y \in \text{Ball}(X)$) instead of a single such rank-1 matrix.

4.2.2 Reducing Bilinear to PSD under Bounded Dual Cotype-2

In this section we show how a C -approximate search algorithm for quadratic maximization over PSD instances can be used to give a $C' = C'(C, \tilde{C}_2(X^*), \tilde{C}_2(Y^*))$ -approximation algorithm for bilinear maximization over X, Y (i.e., $\text{Op}_{X,Y}^{\max}(\cdot)$). Our approach is as follows:

1. We will use the PSD oracles and algorithmic-duality for downward-closed sets ([Theorem 3.14](#)) to construct separation oracles for $\downarrow B_{\wedge}^{\text{Sym}}(X), \downarrow B_{\wedge}^{\text{Sym}}(Y)$.
2. This allows us to approximately solve the following (not exactly computable) convex relaxation $\text{Rlx}(A)$ of $\text{Op}_{X,Y}^{\max}(A)$:

$$\begin{aligned} \text{Rlx}(A) = \max \quad & \langle A, Z \rangle \\ \text{s.t.} \quad & \mathbb{X} \in \downarrow B_{\wedge}^{\text{Sym}}(X), \mathbb{Y} \in \downarrow B_{\wedge}^{\text{Sym}}(Y) \\ & \begin{bmatrix} \mathbb{X} & Z \\ Z^* & \mathbb{Y} \end{bmatrix} \succeq 0. \end{aligned} \tag{38}$$

3. Finally we need a rounding algorithm taking an optimal solution of $\text{Rlx}(A)$ and producing a good solution to $\text{Op}_{X,Y}^{\max}(A)$. A deep factorization theorem of Pisier [[Pis80](#)] states that

$$\|A\|_{Y \rightarrow X^*} \leq \inf_{BC=A} \|C\|_{Y \rightarrow 2} \cdot \|B\|_{2 \rightarrow X^*} \leq C' \cdot \|A\|_{Y \rightarrow X^*}.$$

Pisier's proof can be viewed as having two components. The first component is a bound on the intergrality gap of $\text{Rlx}(A)$:

$$\text{Op}_{X,Y}^{\max}(A) \leq \text{Rlx}(A) \leq C' \cdot \text{Op}_{X,Y}^{\max}(A). \tag{39}$$

Pisier's proof of (39) is non-constructive and does not furnish a rounding algorithm for bilinear maximization. We give such a rounding algorithm below.

The second component of the factorization theorem is a dual characterization of $\text{Rlx}(A)$ as a factorization norm that follows from a duality result of Maurey, i.e., $\text{Rlx}(A) = \inf_{BC=A} \|C\|_{Y \rightarrow 2} \cdot \|B\|_{2 \rightarrow X^*}$. We refer the reader to [Section 5](#) for a detailed discussion of factorization theory and the results of Maurey and Pisier.

Pisier's proof exhibits a (randomized) non-constructive map from any pair of sequences $(y_i), (\bar{y}_i)$ to a unit vector $y \in \text{Ball}(Y)$ such that $\mathbb{E}[\|Ay\|_{X^*}]$ is bounded from below by $(\sum_i \|Ay_i\|_{X^*}^2 / \sum_j \|\bar{y}_j\|_Y^2) / C'$. The analysis of our rounding algorithm borrows heavily from Pisier's proof as we utilize his non-constructive map as a witness of our rounded vector having high expected value.

We proceed with collecting the necessary components of our proof. We require the following fourier analytic decomposition result that is a critical component of Pisier's factorization theorem.

Lemma 4.7 (Pisier [[Pis86](#)]). *Let $g_1, \dots, g_m \sim \mathcal{N}(0, 1)$ be i.i.d. standard Gaussians and let $\varepsilon_1, \dots, \varepsilon_m$ be i.i.d. rademacher (± 1) random variables. If $\sum_{i \in [m]} y_i y_i^* \in \downarrow B_{\wedge}^{\text{Sym}}(Y)$, then for any $0 < \delta < 1$ there exists a vector valued function $\varphi : \mathbb{R}^n \times \{\pm 1\}^m \mapsto \mathbb{R}^n$ such that*

$$\begin{aligned} \mathbb{E}_{g, \varepsilon} [\|\sum_i \varepsilon_i g_i y_i + \varphi(g, \varepsilon)\|_Y^2]^{1/2} &= O(\tilde{C}_2(Y^*) \log 1/\delta) \\ \mathbb{E}_{g, \varepsilon} [\|A\varphi(g, \varepsilon)\|_{X^*}^2]^{1/2} &\leq \delta \cdot \text{Rlx}(A) \quad \text{for any } A \in M_{n,m}(\mathbb{R}). \end{aligned}$$

We will also need the following simple fact about scalar valued random variables.

Fact 4.8. Fix $\alpha, \varepsilon > 0$ and $\beta > 1$. Let x be a random variable supported in the interval $[-\infty, \alpha]$ and let $\mathbb{E}[x] \geq \alpha/\beta$. Then

$$\mathbb{P}\left[x < \frac{\alpha}{(1+\varepsilon)\beta}\right] \leq \frac{(1+\varepsilon)(\beta-1)}{\beta(1+\varepsilon)-1} = 1 - \frac{\varepsilon}{\beta + \varepsilon\beta - 1}.$$

Proof. If the statement is false,

$$\mathbb{E}[x] < \frac{(1+\varepsilon)(\beta-1)}{\beta(1+\varepsilon)-1} \cdot \frac{\alpha}{(1+\varepsilon)\beta} + \left(1 - \frac{(1+\varepsilon)(\beta-1)}{\beta(1+\varepsilon)-1}\right)\alpha = \alpha/\beta,$$

leading to contradiction. \blacksquare

Finally we need the following simple claim about factorization of a 2×2 positive definite block matrix.

Claim 4.9. Consider any invertible matrix $\begin{bmatrix} \mathbb{X} & Z \\ Z^* & \mathbb{Y} \end{bmatrix} \in \text{PSD}^{n+m}$ and fix any factorization $\mathbb{X} = SS^*$ where $S \in M_{n,s}(\mathbb{R})$. Then there exists $T \in M_{m,s}(\mathbb{R})$ such that $Z = ST^*$ and $\mathbb{Y} \succeq TT^*$.

Proof. By assumption there exist matrices $U \in M_{n,n+m}(\mathbb{R})$ and $V \in M_{m,n+m}(\mathbb{R})$ such that $\mathbb{X} = UU^*$, $\mathbb{Y} = VV^*$ and $Z = UV^*$. Let $S^\dagger \stackrel{\text{def}}{=} S^*(SS^*)^{-1} = S^*\mathbb{X}^{-1}$ denote the moore-penrose pseudo-inverse of S (recall S^\dagger satisfies the property $SS^\dagger = I$). Let $T \stackrel{\text{def}}{=} VU^*(S^\dagger)^*$. Then we have $ST^* = SS^\dagger UV^* = UV^* = Z$.

It remains to verify $\mathbb{Y} \succeq TT^*$. To this end note that by definition of pseudo-inverse one has

$$(S^\dagger)^* S^\dagger = (S^\dagger)^* S^* \mathbb{X}^{-1} = \mathbb{X}^{-1}. \quad (40)$$

Let $H \stackrel{\text{def}}{=} U^\dagger U = U^* \mathbb{X}^{-1} U = H^*$. Then $H^2 = U^* \mathbb{X}^{-1} U U^* \mathbb{X}^{-1} U = U^* \mathbb{X}^{-1} U = H$. Thus H is a symmetric orthogonal projector. Now we have

$$TT^* = VU^*(S^\dagger)^* S^\dagger UV^* = VU^* \mathbb{X}^{-1} UV^* = V^* H V \preceq V^* V = \mathbb{Y}$$

where the second equality uses (40), and the (PSD) inequality uses the fact that H is an orthogonal projector. This completes the proof. \blacksquare

We are now ready to prove the main result of this subsection.

Theorem 4.10 (Bilinear Maximization Reduces to PSD Quadratic Maximization).

There is an algorithm $\text{ALG}(A, R, r, \mathcal{O}_X, \mathcal{O}_Y)$ such that if $\|\cdot\|_X$ (resp. $\|\cdot\|_Y$) is an (R, r) -balanced norm over \mathbb{R}^n (resp. \mathbb{R}^m) and \mathcal{O}_X (resp. \mathcal{O}_Y) is an α -approximate search oracle for PSD quadratic maximization over X (resp. Y), then on any input $A \in M_{n,m}(\mathbb{R})$, ALG runs in time $\text{poly}(n, m, \log R, \log 1/r, \text{bit}(A))$ and returns an $\alpha\beta \log \alpha\beta$ -approximate solution to $\text{Op}_{X,Y}^{\max}(A)$ with probability $1 - 2^{-\Omega(n)}$, where $\beta \lesssim C_2(X^*)C_2(Y^*)$.

Proof. Just as in the proof of [Theorem 4.5](#) we obtain $(1 + o(1)) \cdot \alpha$ -separation oracles for $\downarrow B_\lambda^{\text{Sym}}(X)$, $\downarrow B_\lambda^{\text{Sym}}(Y)$. Therefore by [Proposition 3.11](#) we can compute a $(1 + o(1))\alpha$ -approximate solution $(\mathbb{X}, \mathbb{Y}, Z)$ to (38) and apply the following rounding algorithm:

1. Let \mathbb{Y} be a $(1 + o(1))\alpha$ -approximate solution to (38) and fix any decomposition $\mathbb{Y} = \sum_{i \in [m]} y_i y_i^*$.
2. Let $\mathbf{g}_1, \dots, \mathbf{g}_m \sim \mathcal{N}(0, 1)$ be i.i.d. standard Gaussians and let $\varepsilon_1, \dots, \varepsilon_m$ be i.i.d. rademacher (± 1) random variables. Using the exact membership oracle of X (and standard convex optimization), compute $x_{\mathbf{g}} \in \text{Ball}(X)$ satisfying $\langle x_{\mathbf{g}}, \sum_i \varepsilon_i \mathbf{g}_i A y_i \rangle = \|\sum_i \varepsilon_i \mathbf{g}_i A y_i\|_{X^*}$.
3. Using the exact membership oracle of Y (and convex optimization), compute $y_{\mathbf{g}} \in \text{Ball}(Y)$ satisfying $\langle A^* x_{\mathbf{g}}, y_{\mathbf{g}} \rangle = \|A^* x_{\mathbf{g}}\|_{Y^*}$.
4. Output $x_{\mathbf{g}}, y_{\mathbf{g}}$.

We now analyze the rounding algorithm. Since $\wedge_X^{\downarrow \text{Sym}}(\mathbb{X}) \leq 1$, by definition there exists $\mathbb{X}' \succeq \mathbb{X}$ and a finite sequence (x_i) such that $\mathbb{X}' = \sum_i x_i x_i^*$ and $\sum_i \|x_i\|_X^2 \leq 1$. Since $\mathbb{X}' \succeq \mathbb{X}$, we have

$$\begin{bmatrix} \mathbb{X}' & Z \\ Z^* & \mathbb{Y} \end{bmatrix} \succeq 0.$$

Without loss of generality we assume $\begin{bmatrix} \mathbb{X}' & Z \\ Z^* & \mathbb{Y} \end{bmatrix}$ is invertible since otherwise we may add a sufficiently small copy of the identity matrix which affects estimates by only lower order terms. Then by [Claim 4.9](#), there exists a sequence (\bar{y}_i) such that $Z = \sum_i x_i \bar{y}_i^*$ and $\mathbb{Y} \succeq \sum_i \bar{y}_i \bar{y}_i^*$. We have

$$\begin{aligned} & (1 - o(1)) \cdot \text{Rlx}(A) / \alpha \\ & \leq \langle A, Z \rangle \\ & = \sum_i \langle x_i, A \bar{y}_i \rangle \\ & \leq \left(\sum_i \|x_i\|_X^2 \right)^{1/2} \left(\sum_i \|A \bar{y}_i\|_{X^*}^2 \right)^{1/2} && \text{(by Hölder + Cauchy-Schwarz)} \\ & \leq \left(\sum_i \|A \bar{y}_i\|_{X^*}^2 \right)^{1/2} && \text{(by definition of } (x_i)\text{)} \\ & \leq \tilde{C}_2(X^*) \cdot \mathbb{E}_{\mathbf{g}} \left[\|\sum_i \mathbf{g}_i A \bar{y}_i\|_{X^*}^2 \right]^{1/2} && \text{(by definition of cotype-2)} \\ & \leq \tilde{C}_2(X^*) \cdot \mathbb{E}_{\mathbf{g}} \left[\|\sum_i \mathbf{g}_i A y_i\|_{X^*}^2 \right]^{1/2} && \text{(monotonicity - Fact 2.16)} \tag{41} \\ & \text{(since } \sum_i (A y_i)(A y_i)^* = A \mathbb{Y} A^* \succeq \sum_i (A \bar{y}_i)(A \bar{y}_i)^*\text{).} \end{aligned}$$

We are now equipped to relate $\gamma_2(A)$ to the expected value of the rounding algorithm. By (41), we have

$$\begin{aligned} & (1 - o(1)) \cdot \text{Rlx}(A) / (\alpha \cdot \tilde{C}_2(X^*)) \\ & \leq \mathbb{E}_{\mathbf{g}} \left[\|\sum_i \mathbf{g}_i A y_i\|_{X^*}^2 \right]^{1/2} \\ & \lesssim \mathbb{E}_{\mathbf{g}} \left[\|\sum_i \mathbf{g}_i A y_i\|_{X^*} \right] && \text{(Kahane-Khintchine)} \\ & = \mathbb{E}_{\mathbf{g}, \varepsilon} \left[\|\sum_i \varepsilon_i \mathbf{g}_i A y_i\|_{X^*} \right] && \text{(identical distributions)} \\ & = \mathbb{E}_{\mathbf{g}, \varepsilon} \left[\langle x_{\mathbf{g}}, \sum_i \varepsilon_i \mathbf{g}_i A y_i \rangle \right] \\ & = \mathbb{E}_{\mathbf{g}, \varepsilon} \left[\langle A^* x_{\mathbf{g}}, \sum_i \varepsilon_i \mathbf{g}_i y_i \rangle \right] \\ & \leq \mathbb{E}_{\mathbf{g}, \varepsilon} \left[|\langle A^* x_{\mathbf{g}}, \sum_i \varepsilon_i \mathbf{g}_i y_i + \varphi(\mathbf{g}, \varepsilon) \rangle| \right] + \mathbb{E}_{\mathbf{g}, \varepsilon} \left[|\langle A^* x_{\mathbf{g}}, \varphi(\mathbf{g}, \varepsilon) \rangle| \right] && \text{(by triangle inequality)} \end{aligned}$$

$$\begin{aligned}
&\leq \mathbb{E}_{\mathbf{g}, \varepsilon} [|\langle A^* x_{\mathbf{g}}, \sum_i \varepsilon_i \mathbf{g}_i y_i + \varphi(\mathbf{g}, \varepsilon) \rangle|] + \mathbb{E}_{\mathbf{g}, \varepsilon} [\langle A^* x_{\mathbf{g}}, \varphi(\mathbf{g}, \varepsilon) \rangle^2]^{1/2} && \text{(by Jensen's inequality)} \\
&\text{(where } \varphi \text{ is chosen according to Lemma 4.7)} \\
&\leq \mathbb{E}_{\mathbf{g}, \varepsilon} [\|A^* x_{\mathbf{g}}\|_Y \cdot \|\sum_i \varepsilon_i \mathbf{g}_i y_i + \varphi(\mathbf{g}, \varepsilon)\|_{Y^*}] + \mathbb{E}_{\mathbf{g}, \varepsilon} [\|A\varphi(\mathbf{g}, \varepsilon)\|_{X^*}^2]^{1/2} && \text{(Holder's, } x_{\mathbf{g}} \in \text{Ball}(X)) \\
&\leq \mathbb{E}_{\mathbf{g}, \varepsilon} [\|A^* x_{\mathbf{g}}\|_Y \cdot \|\sum_i \varepsilon_i \mathbf{g}_i y_i + \varphi(\mathbf{g}, \varepsilon)\|_{Y^*}] + \delta \cdot \text{Rlx}(A) && \text{(Lemma 4.7)} \\
&\leq \mathbb{E}_{\mathbf{g}, \varepsilon} [\|A^* x_{\mathbf{g}}\|_Y^2]^{1/2} \cdot \mathbb{E}_{\mathbf{g}, \varepsilon} [\|\sum_i \varepsilon_i \mathbf{g}_i y_i + \varphi(\mathbf{g}, \varepsilon)\|_{Y^*}^2]^{1/2} + \delta \cdot \text{Rlx}(A) && \text{(Cauchy-Schwarz)} \\
&\lesssim \mathbb{E}_{\mathbf{g}, \varepsilon} [\|A^* x_{\mathbf{g}}\|_Y^2]^{1/2} \cdot \tilde{C}_2(Y^*) \cdot \log 1/\delta + \delta \cdot \text{Rlx}(A) && \text{(Lemma 4.7)} \\
&= \mathbb{E}_{\mathbf{g}, \varepsilon} [\langle x_{\mathbf{g}}, Ay_{\mathbf{g}} \rangle^2]^{1/2} \cdot \tilde{C}_2(Y^*) \cdot \log 1/\delta + \delta \cdot \text{Rlx}(A) && \text{(definition of } y) \tag{42}
\end{aligned}$$

Finally setting $\delta = 1/(C\alpha\beta)$ for a sufficiently large absolute constant C and rearranging (42) implies that

$$\mathbb{E}_{\mathbf{g}, \varepsilon} [\langle x_{\mathbf{g}}, Ay_{\mathbf{g}} \rangle^2]^{1/2} \geq \gamma_2(A) / O(\alpha\beta \log(\alpha\beta)) \geq \text{Op}_{X,Y}^{\max}(A) / O(\alpha\beta \log(\alpha\beta))$$

Since the random variable $\langle x_{\mathbf{g}}, Ay_{\mathbf{g}} \rangle^2$ lies in the interval $[0, \|A\|_{X \rightarrow Y}^2]$, we may apply [Fact 4.8](#) to it and obtain that with probability at least $\varepsilon / \text{poly}(\alpha\beta)$, $\langle x_{\mathbf{g}}, Ay_{\mathbf{g}} \rangle \geq (1 - \varepsilon) \text{Op}_{X,Y}^{\max}(A) / O(\alpha\beta \log(\alpha\beta))$. Finally, sampling independently for polynomially many rounds and returning the best pair $x_{\mathbf{g}}, y_{\mathbf{g}}$ implies the $1 - 2^{-\Omega(n)}$ success probability claimed in (2). \blacksquare

4.2.3 A Generic Framework for Quadratic/Bilinear Maximization

We are finally equipped to prove our claimed framework theorem.

Theorem 4.11 (Framework: Maximization under Lower Covariance Oracle).

There are algorithms $\text{ALG}_1(A_1, R, r, \mathcal{SO}_X)$ and $\text{ALG}_2(A_2, R, r, \mathcal{SO}_X, \mathcal{SO}_Y)$ such that if $\|\cdot\|_X$ (resp. $\|\cdot\|_Y$) is an (R, r) -balanced norm over \mathbb{R}^n (resp. \mathbb{R}^m) and \mathcal{SO}_X (resp. \mathcal{SO}_Y) is an α -approximate separation oracle for $\mathcal{L}(X)$ (resp. $\mathcal{L}(Y)$) then

- Quadratic: for any $A_1 \in M_n(\mathbb{R})$, ALG_1 runs in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A_1))$ and with probability at least $1 - 2^{-\Omega(n)}$ returns a $(1 + o(1)) \cdot \alpha \cdot T_2(X)^2$ -approximately optimal solution to $Q_X^{\max}(A_1)$.
- Bilinear: for any $A_2 \in M_{n,m}(\mathbb{R})$, ALG_2 runs in time $\text{poly}(n, m, \log R, \log 1/r, \text{bit}(A_2))$ and with probability at least $1 - 2^{-\Omega(n+m)}$, ALG_2 returns a $\beta \log \beta$ -approximately optimal solution to $\text{Op}_{X,Y}^{\max}(A_2)$ where $\beta \lesssim \alpha \cdot \tilde{C}_2(X^*) \cdot \tilde{C}_2(Y^*)$.

Proof. The quadratic (resp. bilinear) claim follows from combining [Theorem 4.5](#) (resp. [Theorem 4.10](#)) with [Proposition 4.4](#). \blacksquare

We note the following useful corollary that illustrates the importance of Khintchine-type inequalities for quadratic/bilinear maximization.

Corollary 4.12 (Khintchine-Type Inequalities Yield Approximation Algorithms).

1. There is an algorithm $\text{ALG}(A, R, r, \mathcal{O}_f)$ such that if $\|\cdot\|_X$ is an (R, r) -balanced norm over \mathbb{R}^n and \mathcal{O}_f is an exact oracle computing a convex function $f : \text{PSID}^n \rightarrow \mathbb{R}$ satisfying

$$\alpha_1^{-1} \cdot f(\mathbb{X}) \leq \mathbb{E}[\|\mathbb{X}^{1/2} \mathbf{g}\|_X^2] \leq \alpha_2 \cdot f(\mathbb{X}) \quad \forall \mathbb{X} \in \text{PSID}^n,$$

then on any input $A \in M_n(\mathbb{R})$, ALG runs in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A))$ and returns an $\alpha_1 \alpha_2 \tilde{T}_2(X)^2$ -approximate solution to $\text{Q}_X^{\max}(A)$ with probability $1 - 2^{-\Omega(n)}$.

2. There are algorithms $\text{ALG}_1(A_1, R, r, \mathcal{O}_f)$ and $\text{ALG}_2(A_2, R, r, \mathcal{O}_f, \mathcal{O}_g)$ such that if $\|\cdot\|_X$ (resp. $\|\cdot\|_Y$) is an (R, r) -balanced norm over \mathbb{R}^n (resp. \mathbb{R}^m) and \mathcal{O}_f (resp. \mathcal{O}_g) is an exact oracle computing a concave function $f : \text{PSID}^n \rightarrow \mathbb{R}$ (resp. $g : \text{PSID}^m \rightarrow \mathbb{R}$) satisfying

$$\begin{aligned} \alpha_1^{-1} \cdot f(\mathbb{W}) &\leq \mathbb{E}[\|\mathbb{W}^{1/2} \mathbf{g}\|_{X^*}^2] \leq \alpha_2 \cdot f(\mathbb{W}) & \forall \mathbb{W} \in \text{PSID}^n \\ \alpha_1^{-1} \cdot g(\mathbb{W}) &\leq \mathbb{E}[\|\mathbb{W}^{1/2} \mathbf{g}\|_{Y^*}^2] \leq \alpha_2 \cdot g(\mathbb{W}) & \forall \mathbb{W} \in \text{PSID}^m, \end{aligned}$$

then

- (A) on any input $A_1 \in M_n(\mathbb{R})$, ALG_1 runs in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A_1))$ and returns an $\alpha_1 \alpha_2 \tilde{T}_2(X)^2$ -approximate solution to $\text{Q}_X^{\max}(A_1)$ with probability $1 - 2^{-\Omega(n)}$,
- (B) on any input $A_2 \in M_{n,m}(\mathbb{R})$, ALG_2 runs in time $\text{poly}(n, m, \log R, \log 1/r, \text{bit}(A_2))$ and returns a $\beta \log \beta$ -approximate solution to $\text{Op}_{X,Y}^{\max}(A_2)$ with probability $1 - 2^{-\Omega(n+m)}$, where $\beta \lesssim \alpha_1 \alpha_2 C_2(X^*) C_2(Y^*)$.

Proof. For claim (1.), we know by [Proposition 4.2](#) that it suffices to give a separation oracle for $\mathcal{U}(X)$. Furthermore, the convex set $\{\mathbb{X} \succeq 0 \mid f(\mathbb{X}) \leq 1\}$ has an exact membership oracle and therefore also a separation oracle. Applying [Observation 3.4](#) yields the desired approximate separation oracle for $\mathcal{U}(X)$.

For claim (2.), we know by [Theorem 4.11](#) that it suffices to give separation oracles for $\mathcal{L}(X), \mathcal{L}(Y)$. Furthermore, the convex set $\{\mathbb{W} \succeq 0 \mid f(\mathbb{W}) \geq 1\}$ has an exact membership oracle and therefore also a separation oracle. Applying [Observation 3.4](#) yields the desired approximate separation oracle for $\mathcal{L}(X)$ and the case for Y is analogous. \blacksquare

4.2.4 Type-2 Equivalence of Quadratic Maximization/Upper Covariance Separation

Below we establish a converse of [Proposition 4.2](#) and thus obtain an algorithm vs. covariance separation oracle duality for Type-2 quadratic maximization.

Proposition 4.13 (Quadratic Maximization Algorithm implies Upper Covariance Separation Oracle). *There is an $(1 + o(1)) \cdot \alpha \cdot \tilde{T}_2(X)^2$ -approximate separation oracle for $\mathcal{U}(X)$ running in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(x))$ where x is the input, assuming access to an α -approximate search oracle \mathcal{O} for quadratic maximization over an (R, r) -balanced norm $(\mathbb{R}^n, \|\cdot\|_X)$.*

Proof. Recall that, by [\(31\)](#), for a PSD matrix A ,

$$\text{Q}_X^{\max}(A) = \max_{\mathbb{W} \in \mathcal{B}_\lambda^{\text{Sym}}(X)} \langle A, \mathbb{W} \rangle = \max_{\mathbb{W} \in \mathcal{B}_\lambda^{\text{Sym}}(X)} \langle A, \mathbb{W} \rangle.$$

Therefore, given an α -approximation search oracle for quadratic maximization, [Theorem 3.14](#) with $K \leftarrow \text{PSID}^n, B \leftarrow \downarrow \mathcal{B}_\lambda^{\text{Sym}}(X)$ implies that $\downarrow \mathcal{B}_\lambda^{\text{Sym}}(X)$ has an $(1 + o(1))\alpha$ -approximate separation oracle. (The balancedness of $\downarrow \mathcal{B}_\lambda^{\text{Sym}}(X)$ follows from [Lemma 2.24](#).)

Then the equivalence between $\downarrow \mathcal{B}_\lambda^{\text{Sym}}(X)$ and $\mathcal{U}(X)$ up to $\tilde{T}_2(X)^2$ ([Observation 2.17](#) and [Observation 3.4](#)) gives a $((1 + o(1)) \cdot \alpha \cdot \tilde{T}_2(X)^2)$ -separation oracle for $\mathcal{U}(X)$ as desired. \blacksquare

4.2.5 Cotype-2 Equivalence of PSD Maximization/Lower Covariance Separation

We next give an analogue of [Proposition 4.13](#) for PSD maximization/lower covariance separation in the presence of dual cotype-2.

Proposition 4.14 (PSD Maximization Algorithm implies Lower Covariance Separation Oracle). *There is a $(1 + o(1)) \cdot \alpha \cdot \tilde{C}_2(X)^2$ -approximate separation oracle for $\mathcal{L}(X)$ running in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(x))$ where x is the input, assuming access to an α -approximate search oracle \mathcal{O} for PSD quadratic maximization over an (R, r) -balanced norm $(\mathbb{R}^n, \|\cdot\|_X)$.*

Proof. We will construct an α -approximate search oracle for $\inf_{W \in B_\lambda^\uparrow(X)} \langle U^*U, W \rangle$ on any input U^*U , and apply [Theorem 3.16](#) with $B \leftarrow B_\lambda^\uparrow(X)$, $K \leftarrow \mathbb{PSD}^n$ to get a $(1 + o(1))\alpha$ -separation oracle for $B_\lambda^\uparrow(X)$. (The inverse balancedness of $B_\lambda^\uparrow(X)$ follows from [Lemma 2.24](#).) Then the equivalence between $B_\lambda^\uparrow(X)$ and $\mathcal{L}(X)$ up to $\tilde{C}_2(X)^2$ ([Observation 2.20](#) and [Observation 3.4](#)) gives a $((1 + o(1)) \cdot \alpha \cdot \tilde{C}_2(X)^2)$ -separation oracle for $\mathcal{L}(X)$ as desired.

Therefore it remains to construct an α -approximate search oracle for $\inf_{W \in B_\lambda^\uparrow(X)} \langle U^*U, W \rangle$ on any input U^*U . To this end observe that

$$\inf_{W \in B_\lambda^\uparrow(X)} \langle U^*U, W \rangle = (\|U\|_{X^* \rightarrow 2}^{\min})^2.$$

We conclude that U is invertible otherwise $\|U\|_{X^* \rightarrow 2}^{\min} = 0$ which contradicts the assumption that $U^*U \in B_\lambda^\uparrow(X)^\circ$. Let v be any vector satisfying $\|v\|_2 \leq \alpha^{1/2} / \|U^{-1}\|_{2 \rightarrow X^*}$ and $\|U^{-1}v\|_{X^*} = 1$.

Such a vector can be found by using \mathcal{O} to obtain u such that $\|u\|_2 = 1$ and $\|U^{-1}u\|_{X^*} \geq \alpha^{-1/2} \cdot \|U^{-1}\|_{2 \rightarrow X^*}$. Then setting $v \stackrel{\text{def}}{=} u / \|U^{-1}u\|_{X^*}$ satisfies the desired properties. Finally taking $y \stackrel{\text{def}}{=} U^{-1}v$ satisfies the conditions $yy^* \in B_\lambda^\uparrow(X)$ and

$$\langle U^*U, yy^* \rangle = \|Uy\|_2^2 = \|v\|_2^2 \leq \alpha \cdot \|U^{-1}\|_{2 \rightarrow X^*}^{-2} = \alpha \cdot (\|U\|_{X^* \rightarrow 2}^{\min})^2.$$

This yields the desired minimization oracle and hence completes our proof. ■

4.2.6 Reducing PSD to Bilinear under Bounded Dual Cotype-2

Proposition 4.15 (PSD reduces to Bilinear under Finite Cotype).

Let $(\mathbb{R}^n, \|\cdot\|_{X^n})_{n \in \mathbb{N}}$ and $(\mathbb{R}^n, \|\cdot\|_{Y^n})_{n \in \mathbb{N}}$ be a sequence of (R, r) -balanced norms satisfying $\sup_n \tilde{C}_2((X^n)^) < \infty$, $\sup_n \tilde{C}_2((Y^n)^*) < \infty$. If there is a family $(\text{ALG}_{n,m})_{n,m \in \mathbb{N}}$ of $\text{poly}(n, m)$ -time search algorithms α -approximating $\text{Op}_{X^n, Y^m}^{\max}(\cdot)$, then there are families of $\text{poly}(n)$ -time search algorithms $4\alpha^2$ -approximating PSD quadratic maximization over (X^n) and over (Y^n) respectively.*

Proof. We give an algorithm for PSD quadratic maximization over (X^n) and the the algorithm for (Y^n) is analogous. It is known through a classical result of Figiel Lindenstrauss Milman [[FLM77](#)] that for any n and $m \gtrsim \tilde{C}_2(Y^m)^2 \cdot n$, there is an $m \times n$ matrix B such that for all $a \in \mathbb{R}^n$,

$$\frac{1}{\sqrt{2}} \cdot \|a\|_{\ell_2^n} \leq \|Ba\|_{(Y^m)^*} \leq \sqrt{2} \cdot \|a\|_{\ell_2^n}. \quad (43)$$

For any matrix B satisfying (43) any $C \in M_n(\mathbb{R})$, we have $\|BC\|_{X^n \rightarrow (Y^m)^*} / \sqrt{2} \leq \|C\|_{X \rightarrow \ell_2^n} \leq \sqrt{2} \|BC\|_{X^n \rightarrow (Y^m)^*}$.

Consider any instance $A = C^*C \in \mathbb{PSID}^n$ of PSD quadratic maximization over X^n . Our algorithm will be run the (Y^m, X^n) bilinear maximization algorithm on BC to obtain a vector $x \in \text{Ball}(X^n)$ satisfying

$$\|BCx\|_{(Y^m)^*} \geq \text{Op}_{Y^m, X^n}^{\max}(BC)/\alpha = \|BC\|_{X^n \rightarrow (Y^m)^*}/\alpha \geq \|C\|_{X \rightarrow \ell_2^n}/(\sqrt{2}\alpha).$$

x is the desired solution since we have

$$\langle x, Ax \rangle = \|Cx\|_{\ell_2^n}^2 \geq \|BCx\|_{(Y^m)^*}^2/2 \geq \|C\|_{X \rightarrow \ell_2^n}^2/(4\alpha^2) = \text{Q}_X^{\max}(A)/(4\alpha^2)$$

where the first inequality follows again from (43). ■

In the next section we discuss how the reductions from quadratic and bilinear maximization to PSD quadratic maximization may be viewed as dual to (algorithmic) factorization theorems, leading to a streamlined exposition of Pisier's abstract factorization theorem (with a new approximation algorithm to obtain such a factorization), as well as a new factorization theorem for quadratic forms. These results are not necessary to obtain the applications promised in [Section 6](#) and [Section 7](#), and readers only interested in applications of [Theorem 4.11](#) may skip [Section 5](#) and proceed to [Section 6](#).

5 Factorization through ℓ_2 via Gaussian Rounding + Convex Programming Duality

For norms $\|\cdot\|_E$ over \mathbb{R}^n , $\|\cdot\|_F$ over \mathbb{R}^m and an linear operator $A : E \rightarrow F$, we define the factorization norm

$$\gamma_2(A) := \inf_{BC=A} \|C\|_{E \rightarrow \ell_2} \cdot \|B\|_{\ell_2 \rightarrow F}. \quad (44)$$

For any factorization $A = BC$ where $C : E \rightarrow \ell_2^d$ and $B : \ell_2^d \rightarrow F$, the distortion of A (i.e., $\|A\|_{E \rightarrow F}$) is trivially upper bounded by the product of distortions (i.e., $\|C\|_{E \rightarrow \ell_2} \cdot \|B\|_{\ell_2 \rightarrow F}$). Thus $\gamma_2(A)$ can be thought of as the best upper bound on the distortion of A when it is viewed as a map from $E \rightarrow \ell_2$ and then from $\ell_2 \rightarrow F$ (there are infinitely many such factorizations). A "factorization theorem" is an inequality of the form

$$\|A\|_{E \rightarrow F} \leq \gamma_2(A) \leq C(E, F) \cdot \|A\|_{E \rightarrow F}.^4$$

i.e., the upper bound on distortion prescribed by the best factorization of A approximates the actual distortion of A . Surprisingly for a wide class of norms, $C(E, F)$ can be taken independent of the dimension. Factorization theory is a fundamental area of research in Banach space theory with many applications.

In this section we discuss how the proofs of powerful theorems on factorization through ℓ_2 (like those of Grothendieck, Maurey and Pisier) can be viewed as having the following components

1. A dual characterization of the factorization norm as a maximization problem (which is a convex but not necessarily computable relaxation of $\|A\|_{E \rightarrow F}$).

⁴Note that the first inequality trivially holds for any E, F .

2. A bound on the “integrality gap” of the maximization problem (i.e., a bound on $\sup_A \gamma_2(A) / \|A\|_{E \rightarrow F}$). Such a bound can be obtained by analyzing an appropriate Gaussian rounding scheme that maps an optimal solution of the maximization problem to a vector $e \in \text{Ball}(E), f \in \text{Ball}(F)$.

While step (1.) above seems to have been proved using sophisticated Hahn-Banach arguments previously⁵, we show how all of these duality results are special cases of conic Lagrangian duality.⁶ We then use this perspective to prove factorization theorems for quadratic maximization which to our knowledge were previously overlooked. En route we also show that the optimal factorization can be computed approximately, assuming access to an oracle for PSD quadratic maximization.

We begin with a discussion of sign-invariant norms (lattices) where factorization takes on a particularly simple form, in order to provide a gentle introduction for a reader encountering these notions for the first time. Readers familiar with factorization theory may wish to skip ahead to [Section 5.2](#).

5.1 Warmup: Sign-Invariant Norms

In this section we discuss factorization in sign-invariant norms (lattices) where factorization takes on a simpler form. We begin by discussing the case of ℓ_∞ and the classical results of Grothendieck.

5.1.1 Grothendieck’s Inequality/Factorization as a Motivating Example

Recall for an $n \times m$ matrix A ,

$$\|A\|_{\infty \rightarrow 1} = \text{Op}_{\infty, \infty}^{\max}(A) = \sup_{|x_i|, |y_j| \leq 1} \sum_{i,j} A_{i,j} \cdot x_i \cdot y_j.$$

Grothendieck’s inequality states that $\|A\|_{\infty \rightarrow 1} \leq \text{SDP}(A) \leq K_G \cdot \|A\|_{\infty \rightarrow 1}$ for an absolute constant K_G ⁷ where the semidefinite programming relaxation $\text{SDP}(A)$ is defined as

$$\begin{aligned} & \text{maximize} && \sum_{i,j} A_{i,j} \cdot \langle u_i, v_j \rangle && \text{s.t.} \\ & \text{subject to} && \|u_i\|_2, \|v_j\|_2 \leq 1 && i \in [n], j \in [m] \\ & && u_1, \dots, u_n, v_1, \dots, v_m \in \mathbb{R}^{m+n} && \end{aligned} \tag{45}$$

and is equivalent to

$$\begin{aligned} & \text{maximize} && \langle A, Z \rangle \\ & \text{subject to} && \mathbb{X}_{i,i}, \mathbb{Y}_{j,j} \leq 1 && i \in [n], j \in [m] \\ & && \begin{bmatrix} \mathbb{X} & Z \\ Z^* & \mathbb{Y} \end{bmatrix} \succeq 0. && \end{aligned} \tag{46}$$

⁵We make no attempt to study the infinite dimensional case in this work as it takes us too far afield of the discussion of algorithms.

⁶also known in the literature as Lagrangian duality of convex programming with generalized inequalities

⁷We use K_G to denote the best such constant.

Grothendieck [Gro53] observed that the above inequality implies a beautiful factorization result; namely for any matrix A there exist $\alpha \in \text{Ball}(\ell_2^n)$, $\beta \in \text{Ball}(\ell_2^m)$ such that $\text{Op}_{\infty, \infty}^{\max}(A) \leq \text{Op}_{2,2}^{\max}(\tilde{A}) \leq K_G \cdot \text{Op}_{\infty, \infty}^{\max}(A)$ where \tilde{A} is obtained by reweighting the rows (resp. columns) of A by $(1/\alpha_i)$ (resp. $(1/\beta_i)$) (Tropp [Tro09] discovered an algorithmic application of such a factorization—specifically he gave an algorithm for column subset selection by using these reweightings as measures of importance of the rows/columns of A).

Stated in a way that better motivates the term “factorization”, $\|A\|_{\infty \rightarrow 1} \leq \gamma_2^{\text{Diag}}(A) \leq K_G \cdot \|A\|_{\infty \rightarrow 1}$ where we define the Grothendieck factorization norm as

$$\gamma_2^{\text{Diag}}(A : \ell_\infty^m \rightarrow \ell_1^n) := \inf_{D_1 B D_2 = A} \|D_2\|_{\infty \rightarrow 2} \cdot \|B\|_{2 \rightarrow 2} \cdot \|D_1\|_{2 \rightarrow 1}$$

where the infimum runs over $n \times m$ matrices B and diagonal matrices D_1, D_2 . Note that the first inequality above is simply a consequence of sub-multiplicativity of the operator norm under composition ($\|BC\|_{Z \rightarrow X} \leq \|C\|_{Z \rightarrow Y} \cdot \|B\|_{Y \rightarrow X}$). Pietsch observed that Grothendieck’s constant is tight for this factorization result by showing the equality $\gamma_2^{\text{Diag}}(A) = \text{SDP}(A)$ using an abstract Hahn-Banach argument (see for e.g. Corollary 23.3 in [Pis12] for an exposition) which in this case involves separation arguments on subsets of $\mathbb{R}^{\mathbb{R}^n}$. Tropp [Tro09] proved this equality by observing that the dual semidefinite program of (46) which is given by

$$\begin{aligned} & \inf (\|t\|_1 + \|s\|_1)/2 \quad \text{s.t.} \\ & \begin{bmatrix} \text{Diag}(s) & -A \\ -A^* & \text{Diag}(t) \end{bmatrix} \succeq 0 \quad s \in \mathbb{R}^n, t \in \mathbb{R}^m \end{aligned} \quad (47)$$

can be easily massaged into the form of $\gamma_2^{\text{Diag}}(A)$. Thus by now Grothendieck’s factorization theorem admits a proof involving only objects of routine occurrence in the optimization literature.

5.1.2 Generalizations to 2-Convex Norms.

Krivine [Kri73] observed that Grothendieck’s inequality/factorization generalizes to a wide class of normed spaces called 2-convex Banach Lattices. For simplicity we specialize the discussion from lattices to the class of sign-invariant norms over \mathbb{R}^n (i.e., norms invariant to flipping signs of the entries) as this class is rich enough to demonstrate the ideas we aim to communicate.

Before stating the factorization theorem, we require a definition. In what follows, for a scalar function $s : \mathbb{R} \rightarrow \mathbb{R}$ and a vector $x \in \mathbb{R}^n$, we use the notation $s(x)$ to denote the vector obtained by entry-wise application of s to x , i.e., $s(x) = (s(x_1), \dots, s(x_n))$. For e.g., $|x|^p$ denotes the vector $(|x_1|^p, \dots, |x_n|^p)$. This notation appears exclusively in Section 7.1.1 and Section 5.1.

Definition 5.1 (2-convexity). *Let X be a sign-invariant norm over \mathbb{R}^n . Then the 2-convexity constant of X , denoted by $M^{(2)}(X)$, is the smallest constant C such that for every finite sequence of vectors (x_i) in X ,*

$$\left\| (\sum_i |x_i|^2)^{1/2} \right\|_X \leq C \cdot (\sum_i \|x_i\|_X^2)^{1/2}$$

where $|\cdot|^2$ and $(\cdot)^{1/2}$ are applied to a vector entry-wise in the left hand expression above. We will say X is exactly 2-convex if $M^{(2)}(X) = 1$.

Theorem 5.2 (Krivine/Grothendieck Factorization Theorem).

For sign-invariant norms X over \mathbb{R}^n and Y over \mathbb{R}^m such that X^ and Y are exactly 2-convex and any operator $A : Y \rightarrow X$, it holds that*

$$\|A\|_{Y \rightarrow X} \leq \gamma_2^{\text{Diag}}(A) \leq K_G \cdot M^{(2)}(X^*) \cdot M^{(2)}(Y) \cdot \|A\|_{Y \rightarrow X}$$

where K_G is Grothendieck's constant and

$$\gamma_2^{\text{Diag}}(A : Y \rightarrow X) := \inf_{D_1 B D_2 = A} \|D_2\|_{Y \rightarrow 2} \cdot \|B\|_{2 \rightarrow 2} \cdot \|D_1\|_{2 \rightarrow X}$$

where the infimum runs over diagonal matrices D_1, D_2 .

Combining this with the known equivalence of 2-convexity and dual cotype-2 (see for e.g. [Pis86]) for Banach lattices yields

$$\gamma_2^{\text{Diag}}(A) \leq O(\tilde{C}_2(X) \tilde{C}_2(Y^*) \log \tilde{C}_2(X) \log \tilde{C}_2(Y^*)) \cdot \|A\|_{Y \rightarrow X}.$$

Here again a proof by Lagrangian duality is available (assuming Grothendieck's inequality as a blackbox). Indeed Nesterov [Nes98] independently reproved Grothendieck's inequality (and Krivine's 2-convex extension) in 1997. In this case one compares the objective $\text{Op}_{X,Y}^{\max}(A)$ to

$$\begin{aligned} & \text{maximize} && \sum_{i,j} A_{i,j} \cdot \langle u_i, v_j \rangle \quad \text{s.t.} \\ & \text{subject to} && \|(\|u_1\|_2, \dots, \|u_n\|_2)\|_X^2 \leq 1 \\ & && \|(\|v_1\|_2, \dots, \|v_m\|_2)\|_Y^2 \leq 1 \\ & && u_1, \dots, u_n, v_1, \dots, v_m \in \mathbb{R}^{m+n} \end{aligned} \quad (48)$$

where the above program is convex precisely when X, Y are exactly 2-convex. Without explicitly noting the connection to factorization theory, Nesterov computed the Lagrangian dual of (48) to be

$$\begin{aligned} & \inf (\text{Op}_{X,X}^{\max}(\text{Diag}(s)) + \text{Op}_{Y,Y}^{\max}(\text{Diag}(t)))/2 \quad \text{s.t.} \\ & \begin{bmatrix} \text{Diag}(s) & -A \\ -A^* & \text{Diag}(t) \end{bmatrix} \succeq 0 \quad s \in \mathbb{R}^n, t \in \mathbb{R}^m. \end{aligned} \quad (49)$$

(49) is readily massaged into the form of $\gamma_2^{\text{Diag}}(A)$ (see lemma A.7. in [BGG⁺19] for a proof).

Simply applying the approximate ellipsoid method to (49) yields a previously overlooked but somewhat surprising result which states that in the case of sign-invariant norms whose duals have bounded cotype-2 constant, diagonal instances are the bottleneck for bilinear maximization:

Proposition 5.3 (Algorithm for Cotype-2 Lattice Duals assuming Oracle for Diagonal Instances). *There is an algorithm $\text{ALG}(A, \mathcal{O}_X, \mathcal{O}_Y)$ running in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A))$ that on any input $A \in M_{n,m}(\mathbb{R})$ returns a C -approximation to $\text{Op}_{X,Y}^{\max}(\cdot)$ assuming access to an α -approximate search oracle \mathcal{O}_X (resp. \mathcal{O}_Y) for quadratic maximization of diagonal instances over X (resp. over Y), where*

$$C \stackrel{\text{def}}{=} \alpha K_G M^{(2)}(X) M^{(2)}(Y) = O(\alpha \cdot C_2(X^*) \log C_2(X^*) \cdot C_2(Y^*) \log C_2(Y^*)).$$

Proof. Using the diagonal-instance oracles, we will exhibit α -approximate separation oracles for the sets

$$S_X \stackrel{\text{def}}{=} \{s \in \mathbb{R}^n \mid s \geq 0, \text{Op}_{X,X}^{\max}(\text{Diag}(s)) \leq 1\}, \quad S_Y \stackrel{\text{def}}{=} \{t \in \mathbb{R}^m \mid t \geq 0, \text{Op}_{Y,Y}^{\max}(\text{Diag}(t)) \leq 1\}.$$

Then Proposition 3.12 with

$$C_1 \leftarrow \left\{ (s, t) : \begin{bmatrix} \text{Diag}(s) & -A \\ -A^* & \text{Diag}(t) \end{bmatrix} \succeq 0 \right\} \text{ and } f(s, t) \leftarrow (\text{Op}_{X,X}^{\max}(\text{Diag}(s)) + \text{Op}_{Y,Y}^{\max}(\text{Diag}(t)))/2$$

implies that we can compute an $(1 + o(1))\alpha$ -approximation to (49), which by [Theorem 5.2](#) implies the desired approximation to $\text{Op}_{X,Y}^{\max}(\cdot)$.

Now we construct α -approximate separation oracles for S_X and S_Y . We will construct an oracle for S_X and the case of S_Y will be analogous. To this end consider any $s \in \mathbb{R}^n$. Let D be the matrix containing s in the diagonal, and let x be the output of $\mathcal{O}_X(D)$. If $\langle D, xx^* \rangle \leq 1$ we return “Inside” and if $\langle D, xx^* \rangle > 1$ we return $\{y \mid \langle y, x \rangle^2 = 1\}$ as a hyperplane separating s from S_X . Indeed since $x \in \text{Ball}(X)$, we have $\langle y, x \rangle^2 \leq 1$ for any $y \in S_X$. Lastly since $\langle D, xx^* \rangle \leq \text{Op}_{X,X}^{\max}(D) \leq \alpha \cdot \langle D, xx^* \rangle$, it is easily checked that the above scheme satisfies all conditions of an α -approximate separation oracle for S_X . ■

It turns out that factorization results with weaker structure (i.e., factorization through ℓ_2) are available for much more general classes of norms (i.e., without any sign-invariance assumptions like in the case of Grothendieck’s factorization). We discuss such factorization results in the next section.

5.2 Factorization Through ℓ_2 without Lattice Assumptions

For norms X over \mathbb{R}^n , Y over \mathbb{R}^m and an operator $A : Y \rightarrow X$, we define the factorization norm

$$\gamma_2(A) := \inf_{BC=A} \|C\|_{Y \rightarrow 2} \cdot \|B\|_{2 \rightarrow X}. \quad (50)$$

Maurey [[Mau74](#)] (extending a result of Kwapien) proved a powerful factorization result

Theorem 5.4 (Maurey Factorization Theorem).

$$\|A\|_{Y \rightarrow X} \leq \gamma_2(A) \leq T_2(Y) \cdot C_2(X) \cdot \|A\|_{Y \rightarrow X}.$$

However in light of Grothendieck’s factorization theorem for $A : \ell_\infty^m \rightarrow \ell_1^n$, one expects a dependence on $\tilde{C}_2(Y^*)$ instead of $\tilde{T}_2(Y)$ above. An important result of Pisier [[Pis80](#)] bridges Grothendieck’s and Maurey’s factorization theorems:

Theorem 5.5 (Pisier’s Factorization Theorem).

$$\|A\|_{Y \rightarrow X} \leq \gamma_2(A) \leq O(C_2(X^*)C_2(Y) \log C_2(X^*) \log C_2(Y)) \cdot \|A\|_{X \rightarrow Y}.$$

Previously constant factor approximation algorithms for bilinear maximization were known only for norms admitting an analogue of Grothendieck’s inequality; this is the case for exactly 2-convex norms which simply extend the classical Grothendieck inequality and this is also the case for Schatten- ∞ (max singular value) where Naor, Regev, and Vidick [[NRV13](#)] algorithmicised Haagerup’s proof of the celebrated non-commutative Grothendieck inequality. Curiously, both the classical and non-commutative Grothendieck inequalities are equivalent to results about “factorization through ℓ_2 ” wherein the factorization satisfies additional structural properties.

We obtain constant factor approximation algorithms for a broader class of norms by utilizing Pisier’s factorization theorem which applies to a much more general class of norms than those for which Grothendieck-type inequalities are available. This demonstrates that for the purpose of algorithm design, the factorizations need not have additional structure.

We next show how $\gamma_2(\cdot)$ can be cast as a convex program. We then apply conic Lagrangian duality to this formulation to show that the best constant in a factorization theorem is equal to the integrality gap of an appropriate convex relaxation of $\|A\|_{E \rightarrow F}$.

5.2.1 Convex Programming Formulation of $\gamma_2(\cdot)$

In this section, we show that $\gamma_2(\cdot)$ can be written as a convex program. As an additional benefit, we can use the approximate ellipsoid method to show that the optimal factorization can be approximately computed using an oracle for quadratic maximization of PSD instances.

Let A be an $n \times m$ matrix. We will show that the following convex program is a reformulation of $\gamma_2(A : Y \rightarrow X^*)$:

$$\begin{aligned} \inf \quad & (Q_X^{\max}(W_1) + Q_Y^{\max}(W_2))/2 \quad \text{s.t.} \\ & \begin{bmatrix} W_1 & -A \\ -A^* & W_2 \end{bmatrix} \succeq 0 \quad W_1 \in M_n(\mathbb{R}), W_2 \in M_m(\mathbb{R}) \end{aligned} \quad (51)$$

which can be equivalently written as

$$\begin{aligned} \inf \quad & (\|U\|_{Y \rightarrow 2}^2 + \|V^*\|_{2 \rightarrow X^*}^2)/2 \quad \text{s.t.} \\ & \begin{bmatrix} UU^* & -A \\ -A^* & VV^* \end{bmatrix} \succeq 0 \quad U \in M_n(\mathbb{R}), V \in M_m(\mathbb{R}). \end{aligned} \quad (52)$$

Lemma 5.6. $\gamma_2(A : Y \rightarrow X^*) = (52)$.

Proof. We begin with the observation that

$$\inf_{BEC=A} \|C\|_{Y \rightarrow 2} \cdot \|E\|_{2 \rightarrow 2} \cdot \|B\|_{2 \rightarrow X^*} = \inf_{BC=A} \|C\|_{Y \rightarrow 2} \cdot \|B\|_{2 \rightarrow X^*}$$

where LHS \geq RHS follows since $\|EC\|_{Y \rightarrow 2} \leq \|C\|_{Y \rightarrow 2} \cdot \|E\|_{2 \rightarrow 2}$ and LHS \leq RHS follows by substituting $E = I$. Consequently we also know

$$\gamma_2(A) = \inf_{\substack{BEC=A \\ \|E\|_{2 \rightarrow 2} \leq 1 \\ \|B\|_{2 \rightarrow X^*} = \|C\|_{Y \rightarrow 2}}} \|C\|_{Y \rightarrow 2} \cdot \|B\|_{2 \rightarrow X^*} \quad (53)$$

since whenever $A = BEC$ we also have $A = B'E'C'$ where

$$E' \stackrel{\text{def}}{=} \frac{E}{\|E\|_{2 \rightarrow 2}}, \quad B' \stackrel{\text{def}}{=} \sqrt{\frac{\|E\|_{2 \rightarrow 2} \|C\|_{Y \rightarrow 2}}{\|B\|_{2 \rightarrow X^*}}} \cdot B, \quad C' \stackrel{\text{def}}{=} \sqrt{\frac{\|E\|_{2 \rightarrow 2} \|B\|_{2 \rightarrow X^*}}{\|C\|_{Y \rightarrow 2}}} \cdot C$$

By AM-GM inequality we have

$$\gamma_2(A) = \inf_{\substack{BEC=A \\ \|E\|_{2 \rightarrow 2} \leq 1}} \|C\|_{Y \rightarrow 2} \cdot \|B\|_{2 \rightarrow X^*} \leq \inf_{\substack{BEC=A \\ \|E\|_{2 \rightarrow 2} \leq 1}} (\|C\|_{Y \rightarrow 2}^2 + \|B\|_{2 \rightarrow X^*}^2)/2. \quad (54)$$

Combining (54) with (53) yields

$$\gamma_2(A) = \inf_{\substack{BEC=A \\ \|E\|_{2 \rightarrow 2} \leq 1}} (\|C\|_{Y \rightarrow 2}^2 + \|B\|_{2 \rightarrow X^*}^2)/2. \quad (55)$$

Consider an optimal solution to (52). We will show (55) \leq (52). We may assume without loss of generality that U and V are invertible since for any $\varepsilon > 0$, adding $(\varepsilon/\max\{\|I\|_{Y \rightarrow 2}, \|I\|_{2 \rightarrow X^*}\}) \cdot I$ to the block matrix in (52) maintains feasibility while increasing the objective value by no more than ε . Thus the infimum in (52) may be taken over invertible U and V without changing the

optimum value. We may now make the following substitution: $B \stackrel{\text{def}}{=} U$, $C \stackrel{\text{def}}{=} V^*$ and $E \stackrel{\text{def}}{=} U^{-1}A(V^{-1})^*$.

Since $BCE = A$, it remains to show that $\|E\|_{2 \rightarrow 2} \leq 1$, and indeed we have,

$$\begin{aligned}
& \begin{bmatrix} UU^* & -A \\ -A^* & VV^* \end{bmatrix} \succeq 0 \\
\Leftrightarrow & \begin{bmatrix} U^{-1} & 0 \\ 0 & V^{-1} \end{bmatrix} \begin{bmatrix} UU^* & -A \\ -A^* & VV^* \end{bmatrix} \begin{bmatrix} (U^{-1})^* & 0 \\ 0 & (V^{-1})^* \end{bmatrix} \succeq 0 \\
\Leftrightarrow & \begin{bmatrix} I & -E \\ -E^* & I \end{bmatrix} \succeq 0 \\
\Leftrightarrow & \|E\|_{2 \rightarrow 2} \leq 1.
\end{aligned} \tag{56}$$

Finally, we need to show (52) \leq (55). So fix any optimal solution to (55) and by the same reasoning as above, we may assume B and C are invertible. We then make the substitution $U \stackrel{\text{def}}{=} B$, and $V \stackrel{\text{def}}{=} C^*$. Since $BEC = A$, we must have $E := U^{-1}A(V^{-1})^*$. Feasibility of U, V then follows from the chain of equivalences in (56). ■

5.2.2 An Alternate Proof of a Dual Characterization of $\gamma_2(A)$

A classical Hahn-Banach result of Lindenstrauss and Pelczynski [LP68] (see for e.g. Theorem 2.4 in [Pis86] or Theorem 7.3.4 in [AK06] for a detailed exposition) gives a dual characterization of $\gamma_2(A : Y \rightarrow X^*)$ as the smallest constant C such that for all finite sequences $(y_i), (\bar{y}_j)$ satisfying $\sum_j \bar{y}_j \bar{y}_j^* \succeq \sum_i y_i y_i^*$, it holds that $\sum_i \|A y_i\|_{X^*}^2 \leq C^2 \cdot \sum_j \|\bar{y}_j\|_Y^2$. A dual reformulation (see chapter 2 in [Pis86]) of this statement is that the dual norm of $\gamma_2(\cdot)$, which we will denote by $\gamma_2^*(\cdot)$, is given by

$$\gamma_2^*(Z) \stackrel{\text{def}}{=} \inf_{UVW^*=Z, \|V\|_{2 \rightarrow 2} \leq 1} \sqrt{\sum_i \|u_i\|_{X^*}^2} \cdot \sqrt{\sum_i \|w_i\|_Y^2} \tag{57}$$

where u_i (resp. w_i) is the i -th column of U (resp. W)

In this section we give an alternate (perhaps more direct) proof of (57) using conic Lagrangian duality.

Lemma 5.7 (Dual Characterization of $\gamma_2(\cdot)$).

For any linear operator $A : Y \rightarrow X^*$, $\gamma_2(A) = (38)$ where (38) is defined as

$$\begin{aligned}
& \max \quad \langle A, Z \rangle \\
& \text{s.t.} \quad \mathbb{X} \in \downarrow B_{\lambda}^{\text{Sym}}(X), \mathbb{Y} \in \downarrow B_{\lambda}^{\text{Sym}}(Y) \\
& \quad \begin{bmatrix} \mathbb{X} & Z \\ Z^* & \mathbb{Y} \end{bmatrix} \succeq 0.
\end{aligned}$$

Proof. We will apply Lagrangian duality to the formulation (51) of $\gamma_2(A)$. Indeed by Lagrangian duality for generalized (conic) inequalities (see theorem 2 on page 224 of [Lue97]; see also section 5.9 in [BV04]), we have

$$\inf_{(W_1, W_2)} \left\{ Q_X^{\max}(W_1) + Q_Y^{\max}(W_2) \mid \begin{bmatrix} W_1 & -A \\ -A^* & W_2 \end{bmatrix} \succeq 0 \right\} \tag{58}$$

$$= \max_{(\mathbb{X}, \mathbb{Y}, Z)} \left\{ \inf_{W_1, W_2 \succeq 0} \{Q_X^{\max}(W_1) + Q_Y^{\max}(W_2) + 2\langle A, Z \rangle - \langle W_1, \mathbb{X} \rangle - \langle W_2, \mathbb{Y} \rangle\} \mid \begin{bmatrix} \mathbb{X} & Z \\ Z^* & \mathbb{Y} \end{bmatrix} \succeq 0 \right\}.$$

where Slater's condition and therefore strong duality can be verified by the substitution $W_1 = \|A\|_{2 \rightarrow 2} \cdot I$, $W_2 = \|A\|_{2 \rightarrow 2} \cdot I$. By (31) we have

$$\inf_{W_1 \succeq 0} \{Q_X^{\max}(W_1) - \langle W_1, \mathbb{X} \rangle\} = \begin{cases} 0 & \text{if } \mathbb{X} \in \downarrow B_{\wedge}^{\text{Sym}}(X) \\ -\infty, & \text{otherwise.} \end{cases}$$

Thus it follows that

$$\begin{aligned} & \inf_{W_1, W_2 \succeq 0} \{Q_X^{\max}(W_1) + Q_Y^{\max}(W_2) + 2\langle A, Z \rangle - \langle W_1, \mathbb{X} \rangle - \langle W_2, \mathbb{Y} \rangle\} \\ &= \begin{cases} 2\langle Z, A \rangle & \text{if } \mathbb{X} \in \downarrow B_{\wedge}^{\text{Sym}}(X), \mathbb{Y} \in \downarrow B_{\wedge}^{\text{Sym}}(Y) \\ -\infty, & \text{otherwise.} \end{cases} \end{aligned}$$

Combining this with (58) yields the claim:

$$\begin{aligned} & \inf_{(W_1, W_2)} \left\{ (Q_X^{\max}(W_1) + Q_Y^{\max}(W_2))/2 \mid \begin{bmatrix} W_1 & -A \\ -A^* & W_2 \end{bmatrix} \succeq 0 \right\} \\ &= \max_{(\mathbb{X}, \mathbb{Y}, Z)} \left\{ \langle A, Z \rangle \mid \mathbb{X} \in \downarrow B_{\wedge}^{\text{Sym}}(X), \mathbb{Y} \in \downarrow B_{\wedge}^{\text{Sym}}(Y), \begin{bmatrix} \mathbb{X} & Z \\ Z^* & \mathbb{Y} \end{bmatrix} \succeq 0 \right\}. \end{aligned}$$

■

Pisier's factorization theorem [Pis80] is now readily deduced from combining Lemma 5.7 with the fact that (38) approximates $\text{Op}_{X,Y}^{\max}(A)$ (which is implicitly shown in Theorem 4.10 via Gaussian rounding).

Theorem 5.8 (Algorithmic Pisier Factorization Theorem).

1. For any norms $(\|\cdot\|_X, \mathbb{R}^n)$, $(\|\cdot\|_Y, \mathbb{R}^m)$ and any linear map $A : Y \rightarrow X^*$, we have

$$\text{Op}_{X,Y}^{\max}(A) \leq \gamma_2(A : Y \rightarrow X^*) \leq O(\tilde{C}_2(X^*)\tilde{C}_2(Y^*) \log \tilde{C}_2(X^*)\tilde{C}_2(Y^*)) \cdot \text{Op}_{X,Y}^{\max}(A).$$

2. There is an algorithm $\text{ALG}(A, R, r, \mathcal{O}_X, \mathcal{O}_Y)$ such that if \mathcal{O}_X (resp. \mathcal{O}_Y) is an α -approximate search oracle for PSD quadratic maximization over an (R, r) -balanced norm $(\|\cdot\|_X, \mathbb{R}^n)$ (resp. $(\|\cdot\|_Y, \mathbb{R}^m)$), then on any input $A \in M_{n,m}(\mathbb{R})$, ALG runs in time $\text{poly}(m, n, \log R, \log 1/r, \text{bit}(A))$ and returns a factorization $A = BC$ satisfying

$$\gamma_2(A : Y \rightarrow X^*) \leq \|C\|_{Y \rightarrow 2} \cdot \|B\|_{2 \rightarrow X^*} \leq (1 + o(1)) \cdot \alpha \cdot \gamma_2(A : Y \rightarrow X^*).$$

Proof. Claim (1.) (i.e., Pisier's factorization theorem) follows from combining Lemma 5.7 (dual characterization of the factorization norm) with the fact that (38) approximates $\text{Op}_{X,Y}^{\max}(A)$ (which is implicitly shown in Theorem 4.10 via Gaussian rounding).

For claim (2.), we will use the search oracles to construct an α -approximate separation oracle for the following sets.

$$S_X \stackrel{\text{def}}{=} \{W \succeq 0 \mid Q_X^{\max}(W) \leq 1\}, \quad S_Y \stackrel{\text{def}}{=} \{W \succeq 0 \mid Q_Y^{\max}(W) \leq 1\}.$$

Then [Proposition 3.12](#) with

$$C_1 \leftarrow \left\{ \begin{bmatrix} W_1 & -A \\ -A^* & W_2 \end{bmatrix} : W \succeq 0 \right\} \text{ and } f \left(\begin{bmatrix} W_1 & -A \\ -A^* & W_2 \end{bmatrix} \right) \leftarrow Q_X^{\max}(W_1) + Q_X^{\max}(W_2)$$

implies that we can compute a $(1 + o(1))\alpha$ -approximation to [\(52\)](#), whose optimal value is equal to $\gamma_2(A : Y \rightarrow X^*)$.

It remains to exhibit approximate separation oracles for S_X and S_Y . We will construct an oracle for S_X and the case of S_Y will be analogous. To this end consider any $W \in M_n(\mathbb{R})$. We may assume W is PSD, since otherwise one can use the separation oracle of the cone of PSD matrices. Let x be the output of $\mathcal{O}_X(W)$. If $\langle W, xx^* \rangle \leq 1$ we return "Inside" and if $\langle W, xx^* \rangle > 1$ we return $\{B \mid \langle B, xx^* \rangle = 1\}$ as a hyperplane separating W from S_X . Indeed since $x \in \text{Ball}(X)$, we have $\langle B, xx^* \rangle \leq 1$ for any $B \in S_X$. Lastly since $\langle W, xx^* \rangle \leq \|W\|_{X \rightarrow X^*} \leq \alpha \cdot \langle W, xx^* \rangle$, it is easily checked that the above scheme satisfies all conditions of an α -approximate separation oracle for S_X . ■

5.3 Factorization Theorem for Quadratic Maximization under Bounded Type-2

In this section we apply the framework of "Gaussian rounding + conic Lagrangian duality" to obtain a new factorization theorem for quadratic maximization under bounded Type-2. Here also we are able to approximately compute the best factorization relative to an oracle for PSD instances.

To do this we first define an appropriate analogue of $\gamma_2(\cdot)$ in the quadratic case and then give a dual characterization of it.

For a norm $(\|\cdot\|_X, \mathbb{R}^n)$, and a symmetric $n \times n$ matrix A , we define the quadratic factorization semi-norm as

$$\begin{aligned} \gamma_2^Q(A) &:= \inf_{B^*CB=A} \|B\|_{X \rightarrow \ell_2^n} \cdot \lambda_{\max}(C) \cdot \|B^*\|_{\ell_2^n \rightarrow X^*} \\ &= \inf_{B^*CB=A} \|B\|_{X \rightarrow \ell_2^n}^2 \cdot \lambda_{\max}(C) \\ &= \inf_W \{Q_X^{\max}(W) \mid W \succeq A, W \succeq 0\} \end{aligned} \tag{59}$$

where the proof of the third equality is a simpler version of the proof of [Lemma 5.6](#). Analogous to the bilinear (operator norm) case, conic Lagrangian duality yields a dual characterization of $\gamma_2^Q(\cdot)$.

Lemma 5.9 (Dual Characterization of $\gamma_2^Q(\cdot)$).

For any norm $(\|\cdot\|_X, \mathbb{R}^n)$ and any symmetric $n \times n$ matrix A , we have

$$\gamma_2^Q(A) = \max_{\mathbb{X}} \{ \langle A, \mathbb{X} \rangle \mid \mathbb{X} \in \downarrow B_{\lambda}^{\text{sym}}(X) \}$$

Proof. By Lagrangian duality for generalized (conic) inequalities (see theorem 2 on page 224 of [\[Lue97\]](#); see also section 5.9 in [\[BV04\]](#)), we have

$$\begin{aligned} &\inf_W \{Q_X^{\max}(W) \mid W \succeq A, W \succeq 0\} \\ &= \max_{(\mathbb{X}, \mathbb{M})} \left\{ \inf_{W \succeq 0} \{Q_X^{\max}(W) + \langle A, \mathbb{X} \rangle - \langle W, \mathbb{X} \rangle - \langle W, \mathbb{M} \rangle\} \mid \mathbb{X}, \mathbb{M} \succeq 0 \right\}. \end{aligned} \tag{60}$$

where Slater's condition and therefore strong duality can be verified by the substitution $W = \lambda_{\max}(A) \cdot I$. By (31) we have

$$\inf_{W \succeq 0} \{Q_X^{\max}(W) - \langle W, X + M \rangle + \langle A, X \rangle\} = \begin{cases} \langle A, X \rangle & \text{if } (X + M) \in \downarrow B_{\wedge}^{\text{Sym}}(X) \\ -\infty, & \text{otherwise.} \end{cases}$$

Substituting this in (60) yields

$$\begin{aligned} & \max_{(X, M)} \left\{ \inf_{W \succeq 0} \{Q_X^{\max}(W) + \langle A, X \rangle - \langle W, X + M \rangle\} \mid X, M \succeq 0 \right\} \\ &= \max_{X, M \succeq 0} \{\langle A, X \rangle \mid X + M \in \downarrow B_{\wedge}^{\text{Sym}}(X)\} \\ &= \max_{X \succeq 0} \{\langle A, X \rangle \mid X \in \downarrow B_{\wedge}^{\text{Sym}}(X)\} \end{aligned}$$

where the final equality follows since $\wedge_X^{\downarrow \text{Sym}}(\cdot)$ is monotone in the Loewner ordering (equivalently $\downarrow B_{\wedge}^{\text{Sym}}(X)$ is downward-closed w.r.t. PSID^n). This completes the proof. \blacksquare

Combining Lemma 5.9 with Theorem 4.5 yields

Theorem 5.10. *For any norm $(\|\cdot\|_X, \mathbb{R}^n)$ we have*

1. *For any symmetric $n \times n$ matrix A ,*

$$Q_X^{\max}(A) \leq \gamma_2^Q(A) \leq \tilde{T}_2(X)^2 \cdot Q_X^{\max}(A).$$

2. *There is an algorithm $\text{ALG}(A, R, r, \mathcal{O})$ such that if \mathcal{O} is an α -approximate search oracle for PSD quadratic maximization over an (R, r) -balanced norm $(\|\cdot\|_X, \mathbb{R}^n)$, then on any input $A \in M_n(\mathbb{R})$, ALG runs in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A))$ and returns a factorization $A = B^*CB$ satisfying*

$$\gamma_2^Q(A) \leq \|B\|_{X \rightarrow 2}^2 \cdot \lambda_{\max}(C) \leq (1 + o(1)) \cdot \alpha \cdot \gamma_2^Q(A).$$

Proof. The first claim follows from combining (60) (dual characterization of the factorization norm) with the fact that (37) is a $\tilde{T}_2(X)^2$ -approximation to $Q_X^{\max}(A)$. The latter fact can be proved via Gaussian rounding (i.e., combining Observation 2.17 and Observation 4.1).

For the second claim, we will use the oracle \mathcal{O} to construct an α -approximate separation oracle for $S_X = \{W \succeq 0 \mid Q_X^{\max}(W) \leq 1\}$ (just as in the proof of Theorem 5.8). Then Proposition 3.12 with

$$C_1 \leftarrow \{W \mid W \succeq A, W \succeq 0\} \text{ and } f(W) = Q_X^{\max}(W)$$

implies that we can compute a $(1 + o(1))\alpha$ -approximation to (59). \blacksquare

6 Algorithmic Closure Properties of Quadratic/Bilinear Maximization

In this section we show that under bounded type-2 (resp. bounded dual cotype-2), constant factor approximability of quadratic (resp. bilinear) maximization over X (resp. X, Y) is preserved under the following operations: (a) Minkowski Sum (b) Intersection (c) Subspaces (d) Quotients (e) Interpolation (resp. (a) Minkowski Sum (b) Quotients). Formally, we define T_2 -Quad-Apx (resp. C_2 -Bi-Apx) as the set of norm sequences $(\mathbb{R}^n, \|\cdot\|_{X^n})$ (resp. $(\mathbb{R}^n, \|\cdot\|_{X^n}), (\mathbb{R}^m, \|\cdot\|_{Y^m})$) satisfying the following properties

(C1) $\sup_n \tilde{T}_2(X^n) < \infty$ (resp. $\sup_n \tilde{C}_2((X^n)^*), \sup_m \tilde{C}_2((Y^m)^*) < \infty$).

(C2) X^n (resp. X^n, Y^n) is $(n^{O(1)}, 1/n^{O(1)})$ -balanced.

(C3) There is a $\text{poly}(n)$ time algorithm computing $\|\cdot\|_{X^n}$ (resp. there is a family of $\text{poly}(n)$ -time algorithms $\|\cdot\|_{X^n}, \|\cdot\|_{Y^n}$) within a factor of $1 + 1/\text{poly}(n)$.

(C4) For an absolute constant $C \geq 1$, there is a $\text{poly}(n, \text{bit}A)$ time randomized algorithm that on any $n \times n$ input matrix A , with probability at least $1 - n^{-\omega(1)}$ returns a C -approximate solution $x \in \text{Ball}(X^n)$ to $\text{Q}_{X^n}^{\max}(A)$ (resp. there is a family $(\text{ALG}_{n,m})_{n,m \in \mathbb{N}}$ of $\text{poly}(n, m, \text{bit}A)$ -time search algorithms returning a C -approximate solution $(x, y) \in \text{Ball}(X^n) \times \text{Ball}(Y^m)$ to $\text{Op}_{X^n, Y^m}^{\max}(A)$ given an $n \times m$ input matrix A).

The main result of this section is the following theorem (notation and definitions are given in the sequel):

Theorem 6.1 (Algorithmic Closure Properties).

T_2 -Quad-Apx is closed under the following operations:

(1a) Minkowski Sum: $(X^n), (\bar{X}^n) \mapsto (X^n + \bar{X}^n)$.

(1b) Intersection: $(X^n), (\bar{X}^n) \mapsto (X^n \vee \bar{X}^n)$.

(1c) Subspaces: $(X^n), (E_n) \mapsto (E_n, \|\cdot\|_{X^n})$.

(1d) Quotients: $(X^n), (E_n) \mapsto (X^n/E_n)$.

(1e) Complex Interpolation: $(\mathbb{C}^n, \|\cdot\|_{X^n}), (\mathbb{C}^n, \|\cdot\|_{\bar{X}^n}) \mapsto ([X^n, \bar{X}^n]_\theta)$.

C_2 -Bi-Apx is closed under the following operations:

(2a) Minkowski Sum: $((X^n), (Y^m)), ((\bar{X}^n), (\bar{Y}^m)) \mapsto (X^n + \bar{X}^n), (Y^m + \bar{Y}^m)$.

(2b) Quotients: $((X^n), (Y^m)), (E_n), (F_m) \mapsto (X^n/E_n), (Y^m/F_m)$.

Remark 6.2.

- Above we assume we are given a polytime algorithm returning a generator matrix for any subspace in the sequence (E_n) (resp. (F_m)).
- An interesting point to note above is that while bilinear maximization admits closure properties under the milder assumption that $\tilde{C}_2(X^*)$ is bounded, quadratic maximization enjoys more closure properties albeit under the stronger assumption of bounded $\tilde{T}_2(X)$. Indeed bounded $\tilde{C}_2(X^*)$ bilinear maximization is only closed under Minkowski-sum/quotients whereas bounded $\tilde{T}_2(X)$ quadratic maximization is closed under Minkowski-sum/quotients as well as the duals of these operations (namely intersection/subspaces). Thus the bounded type-2 assumption is more robust for closure properties.
- (1c) and (2b) above are easy to establish and in fact hold without any type-2/cotype-2 assumptions.

Closure of (C1), (C2) and (C3) under the above operations is immediate (except in the case of complex interpolation where closure of (C2) requires substantial work and was shown in [ANN⁺18]; closure of (C1) under complex interpolation is well known and follows from Riesz-Thorin interpolation). Thus we need only be concerned with closure of (C4) under all of the above operations. We proceed to prove Theorem 6.1 one operation at a time.

6.1 Minkowski Sums

For norms $\|\cdot\|_{X_1}, \|\cdot\|_{X_2}$ over \mathbb{R}^n , we denote by $X_1 + X_2$ and $X_1 \vee X_2$ the norms

$$\|x\|_{X_1+X_2} \stackrel{\text{def}}{=} \inf_{y+z=x} \|y\|_{X_1} + \|z\|_{X_2}, \quad \|x\|_{X_1 \vee X_2} \stackrel{\text{def}}{=} \max\{\|x\|_{X_1}, \|x\|_{X_2}\}$$

which respectively correspond to the Minkowski sum and intersection of $\text{Ball}(X_1), \text{Ball}(X_2)$. It is straightforward to check the duality relation $(X_1 + X_2)^* = X_1^* \vee X_2^*$.

It is easily verified that $\tilde{T}_2(X_1 + X_2) \leq \sqrt{2} \cdot \max\{\tilde{T}_2(X_1), \tilde{T}_2(X_2)\}$ and so boundedness of type-2 is preserved under Minkowski sum. Similarly $\tilde{C}_2((X_1 + X_2)^*) = \tilde{C}_2(X_1^* \vee X_2^*) \leq \sqrt{2} \cdot \max\{\tilde{C}_2(X_1^*), \tilde{C}_2(X_2^*)\}$ and so boundedness of dual cotype-2 is preserved under Minkowski sum. In this section we show that approximation algorithms for quadratic (resp. bilinear) maximization are also preserved under Minkowski sum when type-2 (resp. dual cotype-2) is bounded.

Our approach is to show the claim for PSD maximization first (which is immediate) and then appeal to our reduction theorem from quadratic (resp. bilinear) to PSD maximization to obtain the general case. Indeed the following fact is easy to check.

$$Q_{X_1+X_2}^{\max}(BB^*) = \|B\|_{2 \rightarrow X_1^* \vee X_2^*}^2 = \max\{\|B\|_{2 \rightarrow X_1^*}^2, \|B\|_{2 \rightarrow X_2^*}^2\} = \max\{Q_{X_1}^{\max}(BB^*), Q_{X_2}^{\max}(BB^*)\}. \quad (61)$$

Combining (61) and [Theorem 4.5](#) immediately yields

Proposition 6.3 (Closure of Type-2 Quadratic Maximization under Minkowski Sum).

There is an algorithm $\text{ALG}(A, R, r, \mathcal{O}_{X_1}, \mathcal{O}_{X_2})$ such that if \mathcal{O}_{X_1} (resp. \mathcal{O}_{X_2}) is an α -approximate search oracle for quadratic maximization over an (R, r) -balanced norm $(\mathbb{R}^n, \|\cdot\|_{X_1})$ (resp. $(\mathbb{R}^n, \|\cdot\|_{X_2})$), then on any input $A \in M_n(\mathbb{R})$, ALG runs in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A))$ and returns a β -approximate solution to $Q_{X_1+X_2}^{\max}(A)$ with probability $1 - 2^{-\Omega(n)}$, where $\beta \stackrel{\text{def}}{=} 2\alpha \cdot \max\{\tilde{T}_2(X_1)^2, \tilde{T}_2(X_2)^2\}$.

The bilinear version follows from combining [Proposition 4.15](#), [Theorem 4.11](#), and (61):

Proposition 6.4 (Closure of Dual Cotype-2 Bilinear Maximization under Minkowski Sum).

Consider any pairs of norm sequences $((X^n), (Y^m)) \in C_2\text{-Bi-Apx}$ and $((\bar{X}^n), (\bar{Y}^m)) \in C_2\text{-Bi-Apx}$ (see [Section 6](#) for the definition of $C_2\text{-Bi-Apx}$). Then $((X^n + \bar{X}^n), (Y^m + \bar{Y}^m)) \in C_2\text{-Bi-Apx}$.

6.2 Intersection

Recall for norms $\|\cdot\|_{X_1}, \|\cdot\|_{X_2}$ over \mathbb{R}^n , the intersection $\text{Ball}(X_1) \cap \text{Ball}(X_2)$ is the unit ball of the norm $X_1 \vee X_2$ defined as

$$\|x\|_{X_1 \vee X_2} \stackrel{\text{def}}{=} \max\{\|x\|_{X_1}, \|x\|_{X_2}\}.$$

It is easily verified that $\tilde{T}_2(X_1 \vee X_2) \leq \sqrt{\tilde{T}_2(X_1)^2 + \tilde{T}_2(X_2)^2}$ and so boundedness of type-2 is preserved under intersection. We show that approximability of quadratic maximization is closed under intersection by simply observing that the upper covariance body $\mathcal{U}(X_1 \vee X_2)$ is equivalent to the intersection of the upper covariance bodies $\mathcal{U}(X_1), \mathcal{U}(X_2)$. Formally it is easily checked that for any $\mathbb{X} \in \text{PSD}^n$,

$$\max\{\mathcal{N}_{X_1}(\mathbb{X}), \mathcal{N}_{X_2}(\mathbb{X})\} \leq \mathcal{N}_{X_1 \vee X_2}(\mathbb{X}) \leq \mathcal{N}_{X_1}(\mathbb{X}) + \mathcal{N}_{X_2}(\mathbb{X}) \leq 2 \cdot \max\{\mathcal{N}_{X_1}(\mathbb{X}), \mathcal{N}_{X_2}(\mathbb{X})\} \quad (62)$$

where the first inequality follows immediately from Jensen's inequality and the remaining inequalities are straightforward. Thus we have

$$\frac{1}{2} \cdot \mathcal{U}(X_1) \cap \mathcal{U}(X_2) \subseteq \mathcal{U}(X_1 \vee X_2) \subseteq \mathcal{U}(X_1) \cap \mathcal{U}(X_2). \quad (63)$$

We then obtain

Proposition 6.5 (Closure of Type-2 Quadratic Maximization under Intersection).

There is an algorithm $\text{ALG}(A, R, r, \mathcal{O}_{X_1}, \mathcal{O}_{X_2})$ such that if \mathcal{O}_{X_1} (resp. \mathcal{O}_{X_2}) is an α -approximate search oracle for quadratic maximization over an (R, r) -balanced norm $(\mathbb{R}^n, \|\cdot\|_{X_1})$ (resp. $(\mathbb{R}^n, \|\cdot\|_{X_2})$), then on any input $A \in M_n(\mathbb{R})$, ALG runs in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A))$ and returns a β -approximate solution to $\mathcal{Q}_{X_1 \vee X_2}^{\max}(A)$ with probability $1 - 2^{-\Omega(n)}$, where $\beta \stackrel{\text{def}}{=} 2\alpha \cdot \max\{\tilde{T}_2(X_1)^2, \tilde{T}_2(X_2)^2\}$.

Proof. By [Proposition 4.13](#) we obtain $\tilde{T}_2(X_1)^2$ -approximate and $\tilde{T}_2(X_2)^2$ -approximate separation oracles for $\mathcal{U}(X_1)$ and $\mathcal{U}(X_2)$ respectively. Therefore we obtain a $\max\{\tilde{T}_2(X_1)^2, \tilde{T}_2(X_2)^2\}$ -approximate separation oracle for $\mathcal{U}(X_1) \cap \mathcal{U}(X_2)$. Combining this fact with [Observation 3.4](#) and [\(63\)](#) implies a $2 \cdot \max\{\tilde{T}_2(X_1)^2, \tilde{T}_2(X_2)^2\}$ -approximate separation oracle for $\mathcal{U}(X_1 \vee X_2)$. Finally applying [Proposition 4.2](#) completes the proof. \blacksquare

6.3 Quotients

Let $\|\cdot\|_X$ be a norm over \mathbb{R}^n . For a subspace E of \mathbb{R}^n , the quotient norm $\|X/E\|$ is a norm defined on the space X/E (which can be identified with the orthogonal complement E^\perp) and is given by

$$\|x\|_{X/E} \stackrel{\text{def}}{=} \min_{y \in E} \|x - y\|_X$$

i.e., the distance of x to the subspace E . The dual is the norm $\|\cdot\|_{X^*}$ restricted to the subspace E^\perp , i.e.,

$$\text{for any } x \in E^\perp, \quad \sup_{\substack{\zeta \in E^\perp \\ \|\zeta\|_{X^*} \leq 1}} \langle \zeta, x \rangle = \|x\|_{X/E} \quad \text{and for any } \zeta \in E^\perp, \quad \sup_{\substack{x \in E^\perp \\ \|x\|_{X/E} \leq 1}} \langle x, \zeta \rangle = \|\zeta\|_{X^*}.$$

More generally we say a surjective linear map $T : \mathbb{R}^n \rightarrow \mathbb{R}^k$ (where $k < n$) is an (a, b) -quotient map from X to Q if $\|T\|_{X \rightarrow Q} \leq a$ and $\text{Ball}(Q) \subseteq b \cdot T\text{Ball}(X)$. It is easily checked that for any $x \in \mathbb{R}^n$, $a^{-1} \cdot \|T(x)\|_Q \leq \|x\|_{X/\ker T} \leq b \cdot \|T(x)\|_Q$. An obvious dual transposition implies that for any $x \in \ker T^\perp$,

$$b^{-1} \cdot \|T(x)\|_{Q^*} \leq \|x\|_{X^*} \leq a \cdot \|T(x)\|_{Q^*}. \quad (64)$$

It is easily verified that $\tilde{T}_2(Q) \leq ab \cdot \tilde{T}_2(X)$ and so quotients of X inherit boundedness of type-2. Similarly $\tilde{C}_2(Q^*) \leq ab \cdot \tilde{C}_2(X^*)$ and so quotients of X inherit boundedness of dual cotype-2. In this section we show that approximation algorithms for quadratic (resp. bilinear) maximization are also preserved under quotienting when type-2 (resp. dual cotype-2) is bounded.

Our approach again is to show the claim for PSD maximization first (which is immediate) and then obtain the general case by appealing to our reduction theorem from quadratic (resp. bilinear) to PSD maximization.

Given access to an oracle for T one can compute the linear map $T^\dagger : \mathbb{R}^k \rightarrow \ker T^\perp$ corresponding to the inverse of $T|_{\ker T^\perp}$.

Fact 6.6. Consider any PSD matrix $A = BB^* \in \text{PSID}^k$. Then we have

$$\begin{aligned} \frac{1}{b} \cdot \|B\|_{2 \rightarrow Q^*} &\leq \|T^\dagger B\|_{2 \rightarrow X^*} \leq a \cdot \|B\|_{2 \rightarrow Q^*} && \text{(by (64))} \\ \Rightarrow \frac{1}{b^2} \cdot \mathcal{Q}_Q^{\max}(A) &\leq \mathcal{Q}_X^{\max}(T^\dagger A (T^\dagger)^*) \leq a^2 \cdot \mathcal{Q}_Q^{\max}(A). \end{aligned}$$

Thus if x is an α -approximate solution to $\mathcal{Q}_X^{\max}(T^\dagger A(T^\dagger)^*)$, then $T(x)$ is an $\alpha \cdot a^2 b^2$ -approximate solution to $\mathcal{Q}_Q^{\max}(A)$.

Combining [Fact 6.6](#) and [Theorem 4.5](#) immediately yields

Proposition 6.7 (Closure of Type-2 Quadratic Maximization under Quotienting).

There is an algorithm $\text{ALG}(A, R, r, \mathcal{O}_X, \mathcal{O}_T)$ such that if \mathcal{O}_X is an α -approximate search oracle for quadratic maximization over a norm $(\mathbb{R}^n, \|\cdot\|_X)$, $\|\cdot\|_Q$ is an (R, r) -balanced norm over \mathbb{R}^k and \mathcal{O}_T is an oracle computing an (a, b) -quotient map $T : X \rightarrow Q$, then on any input $A \in M_k(\mathbb{R})$, ALG runs in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A))$ and returns a β -approximate solution to $\mathcal{Q}_Q^{\max}(A)$ with probability $1 - 2^{-\Omega(n)}$, where $\beta = \alpha \cdot a^2 b^2 \cdot \tilde{T}_2(Q)^2 \leq \alpha \cdot a^4 b^4 \cdot \tilde{T}_2(X)^2$.

The bilinear case follows (without any type-2/cotype-2 assumptions) from the following straightforward generalization of [Fact 6.6](#).

Fact 6.8 (Closure of Bilinear Maximization under Quotienting).

Consider norms $(\mathbb{R}^n, \|\cdot\|_X)$, $(\mathbb{R}^m, \|\cdot\|_Y)$, $(\mathbb{R}^k, \|\cdot\|_P)$, $(\mathbb{R}^\ell, \|\cdot\|_Q)$ admitting (a, b) -quotient maps $S : \mathbb{R}^n \rightarrow \mathbb{R}^k$ from X to P and $T : \mathbb{R}^m \rightarrow \mathbb{R}^\ell$ from Y to Q . Consider any $k \times \ell$ matrix A . Then for S^\dagger, T^\dagger defined as above, we have

$$\frac{1}{b} \cdot \|A^*\|_{P \rightarrow Q^*} \leq \|T^\dagger A^*\|_{P \rightarrow Y^*} \leq a \cdot \|A^*\|_{P \rightarrow Q^*} \quad (\text{by (64)})$$

and

$$\frac{1}{b} \cdot \|A(T^\dagger)^*\|_{Y \rightarrow P^*} \leq \|S^\dagger A(T^\dagger)^*\|_{Y \rightarrow X^*} \leq a \cdot \|A(T^\dagger)^*\|_{Y \rightarrow P^*} \quad (\text{by (64)})$$

Combining the above four inequalities (repeatedly) with the fact that $\text{Op}_{E,F}^{\max}(M) = \|M\|_{F \rightarrow E^*} = \|M^*\|_{E \rightarrow F^*}$, yields

$$\frac{1}{b^2} \cdot \text{Op}_{P,Q}^{\max}(A) \leq \text{Op}_{X,Y}^{\max}(S^\dagger A(T^\dagger)^*) \leq a^2 \cdot \text{Op}_{P,Q}^{\max}(A).$$

Thus if (x, y) is an α -approximate solution to $\text{Op}_{X,Y}^{\max}(S^\dagger A(T^\dagger)^*)$, then $(S(x), T(y))$ is an $\alpha \cdot a^2 b^2$ -approximate solution to $\mathcal{Q}_Q^{\max}(A)$.

6.4 Complex Interpolation

6.4.1 Interpolation Preliminaries

Let $\mathcal{S} \stackrel{\text{def}}{=} \{z \in \mathbb{C} \mid \text{Re}(z) \in (0, 1)\}$ be the complex unit open strip, let $\partial\mathcal{S}$ denote its boundary and lastly let $\bar{\mathcal{S}} \stackrel{\text{def}}{=} \mathcal{S} \cup \partial\mathcal{S}$ denote the closed unit strip.

Let \mathcal{F} be the space of bounded continuous functions $f : \bar{\mathcal{S}} \rightarrow \mathbb{C}^n$ that are holomorphic in \mathcal{S} . Given complex norms $(\mathbb{C}^n, \|\cdot\|_{X_0})$ and $(\mathbb{C}^n, \|\cdot\|_{X_1})$, and a parameter $0 < \theta < 1$, the complex interpolant of X_0 and X_1 is defined to be the norm $(\mathbb{C}^n, \|\cdot\|_{[X_0, X_1]_\theta})$ given by

$$\|x\|_{[X_0, X_1]_\theta} \stackrel{\text{def}}{=} \inf_{\substack{f \in \mathcal{F} \\ f(\theta) = x}} \max \left\{ \sup_{\text{Re}(z)=0} \|f(z)\|_{X_0}, \sup_{\text{Re}(z)=1} \|f(z)\|_{X_1} \right\}. \quad (65)$$

Whenever the source and destination are labeled as X_0 and X_1 we will use the shorthand X_θ to denote the interpolant $[X_0, X_1]_\theta$.

In this section we will be interested in quadratic and bilinear maximization over interpolants X_θ . We obtain C -approximation algorithms for a constant C depending on $T_2(X_0)$ and $T_2(X_1)$. We use the shorthand $(X, Y)_\theta$ to denote the interpolant $[(X_0, Y_0), (X_1, Y_1)]_\theta$.

6.4.2 Approximation Algorithms for Interpolants of Type-2 Norms

The main result of this section is that if quadratic maximization is computationally approximable over type-2 norms $(\mathbb{C}^n, \|\cdot\|_{X_0})$ and $(\mathbb{C}^n, \|\cdot\|_{X_1})$, then it is also approximable over any interpolant of X_0 and X_1 :

Theorem 6.9 (Quadratic Maximization over Interpolants of Type-2 Norms.).

Fix any $\theta \in (0, 1)$. There is an algorithm $\text{ALG}(A, R, r, \mathcal{O}_{X_0}, \mathcal{O}_{X_1})$ such that if $\mathcal{O}_{X_0}, \mathcal{O}_{X_1}$ are α -approximate search oracles for quadratic maximization over (R, r) -balanced norms $(\mathbb{C}^n, \|\cdot\|_{X_0}), (\mathbb{C}^n, \|\cdot\|_{X_1})$ respectively, then ALG runs in time $\text{poly}(n, R, 1/r, \text{bit}(A))$ and returns a C -approximate solution to $\mathcal{Q}_{X_\theta}^{\max}(A)$ with probability $1 - 2^{-\Omega(n)}$, where $C = (1 + o(1)) \cdot \alpha \cdot \max_{i \in \{0, 1\}} \{\tilde{T}_2(X_i)^4\}$.

Here again our approach is to reduce quadratic maximization to bilinear maximization and then appeal to a factorization result of Kouba in order to solve the bilinear problem. Indeed [Theorem 6.9](#) follows from combining [Theorem 4.5](#) with the following theorem

Proposition 6.10 (Bilinear Maximization over Interpolants of Type-2 Norms.).

Fix any $\theta \in (0, 1)$. There is an algorithm $\text{ALG}(A, \mathcal{O}_{(X_0, Y_0)}, \mathcal{O}_{(X_1, Y_1)})$ such that if $\mathcal{O}_{(X_0, Y_0)}$ (resp. $\mathcal{O}_{(X_1, Y_1)}$) is and α -approximate search oracle for bilinear maximization over the pair of (R, r) -balanced norms $(\mathbb{C}^n, \|\cdot\|_{X_0}), (\mathbb{C}^m, \|\cdot\|_{Y_0})$ (resp. $(\mathbb{C}^n, \|\cdot\|_{X_1}), (\mathbb{C}^m, \|\cdot\|_{Y_1})$), then ALG runs in time $\text{poly}(n, m, R, 1/r, \text{bit}(A))$ and returns a C -approximation to $\text{Op}_{X_\theta, Y_\theta}^{\max}(A)$, where $C = (1 + o(1)) \cdot \alpha \cdot \max_{i \in \{0, 1\}} \{\tilde{T}_2(X_i)^2 \cdot \tilde{T}_2(Y_i)^2\}$.

Moreover ALG returns a witness $Z \in \text{Ball}(X_\theta \hat{\otimes} Y_\theta)$ satisfying $\langle A, Z \rangle \geq \text{Op}_{X_\theta, Y_\theta}^{\max}(A) / C$.

6.4.3 Proof of [Proposition 6.10](#)

Our approach is to use a result of Kouba [[Kou91](#)] who extended a factorization result of Pisier for matrix valued analytic functions in order to obtain conditions under which $\|\cdot\|_{[(X_0, Y_0), (X_1, Y_1)]_\theta}$ and $\text{Op}_{X_\theta, Y_\theta}^{\max}(\cdot)$ are equivalent.

Theorem 6.11 (Kouba [[Kou91](#)]). Let X_0, X_1 (resp. Y_0, Y_1) be norms over \mathbb{C}^n (resp. \mathbb{C}^m). Then for any $A \in M_{n, m}(\mathbb{C})$ we have

$$\begin{aligned} \|A\|_{(X, Y)_\theta} &\leq \text{Op}_{X_\theta, Y_\theta}^{\max}(A) \leq \max_{i \in \{0, 1\}} \{\tilde{T}_2(X_i)^2 \cdot \tilde{T}_2(Y_i)^2\} \cdot \|A\|_{(X, Y)_\theta} \\ \|A\|_{(X \hat{\otimes} Y)_\theta} &\geq \|A\|_{X_\theta \hat{\otimes} Y_\theta} \geq \min_{i \in \{0, 1\}} \{\tilde{T}_2(X_i)^{-2} \cdot \tilde{T}_2(Y_i)^{-2}\} \cdot \|A\|_{(X \hat{\otimes} Y)_\theta} \end{aligned}$$

Armed with this equivalence, we need only approximate $(X, Y)_\theta \stackrel{\text{def}}{=} [(X_0, Y_0), (X_1, Y_1)]_\theta$ in order to prove [Proposition 6.10](#). Andoni et al. [[ANN⁺18](#)] show that given membership oracles for norms W_0, W_1, W_θ can be computed using the ellipsoid method. We use their result with the substitution $W_0 \stackrel{\text{def}}{=} X_0 \hat{\otimes} Y_0$ and $W_1 \stackrel{\text{def}}{=} X_1 \hat{\otimes} Y_1$ ⁸ with the caveat that we only have approximate separation oracles for W_0, W_1 and so we instead use the approximate ellipsoid method.

To this end note that $\|\cdot\|_{(X \hat{\otimes} Y)_\theta} \stackrel{\text{def}}{=} \|\cdot\|_{[X_0 \hat{\otimes} Y_0, X_1 \hat{\otimes} Y_1]_\theta}$ is the dual norm of $\|\cdot\|_{(X, Y)_\theta}$ and so

$$\|A\|_{(X, Y)_\theta}$$

⁸We interpolate the projective norms along with the injective norms as the reduction from quadratic maximization to psd maximization ([Theorem 4.5](#)) requires a ‘witness’ in order to apply.

$$\begin{aligned}
&= \sup_{\|Z\|_{(X \otimes Y)_\theta} \leq 1} \langle A, Z \rangle \\
&= \sup \left\{ \langle A, Z \rangle \mid \exists f \in \mathcal{F} \text{ s.t. } f(\theta) = Z, \forall t \in \mathbb{R}, \|f(i \cdot t)\|_{X_0 \otimes Y_0}, \|f(1 + i \cdot t)\|_{X_1 \otimes Y_1} \leq 1 \right\} \quad (66)
\end{aligned}$$

where \mathcal{F} is the set of matrix valued functions holomorphic on the strip $\overline{\mathcal{S}}$.

The following result of [ANN⁺18] states that \mathcal{F} can be replaced with a set of functions having only polynomially many non-zero fourier coefficients, thereby enabling us to formulate (66) (approximately) as a convex program in polynomially many variables.

Lemma 6.12 (Lemma 5.11 + Claim 5.15 in Andoni et al.).

Fix any $C \geq 1$. Then there exists $C' \geq 1$ and $M, N, R \leq n^{C'}$ such that for any $Z \in \text{Ball}((X \otimes Y)_\theta)$, there is a sequence of matrices $(v_q)_{q \in \mathbb{Q}_M}$ in $M_{n,m}(\mathbb{C})$ such that

- i. $\forall z \in \mathbb{D}_N^{(0)} \quad \|f_V(z)\|_{X_0 \otimes Y_0} \leq 1 + 1/n^C$
- ii. $\forall z \in \mathbb{D}_N^{(1)} \quad \|f_V(z)\|_{X_1 \otimes Y_1} \leq 1 + 1/n^C$
- iii. $\forall q \in \mathbb{Q}_M \quad \|v_q\|_H \cdot \max\{e^q, 1\} \leq n^{C'}$
- iv. $\|f_V(\theta) - Z\|_H \leq 1/n^C$

$$\text{where } f_V(z) \stackrel{\text{def}}{=} e^{z^2/M} \cdot \sum_{q \in \mathbb{Q}_M} v_q \cdot e^{qz}. \quad (67)$$

Furthermore if for some $Z \in M_n(\mathbb{C})$, a sequence (v_q) satisfying (67) exists, then $\|Z\|_{(X \otimes Y)_\theta} \leq 1 + 1/n$.

We thus apply the approximate ellipsoid method to the following program

maximize $\langle A, Z \rangle$ such that

- i. $\forall z \in \mathbb{D}_N^{(0)} \quad \|f_V(z)\|_{W_0} \leq 1 + 1/n^C$
- ii. $\forall z \in \mathbb{D}_N^{(1)} \quad \|f_V(z)\|_{W_1} \leq 1 + 1/n^C$
- iii. $\forall q \in \mathbb{Q}_M \quad \|v_q\|_H \cdot \max\{e^q, 1\} \leq n^{C'}$
- iv. $\|f_V(\theta) - w\|_H \leq 1/n^C$

$$\text{where } f_V(z) \stackrel{\text{def}}{=} e^{z^2/M} \cdot \sum_{q \in \mathbb{Q}_M} v_q \cdot e^{qz}$$

$$Z, v_q \in M_{n,m}(\mathbb{C}). \quad (68)$$

We are ready to prove our bilinear maximization result.

Proof of Proposition 6.10. By Theorem 3.14, $\mathcal{O}_{(X_0, Y_0)}$ (resp. $\mathcal{O}_{(X_1, Y_1)}$) can be adapted to obtain a poly-time α -approximate separation oracle for $\text{Ball}(X_0 \otimes Y_0)$ (resp. $\text{Ball}(X_1 \otimes Y_1)$). Therefore by applying Proposition 3.11 to the program (68) (where $B \in M_{n,m}(\mathbb{C})$ is treated as an object in $(\mathbb{R}^2)^{n \times m}$), we obtain $Z \in M_{n,m}(\mathbb{C})$ such that $\langle A, Z \rangle = (1 - o(1)) \cdot \|A\|_{(X, Y)_\theta}$ and $(1 - o(1))Z/\alpha$ satisfies the conditions of (67). Thus by Lemma 6.12 $\|Z\|_{(X \otimes Y)_\theta} \leq \alpha(1 + o(1))$ and furthermore by Theorem 6.11 (the trivial direction), $\|Z\|_{X_\theta \otimes Y_\theta} \leq \alpha(1 + o(1))$. Again by Theorem 6.11 (the non-trivial direction), $\langle A, Z \rangle \geq (1 - o(1)) \cdot \min_{i \in \{0,1\}} \{\tilde{T}_2(X_i)^{-2} \cdot \tilde{T}_2(Y_i)^{-2}\} \cdot \text{Op}_{X_\theta, Y_\theta}^{\max}(A)$. This completes the proof since $(1 - o(1))Z/\alpha$ is the desired witness. \blacksquare

6.5 Sections (Subspaces)

A section of a origin-symmetric convex body C is defined as the intersection of C with a subspace V . Let X be the norm whose ball is C and let E denote the norm on the space V whose unit ball is $C \cap V$. The corresponding quadratic maximization problem over E is to compute given a linear map $A : V \rightarrow V$, the following quantity

$$Q_E^{\max}(A) \stackrel{\text{def}}{=} \sup_{x \in C \cap V} \langle x, A(x) \rangle.$$

It is easy to reduce $Q_E^{\max}(\cdot)$ to $Q_X^{\max}(\cdot)$. In fact, this holds without any type-2 assumption on X – a fact that will come in handy for hardness reductions in [Section 8](#). Indeed we have

Observation 6.13 (Enforcing a Subspace Constraint).

Let A be an $n \times n$ matrix, $V \subseteq \mathbb{R}^n$ be a subspace and X be a norm over \mathbb{R}^n such that $r \cdot \text{Ball}(\ell_2^n) \subseteq \text{Ball}(X)$.

Let Π be the projector to V^\perp and let $\alpha \stackrel{\text{def}}{=} Q_X^{\max}(A') / (n \|A'\|_{X \rightarrow X^*})$. Let $A' \stackrel{\text{def}}{=} A - (r\alpha)^{-2} \cdot Q_X^{\max}(A) \cdot \Pi$. Then we have

$$\sup_{x \in \text{Ball}(X) \cap V} \langle x, Ax \rangle \leq Q_X^{\max}(A') \leq (1 + O(1/n)) \cdot \sup_{x \in \text{Ball}(X) \cap V} \langle x, Ax \rangle$$

Proof. The first inequality is immediate. For the second inequality, consider any $x \in \text{Ball}(X)$ that maximizes $\langle x, A'x \rangle$. Decompose x as $x^\parallel + x^\perp$ where x^\perp is the component of x in V^\perp and x^\parallel is the component in V . We may assume $\|x^\perp\|_2 \leq r\alpha$ since otherwise $\langle x, A'x \rangle < 0$ which would contradict optimality of x . This implies that $\|x^\perp\|_X \leq \alpha$ and therefore $\|x^\parallel\|_X \leq 1 + \alpha \leq 1 + 1/n$. Finally we have

$$\begin{aligned} Q_X^{\max}(A') &= \langle x, A'x \rangle = \langle x^\parallel, A'x^\parallel \rangle + \langle x^\parallel, A'x^\perp \rangle + \langle x^\perp, A'x^\parallel \rangle + \langle x^\perp, A'x^\perp \rangle \\ &\leq \langle x^\parallel, A'x^\parallel \rangle + (2(1 + \alpha)\alpha + \alpha^2) \cdot \|A'\|_{X \rightarrow X^*} \\ &\leq \langle x^\parallel, A'x^\parallel \rangle + 4 \cdot Q_X^{\max}(A')/n \\ &= (1 + O(1/n)) \langle x^\parallel, A'x^\parallel \rangle / \|x^\parallel\|_X^2 + 4 \cdot Q_X^{\max}(A')/n \\ &= (1 + O(1/n)) \langle x^\parallel, Ax^\parallel \rangle / \|x^\parallel\|_X^2 + 4 \cdot Q_X^{\max}(A')/n \end{aligned}$$

as desired. ■

7 Unconditional Algorithms for Special Families

In this section we study various special families of norms for which one can design separation oracles for the lower covariance region $\mathcal{L}(X)$ using only a membership oracle for $\text{Ball}(X)$. Thus using [Theorem 4.11](#) we obtain approximation algorithms for quadratic/bilinear maximization that require only a membership oracle for $\text{Ball}(X)$. In doing so we recover constant factor approximation algorithms for the cases studied previously in the literature and also obtain new results. We obtain constant factor algorithms for the following norm families (see [Section 7.1.1](#) and [Section 7.3.1](#) for definitions):

1. Quadratic maximization over exactly 2-convex sign-invariant norms with bounded q -concavity for finite q . This recovers a result of Naor and Schechtman [[NS09](#)].

2. Bilinear maximization over exactly 2-convex sign-invariant norms. This recovers (with a worse constant) a result implicit in the work of Krivine [Kri73] and independently rediscovered by Nesterov [Nes98].
3. Quadratic (resp. bilinear) maximization over symmetric norms (i.e., invariant to permutation and sign changes) that have bounded type-2 (resp. bounded dual cotype-2). This includes many new examples not covered in the exactly 2-convex case.
4. Quadratic (resp. bilinear) maximization over unitarily invariant matrix norms that have bounded type-2 (resp. bounded dual cotype-2). This recovers as a special case (with a worse constant) a result of Naor, Regev and Vidick [NRV13] for bilinear maximization over Schatten- ∞ .

7.1 Approximation Algorithms for Sign Invariant Norms

We require some preliminaries.

7.1.1 p -convexity and q -concavity Preliminaries

The notions of p -convexity and q -concavity are well defined for a wide class of normed spaces known as Banach lattices. In this document we only define these notions for finite dimensional norms that are 1-unconditional in the elementary basis (i.e., those norms $\|\cdot\|_X$ for which flipping the sign of an entry of x does not change the norm. We shall refer to such norms as *sign-invariant norms*). Most of the statements we make in this context can be readily extended to the case of norms admitting some 1-unconditional basis, but we choose to fix the elementary basis in the interest of simplicity.

In what follows, for a scalar function $s : \mathbb{R} \rightarrow \mathbb{R}$ and a vector $x \in \mathbb{R}^n$, we use the notation $s(x)$ to denote the vector obtained by entry-wise application of s to x , i.e., $s(x) = (s(x_1), \dots, s(x_n))$. For e.g., $|x|^p$ denotes the vector $(|x_1|^p, \dots, |x_n|^p)$. This notation appears exclusively in [Section 7.1.1](#) and [Section 5](#).

Definition 7.1 (p -convexity/ q -concavity). *Let X be a sign-invariant norm over \mathbb{R}^n . Then for $1 \leq p \leq \infty$ the p -convexity constant of X , denoted by $M^{(p)}(X)$, is the smallest constant C such that for every finite sequence of vectors (x_i) in X ,*

$$\left\| \left(\sum_i |x_i|^p \right)^{1/p} \right\|_X \leq C \cdot \left(\sum_i \|x_i\|_X^p \right)^{1/p}$$

We will say X is exactly p -convex if $M^{(p)}(X) = 1$.

For $1 \leq q \leq \infty$, the q -concavity constant of X , denoted by $M_{(q)}(X)$, is the smallest constant C such that for every finite sequence of vectors $\{x_i\}$ in X ,

$$\left\| \left(\sum_i |x_i|^q \right)^{1/q} \right\|_X \geq \frac{1}{C} \cdot \left(\sum_i \|x_i\|_X^q \right)^{1/q}.$$

We will say X is exactly q -concave if $M_{(q)}(X) = 1$.

Remark 7.2. *Every sign-invariant norm is exactly 1-convex and exactly ∞ -concave. It is also known (see Lindenstrauss-Tzafriri) that any sign-invariant norm X is C -equivalent to an exactly p -convex norm (resp. an exactly q -concave norm) for C depending only on $M^{(p)}(X)$ (resp. $M_{(q)}(X)$).*

Fact 7.3. Let X be a sign-invariant norm over \mathbb{R}^n . Then whenever $\|[x]\| \leq \|[y]\|$ (entry-wise), it must be that $\|x\|_X \leq \|y\|_X$.

For a sign-invariant norm $\|\cdot\|_X$ over \mathbb{R}^n , and any $0 < p < \infty$ let $\|\cdot\|_{X^{(p)}}$ denote the function $\| |x|^p \|_X^{1/p}$. $X^{(p)}$ is referred to as the p -convexification of X . It is easily verified that $M^{(p)}(X^{(p)}) = M^{(1)}(X)$ and further that $\|\cdot\|_{X^{(p)}}$ is an exactly p -convex sign-invariant norm if and only if $\|\cdot\|_X$ is an exactly 1-convex sign-invariant norm.

We will appeal to a known equivalence between (2-convexity + $(q < \infty)$ -concavity) and Type-2 for Banach lattices that is implicit in the following Khintchine-type inequality.

Theorem 7.4 (Banach Lattice Khintchine [Mau73]).

Let X be a sign-invariant norm over \mathbb{R}^n . Then for any $1 \leq p \leq 2 \leq q < \infty$ and any finite sequence (x_i) and i.i.d. standard Gaussians (g_i) ,

$$\frac{\gamma_p}{M^{(p)}(X)} \cdot \left\| \left(\sum_i |x_i|^2 \right)^{1/2} \right\|_X \leq \sqrt{\mathbb{E} \left[\left\| \sum_i g_i \cdot x_i \right\|_X^2 \right]} \leq \gamma_q \cdot M^{(q)}(X) \cdot \left\| \left(\sum_i |x_i|^2 \right)^{1/2} \right\|_X$$

where $\gamma_p \stackrel{\text{def}}{=} \mathbb{E}_{g \sim \mathcal{N}(0,1)} [|g|^p]^{1/p}$.

Remark 7.5.

1. If X is exactly 2-convex then the constant in the left inequality can be taken as 1 since in this case $\gamma_2 / M^{(2)}(X) = 1$.
2. If X is exactly 2-concave then the constant in the right inequality can be taken as 1 since in this case $\gamma_2 \cdot M^{(2)}(X) = 1$.
3. Since X is a sign-invariant norm, $M^{(1)}(X) = 1$ and so the constant in the left inequality can always be taken as $\gamma_1 = \sqrt{2/\pi}$.
4. Since X is a norm, $M^{(\infty)}(X) = 1$ and so $M^{(\log n)}(X) \leq e$. Thus the constant in the right inequality can always be taken as $e \cdot \gamma_{\log n} = (1 + o(1)) \sqrt{e \log n}$.

Finally, we record the well known equivalence between 2-convexity and dual cotype-2 (see for e.g. [Pis86]) for Banach lattices:

$$\tilde{C}_2(X^*) \lesssim M^{(2)}(X) \lesssim \tilde{C}_2(X^*) \log \tilde{C}_2(X^*).$$

7.1.2 Approximation Algorithms for Maximization over Exactly 2-Convex Norms

Naor and Schechtman [NS09] gave an $M^{(q)}(X)^2 \cdot \gamma_q^2$ -approximation algorithm for quadratic maximization over an exactly 2-convex norm X . We show how our framework for quadratic maximization (Proposition 4.2) recovers their result.

Theorem 7.6 (Quadratic Maximization over Exactly 2-convex Norms).

There is an algorithm $\text{ALG}(A, R, r, \mathcal{O}_X)$ such that if \mathcal{O}_X is an oracle for exactly computing a sign-invariant, exactly 2-convex, (R, r) -balanced norm $(\mathbb{R}^n, \|\cdot\|_{X_1})$, then on any input $A \in M_n(\mathbb{R})$, ALG runs in time $\text{poly}(n, \log R, \log 1/r, \text{bit}(A))$ and returns a β -approximate solution to $\text{Q}_X^{\max}(A)$ with probability $1 - 2^{-\Omega(n)}$, where $\beta \stackrel{\text{def}}{=} \inf_{2 \leq q < \infty} M^{(q)}(X)^2 \cdot \gamma_q^2$.

Proof. For the domain $\mathbb{X} \in M_n(\mathbb{R})$, $\mathbb{X} \succeq 0$, we define a computable convex function f as

$$f(\mathbb{X}) \stackrel{\text{def}}{=} \|\text{diag}(\mathbb{X})\|_{X^{(1/2)}} = \|\text{diag}(\mathbb{X})^{1/2}\|_X^2 = \left\| \left(\sum_i |x_i|^2 \right)^{1/2} \right\|_X^2 \quad (69)$$

for any decomposition $\mathbb{X} = \sum_{k \in [K]} x_k(x_k)^*$.

Convexity is evident from the first equality above since exact 2-convexity of X implies $X^{(1/2)}$ is a norm. Computability (using \mathcal{O}_X) is evident from the second equality above. [Theorem 7.4](#) implies that $\mathcal{N}_X(\mathbb{X})$ and $f(\mathbb{X})$ are β -equivalent for the claimed β . Thus combining part (1.) of [Corollary 4.12](#) and [Theorem 7.4](#) yields the claim. ■

Krivine [[Kri73](#)] observed that Grothendieck's inequality extends (with the same constant K_G) to the case of bilinear maximization over exactly 2-convex lattices. Nesterov [[Nes98](#)] independently rediscovered Grothendieck's inequality (with constant ~ 2.29) and also noted its extension to the exactly 2-convex case.

We give an alternate proof of this result (with a worse constant) using our framework ([Theorem 4.11](#)).

Theorem 7.7 (Bilinear Maximization over Exactly 2-convex Norms).

There is an absolute constant $\beta > 1$ and an algorithm $\text{ALG}(A, R, r, \mathcal{O}_X, \mathcal{O}_Y)$ such that if \mathcal{O}_X (resp. \mathcal{O}_Y) is an oracle for exactly computing a sign-invariant, exactly 2-convex, (R, r) -balanced norm $(\mathbb{R}^n, \|\cdot\|_X)$ (resp. $(\mathbb{R}^m, \|\cdot\|_Y)$), then on any input $A \in M_{n,m}(\mathbb{R})$, ALG runs in time $\text{poly}(n, m, \log R, \log 1/r, \text{bit}(A))$ and returns a β -approximate solution to $\text{Op}_{X,Y}^{\max}(A)$ with probability $1 - 2^{-\Omega(n+m)}$.

Proof. For the domain $\mathbb{W} \in M_n(\mathbb{R})$, $\mathbb{W} \succeq 0$, we define a computable concave function f as

$$f(\mathbb{W}) \stackrel{\text{def}}{=} \|\text{diag}(\mathbb{W})\|_{(X^*)^{(1/2)}} = \|\text{diag}(\mathbb{W})^{1/2}\|_{X^*}^2 = \left\| \left(\sum_i |x_i|^2 \right)^{1/2} \right\|_{X^*}^2 \quad (70)$$

for any decomposition $\mathbb{W} = \sum_{k \in [K]} x_k(x_k)^*$.

Concavity is evident from the first equality above since exact 2-convexity of X implies $(X^*)^{(1/2)}$ is concave. Computability (using \mathcal{O}_X) is evident from the second equality above. [Theorem 7.4](#) implies that $\mathcal{N}_{X^*}(\mathbb{W})$ and $f(\mathbb{W})$ are equivalent within a universal constant. We proceed similarly for Y . Thus combining [Theorem 7.4](#) and part (B) of part (2.) of [Corollary 4.12](#) yields the claim. ■

7.2 Approximation Algorithms for Symmetric Norms

In this section we will use [Theorem 4.11](#) to give constant factor approximation algorithms for quadratic (resp. bilinear) maximization over symmetric norms with bounded type-2 (resp. dual cotype-2) constant (assuming only oracles computing the norms being optimized over).

To do so, we will design a separation oracle for the lower covariance region. We begin with a technical ingredient common to all our proofs - namely that linear optimization over symmetric downward/upward-closed subsets of the non-negative orthant can be done efficiently (assuming only an approximate membership oracle for the set).

7.2.1 Preliminaries: Linear Optimization over Symmetric Upward/Downward-closed Sets

The following lemma shows that one can efficiently minimize linear functions over a symmetric upward closed subset of the non-negative orthant. This will be used in constructing a separation oracle for the lower covariance region of a symmetric norm.

Lemma 7.8. *Let $K = (\mathbb{R}_{\geq 0}^n)$ be the positive orthant, and let $B \subseteq K$ be a, permutation symmetric, upward-closed, inverse (R, r, K) -balanced, and exactly convex body with a C -approximate membership oracle \mathcal{O} . Then for any $\varepsilon > 0$, there is a $\text{poly}(n, R/r, 1/\log \varepsilon, \text{bit}(z))$ -time algorithm ALG that takes a vector $z \in \mathbb{R}_{\geq 0}^n$ as input and outputs y' with $y' \in B$ such that*

$$\inf_{y \in B} \langle z, y \rangle \leq \langle z, y' \rangle \leq (C \cdot (1 + \varepsilon)) \inf_{y \in B} \langle z, y \rangle$$

Proof. Without loss of generality, assume $r = 1$. If $z_i = 0$ for some $i \in [n]$, the optimum is 0 since $R \cdot e_i \in B$, and we are done. Otherwise, by sign-invariance, and scaling, we may assume z and the optimal vector $y^* \in B$ satisfy $1 = z_1 \leq \dots \leq z_n \leq n^c$, and $y_1^* \geq \dots \geq y_{n'}^* = 1 > y_{n'+1}^* = \dots = y_n^* = 0$. (I.e., the smallest nonzero entry of y^* is 1.) Then $\text{OPT} = \langle z, y^* \rangle \geq 1$, and since $R \cdot e_1 \in B$, $\text{OPT} \leq R$. This implies that $y_1^* \leq R$ and $z_{n'} \leq R$. Let z and y^* be their projections to the first n' coordinates. (We do not know n' in advance, but one can guess it with n tries.) We perform the following operations to covert z and y^* to nicer forms while approximately preserving $\langle z, y^* \rangle$ and $y^* \in B$.

- Round up each entry of z_i to the smallest $(1 + \varepsilon)^t$ for some integer $t \geq 0$. The value $\langle z, y^* \rangle$ increases multiplicatively by a factor at most $(1 + \varepsilon)$. Let $L = \lceil \log_{1+\varepsilon} R \rceil$ and Let $S_0, \dots, S_L \subseteq [n']$ be the set of coordinates such that $S_i := \{j \in [n'] : z_j = (1 + \varepsilon)^i\}$. Note that S_0, \dots, S_L partition $[n']$ and they are “consecutive” in the sense that S_0 contains the first $|S_0|$ coordinates, S_1 contains the next $|S_1|$ coordinates, \dots , and S_L contains the last $|S_L|$ coordinates.
- For $i \in \{0, \dots, L\}$, let $a_i := (\sum_{j \in S_i} y_j^*) / |S_i|$, and change y^* by setting $y_j^* = a_i$ for all $j \in S_i$ for all $i \in \{0, \dots, L\}$. Since the new y^* can be written as a convex combination of coordinate permutations of the old y^* , the convexity of B implies the new y^* satisfies $y^* \in B$. The value $\langle z, y^* \rangle$ does not change.
- Similarly to z , round up each entry of y_i^* to the largest power of $(1 + \varepsilon)^t$ for some $t \geq 0$. The value $\langle z, y^* \rangle$ increases multiplicatively by a factor at most $(1 + \varepsilon)$. Since we only increased coordinates of y^* , by upward-closedness of B , y^* is still in B .

Therefore, for both z and y^* , the coordinates in the same S_i have the same value, which is of the form $(1 + \varepsilon)^t$ for some $t \in \{1, \dots, L\}$. The objective function $\langle z, y^* \rangle$ is at most $(1 + \varepsilon)^2 \text{OPT}$, and $y \in B$. Note that we do not know y^* , but we can compute the new z and S_0, \dots, S_L given the original z and correctly guessed n' .

Finally, we exhaustively search the coordinate values for y^* . Since we search for at most $L + 1$ different values and the number of possible values is also at most $L + 1$, the number of possible choices is bounded by $2^{O(L)} = \text{poly}(R)$. Once we find the correct y^* , the C -approximate membership oracle will at least correctly certify that $Cy^* \in B$. The returned value is $\langle Cy^*, z \rangle \leq C(1 + O(\varepsilon))\text{OPT}$. ■

The next lemma states that one can efficiently maximize a linear function over a downward-closed subset of the non-negative orthant.

Lemma 7.9. *Let Y be a sign-invariant, symmetric (under permutation), quasi-norm over \mathbb{R}^n with quasi-triangle constant C and an exact membership oracle \mathcal{O} . Then for any $\varepsilon > 0$, there is an $\text{poly}(n^{O(1/\varepsilon)}, \text{bit}(z))$ -time algorithm ALG that takes a vector $z \in \mathbb{R}^n$ as input and outputs $y' \in \text{Ball}(Y)$ such that*

$$\frac{\sup_{y \in \text{Ball}(Y)} \langle z, y \rangle}{(1 + o(1)) \cdot C \cdot (1 + \varepsilon)} \leq \langle z, y' \rangle \leq \sup_{y \in \text{Ball}(Y)} \langle z, y \rangle$$

Proof. By symmetry and sign-invariance we may assume z as well as the optimal vector y^* have non-negative entries sorted in descending order. Without loss of generality we may assume $z_1 = 1$ and $\|e_1\|_Y = 1$, which implies that $y_1^* \leq 1$. Let $\text{OPT} := \langle y^*, z \rangle$. Since $y = e_1$ is a feasible solution with $\langle y, z \rangle = 1$, it also implies that $\text{OPT} \geq 1$. We perform the following operations to covert z and y^* to nicer forms while approximately preserving $\langle z, y^* \rangle$.

- For any $i \in [n]$, if z_i or y_i^* is less than $1/n^3$, make it 0. The value $\langle z, y^* \rangle$ decreases additively by at most $n^2/n^3 = 1/n \leq \text{OPT}/n$.
- Round down each entry of z_i to the largest power of $(1 - \varepsilon)^t$ for some $t \geq 0$ while keeping 0 as 0. The value $\langle z, y^* \rangle$ decreases multiplicatively by a factor at most $1/(1 - \varepsilon)$. Let $L = \lceil \log_{1/(1-\varepsilon)} n^3 \rceil$ and Let $S_0, \dots, S_L \subseteq [n]$ be the set of coordinates such that $S_i := \{j \in [n] : z_j = (1 - \varepsilon)^i\}$. Let $S_z := \{j \in [n] : z_j = 0\}$. Note that S_0, \dots, S_L, S_z partition $[n]$ and they are “consecutive” in the sense that S_0 contains the first $|S_0|$ coordinates, S_1 contains the next $|S_1|$ coordinates, \dots , and S_z contains the last $|S_z|$ coordinates.
- For $i \in \{0, \dots, L\}$, let $a_i := (\sum_{j \in S_i} y_j^*)/|S_i|$, and change y^* by setting $y_j^* = a_i$ for all $j \in S_i$ for all $i \in \{0, \dots, L\}$. Since the new y^* can be written as a convex combination of coordinate permutations of the old y^* , the quasi-triangle inequality implies that $\|y^*\|_Y$ is increased by a factor at most C . The value $\langle z, y^* \rangle$ does not change.
- Similarly to z , round down each entry of y_i^* to the largest power of $(1 - \varepsilon)^t$ for some $t \geq 0$ while keeping 0 as 0. The value $\langle z, y^* \rangle$ decreases multiplicatively by a factor at most $1/(1 - \varepsilon)$, and $\|y^*\|_Y$ does not increase.

Therefore, for both z and y^* , the coordinates in the same S_i have the same value, which is of the form $(1 - \varepsilon)^t$ for some $t \in \{0, \dots, L\}$ or 0. The objective function $\langle z, y^* \rangle$ is at least $(1 - 1/n)(1 - \varepsilon)^2 \text{OPT}$, and $\|y^*\|_Y \leq C$. Note that we do not know y^* , but we can compute the new z and S_0, \dots, S_L, S_z given the original z .

Finally, we exhaustively search the coordinate values for y^* . Since we search for at most $L + 2$ different values and the number of possible values is also at most $L + 2$, the number of possible choices is bounded by $2^{O(L)} = n^{O(1/\varepsilon)}$. Once we find the correct y^* , we can divide by C so that the $\|y^*\|_Y \leq 1$. The returned value of $\langle y^*, z \rangle \geq \frac{(1-1/n)(1-2\varepsilon)}{C} \text{OPT}$. ■

The following lemma shows that one can reduce linear function maximization over a unitarily invariant subset of the positive semidefinite cone to an instance of linear function maximization over a symmetric subset of the non-negative orthant. This will be used in constructing a separation oracle for the lower covariance region of a unitarily invariant norm.

Lemma 7.10. *Let $\Lambda \in (\mathbb{R}_{\geq 0})^n$ be a permutation symmetric set, and let \mathcal{B} be the set of PSD matrices whose eigenvalue sequence (y_1, \dots, y_n) belongs to Λ . Let A be an $n \times n$ PSD matrix, and let $z_1 \geq \dots \geq z_n \geq 0$*

be its eigenvalues. Then

$$\inf_{B \in \mathcal{B}} \langle A, B \rangle = \inf_{y \in \Lambda} \langle z, y \rangle, \quad \text{and} \quad \sup_{B \in \mathcal{B}} \langle A, B \rangle = \sup_{y \in \Lambda} \langle z, y \rangle,$$

and the optimal $B \in \mathcal{B}$ can be efficiently computed given the optimal solution $y \in \Lambda$.

Proof. Let $A = \sum z_i u_i u_i^T$ be its eigendecomposition for some orthonormal basis $\{u_i\}$. For any $y \in \Lambda$, $B = \sum_i y_i u_i u_i^T \in \mathcal{B}$ satisfies $\langle A, B \rangle = \sum z_i y_i$, so $\inf_B \langle A, B \rangle \leq \inf_y \langle z, y \rangle$ and $\sup_B \langle A, B \rangle \geq \sup_y \langle z, y \rangle$. To prove the other direction, take any $B \in \mathcal{B}$ and let $B = \sum_i y_i u_i u_i^T$ be its eigendecomposition with $y \in \Lambda$ with $y_1 \geq \dots \geq y_n \geq 0$. Since $\langle A, B \rangle = \sum_{i,j} z_i y_j \langle v_i, u_j \rangle^2$ and both $\{u_i\}$ and $\{v_i\}$ are orthonormal bases, $\sum_j \langle v_i, u_j \rangle^2 = 1$ and $\sum_j \langle v_j, u_i \rangle^2 = 1$. Therefore, $\langle A, B \rangle$ is maximized when $v_i = u_i$ and minimized $v_i = u_{n+1-i}$. Therefore, $\langle A, B \rangle \in [\sum_i z_i y_{n+1-i}, \sum_i z_i y_i] \in [\inf_y \langle z, y \rangle, \sup_y \langle z, y \rangle]$. ■

7.2.2 Symmetric Type-2 Norms

In this section we will give an unconditional constant factor approximation algorithm for quadratic (resp. bilinear) maximization over symmetric norms with bounded type-2 (resp. dual cotype-2) constant.

We begin by constructing a separation oracle for the lower covariance region $\mathcal{L}(X)$ of a symmetric norm $\|\cdot\|_X$.

Lemma 7.11 (Lower Covariance Separation Oracle for a Symmetric Norm with Bounded Dual Cotype-2). *For any (R, r) -balanced symmetric norm $\|\cdot\|_X$ over \mathbb{R}^n the lower covariance region $\mathcal{L}(S_X)$ has a $\text{poly}(n, \log R, \log 1/r, x)$ time α -approximate separation oracle with input x , assuming access to an oracle \mathcal{O}_X computing $\|\cdot\|_X$, where $\alpha \lesssim M^{(2)}(X)^4 = O(\tilde{C}_2(X^*))^4 \log^4(\tilde{C}_2(X^*))$.*

Proof. Let F be the renorming of X that is symmetric, exactly 2-convex and satisfies $\|x\|_F \leq \|x\|_X \leq M^{(2)}(X) \cdot \|x\|_X$. We know F^* is exactly 2-concave and therefore $\|\cdot\|_{(F^*)^{(1/2)}}$ is a concave function on $\mathbb{R}_{\geq 0}^n$ (we abuse notation and use $\|\cdot\|_{(F^*)^{(1/2)}}$ even though it is not a norm). Similar to the proof of [Theorem 7.7](#) we define a concave function $g : \text{PSID}^n \rightarrow \mathbb{R}_{\geq 0}$ and its associated convex level set L_g as follows

$$g(\mathbb{W}) \stackrel{\text{def}}{=} \|\text{diag}(\mathbb{W})^{1/2}\|_{F^*}^2 = \|\text{diag}(\mathbb{W})\|_{(F^*)^{(1/2)}}$$

$$L_g \stackrel{\text{def}}{=} \{\mathbb{W} \succeq 0 \mid g(\mathbb{W}) \geq 1\}.$$

[Theorem 7.4](#) implies that $\mathcal{N}_{F^*}(\mathbb{W})$ and $g(\mathbb{W})$ are equivalent within a factor of $\pi/2$. Thus by [Observation 3.4](#) it suffices to give a C -separation oracle for L_g since this would imply a $C \cdot \pi/2$ -approximate separation oracle for $\mathcal{L}(F)$ and therefore a $C \cdot M^{(2)}(X)^2 \cdot \pi/2$ -approximate separation oracle for $\mathcal{L}(X)$.

Observe further that it suffices to design a separation oracle for the convex set $L_{(F^*)^{(1/2)}}$ $\stackrel{\text{def}}{=} \{\xi \in \mathbb{R}_{\geq 0}^n \mid \|\xi\|_{(F^*)^{(1/2)}} \geq 1\}$. Indeed consider any $\mathbb{W} \in \text{PSID}^n \setminus L_g/C$ (if $\mathbb{W} \not\succeq 0$ then we may use the separation oracle of PSID^n). By definition $\text{diag}(\mathbb{W}) \in \mathbb{R}_{\geq 0}^n \setminus L_{(F^*)^{(1/2)}}$ and so there exists a hyperplane $\{\xi \mid \langle h, \xi \rangle = 1\}$ that separates $\text{diag}(\mathbb{W})$ from $L_{(F^*)^{(1/2)}}$. Thus the hyperplane $\{M \mid \langle \text{Diag}(h), M \rangle = 1\}$ separates \mathbb{W} from L_g .

To obtain a separation oracle for $L_{(F^*)^{(1/2)}}$, we combine [Lemma 7.8](#) (i.e., $(1 + \varepsilon)M^{(2)}(X)^2$ -approximate linear minimization can be done over $L_{(F^*)^{(1/2)}}$) with [Theorem 3.16](#) (i.e., linear minimization over an upward-closed set implies a separation oracle). In order to ensure linear minimization runs in $\text{poly}(n)$ time, we need to ensure that $L_{(F^*)^{(1/2)}}$ is inverse $(R, r, \mathbb{R}_{\geq 0}^n)$ -balanced for $R/r = \text{poly}(n)$. This can be verified from the following simple inequalities:

$$\|\tilde{\zeta}\|_{\infty} \leq \frac{\|\tilde{\zeta}\|_{(F^*)^{(1/2)}}}{\|e_1\|_{F^*}} \leq n^2 \cdot \|\zeta\|_{\infty} \quad \forall \zeta \in \mathbb{R}_{\geq 0}^n$$

where the first inequality uses monotonicity of $\|\cdot\|_{F^*}$ (in the entry-wise ordering) and the second inequality follows from triangle inequality. Thus we obtain an $O(M^{(2)}(X)^2)$ -approximate separation oracle for $L_{(F^*)^{(1/2)}}$ running in time $\text{poly}(n)$. \blacksquare

Finally we obtain

Theorem 7.12 (Maximization over Symmetric Norms under Type-2/Dual Cotype-2).

There are algorithms $\text{ALG}_1(A_1, \mathcal{O}_X)$ and $\text{ALG}_2(A_2, \mathcal{O}_X, \mathcal{O}_Y)$ such that if \mathcal{O}_X (resp. \mathcal{O}_Y) is an oracle computing a symmetric norm $(\mathbb{R}^n, \|\cdot\|_X)$ (resp. $(\mathbb{R}^m, \|\cdot\|_Y)$), then

- Quadratic: on any input $A_1 \in M_n(\mathbb{R})$, ALG_1 runs in time $\text{poly}(n, \text{bit}(A_1))$ and returns an α -approximate solution to $\text{Q}_{S_X}^{\max}(A_1)$ with probability $1 - 2^{-\Omega(n)}$, where

$$\alpha \lesssim T_2(X)^2 \cdot \tilde{C}_2(X^*)^4 \cdot \log^4 \tilde{C}_2(X^*) \leq \tilde{T}_2(X)^6 \log^4 \tilde{T}_2(X).$$

- Bilinear: on any input $A_2 \in M_{nd, mh}(\mathbb{R})$, ALG_2 runs in time $\text{poly}(n, m, \text{bit}(A_2))$ and returns a β -approximate solution to $\text{Op}_{S_X, S_Y}^{\max}(A_2)$ with probability $1 - 2^{-\Omega(n)}$, where

$$\beta \lesssim \max\{\tilde{C}_2(X^*)^6 \cdot \log^5 \tilde{C}_2(X^*), \tilde{C}_2(Y^*)^6 \cdot \log^5 \tilde{C}_2(Y^*)\}.$$

Proof. Without loss of generality we may assume $\|e_1\|_X = \|e_1\|_Y = 1$ since otherwise we can just renormalize. $\|\cdot\|_X$ (resp. $\|\cdot\|_Y$) is always $(n^{O(1)}, n^{-\Omega(1)})$ -balanced (resp. $(m^{O(1)}, m^{-\Omega(1)})$ -balanced) since

$$\|x\|_{\infty} \leq \frac{\|x\|_X}{\|e_1\|_X} \leq n \cdot \|x\|_{\infty} \quad \forall x \in \mathbb{R}^n$$

and $\|e_1\|_X = 1$ by assumption.

Combining [Theorem 4.11](#) with [Lemma 7.11](#) yields the desired theorem. \blacksquare

Remark 7.13. Combining [Proposition 5.3](#) and [Lemma 7.9](#) yields an alternate proof of [Theorem 7.12](#), achieving the approximation factor $(1 + \varepsilon) \cdot K_G \cdot M^{(2)}(X) \cdot M^{(2)}(Y) \cdot \max\{M^{(2)}(X)^2, M^{(2)}(Y)^2\}$. In fact the factor can be improved to $(1 + \varepsilon) \cdot K_G \cdot M^{(2)}(X) \cdot M^{(2)}(Y)$ if one allows non-uniform algorithms with quasipolynomial runtime.

We chose to present the above proof via [Theorem 4.11](#) to highlight the generality of our approach and furthermore because it generalizes well to the unitarily invariant case.

7.3 Approximation Algorithms for Unitarily Invariant Norms

In this section we will use [Theorem 4.11](#) to give constant factor approximation algorithms for quadratic (resp. bilinear) maximization over unitarily-invariant norms with bounded type-2 (resp. dual cotype-2) constant (assuming only oracles computing the norms being optimized over). To do so, we will design a separation oracle for the lower covariance region, for which we first require some additional preliminaries.

7.3.1 Preliminaries: Non-Commutative Khintchine Inequality

Let E be a symmetric norm over \mathbb{R}^n (i.e., invariant to permutation and flipping signs of entries). For a matrix $M \in M_{n,d}(\mathbb{R})$ we define the norm $\|M\|_{S_E} \stackrel{\text{def}}{=} \|\sigma(M)\|_E$ where $\sigma(M)$ denotes the vector of singular values of M (when $d \geq n$ we assume $\sigma(M)$ returns only the n largest singular values of M – the rest are of course 0's). We use the shorthand C_p for $C_{\ell_p^n}$.

For a finite sequence of matrices $(M_k)_{k=1}^K$ in $M_{n,d}(\mathbb{R})$, we define the norms (we refer the reader to [LPP91] for a more detailed discussion of the these norms and their associated properties quoted here):

$$\begin{aligned} \|(M_k)\|_{S_E(\ell_R^2)} &\stackrel{\text{def}}{=} \left\| \sqrt{\sum_k M_k^* M_k} \right\|_{S_E} \\ \|(M_k)\|_{S_E(\ell_L^2)} &\stackrel{\text{def}}{=} \left\| \sqrt{\sum_k M_k M_k^*} \right\|_{S_E} \\ \|(M_k)\|_{S_E(\ell_R^2) \vee S_E(\ell_L^2)} &\stackrel{\text{def}}{=} \max\{\|(M'_k)\|_{S_E(\ell_R^2)}, \|(M''_k)\|_{S_E(\ell_L^2)}\} \\ \|(M_k)\|_{S_E(\ell_R^2) + S_E(\ell_L^2)} &\stackrel{\text{def}}{=} \inf_{(M_k) = (M'_k) + (M''_k)} \|(M'_k)\|_{S_E(\ell_R^2)} + \|(M''_k)\|_{S_E(\ell_L^2)}. \end{aligned}$$

For a finite sequence of matrices $(\overline{M}_k)_{k=1}^K$ in $M_{n,d}(\mathbb{R})$, we define the inner product $\langle (M_k), (\overline{M}_k) \rangle \stackrel{\text{def}}{=} \sum_k \langle M_k, \overline{M}_k \rangle = \sum_k \text{Tr}(M_k \overline{M}_k^*)$. It is a standard fact that $S_E(\ell_R^2) + S_E(\ell_L^2)$ and $(S_{E^*}(\ell_R^2) \vee S_{E^*}(\ell_L^2))$ are dual to each other, i.e.,

$$\|(\overline{M}_k)\|_{(S_E(\ell_R^2) + S_E(\ell_L^2))^*} = \sup_{\|(M_k)\|_{S_E(\ell_R^2) + S_E(\ell_L^2)} \leq 1} \langle (M_k), (\overline{M}_k) \rangle = \|(\overline{M}_k)\|_{S_{E^*}(\ell_R^2) \vee S_{E^*}(\ell_L^2)}. \quad (71)$$

We require the following matrix khintchine inequality due to Lust-Piquard and Xu [LPX07]

Theorem 7.14 (2-concave Matrix Khintchine Inequality).

There is an absolute constant $C > 1$ such that for any exactly 2-concave symmetric norm E and any finite sequence of matrices $(M_k)_{k=1}^K$ in $M_{n,d}(\mathbb{R})$ we have

$$\frac{1}{C} \cdot \|(M_k)\|_{S_E(\ell_R^2) + S_E(\ell_L^2)} \leq \sqrt{\mathbb{E}[\|\sum_k \mathbf{g}_k \cdot M_k\|_{S_E}^2]} \leq \|(M_k)\|_{S_E(\ell_R^2) + S_E(\ell_L^2)},$$

where $(\mathbf{g}_k)_{k=1}^K$ is a sequence of i.i.d. standard Gaussians.

We define linear maps $T_L : \text{PSID}^{n \cdot d} \rightarrow M_n(\mathbb{R})$ (resp. $T_R : \text{PSID}^{n \cdot d} \rightarrow M_d(\mathbb{R})$) as

$$\begin{aligned} (T_R(\mathbb{X}))_{[i,j]} &\stackrel{\text{def}}{=} \sum_{\ell \in [n]} \mathbb{X}[(\ell, i), (\ell, j)] \\ (T_L(\mathbb{X}))_{[i,j]} &\stackrel{\text{def}}{=} \sum_{\ell \in [n]} \mathbb{X}[(i, \ell), (j, \ell)] \end{aligned}$$

Note that for any decomposition $\mathbb{X} = \sum_{k \in [K]} \text{vec}(M_k) \text{vec}(M_k)^*$ one has

$$\begin{aligned} T_R(\mathbb{X}) &= \sum_{k \in [K]} M_k^* M_k \\ T_L(\mathbb{X}) &= \sum_{k \in [K]} M_k M_k^* \end{aligned}$$

where $\text{vec}(H)$ denotes the (row-wise) nd -dimensional vector associated with the $n \times d$ matrix H .

7.3.2 Unitarily Invariant Type-2 Matrix Norms

In this section we will give an unconditional constant factor approximation algorithm for quadratic (resp. bilinear) maximization over unitarily invariant matrix norms with bounded type-2 (resp. dual cotype-2) constant.

We first construct a separation oracle for the lower covariance region of a unitarily invariant norm.

Lemma 7.15 (Separation Oracle for $\mathcal{L}(S_E)$).

For any (R, r) -balanced unitarily invariant matrix norm $\|\cdot\|_{S_E}$ over $M_{n,d}(\mathbb{R})$ the lower covariance region $\mathcal{L}(S_E)$ has a poly($n, d, \log R, \log 1/r, \text{bit}(x)$) time α -approximate separation oracle on input x , assuming access to an oracle \mathcal{O}_{S_E} computing $\|\cdot\|_{S_E}$, where $\alpha = O(M^{(2)}(E)^6) = O(\tilde{C}_2(E^*)^6 \log^6(\tilde{C}_2(E^*)))$.

Proof. Let F be the renorming of E that is symmetric, exactly 2-convex and satisfies $\|x\|_F \leq \|x\|_E \leq M^{(2)}(E) \cdot \|x\|_F$. For $\mathbb{W} \in \text{PSID}^{n \cdot d}$, we define a concave function g non-decreasing in the Loewner ordering as follows

$$g(\mathbb{W}) \stackrel{\text{def}}{=} \sup_{(M_k)} \|(M_k)\|_{S_{F^*}(\ell_R^2) + S_{F^*}(\ell_L^2)} \quad (72)$$

where the supremum runs over all decompositions

$$\mathbb{W} \succeq \sum_{k \in [2n^2]} \text{vec}(M_k) \text{vec}(M_k)^* .$$

By [Theorem 7.14](#) and monotonicity in the Loewner ordering of $\mathcal{N}_{S_{F^*}}(\cdot)$, we conclude $\mathcal{N}_{S_{F^*}}(\mathbb{W})^{1/2}$ and $g(\mathbb{W})$ are equal within an absolute constant. $g(\mathbb{W})$ is alternatively given by the following (not exactly computable) convex program

$$\begin{aligned} & \sup \text{Tr}(Z) \quad \text{s.t.} \\ & \|T_L(\mathbb{X})\|_{S_{F(1/2)}} , \|T_R(\mathbb{X})\|_{S_{F(1/2)}} \leq 1 \\ & \begin{bmatrix} \mathbb{X} & Z \\ Z^* & \mathbb{W}' \end{bmatrix} \succeq 0 \\ & \mathbb{W} \succeq \mathbb{W}' \\ & \mathbb{X}, Z \in M_{n \cdot d}(\mathbb{R}) . \end{aligned} \quad (73)$$

Before verifying equivalence of the two definitions, we note that concavity of g follows easily from (73) since whenever \mathbb{X}_1, Z_1 (resp. \mathbb{X}_2, Z_2) are feasible for \mathbb{W}_1 (resp. \mathbb{W}_2), $\lambda\mathbb{X}_1 + (1-\lambda)\mathbb{X}_2, \lambda Z_1 + (1-\lambda)Z_2$ are feasible for $\lambda\mathbb{W}_1 + (1-\lambda)\mathbb{W}_2$. It follows that $g(\lambda\mathbb{W}_1 + (1-\lambda)\mathbb{W}_2) \geq \lambda g(\mathbb{W}_1) + (1-\lambda)g(\mathbb{W}_2)$.

Claim 7.16. (72) = (73).

Proof. To show (72) \leq (73), consider any sequence (M_k) and let (\bar{M}_k) be any sequence satisfying (such a sequence exists by (71)),

$$\langle (\bar{M}_k), (M_k) \rangle = \|(M_k)\|_{S_{F^*}(\ell_R^2) + S_{F^*}(\ell_L^2)} \quad \text{where } \|(\bar{M}_k)\|_{S_F(\ell_L^2)}, \|(\bar{M}_k)\|_{S_F(\ell_R^2)} \leq 1 .$$

The claim then follows by considering the substitution

$$\mathbb{X} \stackrel{\text{def}}{=} \sum_k \text{vec}(\bar{M}_k) \text{vec}(\bar{M}_k)^* , \quad Z \stackrel{\text{def}}{=} \sum_k \text{vec}(\bar{M}_k) \text{vec}(M_k)^* , \quad \mathbb{W}' \stackrel{\text{def}}{=} \sum_k \text{vec}(M_k) \text{vec}(M_k)^* .$$

For the other direction, consider any $\mathbb{X}, Z, \mathbb{W}'$ feasible for \mathbb{W} . By the spectral theorem there exist sequences $(M_k)_{k \in [K]}, (\overline{M}_k)_{k \in [K]}$ for $K = 2n^2$, such that

$$\mathbb{X} = \sum_{k \in [K]} \text{vec}(\overline{M}_k) \text{vec}(\overline{M}_k)^*, \quad Z = \sum_{k \in [K]} \text{vec}(\overline{M}_k) \text{vec}(M_k)^*, \quad \mathbb{W}' = \sum_{k \in [K]} \text{vec}(M_k) \text{vec}(M_k)^*.$$

Since $\|(\overline{M}_k)\|_{\mathcal{S}_F(\ell_L^2)}, \|(\overline{M}_k)\|_{\mathcal{S}_F(\ell_R^2)} \leq 1$, we have

$$\text{Tr}(Z) = \langle (\overline{M}_k), (M_k) \rangle \leq \|(M_k)\|_{\mathcal{S}_{F^*}(\ell_R^2) + \mathcal{S}_{F^*}(\ell_L^2)}.$$

Taking (M_k) as the sequence in (72) then implies (72) \geq (73). \blacksquare

We define an upwards closed convex set $B \stackrel{\text{def}}{=} \{\mathbb{W} \succeq 0 \mid g(\mathbb{W}) \geq 1\}$ where convexity of B follows from concavity of g and upward-closure follows from monotonicity of g (in the Loewner ordering). By [Claim 7.16](#), we know that $\inf_{\mathbb{W} \in B} \langle A, \mathbb{W} \rangle$ is alternatively given by

$$\begin{aligned} & \inf \langle A, \mathbb{W} \rangle \quad \text{s.t.} \\ & \text{Tr}(Z) \geq 1 \\ & \|T_L(\mathbb{X})\|_{\mathcal{S}_{F(1/2)}}, \|T_R(\mathbb{X})\|_{\mathcal{S}_{F(1/2)}} \leq 1 \\ & \begin{bmatrix} \mathbb{X} & Z \\ Z^* & \mathbb{W}' \end{bmatrix} \succeq 0 \\ & \mathbb{W} \succeq \mathbb{W}' \\ & \mathbb{X}, Z \in M_{n,d}(\mathbb{R}). \end{aligned} \tag{74}$$

By [Proposition 3.12](#), it suffices to give an approximate separation oracle for

$$\{\mathbb{X} \in \text{PSID}^{n,d} \mid \|T_L(\mathbb{X})\|_{\mathcal{S}_{F(1/2)}} \leq 1\} \quad (\text{resp. } \|T_R(\mathbb{X})\|_{\mathcal{S}_{F(1/2)}}).$$

For this it suffices to give a separation oracle for the set $S \stackrel{\text{def}}{=} \{H \in \text{PSID}^n \mid H \in \text{Ball}(\mathcal{S}_{F(1/2)})\}$. By [Lemma 7.10](#) and [Lemma 7.9](#) for any fixed $\varepsilon > 0$, one can compute in oracle-polytime an $M^{(2)}(E)^4 \cdot (1 + \varepsilon)$ -approximate solution to $\sup_{H \in S} \langle M, H \rangle$ for any $M \in \text{PSID}^n$. So by linear function maximization duality for downward closed sets ([Theorem 3.14](#)) there is an oracle-polytime algorithm to find an $M^{(2)}(E)^4 \cdot (1 + \varepsilon)$ -approximate separation oracle for S . Then by [Proposition 3.12](#), we obtain an oracle-polytime algorithm to find an $O(M^{(2)}(E)^4)$ -approximate solution to (74) (i.e., linear function minimization over B).

Thus by linear function minimization duality for upward closed sets ([Theorem 3.16](#)), B admits an $M^{(2)}(E)^4 \cdot (1 + \varepsilon)$ -approximate separation oracle. Recalling that by [Theorem 7.14](#) $\mathcal{L}(S_E)$ and B are equivalent within $O(M^{(2)}(E))^2$ and applying [Observation 3.4](#) yields the claimed separation oracle for $\mathcal{L}(S_E)$. \blacksquare

We are now ready to prove the main result of this section.

Theorem 7.17 (Maximization over Unitarily Invariant Norms under Type-2/Dual Cotyep-2).

There are algorithms $\text{ALG}_1(A_1, \mathcal{O}_{S_X})$ and $\text{ALG}_2(A_2, \mathcal{O}_{S_X}, \mathcal{O}_{S_Y})$ such that if \mathcal{O}_{S_X} (resp. \mathcal{O}_{S_Y}) is an oracle computing a unitarily invariant matrix norm $(M_{n,d}(\mathbb{R}), \|\cdot\|_{S_X})$ (resp. $(M_{m,h}(\mathbb{R}), \|\cdot\|_{S_Y})$), then

- Quadratic: on any input $A_1 \in M_{nd}(\mathbb{R})$, ALG_1 runs in time $\text{poly}(n, d, \text{bit}(A_1))$ and returns an α -approximate solution to $\text{Q}_{S_X}^{\max}(A_1)$ with probability $1 - 2^{-\Omega(n)}$, where

$$\alpha \lesssim T_2(X)^2 \cdot \tilde{C}_2(X^*)^6 \cdot \log^6 \tilde{C}_2(X^*) \leq \tilde{T}_2(X)^8 \log^6 \tilde{T}_2(X).$$

- *Bilinear*: on any input $A_2 \in M_{nd, mh}(\mathbb{R})$, ALG_2 runs in time $\text{poly}(n, d, m, h, \text{bit}(A_2))$ and returns a β -approximate solution to $\text{Op}_{S_X, S_Y}^{\max}(A_2)$ with probability $1 - 2^{-\Omega(n)}$, where

$$\beta \lesssim \max\{\tilde{C}_2(X^*)^8 \cdot \log^7 \tilde{C}_2(X^*), \tilde{C}_2(Y^*)^8 \cdot \log^7 \tilde{C}_2(Y^*)\}.$$

Proof. Without loss of generality we may assume $\|e_1 e_1^*\|_{S_X} = \|e_1 e_1^*\|_{S_Y} = 1$ since otherwise we can just renormalize. $\|\cdot\|_{S_X}$ (resp. $\|\cdot\|_{S_Y}$) is always $(n^{O(1)}, n^{-\Omega(1)})$ -balanced (resp. $(m^{O(1)}, m^{-\Omega(1)})$ -balanced) since

$$\|M\|_{C_\infty} \leq \frac{\|M\|_{S_E}}{\|e_1 e_1^*\|_{S_E}} \leq \max\{n, d\} \cdot \|M\|_{C_\infty} \quad \forall M \in M_{n,d}(\mathbb{R})$$

and $\|e_1 e_1^*\|_{S_E} = 1$ by assumption.

Combining [Theorem 4.11](#) with [Lemma 7.15](#) yields the desired theorem. ■

8 Hardness in the Absence of Type-2

In this section, we prove that if a sequence of norms $(\|\cdot\|_{X^n}, \mathbb{R}^n)_n$ has $\tilde{T}_2(X^n) = n^{\Omega(1)}$, then a constant-factor approximation algorithm for $\text{Q}_{X^n}^{\max}(\cdot)$ will refute the (randomized) Small Set Expansion Hypothesis (SSEH) defined below. Given a d -regular graph $G = (V, E)$ with n vertices and $\delta \in [0, 1]$, let

$$\Phi(\delta) := \min_{S \subseteq V, |S| = \delta n} \frac{|E(S, V \setminus S)|}{d|S|}.$$

Definition 8.1 (Gap Small Set Expansion (Gap-SSE) [[RS10](#)]). *Given a regular graph $G = (V, E)$ and $\varepsilon, \delta > 0$, the Gap Small Set Expansion problem $\text{Gap-SSE}(\varepsilon, \delta)$ asks to distinguish between*

- YES: $\Phi(\delta) \leq \varepsilon$.
- NO: $\Phi(\delta) \geq 1 - \varepsilon$.

Hypothesis 8.2 (Small Set Expansion Hypothesis [[RS10](#)]). *For any $\varepsilon > 0$, there exists $\delta > 0$ such that $\text{Gap-SSE}(\varepsilon, \delta)$ is NP-hard.*

Similarly, in the sequel we will say “Gap-SSE does not admit an algorithm of runtime $T(n)$ ” to refer to the assumption that for every constant $\varepsilon > 0$, there exists $\delta > 0$ such that no algorithm of runtime $T(n)$ solves the distinguishing problem $\text{Gap-SSE}(\varepsilon, \delta)$.

8.1 SSE-Hardness of ℓ_p -Quadratic Maximization when $p < 2$

Barak et al. [[BBH⁺12](#)] proved hardness of approximating $2 \rightarrow 4$ norm assuming the SSEH.

We will require a hardness result for $2 \rightarrow q$ norm when q approaches 2 sufficiently slowly (i.e., $q - 2$ is subconstant in n). To do this we exploit the fact that in the reduction of [[BBH⁺12](#)] the optimal vector in the completeness case does not depend on q . We then use interpolation to argue that the soundness decays sufficiently slowly – here we make use of another special property of the reduction, namely the instance is a projector and therefore the $2 \rightarrow 2$ norm is bounded. We begin by stating the result of [[BBH⁺12](#)] (also see theorem 21 in [[BBB⁺19](#)] for details) with the desired additional properties.

Theorem 8.3 ([BBH⁺12]). *Assuming SSEH, for any sufficiently small $\delta > 0$, no polynomial time algorithm can distinguish between the following two cases given an input matrix $P \in M_n(\mathbb{R})$ satisfying $\|P\|_{2 \rightarrow 2} \leq 1$:*

$$\text{YES: } \|P\|_{2 \rightarrow q} \geq (n/(10 \cdot \delta))^{1/2-1/q} \quad \forall q \in (2, \infty).$$

$$\text{NO: } \|P\|_{2 \rightarrow 4} \leq 2 \cdot n^{1/4}/\delta^{1/8}.$$

By interpolation (Holder's inequality in this case) we obtain the following hardness result for q approaching 2.

Corollary 8.4. *Fix any $\theta \in (0, 1)$ and let $q_\theta \stackrel{\text{def}}{=} ((1 - \theta)/2 + \theta/4)^{-1} = 4/(2 - \theta)$. Assuming SSEH, for any $C > 1$, no polynomial time algorithm can approximate $\|\cdot\|_{2 \rightarrow q_\theta}$ within a factor better than $1 + C\theta$.*

Proof. By [Theorem 8.3](#) in the YES case we have

$$\|P\|_{2 \rightarrow q_\theta} \geq n^{1/2-1/q_\theta} / (10 \cdot \delta^{1/2-1/q_\theta}) = (n/(10 \cdot \delta))^{\theta/4}.$$

In the NO case by Holder's inequality we have

$$\|P\|_{2 \rightarrow q_\theta} \leq \|P\|_{2 \rightarrow 2}^{1-\theta} \cdot \|P\|_{2 \rightarrow 4}^\theta \leq (2n)^{\theta/4} / \delta^{\theta/8}.$$

Thus we obtain a gap of $1/(20 \cdot \delta)^{\theta/8} \geq 1 + (1/100) \cdot \theta \cdot \log 1/\delta$. Taking δ sufficiently small yields the claim. \blacksquare

We next amplify the gap in the previous hardness result, and also adapt it to obtain hardness results for ℓ_p -quadratic maximization when $1 < p < 2$ (even for the case when $p = p(n)$ approaches 2 as $n \rightarrow \infty$).

Proposition 8.5 (ℓ_p -Quadratic Maximization Hardness for $p < 2$).

- (1) *Consider any fixed $q > 2$. Assuming the SSEH and that $P \neq NP$, there is no polynomial time constant factor approximation algorithm for $\|\cdot\|_{2 \rightarrow q}$.*
- (2) *Consider any fixed $q > 2, \varepsilon > 0$. Assuming that Gap-SSE does not admit a quasi-polynomial time algorithm, there is no polynomial time $2^{\log^{1-\varepsilon} n}$ -approximation algorithm for $\|\cdot\|_{2 \rightarrow q}$.*
- (3) *Consider any increasing function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ satisfying $f(n) \leq n$, and let $g \stackrel{\text{def}}{=} f^{-1}$. For any constant $C > 1$ there is a reduction running in time $\text{poly}(g(\bar{m}^{O(1)}))$ from a size \bar{m} instance of Gap-SSE to the problem of obtaining a C -approximation to $\|\cdot\|_{2 \rightarrow 2 + \log f(n)/\log n}$.*
- (4) *All results above extend to the case of $q^* \rightarrow q$ norm.*
- (5) *All results above extend to the case of $Q_{q^*}^{\max}(B)$ (even for instances B with 0's on the diagonal).*

Proof. Since for any fixed $q > 2$ arbitrary constant SSEH-hardness of $2 \rightarrow q$ norm was shown in [BBH⁺12], (1) and (2) follow from amplifying the gap using the fact that $\|A^{\otimes t}\|_{2 \rightarrow q} = \|A\|_{2 \rightarrow q}^t$ (see [BGG⁺19]).

The proof of (3) proceeds again by amplifying the gap obtained in [Corollary 8.4](#), the caveat being that q now depends on the size of the matrix and so care must be taken in the calculations. To this end, consider any constant $C > 1$. Let $q(n) \stackrel{\text{def}}{=} 2 + \log f(n)/\log n$, $\bar{q}(m) \stackrel{\text{def}}{=} q(g(m)) =$

$2 + \log m / \log g(m)$ and let $t \stackrel{\text{def}}{=} \log C \cdot \log g(m) / \log m$. We set $m = m(n)$ so that $m^t = n$. By [Corollary 8.4](#) there is a reduction from Gap-SSE of size $\bar{m} = m^{\Omega(1)}$ to an instance A of $\ell_2^m \rightarrow \ell_{\bar{q}(m)}^m$ with gap $1 + 10 \log m / \log g(m)$. We then reduce to $\ell_2^n \rightarrow \ell_{\bar{q}(m)}^n$ by considering $A^{\otimes t}$ and as claimed we obtain a gap of at least

$$(1 + 10 \log m / \log g(m))^t \geq e^{t \log m / \log g(m)} = C.$$

Since $m^t = n$, we have $n = \text{poly}(g(m))$. Therefore the runtime of the reduction is $\text{poly}(g(\bar{m}^{O(1)}))$ as desired (where \bar{m} is the size of the Gap-SSE instance).

(4) follows from combining (1), (2) and (3) with the observation that $\|AA^*\|_{q^* \rightarrow q} = \|A\|_{2 \rightarrow q}^2$. Lastly, (5) follows from combining (4) with the following claim:

Claim 8.6. *For any $p \in [1, \infty]$, there is a polytime computable matrix $B = B(A)$ (with 0's on the diagonal) such that $Q_p^{\max}(B) \leq \|A\|_{p \rightarrow p^*} \leq 2 \cdot Q_p^{\max}(B)$.*

Proof. Consider any $A \in M_{m,n}(\mathbb{R})$ and let $B \in M_{m+n}(\mathbb{R})$ be given by

$$B \stackrel{\text{def}}{=} \frac{1}{2} \cdot \begin{pmatrix} 0 & A \\ A^* & 0 \end{pmatrix}$$

Let $p \oplus_\infty p$ be the norm over \mathbb{R}^{m+n} defined as $\|x \oplus y\|_{p \oplus_\infty p} \stackrel{\text{def}}{=} \max\{\|x\|_p, \|y\|_p\}$. Observe that $Q_{p \oplus_\infty p}^{\max}(B) = \|A\|_{p \rightarrow p^*}$. Further, it is easily checked that $\|\cdot\|_{p \oplus_\infty p}$ and ℓ_p^{m+n} are equivalent within a factor of 2. Thus $Q_p^{\max}(B)$ and $\|A\|_{p \rightarrow p^*}$ are within a factor of 2 as desired. ■

This completes the proof of [Proposition 8.5](#). ■

Remark 8.7. *Observe that if $q(n) = 2 + O(1/\log n)$, then $\|\cdot\|_{2 \rightarrow q(n)}$ is equivalent to $\|\cdot\|_{2 \rightarrow 2}$ within a constant and therefore admits a polytime constant factor approximation algorithm. Part (2) of [Proposition 8.5](#) provides an almost matching hardness result, i.e., assuming SSEH and ETH, for any $q(n) = 2 + \omega(\log \log n / \log n)$, there is no polytime constant factor approximation algorithm for $\|\cdot\|_{2 \rightarrow q(n)}$. We speculate that one can bridge this small gap with a hardness result for the case of $q(n) = 2 + \omega(1/\log n)$. One approach to establishing such a result would be to employ a more size-efficient gap-amplification procedure than tensoring (for example using random walks).*

8.2 SSE-Hardness of Approximation when Type-2 Fails

We begin with some preliminaries.

8.2.1 Preliminaries: Embedding Copies of ℓ_p^k in X

To obtain more general hardness results we will require the notion of *stable type*. For $1 \leq p \leq 2$, the stable type- p constant $ST_p(X)$ is defined analogously to $T_p(X)$ with p -stable random variables replacing rademachers. We also define a constant c_p which will appear in the sequel. Let $(E_i)_{i \in \mathbb{N}}$ be an independent sequence of exponential random variables and let $\Gamma_j \stackrel{\text{def}}{=} \sum_{i=1}^j E_i$. Then we define

$$c_p \stackrel{\text{def}}{=} \mathbb{E} \left[\left\| \sum_{j \in \mathbb{N}} \Gamma_j^{-1/p} \cdot \varepsilon_j \right\| \right]$$

where $(\varepsilon_j)_{j \in \mathbb{N}}$ is an independent sequence of rademacher random variables. It's not hard to check that $c_p \gtrsim 1$ (a fact we will use later).

We now state an important result of Pisier [Pis83] which states that for any $1 < p \leq 2$, norms with large stable type- p constant contain large-dimensional isomorphic copies of ℓ_p^k .

Theorem 8.8 (Pisier's Stable Type Theorem [Pis83]). *Let X be a norm over \mathbb{R}^n . Fix any $1 < p < 2$ and $\varepsilon > 0$, and let $\delta(\varepsilon, p) \stackrel{\text{def}}{=} \frac{2-p}{p} \left(\frac{\varepsilon c_p}{2^{p^*+2}} \right)^{p^*}$. Then for any $k < \delta(\varepsilon, p) \cdot (\text{ST}_p(X))^{1/p^*}$, there is an $n \times k$ matrix B such that for all $a \in \mathbb{R}^k$,*

$$(1 - \varepsilon) \cdot \|a\|_p \leq \|Ba\|_X \leq (1 + \varepsilon) \cdot \|a\|_p.$$

The existence of the matrix B above is shown by a probabilistic construction. Below we check that it can be efficiently sampled.

Corollary 8.9 (Efficient Sampling of [Pis83]).

Fix any $1 < p < 2$ and $\varepsilon > 0$, and let $\delta(\varepsilon, p) \stackrel{\text{def}}{=} \frac{2-p}{p} \left(\frac{\varepsilon c_p}{2^{p^+2}} \right)^{p^*}$. Let $X = (\mathbb{R}^n, \|\cdot\|_X)$ be a normed space and assume we are given a sequence of vectors $x_1, \dots, x_t \in \text{Ball}(X)$ with $t = n^{O(1)}$, and such that $\mathbb{E}[\|\sum_i \varepsilon_i x_i\|_X^2]^{1/2} \geq f(n) \cdot \sqrt{t}$. Then for $k(n) = \frac{2-p}{100p} \left(\frac{c_p \cdot f(n)}{(p^*+1)n^{1/p-1/2}} \right)^{p^*}$, there is an $n \times k(n)$ random matrix B that can be sampled in time $n^{O(1/(2-p))}$, and a scalar C , such that with probability $1 - 1/n^4$ it holds that for all $a \in \mathbb{R}^k$,*

$$0.9 \cdot C \cdot \|a\|_p \leq \|Ba\|_X \leq 1.1 \cdot C \cdot \|a\|_p.$$

Proof. We describe the distribution of an $n \times k$ random matrix B' as specified in a presentation of Theorem 8.8 in [MS86]. Let v be the \mathbb{R}^n valued random variable given by choosing uniformly at random from the set $\{x_1, \dots, x_t\}$. Let $(v_{i,j})_{i \in [k], j \in \mathbb{N}}$ be a double sequence of i.i.d. \mathbb{R}^n -valued random variables having the same distribution as v and similarly let $(\varepsilon_{i,j})_{i \in [k], j \in \mathbb{N}}$ be a double sequence of i.i.d. Rademacher random variables. Let B'_i denote the i -th column of B' . Then each B'_i (for $i \in [k]$) is given by

$$B'_i \stackrel{\text{def}}{=} \sum_{j=1}^{\infty} j^{-1/p} \cdot \varepsilon_{i,j} \cdot v_{i,j}.$$

Let $\tau \stackrel{\text{def}}{=} n^{50(2-p)}$ and let B be an $n \times k$ random matrix whose i -th column B_i is given by

$$B_i \stackrel{\text{def}}{=} \sum_{1 \leq j \leq \tau} j^{-1/p} \cdot \varepsilon_{i,j} \cdot v_{i,j}.$$

It is shown in [MS86] (see Theorem 13.12 and Proposition 13.14) that with probability $1 - 2^{-\Omega(k)}$, it holds that for all $a \in \mathbb{R}^k$,

$$0.99 \cdot C \cdot \|a\|_p \leq \|B'a\|_X \leq 1.01 \cdot C \cdot \|a\|_p.$$

where $C \stackrel{\text{def}}{=} \mathbb{E}[\|B'_1\|_X]$. We will show that w.h.p. $B'' \stackrel{\text{def}}{=} B' - B$ has $p \rightarrow X$ operator norm bounded by C/n . Since $\|B'a\|_X - \|B''a\|_X \leq \|B''a\|_X \leq \|B''\|_{p \rightarrow X} \cdot \|a\|_p \leq \|B'a\|_X + \|B''\|_{p \rightarrow X} \cdot \|a\|_p$, this would imply (for sufficiently large n) that

$$0.9 \cdot C \cdot \|a\|_p \leq \|Ba\|_X \leq 1.1 \cdot C \cdot \|a\|_p.$$

In what follows we assume WLOG that $k \leq n \leq t$. We proceed with showing $\|B''\|_{p \rightarrow X} \leq C/t^{3/2}$ with probability at least $1 - \Omega(1/t^4)$. It suffices to show that $\|B'_i\|_X \leq C/t^{5/2}$ with probability $1 - \Omega(1/t^5)$, since by triangle inequality we have $\|B''a\|_X \leq \|a\|_1 \cdot \max_{i \in [k]} \|B'_i\|_X \leq k^{1-1/p} \cdot \|a\|_p \cdot \max_{i \in [k]} \|B'_i\|_X$. To this end we have,

$$\begin{aligned}
& \mathbb{E}_{(v_{i,j})} \left[\mathbb{E}_{(\varepsilon_{i,j})} [\|B''_i\|_X^2] \right] \\
&= \mathbb{E}_{(v_{i,j})} \left[\mathbb{E}_{(\varepsilon_{i,j})} [\|\sum_{1 \leq j \leq \tau} j^{-1/p} \cdot \varepsilon_{i,j} \cdot v_{i,j}\|_X^2] \right] \\
&\leq n \cdot \mathbb{E}_{(v_{i,j})} [\sum_{\tau \leq j \leq \infty} j^{-2/p} \cdot \|v_{i,j}\|_X^2] && \text{(since } \tilde{T}_2(X) \leq \sqrt{n}\text{)} \\
&= n \cdot \sum_{\tau \leq j \leq \infty} j^{-2/p} \cdot \mathbb{E}_{(v_{i,j})} [\|v_{i,j}\|_X^2] \\
&\leq n \cdot (\sum_{\ell \in [t]} \|x_\ell\|_X^2 / t) \cdot \sum_{\tau \leq j \leq \infty} j^{-2/p} && \text{(by definition of } v\text{)} \\
&\leq n \cdot \max_{\ell \in [t]} \|x_\ell\|_X \cdot \sum_{\tau \leq j \leq \infty} j^{-2/p} \\
&\lesssim \frac{n}{\tau^{2/p-1}} \cdot \max_{\ell \in [t]} \|x_\ell\|_X && \text{(integration).}
\end{aligned}$$

Taking $\tau \stackrel{\text{def}}{=} t^{20/(2-p)}$ and applying Markov's inequality implies that with probability $1 - \Omega(1/t^5)$ we have $\|B'_i\|_X \leq \max_{\ell \in [t]} \|x_\ell\|_X / t^4$.

By an analogous argument to above (this time using the cotype inequality $\tilde{C}_2(X) \leq \sqrt{n}$), we conclude that

$$C = \mathbb{E} [\|B'_1\|_X] \gtrsim n^{-1/2} \cdot \mathbb{E}_{(v)} [\|v\|_X^2]^{1/2} \geq \max_{i \in [t]} \|x_i\|_X / (\sqrt{nt}).$$

This implies $\|B''_i\|_X \leq C/t^{5/2}$ as desired.

The truncation B of B' can be computed in time $t^{O(1/(2-p))} = n^{O(1/(2-p))}$ and this completes the proof. \blacksquare

8.2.2 The Final Reduction

We are finally ready to prove our main hardness result. The core idea is to use Pisier's stable type theorem in order to encode an ℓ_p -quadratic maximization instance (for $p < 2$) as an instance of quadratic maximization over a general norm X with large type-2. The proof is then completed by appealing to hardness results for ℓ_p -quadratic maximization.

Theorem 8.10 (Hardness of Quadratic Maximization when Type-2 Grows Sufficiently Fast).

Let $(\|\cdot\|_{X^n}, \mathbb{R}^n)_{n \in \mathbb{N}}$ be a sequence of normed spaces and let $f(n) \stackrel{\text{def}}{=} \tilde{T}_2(X^n)$. Assume there is a poly(n) time algorithm computing $\|\cdot\|_{X^n}$, and an algorithm that on input $n \in \mathbb{N}$ returns in poly(n) time a sequence of vectors $x_1, \dots, x_t \in \text{Ball}(X^n)$ satisfying

$$\mathbb{E} [\|\sum_{i \in [m]} \varepsilon_i \cdot x_i\|_{X^n}^2] \gtrsim f(n) \cdot \sqrt{t}. \quad (75)$$

Then we have

- (1) If $f(n) = n^{\Omega(1)}$, then assuming the SSEH and that $\text{NP} \not\subseteq \text{BPP}$, there is no polynomial time algorithm approximating $Q_{X^n}^{\max}(\cdot)$ within a constant factor.

(2) If $f(n) = n^{\Omega(1)}$, then assuming Gap-SSE cannot be solved in randomized quasi-polynomial time, there is no polynomial time algorithm approximating $Q_{X^n}^{\max}(\cdot)$ within a factor of $2^{\log^{1-\epsilon} n}$.

(3) If for every constant $C > 1$, $f(f^{1/C}(n^{1/C})) = \log^{\omega(1)} n$, then assuming Gap-SSE $\notin \bigcap_{\epsilon > 0} \text{RTIME}(2^{n^\epsilon})$, there is no polynomial time algorithm approximating $Q_{X^n}^{\max}(\cdot)$ within a constant factor.

Proof. We proceed by reducing $Q_{\mathbb{R}^k}^{\max}(\cdot)$ to quadratic maximization over a subspace of X^n . Applying [Observation 6.13](#) then completes the proof. We conclude (1) and (2) by appropriately choosing $p = p(n)$ and $k = k(n)$ at the end.

By [Corollary 8.9](#), setting $k(n) \stackrel{\text{def}}{=} \frac{2-p}{100p} \left(\frac{c_p f(n)}{(p^*+1)n^{1/p-1/2}} \right)^{p^*}$, there is an $n \times k(n)$ random matrix B that can be sampled in time $n^{O(1/(2-p))}$, and a scalar C such that with probability $1 - o(1)$, it holds that for all $a \in \mathbb{R}^k$,

$$0.9 \cdot C \cdot \|a\|_p \leq \|Ba\|_{X^n} \leq 1.1 \cdot C \cdot \|a\|_p. \quad (76)$$

Let B^\dagger be the $k \times n$ matrix defined as the linear map that maps $x \in \mathbb{R}^n$ to $a \in \mathbb{R}^k$ where a is such that Ba is the projection of x onto the column span of B (note that a is unique since the columns of B must be linearly independent in order to satisfy (76)). Let E be the column span of B . Then we have for any $H \in M_k(\mathbb{R})$

$$\sup_{x \in E \setminus \{0\}} \frac{\langle x, (B^\dagger)^* H B^\dagger x \rangle}{\|x\|_{X^n}^2} = \sup_{a \in \mathbb{R}^k \setminus \{0\}} \frac{\langle a, Ha \rangle}{\|Ba\|_{X^n}^2} \asymp \frac{1}{C^2} \cdot \sup_{a \in \mathbb{R}^k \setminus \{0\}} \frac{\langle a, Ha \rangle}{\|a\|_p^2}$$

where the final step follows from (76). Thus $H \mapsto (B^\dagger)^* H B^\dagger$ is a valid reduction from quadratic maximization over ℓ_p^k to quadratic maximization over the subspace E of X^n .

We are now ready to set parameters. Let $p(n) \stackrel{\text{def}}{=} 2(1 + \log f(n)/(2 \log n))^{-1}$ (so that $f(n)/n^{1/p-1/2} = \sqrt{f(n)}$). We then have the lower estimate

$$k(n) \geq f^{1/5}(n) / \log n \quad (77)$$

where the choice of $1/5$ above is not special - any sufficiently small constant will do.

For (1), it follows from (77) that if $f(n) = n^{\Omega(1)}$, then $k(n) = n^{\Omega(1)}$ and thus composing with [Proposition 8.5](#) part (1), we obtain a reduction from a size m instance of Gap-SSE to $Q_{X^n}^{\max}(\cdot)$ where $n = m^{O(1)}$. The runtime of the reduction is $n^{O(1/(2-p(n)))} = n^{O(1)} = m^{O(1)}$. This completes the proof of (1).

For (2), it follows from (77) that if $f(n) = n^{\Omega(1)}$, then $k(n) = n^{\Omega(1)}$ and thus composing with [Proposition 8.5](#) part (2), we obtain a reduction from a size m instance of Gap-SSE to $Q_{X^n}^{\max}(\cdot)$ where $n = m^{\log^{O(1)} m}$. The runtime of the reduction is $n^{O(1/(2-p(n)))} = n^{O(1)} = m^{\log^{O(1)} m}$. This completes the proof of (2).

For (3), it follows from (77) that if for every constant $C > 1$, $f(f^{1/C}(n^{1/C})) = \log^{\omega(1)} n$, then $k(n) \geq f^{1/5}(n) / \log n \geq f^{1/6}(n)$ (where we use the fact that $c_p \gtrsim 1$) and thus composing with [Proposition 8.5](#) part (3), we obtain a reduction from a size m instance of Gap-SSE to $Q_{X^n}^{\max}(\cdot)$ where $n = g(g(m^{O(1)})^{O(1)})$ and $g = f^{-1}$. Finally since for every $C > 1$, $f(f^{1/C}(n^{1/C})) = \log^{\omega(1)} n$, it follows that $n = g(g(m^{O(1)})^{O(1)}) \leq 2^{m^{o(1)}}$. The runtime of the reduction is $n^{O(1/(2-p(n)))} = n^{O(1/\log n)} = (2^{m^{o(1)}})^{m^{o(1)}} = 2^{m^{o(1)}}$, which completes the proof of (3). ■

Remark 8.11.

1. *Theorem 1.3* in the introduction states that if $\tilde{T}_2(X) = n^{\Omega(1)}$ then $Q_X^{\max}(\cdot)$ cannot be approximated within a constant in polynomial time assuming the (randomized) small set expansion hypothesis. This is implied by the above proof (with parameters set differently). Above, we chose to highlight two settings of parameters we believed to be most interesting.
2. If $Q_{X^n}^{\max}(\cdot)$ admits a polytime constant factor approximation, then (3) above necessitates in particular that $\tilde{T}_2(X^n)$ grows slower than every function in the following series: $2^{\log^\varepsilon n}, 2^{2^{\log \log^\varepsilon n}}, \dots$.
3. The assumption (75) of *Theorem 8.10* may be removed in any of (1), (2), (3) provided one replaces the assumption “there is no algorithm for Gap-SSE with runtime $T(n)$ ” by the stronger hypothesis that “Gap-SSE cannot be solved by circuits of size $T(n)$ ”. In (1) for instance, this is equivalent to assuming the SSEH and that $\text{NP} \not\subseteq \text{P}/\text{poly}$.
4. By [FLM77] the type-2 constant of an n -dimensional normed space is attained up to a universal constant factor for $t = O(n^2)$ vectors (see also [JN09] for the proof of this). In fact, by [TJ79] we can take $t = O(n)$. The assumption (75) of *Theorem 8.10* states that such vectors can be found efficiently. It is open whether or not it is possible to find such vectors using only polynomially many calls to the assumed membership oracle for X ; this would amount to making the proofs in [FLM77] or [TJ79] algorithmic. If this were possible, then we could remove assumption (75). This nuance is secondary to the main content of *Theorem 8.10*, namely to demonstrate that type-2 is inherently linked to the computational complexity of quadratic maximization. We note that if one were only interested in removing the assumption that vectors satisfying (75) could be found efficiently, then it would suffice to do so using polynomially many calls to an oracle that approximates $Q_{X^n}^{\max}(\cdot)$. But, it is an independently interesting question to understand when it is possible to approximate type and cotype constants efficiently.

9 Oracle Lower Bound for General Type-2 Norms

Our construction is identical to that of Brieden et al. [BGK⁺01] who gave a query lower bound for approximating the ℓ_2 -diameter of a convex body in the membership oracle model with the caveat that a constant bound on Type-2 needs to be verified. We include the full analysis for completeness.

9.1 The Construction

Fix any constant $\delta < 1/2$ and let $r \stackrel{\text{def}}{=} g/n^{1/2-\delta} \in \mathbb{R}^n$ where g is a standard Gaussian random vector (so that r has length n^δ with probability $1 - o(1)$). Let $\|\cdot\|_X$ be the norm whose unit ball is $\text{conv}(S^{n-1} \cup \{\pm r\})$.

9.2 Oracle Lower Bound

We now show that PSD quadratic maximization over general norms with bounded type-2 (and therefore also bilinear maximization) cannot be approximated (within even $n^{o(1)}$).

Theorem 9.1 (Bounded Type-2 Hardness).

Consider constants $\delta \in (0, 1/2)$ and $\varepsilon < 1 - 2\delta$. Then for the random norm X over \mathbb{R}^n defined as above

(1) $\tilde{T}_2(\mathbf{X}) \lesssim 1$.

(2) For any fixed set $S \subseteq \mathbb{R}^n$ of size at most 2^{n^ε} , $\text{Ball}(\mathbf{X}) \cap S = \text{Ball}(\ell_2^n) \cap S$ with probability $1 - o(1)$.

(3) $\text{Op}_{\mathbf{X}, \mathbf{X}}^{\max}(\mathbf{I}) = \text{Q}_{\mathbf{X}}^{\max}(\mathbf{I}) = \Omega(n^{2\delta})$ with probability $1 - o(1)$, while on the other hand $\text{Op}_{\ell_2^n, \ell_2^n}^{\max}(\mathbf{I}) = \text{Q}_{\ell_2^n}^{\max}(\mathbf{I}) = 1$.

Proof. The squared ℓ_2 -radius of $\text{Ball}(\mathbf{X})$ is $\Omega(n^{2\delta})$ w.h.p. by construction. (3) then follows from observing that $\text{Q}_{\mathbf{X}}^{\max}(\mathbf{I})$ is precisely the squared ℓ_2 -radius of $\text{Ball}(\mathbf{X})$.

We next establish (1). Let $x_1, \dots, x_m \in \mathbb{R}^n$ be a sequence of vectors. We may write $x_i / \|x_i\|_{\mathbf{X}}$ as a convex combination $x_i / \|x_i\|_{\mathbf{X}} = \lambda_i \cdot r + (1 - \lambda_i) \cdot y_i$ where $y_i \in \text{Ball}(\ell_2^n)$. Then for a sequence $(\mathbf{g}_i)_{i \in [m]}$ of i.i.d. standard Gaussians we have

$$\begin{aligned}
& \mathbb{E} \left[\left\| \sum_i \mathbf{g}_i x_i \right\|_{\mathbf{X}} \right] \\
& \leq \mathbb{E} \left[\left\| \sum_i \mathbf{g}_i \cdot \|x_i\|_{\mathbf{X}} \cdot (1 - \lambda_i) \cdot y_i \right\|_{\mathbf{X}} \right] + \mathbb{E} \left[\left\| \sum_i \mathbf{g}_i \cdot \|x_i\|_{\mathbf{X}} \cdot \lambda_i \cdot r \right\|_{\mathbf{X}} \right] \\
& \leq \mathbb{E} \left[\left\| \sum_i \mathbf{g}_i \cdot \|x_i\|_{\mathbf{X}} \cdot (1 - \lambda_i) \cdot y_i \right\|_2 \right] + \mathbb{E} \left[\left\| \sum_i \mathbf{g}_i \cdot \|x_i\|_{\mathbf{X}} \cdot \lambda_i \cdot r \right\|_{\mathbf{X}} \right] \\
& \leq \sqrt{\sum_i \|x_i\|_{\mathbf{X}}^2} + \mathbb{E} \left[\left\| \sum_i \mathbf{g}_i \cdot \|x_i\|_{\mathbf{X}} \cdot \lambda_i \cdot r \right\|_{\mathbf{X}} \right] \\
& = \sqrt{\sum_i \|x_i\|_{\mathbf{X}}^2} + \|r\|_{\mathbf{X}} \cdot \mathbb{E} \left[\left| \sum_i \mathbf{g}_i \cdot \|x_i\|_{\mathbf{X}} \cdot \lambda_i \right| \right] \\
& = \sqrt{\sum_i \|x_i\|_{\mathbf{X}}^2} + \mathbb{E} \left[\left| \sum_i \mathbf{g}_i \cdot \|x_i\|_{\mathbf{X}} \cdot \lambda_i \right| \right] \\
& \leq 2 \cdot \sqrt{\sum_i \|x_i\|_{\mathbf{X}}^2}.
\end{aligned}$$

Lastly by Kahane-Khintchine inequality [Theorem 2.1](#) we have $\mathbb{E} \left[\left\| \sum_i \mathbf{g}_i x_i \right\|_{\mathbf{X}} \right] \gtrsim \mathbb{E} \left[\left\| \sum_i \mathbf{g}_i x_i \right\|_{\mathbf{X}}^2 \right]^{1/2}$. This completes the proof of (1).

Finally we shall establish (2). Consider any any fixed set $S \subseteq \mathbb{R}^n$ of size at most 2^{n^ε} (ε to be chosen later) and let $S' \stackrel{\text{def}}{=} \{v \in S \mid \|v\|_2 > 1\}$. Since $\text{Ball}(\ell_2^n) \subseteq \text{Ball}(\mathbf{X})$, we need only show that $S' \cap \text{Ball}(\mathbf{X}) = \emptyset$ with probability $1 - o(1)$. We do so by exhibiting dual witnesses. Note that for any $v \in S'$, we have $\langle v, v / \|v\|_2 \rangle > 1$. So if $v / \|v\|_2 \in \text{Ball}(\mathbf{X}^*)$, then we conclude $v \notin \text{Ball}(\mathbf{X})$. Thus in order to prove (2), it suffices to show that with probability $1 - o(1)$, the set $\widehat{S} \stackrel{\text{def}}{=} \{v / \|v\|_2 \mid v \in S\}$ is contained in $\text{Ball}(\mathbf{X}^*)$.

To this end note that $\text{Ball}(\mathbf{X}^*)$ is simply given by

$$\{\zeta \in \text{Ball}(\ell_2^n) \mid |\langle \zeta, r \rangle| \leq 1\} = \{\zeta \in \text{Ball}(\ell_2^n) \mid |\langle \zeta, \mathbf{g} \rangle| \leq n^{1/2-\delta}\}.$$

Thus showing that $\widehat{S} \subseteq \text{Ball}(\mathbf{X}^*)$ is equivalent to showing that $\sup_{\zeta \in \widehat{S}} |\langle \mathbf{g}, \zeta \rangle| \leq n^{1/2-\delta}$. Since for each $\zeta \in \widehat{S}$, $\langle \mathbf{g}, \zeta \rangle$ is a standard Gaussian, a standard union bound argument yields that with probability $1 - 1/|\widehat{S}|$,

$$\sup_{\zeta \in \widehat{S}} |\langle \mathbf{g}, \zeta \rangle| \lesssim \sqrt{\log |\widehat{S}|} \lesssim n^{\varepsilon/2}.$$

Thus it suffices to take $\varepsilon < 1 - 2\delta$. This completes the proof. \blacksquare

Remark 9.2. Note that the above theorem rules out the possibility of any approximation algorithm for PSD quadratic maximization assuming only membership access to $\text{Ball}(X)$. This follows from taking the set S to be the set of queries made by the algorithm.

We conclude by noting that the above hardness result applies even to quadratic maximization of matrices with 0's on the diagonal (this is a special case considered frequently in the literature). Let

$$A \stackrel{\text{def}}{=} \frac{1}{2} \cdot \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix}.$$

Further let $\|\cdot\|_{X \oplus_\infty X}$ and $\|\cdot\|_{\ell_2^n \oplus_\infty \ell_2^n}$ be the norms over \mathbb{R}^{2n} given by $\|(x, y)\|_{X \oplus_\infty X} = \max\{\|x\|_X, \|y\|_X\}$ and $\|(x, y)\|_{\ell_2^n \oplus_\infty \ell_2^n} = \max\{\|x\|_2, \|y\|_2\}$. From the simple observation that $\sup_{x, y \in \text{Ball}(X)} \langle x, y \rangle = Q_{X \oplus_\infty X}^{\max}(A)$, we obtain the following corollary.

Corollary 9.3 (Bounded Type-2 Hardness even with 0's on the Diagonal).

Consider constants $\delta \in (0, 1/2)$ and $\varepsilon < 1 - 2\delta$. Then for the matrix A and the norms $X \oplus_\infty X, \ell_2^n \oplus_\infty \ell_2^n$ over \mathbb{R}^{2n} defined as above, we have

- (1) All diagonal entries of A are 0.
- (2) $\tilde{T}_2(X \oplus_\infty X), \tilde{T}_2(\ell_2^n \oplus_\infty \ell_2^n) \lesssim 1$.
- (3) For any fixed set $S \subseteq \mathbb{R}^n$ of size at most 2^{n^ε} , $\text{Ball}(X \oplus_\infty X) \cap S = \text{Ball}(\ell_2^n \oplus_\infty \ell_2^n) \cap S$ with probability $1 - o(1)$.
- (4) $Q_{X \oplus_\infty X}^{\max}(A) = \Omega(n^{2\delta})$ w.h.p., while on the other hand $Q_{\ell_2^n \oplus_\infty \ell_2^n}^{\max}(A) = 1$.

References

- [ABH⁺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. [6, 15](#)
- [ABS98] Noga Alon, Ravi Boppana, and Joel Spencer. An asymptotic isoperimetric inequality. *Geometric & Functional Analysis GAFA*, 8(3):411–436, 1998. [16](#)
- [AK06] Fernando Albiac and Nigel J Kalton. *Topics in Banach space theory, volume 233 of Graduate Texts in Mathematics*. Springer, New York, 2006. [53](#)
- [Alo06] Noga Alon. Maximizing a quadratic form on the ℓ_1^n unit ball. Unpublished manuscript, 2006. [6](#)
- [ALS⁺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. [12](#)
- [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. [15](#)
- [AN06] Noga Alon and Assaf Naor. Approximating the cut-norm via Grothendieck's inequality. *SIAM J. Comput.*, 35(4):787–803, 2006. [15](#)

- [ANN⁺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. [12](#)
- [ANN⁺18] Alexandr Andoni, Assaf Naor, Aleksandar Nikolov, Ilya Razenshteyn, and Erik Waingarten. Hölder homeomorphisms and approximate nearest neighbors. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 159–169. IEEE, 2018. [57](#), [61](#), [62](#)
- [BBB⁺19] Frank Ban, Vijay Bhattiprolu, Karl Bringmann, Pavel Kolev, Euiwoong Lee, and David P Woodruff. A ptas for lp-low rank approximation. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 747–766. SIAM, 2019. [74](#)
- [BBC⁺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. [12](#)
- [BBH⁺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. [16](#), [74](#), [75](#)
- [BGG⁺19] Vijay Bhattiprolu, Mrinalkanti Ghosh, Venkatesan Guruswami, Euiwoong Lee, and Madhur Tulsiani. Approximability of $p \rightarrow q$ matrix norms: generalized krivine rounding and hypercontractive hardness. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1358–1368. SIAM, 2019. [50](#), [75](#)
- [BGK⁺01] Andreas Brieden, Peter Gritzmann, Ravindran Kannan, Victor Klee, László Lovász, and Miklós Simonovits. Deterministic and randomized polynomial-time approximation of radii. *Mathematika*, 48(1-2):63–105, 2001. [80](#)
- [Bha97] Rajendra Bhatia. *Matrix analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997. [14](#)
- [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. [16](#)
- [BL76] Jöran Bergh and Jörgen Löfström. *Interpolation spaces. An introduction*. Springer-Verlag, Berlin-New York, 1976. Grundlehren der Mathematischen Wissenschaften, No. 223. [11](#)
- [BMMN13] Mark Braverman, Konstantin Makarychev, Yury 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. [15](#)
- [BR95] Mihir Bellare and Phillip Rogaway. The complexity of approximating a nonlinear program. *Math. Programming*, 69(3, Ser. A):429–441, 1995. [15](#)

- [Bri02] Andreas Brieden. Geometric optimization problems likely not contained in \mathbb{APX} . *Discrete Comput. Geom.*, 28(2):201–209, 2002. [15](#)
- [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. [16](#)
- [BSS18] Jarosław 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. [13](#)
- [BV04] Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004. [53](#), [55](#)
- [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. [13](#)
- [CS19b] Deeparnab Chakrabarty and Chaitanya Swamy. Simpler and better algorithms for minimum-norm load balancing. *arXiv preprint arXiv:1905.00044*, 2019. [13](#)
- [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. [15](#)
- [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. [12](#)
- [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. [16](#)
- [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. [15](#)
- [DPR72] Ed Dubinsky, A. Pełczyński, and H. P. Rosenthal. On Banach spaces X for which $\Pi_2(\mathcal{L}_\infty, X) = B(\mathcal{L}_\infty, X)$. *Studia Math.*, 44:617–648, 1972. [3](#)
- [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. [14](#)
- [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. [15](#)
- [FLM77] Tadeusz Figiel, Joram Lindenstrauss, and Vitali D Milman. The dimension of almost spherical sections of convex bodies. *Acta Mathematica*, 139(1):53–94, 1977. [46](#), [80](#)

- [GLS93] M Grotscchel, L Lovász, and A Schrijver. Geometric algorithms and combinatorial optimization, 1993. 10, 29
- [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. 16
- [Gro53] Alexander Grothendieck. Résumé de la théorie métrique des produits tensoriels topologiques. *Bol. Soc. Mat. São Paulo* 8, 1953. 15, 16, 48
- [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. 2
- [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. 14
- [Haa85] Uffe Haagerup. The grothendieck inequality for bilinear forms on C*-algebras. *Advances in Mathematics*, 56(2):93–116, 1985. 16
- [HH06] Eran Halperin and Elad Hazan. HAPLOFREQ—estimating haplotype frequencies efficiently. *J. Comput. Biol.*, 13(2):481–500 (electronic), 2006. 15
- [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. 16
- [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. 16
- [JN09] William B. Johnson and Assaf Naor. The Johnson-Lindenstrauss lemma almost characterizes Hilbert space, but not quite. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 885–891. SIAM, Philadelphia, PA, 2009. 80
- [Kah64] Jean-Pierre Kahane. Sur les sommes vectorielles sigma plus minus un. *COMPTEs RENDUS HEBDOMADAIRES DES SEANCES DE L ACADEMIE DES SCIENCES*, 259(16):2577, 1964. 17
- [Kat98] Eleni Katirtzoglou. Type and cotype of musielak-orlicz sequence spaces. *Journal of mathematical analysis and applications*, 226(2):431–455, 1998. 12
- [KN12] Subhash Khot and Assaf Naor. Grothendieck-type inequalities in combinatorial optimization. *Communications on Pure and Applied Mathematics*, 65(7):992–1035, 2012. 6, 15
- [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. 2, 6
- [KO09] Subhash Khot and Ryan O’Donnell. SDP gaps and UGC-hardness for Max-Cut-Gain. *Theory OF Computing*, 5:83–117, 2009. 15

- [Kou91] Omran Kouba. On the interpolation of injective or projective tensor products of banach spaces. *Journal of Functional Analysis*, 96(1):38–61, 1991. [61](#)
- [Kri73] JL Krivine. Théoremes de factorisation dans les espaces réticulés. *Séminaire Analyse fonctionnelle (dit "Maurey-Schwartz")*, pages 1–22, 1973. [15](#), [49](#), [64](#), [66](#)
- [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. [12](#)
- [LP68] Joram Lindenstrauss and Aleksander Pełczyński. Absolutely summing operators in \mathcal{L}_p -spaces and their applications. *Studia Mathematica*, 29(3):275–326, 1968. [15](#), [53](#)
- [LPP91] Françoise Lust-Piquard and Gilles Pisier. Noncommutative Khintchine and Paley inequalities. *Ark. Mat.*, 29(2):241–260, 1991. [16](#), [71](#)
- [LPX07] Françoise Lust-Piquard and Quanhua Xu. The little grothendieck theorem and khintchine inequalities for symmetric spaces of measurable operators. *Journal of Functional Analysis*, 244(2):488–503, 2007. [71](#)
- [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. [16](#)
- [Lue97] David G Luenberger. *Optimization by vector space methods*. John Wiley & Sons, 1997. [53](#), [55](#)
- [Mau73] B Maurey. Type et cotype dans les espaces munis de structures locales inconditionnelles. *Séminaire Analyse fonctionnelle (dit "Maurey-Schwartz")*, pages 1–25, 1973. [65](#)
- [Mau74] Bernard Maurey. Théorèmes de factorisation pour les opérateurs à valeurs dans un espace l_p . *Séminaire Analyse fonctionnelle (dit)*, pages 1–5, 1974. [51](#)
- [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. [3](#)
- [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. [4](#), [18](#)
- [MR14] Andrea Montanari and Emile Richard. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems*, pages 2897–2905, 2014. [16](#)
- [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. [3](#), [77](#)
- [Nes98] Yurii Nesterov. Semidefinite relaxation and nonconvex quadratic optimization. *Optimization methods and software*, 9(1-3):141–160, 1998. [2](#), [16](#), [50](#), [64](#), [66](#)
- [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. [14](#), [16](#), [51](#), [64](#)

- [NS09] Assaf Naor and Gideon Schechtman. An approximation scheme for quadratic form maximization on convex bodies. 2009. [2](#), [13](#), [16](#), [63](#), [65](#)
- [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. [16](#)
- [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. [4](#), [10](#), [41](#), [51](#), [54](#)
- [Pis83] Gilles Pisier. On the dimension of the ℓ_p^n -subspaces of banach spaces, for $1 \leq p < 2$. *Transactions of the American Mathematical Society*, 276(1):201–211, 1983. [77](#)
- [Pis86] Gilles Pisier. *Factorization of linear operators and geometry of Banach spaces*. American Mathematical Soc., 1986. [7](#), [10](#), [41](#), [50](#), [53](#), [65](#)
- [Pis12] Gilles Pisier. Grothendieck’s theorem, past and present. *Bulletin of the American Mathematical Society*, 49(2):237–323, 2012. [15](#), [49](#)
- [RRS16] Prasad Raghavendra, Satish Rao, and Tselil Schramm. Strongly refuting random csp’s below the spectral threshold. *arXiv preprint arXiv:1605.00058*, 2016. [16](#)
- [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. [15](#)
- [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. [6](#), [74](#)
- [Rud73] Walter Rudin. *Functional analysis*. McGraw-Hill Book Co., New York-Düsseldorf-Johannesburg, 1973. McGraw-Hill Series in Higher Mathematics. [3](#)
- [RV15] Oded Regev and Thomas Vidick. Quantum xor games. *ACM Transactions on Computation Theory (ToCT)*, 7(4):15, 2015. [16](#)
- [SV15] Yaron Singer and Jan Vondrák. Information-theoretic lower bounds for convex optimization with erroneous oracles. In *Advances in Neural Information Processing Systems*, pages 3204–3212, 2015. [33](#)
- [SWY⁺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. [12](#)
- [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. [12](#)
- [TJ79] Nicole Tomczak-Jaegermann. Computing 2-summing norm with few vectors. *Ark. Mat.*, 17(2):273–277, 1979. [80](#)

- [Tro09] Joel A Tropp. Column subset selection, matrix factorization, and eigenvalue optimization. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 978–986. Society for Industrial and Applied Mathematics, 2009. [48](#), [49](#)