

Algorithmic and optimization aspects of Brascamp-Lieb inequalities, via Operator Scaling

Ankit Garg ^{*} Leonid Gurvits[†] Rafael Oliveira [‡] Avi Wigderson [§]

December 17, 2017

Abstract

The celebrated Brascamp-Lieb (BL) inequalities [BL76, Lie90], and their reverse form of Barthe [Bar98], are an important mathematical tool, unifying and generalizing numerous inequalities in analysis, convex geometry and information theory, with many used in computer science. While their structural theory is very well understood, far less is known about computing their main parameters (which we later define below). Prior to this work, the best known algorithms for any of these optimization tasks required at least exponential time. In this work, we give polynomial time algorithms to compute: (1) Feasibility of BL-*datum*, (2) Optimal BL-*constant*, (3) Weak separation oracle for BL-*polytopes*.

What is particularly exciting about this progress, beyond the better understanding of BL-inequalities, is that the objects above naturally encode rich families of optimization problems which had no prior efficient algorithms. In particular, the BL-constants (which we efficiently compute) are solutions to *non-convex* optimization problems, and the BL-polytopes (for which we provide efficient membership and separation oracles) are linear programs with *exponentially many* facets. Thus we hope that new combinatorial optimization problems can be solved via reductions to the ones above, and make modest initial steps in exploring this possibility.

Our algorithms are obtained by a simple efficient reduction of a given BL-*datum* to an instance of the *Operator Scaling* problem defined by [Gur04]. To obtain the results above, we utilize the two (very recent and different) algorithms for the operator scaling problem [GGOW16, IQS15]. Our reduction implies algorithmic versions of many of the known structural results on BL-inequalities, and in some cases provide proofs that are different or simpler than existing ones. Further, the analytic properties of the [GGOW16] algorithm provide new, effective bounds on the magnitude and continuity of BL-constants; prior work relied on compactness, and thus provided no bounds.

On a higher level, our application of operator scaling algorithm to BL-inequalities further connects analysis and optimization with the diverse mathematical areas used so far to motivate and solve the operator scaling problem, which include commutative invariant theory, non-commutative algebra, computational complexity and quantum information theory.

^{*}Microsoft Research New England, email: garga@microsoft.com. Part of research was done when the author was a student at Princeton University. Research supported by Mark Braverman's NSF grant CCF-1149888, Simons Collaboration on Algorithms and Geometry, Simons Fellowship in Theoretical Computer Science and Siebel Scholarship.

[†]Department of Computer Science, The City College of New York, email: l.n.gurvits@gmail.com

[‡]Department of Computer Science, Princeton University, email: rmo@cs.princeton.edu. Research supported by NSF CAREER award DMS-1451191, NSF grant CCF-1523816 and Siebel Scholarship.

[§]Institute for Advanced Study, Princeton, email: avi@math.ias.edu. Research supported by NSF grant CCF-1412958.

Contents

1	Introduction	1
1.1	The Brascamp-Lieb inequalities: basic notions	1
1.2	The Brascamp-Lieb inequalities: known and new results	3
1.3	Efficient computation of the BL-constant via scaling	4
1.4	Applications of Brascamp-Lieb in Computer Science, Optimization and Beyond	9
2	Square completely positive operators	12
3	Rectangular operators	14
4	Reduction from Brascamp-Lieb to completely positive operators	20
4.1	Connection to quiver semi-invariants	24
5	Upper bound on the Brascamp-Lieb constant	25
6	Computing the Brascamp-Lieb constant	27
7	Continuity of the Brascamp-Lieb constant	28
8	Interesting special cases of BL polytopes	29
9	Additional properties of the Brascamp-Lieb constant	31
10	Iteration analysis for Algorithm 1	34
11	Conclusion and open questions	35

1 Introduction

This long introduction summarizes the main contributions of the paper and the intuitions behind many of the definitions, results and proofs. We start with the main object of this paper: introducing the Brascamp-Lieb inequalities, and surveying known structural results and our new algorithmic results. We then describe how to instantiate the operator scaling algorithm of [GGOW16] directly in the context of BL-inequalities, in a way that does not require previous familiarity with it. The properties of this algorithm lead to further structural results, as well as highlight it as a provably efficient instance of the general *alternate minimization*¹ heuristic. Finally, we discuss several known and potential applications of Brascamp-Lieb inequalities in computer science and optimization, which further motivate this and future studies.

1.1 The Brascamp-Lieb inequalities: basic notions

Many important inequalities, including Hölder's inequality, Loomis-Whitney inequality, Young's convolution inequality, hypercontractivity inequalities and many others are all special cases of the extremely general Brascamp-Lieb inequalities, introduced by these authors in [BL76, Lie90]. Yet many others, including Prekopa-Lindler inequality, versions of Brunn-Minkowski inequality and others are special cases of Barthe's reverse form of Brascamp-Lieb [Bar98]. We discuss below only the original form, and only for Euclidean spaces. The notation we use is taken mostly from [BCCT08a], which is an excellent source both for background and motivation on this topic, as well as the state-of-art on the basic questions in this field. As is common, we often use BL as abbreviation for Brascamp-Lieb.

A *Brascamp-Lieb datum* is specified by two m -tuples: one of linear transformations $\mathbf{B} = (B_1, B_2, \dots, B_m)$, with $B_j : \mathbb{R}^n \rightarrow \mathbb{R}^{n_j}$ (along with (n, n_1, \dots, n_m) , the vector of *dimensions*), and another of non-negative reals $\mathbf{p} = (p_1, p_2, \dots, p_m)$ (which is a vector of *exponents*). This combined information is denoted by (\mathbf{B}, \mathbf{p}) . A *Brascamp-Lieb inequality*² with the datum above states that for *every* tuple of m non-negative functions, $f = (f_1, f_2, \dots, f_m)$, which are integrable according to the Lebesgue measure, we have the following inequality.

$$\int_{x \in \mathbb{R}^n} \prod_{j=1}^m (f_j(B_j x))^{p_j} dx \leq C \prod_{j=1}^m \left(\int_{x_j \in \mathbb{R}^{n_j}} f_j(x_j) dx_j \right)^{p_j}$$

where $C \in (0, \infty]$ is independent of the functions f . Of course, we really get an interesting inequality if C is finite. In that case, we call the datum (\mathbf{B}, \mathbf{p}) *feasible*, and we denote smallest C for which this inequality holds by $\text{BL}(\mathbf{B}, \mathbf{p})$, called the *BL-constant*.

Many familiar special cases of this inequality are listed in the introduction of [BCCT08a]; we describe only a very simple one here, which hopefully makes concrete and intuitive the complicated formal expression above and its constituents. A special case of the Loomis-Whitney inequality asserts that *the volume of every measurable set S in the unit cube is at most the square root of the product of the areas of its projections on the three coordinate planes*. In this theorem, the linear transformations $B_j : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ are defined by the simple projections $B_1(x, y, z) = (y, z)$, $B_2(x, y, z) = (x, z)$ and $B_3(x, y, z) = (x, y)$, the functions f_j are the indicator functions of these

¹sometimes also called alternating minimization.

²While this is the common name of the inequality below found in the literature, we note that only a restricted form appears in the original [BL76] paper, and this general form appeared only in the later paper of Lieb [Lie90].

three projections of the set S , and the exponent vector is $p = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. The corresponding BL-constant in this case happens to be 1 (this case is an important situation which will be important later). The reader may be familiar with an entropy³ version of the inequality above, namely that $H(X, Y, Z) \leq \frac{1}{2}(H(X, Y) + H(Y, Z) + H(X, Z))$ for every random variable (X, Y, Z) . Indeed, BL-inequalities may be viewed in general as entropy subadditivity inequalities, as developed in [CCE09, LCCV16, LCCV17].

There has been extensive work on precisely understanding (and computing) when a given datum (\mathbf{B}, \mathbf{p}) is feasible, and if it is, determining the BL-constant $\text{BL}(\mathbf{B}, \mathbf{p})$. Clean characterizations exist for both questions, which clarify that they are both decidable (have a finite algorithm). Let us overview these characterizations and then turn to the computational complexity of the existing algorithms and our new ones. The first, regarding feasibility, is from the paper [BCCT08a], and the second, on the optimal constant, is from [Lie90]. Before doing so, we make two general comments, one about the “reverse BL-inequalities” and the second regarding how we measure input size when analyzing algorithms.

Reverse BL inequalities This informal comment simply clarifies that all results in this paper stated for BL-inequalities hold for their reverse form. In [Bar98], Barthe introduced a reverse form of the Brascamp-Lieb inequalities, sometimes called RBL inequalities, which turn out to generalize some known inequalities not captured by the the original BL inequalities. These RBL inequalities take the *same* data as standard BL inequalities, and have an optimal RBL constant associated with each. All questions raised above for the BL inequalities, like feasibility, computation of the constant and its continuity properties are relevant in this reverse setting. However, they are actually equivalent to the original ones for a simple reason: Barthe [Bar98] proved that for any feasible datum the optimal BL constant and the optimal RBL constant multiply to 1, and when the BL datum is infeasible the RBL constant is 0. In short, these two optimal constants determine each other. Thus all structural results above translate to the reverse setting, and so do all our algorithmic results.

Input size conventions Before we start, let us formalize the input conventions to all algorithms, and the parameters we use to measure their complexity. The input to all algorithms will be a BL datum (\mathbf{B}, \mathbf{p}) , and the “size” of each part will be measured differently. The entries of the matrices B_j will be rational numbers, given in binary. We will let $b = b(\mathbf{B})$ denote their total binary length (note that in particular $b \geq nm$, and so lower bounds the “combinatorial size” of the input). The vector \mathbf{p} will be given as a sequence of rationals with a common denominator, namely $p_j = c_j/d$ with c_j, d integers. We use this convention for two reasons. First, our algorithms will use this integer representation, and their complexity will depend on $d = d(p)$ (while in many cases d is only polynomial in the dimensions n, m of the problem, it can definitely be as large as exponential in other cases). Second, in the context of *quiver representations*, which is very relevant to this study as we shall later see⁴, it is natural to use these integer “weights” c_j and d .

Summarizing, the two size parameters for a BL datum (\mathbf{B}, \mathbf{p}) will be b and d as above.

³Appropriately defined for these continuous variables.

⁴With hindsight, our reduction from BL data to operator scaling data may be viewed as a special case of the reduction of Derksen and Makam [DM15] of general quivers to the Left-Right quiver. More on this in Section 4.

1.2 The Brascamp-Lieb inequalities: known and new results

Testing feasibility (and more) The following theorem of Bennett et al [BCCT08a, BCCT08b] precisely characterizes when a given BL datum is feasible. This work will provide a different proof of this important theorem (see Corollary 4.5).

Theorem 1.1 ([BCCT08a]). *The datum (\mathbf{B}, \mathbf{p}) is feasible if and only if the following conditions hold:*

1. $n = \sum_j p_j n_j$.
2. $\dim(V) \leq \sum_j p_j \dim(B_j V)$ holds for all subspaces V of \mathbb{R}^n .

Note that these are simply linear conditions on the exponent vector \mathbf{p} , albeit defined by infinitely many subspaces V . However, as the coefficients are integers in $[n] = \{1, 2, \dots, n\}$, there are at most n^m different inequalities, and so the linear maps \mathbf{B} define a polytope $P_{\mathbf{B}}$ (sometimes called the *BL polytope*) in \mathbb{R}^n , such that (\mathbf{B}, \mathbf{p}) is feasible iff $\mathbf{p} \in P_{\mathbf{B}}$. The BL polytopes $P_{\mathbf{B}}$ have received a lot of attention. For example, their vertices were characterized for the “rank-1 case” (namely, when all dimensions $n_j = 1$) by Barthe [Bar98], and this was extended to other cases by Valdimarsson [Val10]. In the same paper, Valdimarsson considers the question of generating the inequalities defining $P_{\mathbf{B}}$, and gives a finite algorithm to do so (needless to say, after they are given, feasibility testing becomes a linear programming problem). No explicit upper bound on the complexity of this algorithm is given, but it is at least exponential in m . The same holds for the algorithm in [CDK⁺13] for generating the inequalities defining $P_{\mathbf{B}}$.

We give a polynomial time algorithm for the feasibility problem, and much more. Our algorithm also gives a “separation oracle” (namely finds a violated inequality when infeasible). Recall again that the exponents \mathbf{p} in the BL datum (\mathbf{B}, \mathbf{p}) are given by $p_j = c_j/d$ where c_j, d are integers.

Theorem 1.2 (Corollary 4.6 rephrased). *There is an algorithm that on input (\mathbf{B}, \mathbf{p}) of binary length b and common denominator d runs in time $\text{poly}(b, d)$ and provides the following information:*

Membership oracle: Tests if $\mathbf{p} \in P_{\mathbf{B}}$.

Separation oracle: When $\mathbf{p} \notin P_{\mathbf{B}}$, provides a violated inequality, namely a basis for a vector space V in \mathbb{R}^n such that $\dim(V) > \sum_j p_j \dim(B_j V)$.

We believe that the ability to efficiently optimize over such a wide family of polytopes with exponentially many facets should be useful for different optimization problems via appropriate reductions to this setting. We will give some simple examples of such reductions in Section 8. One concrete challenge is e.g. trying to find such a reduction which embeds the Edmonds polytope [Edm65] of perfect matchings in a general (non-bipartite) graph into some BL polytope.

It is an interesting open problem if one can improve the algorithm to depend polynomially on $\log d$ instead of d . This would allow for optimization over the BL polytopes in polynomial time via the ellipsoid algorithm. We believe even the current separation oracle for BL polytopes with polynomial dependence on d should allow for approximate optimization over BL polytopes but we haven’t been able to prove this yet.

Characterizing the BL constant The following theorem of Lieb [Lie90] characterizes the BL constant $\text{BL}(\mathbf{B}, \mathbf{p})$ in all cases where the BL datum is feasible. The heart of the proof establishes that the optimal constant in any BL inequality is attained (or approached) by plugging in density functions of appropriate centered Gaussians. For such densities, the BL constant has a nice expression via the following identity:

$$\int_{\mathbb{R}^n} \exp(-\pi x^T A x) dx = \text{Det}(A)^{-1/2}$$

Theorem 1.3 ([Lie90]). *Assume (\mathbf{B}, \mathbf{p}) is feasible. Then*

$$\text{BL}(\mathbf{B}, \mathbf{p}) = \left[\sup \frac{\prod_j (\det X_j)^{p_j}}{\det \left(\sum_j p_j B_j^\dagger X_j B_j \right)} \right]^{1/2}$$

where the supremum is taken over all choices of positive definite matrices X_j in dimension n_j , and B_j^\dagger denotes the adjoint map corresponding to B_j .

Thus the BL constant for given BL datum is a solution to an optimization problem. However, as defined, it is not even convex. We are not aware of any algorithms to compute the BL constant in general. Moreover, to the best of our knowledge no explicit bounds on the BL constant (when finite) in terms of the BL datum were known. We resolve both issues: provide an explicit bound and give a polynomial time algorithm for computing the BL constant to any accuracy.

Theorem 1.4 (Corollary 5.5 rephrased). *For any feasible BL datum (\mathbf{B}, \mathbf{p}) that has binary length b and common denominator d , it holds that $\text{BL}(\mathbf{B}, \mathbf{p}) \leq \exp(O(b \log(bd)))$.*

Theorem 1.5 (Theorem 6.2 rephrased). *There is an algorithm that on input (\mathbf{B}, \mathbf{p}) of binary length b and common denominator d , and an accuracy parameter $\epsilon > 0$, runs in time $\text{poly}(b, d, 1/\epsilon)$ and computes a factor $(1 + \epsilon)$ -approximation of $\text{BL}(\mathbf{B}, \mathbf{p})$. Furthermore, the algorithm outputs a scaling \mathbf{B}' which is almost geometric i.e. $\text{BL}(\mathbf{B}', \mathbf{p}) \leq 1 + \epsilon$.*

1.3 Efficient computation of the BL-constant via scaling

The algorithm underlying the proof of Theorem 1.5 will be shown (via a reduction) to be a special case of the operator scaling algorithm of Gurvits [Gur04], which the current authors analyzed and proved to be polynomial time in [GGOW16].

In this subsection, we explicitly describe the algorithm in this special case of BL-inequalities only, without referring to [GGOW16]. It will be instructive to see how notion of scaling, which naturally exists in the theory of BL inequalities, can be used algorithmically. Moreover, it will become clear how the algorithm implies (known and new) structural consequences to the BL theory. This description will also help to motivate our reduction in the technical chapters which follow (which will provide the proof for its run-time).

Let us return to the BL constant, and to an important family of BL data called *geometric*⁵ defined below. It was introduced by Ball [Bal89] and extended by Barthe [Bar98].

Definition 1.6. A BL datum (\mathbf{B}, \mathbf{p}) is called *geometric* if it satisfies the following conditions (with I_k denoting the $k \times k$ identity matrix).

⁵The analogous term in the operator scaling setting will be called *doubly stochastic*.

Projection: For every $j \in [m]$, B_j is a projection, namely $B_j B_j^\dagger = I_{n_j}$.

Isotropy: $\sum_j p_j B_j^\dagger B_j = I_n$

Ball proved that for the special case of geometric BL datum with $n_j = 1$, the constant is always one. Barthe extended this to the general case⁶.

Theorem 1.7 ([Bal89, Bar98]). *For every geometric BL datum (\mathbf{B}, \mathbf{p}) we have $\text{BL}(\mathbf{B}, \mathbf{p}) = 1$.*

A simple and natural action of the general linear groups on the Euclidean spaces involved (which simply performs a basis change in each space) yields an equivalence relation on BL data⁷. Specifically, we say that (\mathbf{B}, \mathbf{p}) and $(\mathbf{B}', \mathbf{p}')$ are *equivalent* if there exist matrices $C \in GL_n(\mathbb{R})$, $C_j \in GL_{n_j}(\mathbb{R})$ (which are called *intertwining transformations*⁸ in [BCCT08a]) such that $B'_j = C_j^{-1} B_j C$ and $p'_j = p_j$ for all j . It is easy to compute the effect of such action on the BL constant.

Proposition 1.8 ([BCCT08a]). *Assume (\mathbf{B}, \mathbf{p}) and $(\mathbf{B}', \mathbf{p}')$ are equivalent via C, C_j as above. Then*

$$\text{BL}(\mathbf{B}', \mathbf{p}') = \frac{\prod_j (\det(C_j))^{p_j}}{\det(C)} \text{BL}(\mathbf{B}, \mathbf{p}).$$

where \det is the determinant polynomial on matrices of the appropriate dimension.

Equivalence and this simple, efficiently computable formula suggests a natural path to computing BL constants. Given BL datum (\mathbf{B}, \mathbf{p}) , compute an equivalent *geometric* datum $(\mathbf{B}', \mathbf{p}')$ (if one exists), and the intertwining transformations relating the two. The main questions are, is there such a geometric equivalent datum, and how to compute these transformations. The theory of (quantum) operator scaling [Gur04, GGOW16] suggests that a simple, greedy procedure will work. This algorithm will underly (most of) the statements in Theorem 1.2 and Theorem 1.5. As mentioned, we will explain the connection and reduction to operator scaling in the next sections, and here describe informally how it is applied to the BL setting.

The scaling algorithm The key idea of *greedy, iterative* scaling, which goes back to Sinkhorn [Sin64] in the (classical) matrix scaling, is that when trying to achieve a pair of conditions as in the definition of geometric BL datum, we should try to satisfy them one at a time! Say we are given (\mathbf{B}, \mathbf{p}) and assume for example that the isotropy condition is not met. A non-triviality condition (otherwise (\mathbf{B}, \mathbf{p}) is not feasible) implies that $C \triangleq \sum_j p_j B_j^\dagger B_j$ is invertible, and so we can set $B_j \leftarrow B_j C^{-1/2}$ to get a new datum which satisfies isotropy. So, we “only” have to fix the projection property; lets do it. Again, non-triviality implies that each of $C_j \triangleq B_j B_j^\dagger$ is invertible, and we can now set $B_j \leftarrow C_j^{-1/2} B_j$, satisfying the projection property. Of course, we may now have ruined isotropy. No problem: lets fix it again in the same way, and repeat alternately fixing the unsatisfied property. The magic is that this sequence of *normalization* steps converges, and

⁶The analogous theorem for operator scaling is that the capacity of doubly stochastic operators is always 1. Another family of inequalities where the constant is always 1 are the “discrete” BL inequalities of [CDK⁺15].

⁷This group action naturally calls for a study of BL from an invariant theory viewpoint, which indeed exists in a much more general context that will be relevant to us in several ways, namely that of *quiver representations*. An extensive survey is [DW07].

⁸They will be called *scaling matrices* in the operator scaling setting.

moreover, converges in polynomial time, to a geometric datum, whenever the original datum (\mathbf{B}, \mathbf{p}) is feasible!⁹ Let us describe the BL algorithm more precisely and then state its properties formally.

Input: A BL datum (\mathbf{B}, \mathbf{p})

Set $\mathbf{B}^0 = \mathbf{B}$.

Repeat for $i = 1$ to t normalization steps.

- if i is odd, normalize \mathbf{B}^i from \mathbf{B}^{i-1} to satisfy isotropy.
- if i is even, normalize \mathbf{B}^i from \mathbf{B}^{i-1} to satisfy projection.

Output: \mathbf{B}^t .

Algorithm 1: The BL scaling algorithm

The BL scaling algorithm defines a dynamics in the space of BL data which are restricted to stay in one equivalence class. If (\mathbf{B}, \mathbf{p}) is the original datum which is not already geometric, it defines a sequence $\mathbf{B}^0 = \mathbf{B}, \mathbf{B}^1, \mathbf{B}^2, \dots$ which alternately satisfies the isotropy or projection properties (note that the exponent vector p is used, but not changed in this process). We now state the main three properties which underlie the analysis.

Theorem 1.9. *For every (\mathbf{B}, \mathbf{p}) the following holds:*

1. Upper bound: *If (\mathbf{B}, \mathbf{p}) is feasible, then $\text{BL}(\mathbf{B}, \mathbf{p}) \leq \exp(O(b \log(bd)))$.*
2. Lower bound: *If the datum (\mathbf{B}, \mathbf{p}) is either isotropy-normalized or projection-normalized, then $\text{BL}(\mathbf{B}, \mathbf{p}) \geq 1$.*
3. Progress per step: *Let $(\mathbf{B}', \mathbf{p})$ denote the result of applying either isotropy-normalization or projection-normalization to (\mathbf{B}, \mathbf{p}) . Then (as long as $\text{BL}(\mathbf{B}, \mathbf{p}) \geq 1 + \epsilon$), we have*

$$\text{BL}(\mathbf{B}', \mathbf{p}) \leq (1 - \text{poly}(\epsilon/nd))\text{BL}(\mathbf{B}, \mathbf{p}).$$

We will not need to prove this theorem, as it will follow via our reduction¹⁰ from the analogous theorem in the more general operator scaling setting¹¹. However, let us say a few words here about what goes into each item above using BL language. Property (3) on progress per step is the simplest: it follows from a robust version of the AM-GM inequality. Property (2) follows from plugging in densities of spherical Gaussians (it was proved in [Val11], also see Section 9). Property (1) is the hardest, in that it requires degree bounds on the generating polynomial invariants in the null-cone of the quiver associated to the group action above. We will briefly explain this mouthful in Section 5, and refer the reader to [GGOW16] for details. In Section 5, we will also give an improved and simplified analysis of the main theorem from [GGOW16] and derive the desired upper bounds on the BL constant via our main reduction. All in all, it is worth stressing again that beyond this

⁹Note that the sequence might only converge to a geometric datum and it is possible that no element of the sequence is geometric.

¹⁰In the case when the exponent vector \mathbf{p} is irrational, the reduction to operators does not exist and so in this case, we provide an analysis of the number of iterations of Algorithm 1 in Section 10.

¹¹Where statements made here regarding the BL constants will be replaced by statements about a related notion called the *capacity* of an operator.

non-trivial use of algebra, on the analytic side the most general BL inequalities follow from the most basic one, AM-GM.

This analysis above implies that we can choose the number t of steps in the BL scaling algorithm to be $\text{poly}(b, d, 1/\epsilon)$ so that $\text{BL}(\mathbf{B}^i, \mathbf{p}) \leq 1 + \epsilon$ for some $1 \leq i \leq t$ (whenever the initial datum $(\mathbf{B}^0, \mathbf{p})$ is feasible). Testing if the lower bound is met ($\text{BL}(\mathbf{B}^i, \mathbf{p}) \leq 1 + \epsilon$) decides feasibility¹², thus proving the first item in Theorem 1.2. For proving the second item we will employ a very different, beautiful combinatorial algorithm for operator scaling (only the decision version of it), that was developed after ours by Ivanyos, Qiao and Subrahmanyam [IQS15]. Their algorithm relies on even tighter degree bounds from invariant theory which were proved very recently [DM15], and has several advantages over the algorithm in [GGOW16]; in particular it can find violated dimension inequalities. Going back to Theorem 1.5 on approximating the BL constant, note that it suffices to multiply together the (respective) normalizing matrices used in all steps of the BL scaling algorithm, using the formula in Proposition 1.8.

We now return to the scaling algorithm above, and discuss connections and applications of its structure and properties in the next two bullets.

The BL constant, scaling and alternate minimization The problem of computing the BL constant, as it is formulated here, is a non-convex problem. We presented an efficient algorithm for this optimization problem which may be viewed as an instance of a general heuristic called *alternate minimization*. Let us describe and discuss this method.

In general, alternate minimization is a very general technique introduced by [CT84], which was devised to solve problems of the following kind. Given some universe U , a (distance) function $d : U \times U \rightarrow \mathbb{R}_+$ and two subsets $X, Y \subseteq U$, find $(\mathbf{x}^*, \mathbf{y}^*) = \arg \min_{\mathbf{x} \in X, \mathbf{y} \in Y} d(\mathbf{x}, \mathbf{y})$. We note that such problems are in general NP-hard, even over convex sets X, Y and natural distances d . Nevertheless, alternate minimization is widely used to solve special cases of such problems in practice.¹³

An alternate minimization algorithm for the problem above starts with a arbitrary point $(x_0, y_0) \in X \times Y$ and generates a sequence of points (x_i, y_i) as follows. In even steps i , $x_{i+1} = x_i$, and y_{i+1} is chosen to be the value of y which minimizes $d(x_i, y)$ over the second coordinate alone (an optimization problem that is often far simpler). In odd steps the roles of x and y are reversed. It is easy to cast the scaling algorithm above as a process of this nature. It is also easy to see that it enjoys some nice properties in general situations. For example, if both X and Y are strictly convex, and d is a metric, then this process converges to the optimum. Of course, the speed of convergence is the main question. We list below some other, possibly familiar instances of alternate minimization for which such convergence is not known.

The famous Lemke-Howson algorithm for finding a Nash equilibrium for 2-player games [LH64] is of that nature. Recall that in that algorithm one starts with an arbitrary strategy for both players, and then one alternates fixing the strategy of one, and finding the best response for the other via linear programming. Another such example is the work of Zafeiriou and Petrou [ZP11] on computing non-negative tensor factorization. As above, the analysis of such alternating minimization algorithms just provide convergence of the procedure, without proving *how fast* it converges

¹²It is not immediately clear how to test $\text{BL}(\mathbf{B}^i, \mathbf{p}) \leq 1 + \epsilon$. However this can be done by checking closeness to the geometric position in an appropriate metric, similar to the distance to double stochasticity for operators that will be defined later.

¹³Note that such problems can be defined with more than two components, and the alternate minimization approach below can be extended to such problems, but we stick here with two.

to the minimum, with the hope that this heuristic will lead to quick convergence on instances arising in practice. Of course, there are examples, e.g. the alternate minimization algorithms for matrix completion in [JNS13, Har14] which prove rapid convergence under certain conditions on the inputs.

In this work, as well as in [GGOW16], we prove that our alternate minimization algorithm not only converges, but also converges to the infimum in polynomial time. Our proof is based on the introduction of a potential function which measures how much progress we make in every step of the alternating minimization algorithm. In our case the potential function *is* the very BL constant we are trying to optimize, which is a special case of the so-called *capacity* we use as potential in our operator scaling algorithm [GGOW16]. As it happens, the analysis of convergence relies on a combination of algebraic tools, in particular from (the commutative) invariant theory and (the non-commutative) theory of skew fields, as explained in that paper. We hope that such potential functions as well as the methods to analyze their progress, which we call *capacity methods*, will be of use in the analyses of other problems which use alternate minimization, especially when the result of optimizing each of the solution components in individual optimization steps may be viewed as the action of a group. We are currently exploring this direction for alternate minimization with more than two components, and these lead to interesting algebraic questions related to questions that arise in invariant theory as well as in Geometric Complexity Theory of Mulmuley and Sohoni [MS06, Mul12].

Bounds and continuity of the BL constant, and non-linear BL inequalities Another important property of the BL scaling algorithm above, evident from its form, is that it is smooth! Namely, the orbits under the algorithm of two sufficiently close BL data will remain close to each other for the duration of the algorithm. This implies that the BL constants of both will be close as well, and moreover this can be quantified! This simple observation gives a much stronger result than existing qualitative results on the continuity of BL constants which has recently received significant attention. This is yet another example highlighting the usefulness of algorithmic results in mathematics. We briefly overview the motivation and known results before stating ours.

How smooth (or regular) the BL constant is (in terms of the BL data) is of course a natural question, especially as the expression in Lieb’s theorem 1.3 is not even convex. Further motivation to study this question arises from a variety of non-linear variants of the BL inequalities, in which the maps B_j are non-linear, but at least smooth enough so as to be approximable by linear ones in a small ball. Some such variants, needed in specific applications, were proved directly. The papers [BB10, BBFL17] ask this question in full generality, and demonstrate the importance of such smoothness conditions of the BL constant for general non-linear extensions of BL inequalities. These papers also exposit the many diverse applications of such generalized BL inequalities in analysis, number theory and other areas (possibly the most impressive recent ones are the applications of “Kakeya-type” BL inequalities to number theory in [BD16, BDG16], resolving long standing open problems including the “Vinogradov mean-value conjecture”).

Local boundedness, a condition weaker than continuity, is established in [BBFL17].

Theorem 1.10 ([BBFL17]). *If (\mathbf{B}, \mathbf{p}) is feasible, then there are $\delta > 0$ and $C < \infty$ such that for all $(\mathbf{B}', \mathbf{p})$ such that $\|\mathbf{B}' - \mathbf{B}\|_2 \leq \delta$ we have $\text{BL}(\mathbf{B}', \mathbf{p}) \leq C$.*

This was followed recently by the paper [BBCF17], which proves continuity of the BL constant.

Theorem 1.11 ([BBCF17]). *The BL constant $\text{BL}(\mathbf{B}, \mathbf{p})$ is continuous in \mathbf{B} . Namely, for every feasible (\mathbf{B}, \mathbf{p}) and $\epsilon > 0$, there exists a $\delta > 0$ such that if $\|\mathbf{B}' - \mathbf{B}\|_2 \leq \delta$, then we have that $|\text{BL}(\mathbf{B}', \mathbf{p}) - \text{BL}(\mathbf{B}, \mathbf{p})| \leq \epsilon$.*

In both theorems above, no quantitative bounds are given. This is not surprising, as the proofs of both use compactness. On the other hand, the availability of a smooth algorithm, together with the quantitative analysis of its continuity provided in [GGOW16], imply via our reduction the following quantitative bound on the continuity¹⁴ of the BL constant (giving a multiplicative approximation).

Theorem 1.12 (Theorem 7.2 rephrased). *For every parameters $b, d, \epsilon > 0$ there exists a*

$$\delta = \delta(b, d, \epsilon) \leq \exp\left(-\frac{\text{poly}(b, d)}{\epsilon^3}\right)$$

such that for every BL data (\mathbf{B}, \mathbf{p}) of size b, d and every BL data $(\mathbf{B}', \mathbf{p})$ such that $\|\mathbf{B}' - \mathbf{B}\|_2 \leq \delta$, we have that if $\text{BL}(\mathbf{B}, \mathbf{p})$ is finite, then $\text{BL}(\mathbf{B}', \mathbf{p})$ is finite. Furthermore,

$$(1 - \epsilon)\text{BL}(\mathbf{B}, \mathbf{p}) \leq \text{BL}(\mathbf{B}', \mathbf{p}) \leq (1 + \epsilon)\text{BL}(\mathbf{B}, \mathbf{p})$$

We hope that this explicit bound will be useful to making non-linear BL inequalities quantitative, and to their applications.

1.4 Applications of Brascamp-Lieb in Computer Science, Optimization and Beyond

Besides the applications to other areas of mathematics as we mentioned above, the structural theory of BL inequalities (sometimes with the connection to operator scaling) have encountered several applications in computer science and optimization. We give below a (partial) survey of these applications.

Many diverse applications exist even for the special case of rank-1 BL setting (when all the n_i 's are 1), which has been studied extensively. In this case, the BL polytopes are exactly the *basis polytopes* [Bar98]. More specifically, if the maps B_i are given by $B_i(x) = \langle v_i, x \rangle$, then the BL polytope is the convex hull of the following set:

$$\{1_I : I \subset [m], |I| = n, (v_i)_{i \in I} \text{ forms a basis for } \mathbb{R}^n\}$$

Hence Theorem 1.2 for this special case is well known. Furthermore, computing the BL constant in this special case can be formulated as a convex problem and hence Theorem 1.5 in this special case can be obtained by the ellipsoid algorithm [GS00, HM13] ([GS00] in fact solves a more general problem). We list some of the applications of this setting.

- In [HM13], Hardt and Moitra provide an application of scaling the rank-1 BL datum to a geometric position (which they call isotropic) to the problem of *robust subspace recovery*.

¹⁴While Theorem 1.12 only implies continuity of $\text{BL}(\mathbf{B}, \mathbf{p})$ for rational \mathbf{B} and \mathbf{p} , we describe in Section 7 how our proof can be extended to the case of irrational \mathbf{B} and speculate how this might also extend to irrational \mathbf{p} .

- Forster [For02] used the existence of a scaling to geometric position for rank-1 BL datum (sometimes called Barthe’s theorem) to give the first nontrivial lower bounds for the *sign rank* of an explicit matrix (in particular the Hadamard matrix), and thereby prove the first nontrivial *unbounded error communication complexity* lower bounds for the inner product function.
- In [DSW14a], Dvir et al. use Barthe’s theorem as part of the first super-quadratic lower bounds on the size of 3-query *LCCs (Locally Correctable Codes)* over the Real numbers.
- Recently, Nikolov and Singh [NS16] use a optimization problem similar to the rank-1 BL constant to give a polynomial time algorithm for approximating the maximum value of the determinant of a submatrix of a psd matrix under partition constraints.

The class of BL polytopes is a rich class of polytopes. As mentioned above, the rank-1 case gives rise to basis polytopes. The rank-2 case (when all n_i ’s are 1 or 2) has been completely characterized by Valdimarsson [Val10]. An important special case is given by the *linear matroid intersection* polytopes which we describe next. Suppose (v_1, \dots, v_m) and (w_1, \dots, w_m) are two collection of vectors in \mathbb{R}^n . Then the matroid intersection polytope given by these collections is the convex hull of the following set:

$$\{1_I : I \subset [m], |I| = n, (v_i)_{i \in I}, (w_i)_{i \in I} \text{ both form a basis for } \mathbb{R}^n\}$$

Let us define the maps $B_i : \mathbb{R}^{2n} \rightarrow \mathbb{R}^2$ in the following way:

$$B_i(x, y) = (\langle v_i, x \rangle, \langle w_i, y \rangle)$$

Then the BL polytope $P_{\mathbf{B}}$ is exactly the corresponding matroid intersection polytope. In particular, the bipartite matching polytopes are a special case of rank-2 BL polytopes.

We will give proofs of these statements in Section 8, with the hope of demonstrating the simplicity by which such exponential size linear programs of this form capture (simple) combinatorial optimization problems. As mentioned before, it remains a very interesting challenge to see if the general matching polytopes (or in general the linear matroid matching polytopes) are a special case of BL polytopes, for which it may still suffice to stay in the rank-2 case. We note a peculiar discrepancy between operator scaling and BL-inequalities when encoding the matroid intersection problem. In operator scaling, this problem was one of the first special cases considered in Gurvits’ paper [Gur04], and he showed that one could encode it by an operator scaling problem with matrices of rank 1. Here, when encoded as a BL polytope, we use rank-2 projectors, and believe it is impossible to do in rank 1. Even though BL is a special case of operator scaling (as our main reduction shows) the relative power of these two formulations with respect to rank is not clear.

Characterizing the vertices of BL-polytopes for projectors of rank 3 and higher remains a very interesting question, as is the case of rank-2 matrices in the formulation of operator scaling.

BL-inequalities (and operator scaling) have recently found applications in computational and combinatorial geometry in the paper of Dvir et al. [DGOS16]. They generalize the rank lower bounds for design matrices given in [BDWY12, DSW14b] to the setting of block matrices. This result allows them to obtain sharp bounds on Sylvester-Gallai type theorems for arrangements of subspaces, to obtain new incidence bounds for high-dimensional line/curve arrangements, as well as to prove structural rigidity results in the projective setting.

Beyond *applying* BL-inequalities, we note that in recent years computer scientists started getting interested in *proving* them efficiently. Efficiency here here should be taken to mean in the Sum-of-Squares (SoS) framework, a hierarchy of semi-definite programs that is perhaps the most powerful general algorithmic technique for a variety of optimization problems. The paper [BKS14], generalizing from many particular examples on the power of SoS algorithms, suggests a general framework for rounding solutions in this system, in which a crucial element is efficient SoS proofs of inequalities (which in all known applications are special cases of BL-inequalities). This viewpoint has lead to many new recent SoS algorithms [BKS15, MSS16, HSS15], and naturally also lead to the question of efficiently proving (in this sense) all BL-inequalities. This was taken up by Lei and Sheng [LS17], who have (in a very general setting) SoS proofs of degree $\text{poly}(d)$ for BL-inequalities that have denominator d in its exponent \mathbf{p} -vector. This yields an algorithm for testing feasibility that is *exponential* in d . Our algorithms depend polynomially in d , and raise the question of relating operator scaling and SoS algorithms, and whether such a relation can help improve the bounds above. Needless to say, relating the power of these two optimization methods is motivated for many other reasons!

Dictionary We give here a basic dictionary describing translation between central notions in the theory of Brascamp-Lieb inequalities and completely positive operators. This will be useful in reading the technical sections below. The basic notions concerning completely positive operators will be defined in Sections 2 and 3.

Brascamp-Lieb inequalities	Completely positive operators
BL datum	Kraus operators of completely positive operators
Intertwining transformations	Scaling operations
Geometric BL datum	Doubly stochastic operator
BL constant	Capacity

Organization of the paper All but the last two sections are devoted to providing background and then proving the main technical results of this paper. Preliminaries and results from [Gur04, GGOW16] about (square) completely positive operators, their scaling, capacity and more will be reviewed in Section 2. In Section 3, we will extend these results to *rectangular* completely positive operators via a simple reduction. In Section 4, we describe our main reduction, from BL data to (rectangular) completely positive operators, and prove Theorem 1.2. We also show how this reduction yields, as a corollary, a simple proof of the characterization (Theorem 1.1) of [BCCT08a]. In Section 5, we describe an upper bound on the Brascamp-Lieb constants and prove Theorem 1.4. In Section 6, we show how to approximate Brascamp-Lieb constants by using the main reduction above and our operator scaling algorithm (proving Theorem 1.5). In Section 7, we use the smoothness of the scaling algorithm to give explicit bounds for the continuity of Brascamp-Lieb constants and prove Theorem 1.12. In Section 8 (which is independent from all others) we exemplify the power of exponentially large linear programs arising as BL-polytopes to capture certain simple combinatorial optimization problems. Section 9 contains some additional properties of the Brascamp-Lieb constant, strengthening the results in [Val11]. Section 10 contains an analysis of the number of iterations required in Algorithm 1 to get close to geometric position. Finally, we conclude with some open problems in Section 11.

2 Square completely positive operators

In this section, we introduce (square) completely positive operators and important facts and theorems about them which we will need [Gur04, GGOW16].

Given a complex matrix A , we will use A^\dagger to denote the conjugate-transpose of A . For matrices with real entries this will just be A^T . \succeq will be used to denote the Loewner order defined by positive semidefinite matrices. So we will write $A \succeq B$ if $A - B$ is positive semidefinite and $A \succ B$ if $A - B$ is positive definite.

Definition 2.1 (Completely Positive Operators). An operator (or map) $T : M_n(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ is called completely positive if there are $n \times n$ complex matrices A_1, \dots, A_m s.t. $T(X) = \sum_{i=1}^m A_i X A_i^\dagger$. The matrices A_1, \dots, A_m are called Kraus operators of T (they need not be unique). T is called completely positive trace preserving (cptp) if in addition, $\text{tr}(T(X)) = \text{tr}(X)$ for all X . This is equivalent to the condition $\sum_{i=1}^m A_i^\dagger A_i = I$.

Remark 2.2. *The above is actually not the usual definition of completely positive operators. T is defined to be positive if $T(X) \succeq 0$ whenever $X \succeq 0$. T is completely positive if $I_n \otimes T$ is positive for all $n \geq 1$. Choi [Cho75] proved that an operator is completely positive iff it is of the form stated above.*

Definition 2.3 (Tensor products of operators). Given operators $T_1 : M_{d_1}(\mathbb{C}) \rightarrow M_{d_1}(\mathbb{C})$ and $T_2 : M_{d_2}(\mathbb{C}) \rightarrow M_{d_2}(\mathbb{C})$, we define their tensor product $T_1 \otimes T_2 : M_{d_1 d_2}(\mathbb{C}) \rightarrow M_{d_1 d_2}(\mathbb{C})$ in the natural way

$$(T_1 \otimes T_2)(X \otimes Y) = T_1(X) \otimes T_2(Y)$$

and extend by linearity to the whole of $M_{d_1 d_2}(\mathbb{C})$.

Definition 2.4. If $T(X) = \sum_{i=1}^m A_i X A_i^\dagger$ is a completely positive operator, we define its dual T^* by $T^*(X) = \sum_{i=1}^m A_i^\dagger X A_i$. If both T and T^* are trace preserving, namely $T(I) = T^*(I) = I$ then we call T (and T^*) *doubly stochastic*.

Definition 2.5 (Rank Decreasing Operators, [Gur04]). A completely positive operator T is said to be rank-decreasing if there exists an $X \succeq 0$ s.t. $\text{rank}(T(X)) < \text{rank}(X)$.

Now that we defined completely positive operators, we define their *capacity*, which is a very important complexity measure of such operators suggested in [Gur04]. For a special subset of operators, their capacity will be related to the Brascamp-Lieb constants.

Definition 2.6 (Capacity, [Gur04]). The capacity of a completely positive operator T , denoted by $\text{cap}(T)$, is defined as

$$\text{cap}(T) = \inf\{\text{Det}(T(X)) : X \succ 0, \text{Det}(X) = 1\}$$

Remark 2.7. *In the definition of capacity, we didn't specify whether to optimize over complex psd matrices or real psd matrices. However, as long as T is defined by real Kraus operators, the value of capacity doesn't change if we optimize over complex or real psd matrices. This can be seen by the fact that an almost optimizing solution to capacity can be obtained by a natural iterative sequence (Theorem 2.13).*

Next we define the notion of operator scaling.

Definition 2.8 (Operator Scaling, [Gur04]). An operator T' is called an operator scaling of T if there exist invertible matrices B, C s.t.

$$T'(X) = B \cdot T \left(C \cdot X \cdot C^\dagger \right) \cdot B^\dagger$$

Alternatively, if A_1, \dots, A_m are Kraus operators for T , then BA_1C, \dots, BA_mC are Kraus operators for T' .

Next we define a “distance” of an operator from being doubly stochastic.

Definition 2.9 (Distance to double stochasticity).

$$\text{ds}(T) = \text{tr} \left[(T(I) - I)^2 \right] + \text{tr} \left[(T^*(I) - I)^2 \right]$$

Next we define two operations on operators, which enforce one of the conditions of being doubly stochastic.

Definition 2.10 (Right Normalization). Given a completely positive operator T , define its *right normalization* T_R as follows:

$$T_R(X) = T \left(T^*(I)^{-1/2} \cdot X \cdot T^*(I)^{-1/2} \right)$$

Note that $T_R^*(I) = I$. This is because $T_R^*(X) = T^*(I)^{-1/2} \cdot T^*(X) \cdot T^*(I)^{-1/2}$.

Definition 2.11 (Left Normalization). Given a completely positive operator T , define its *left normalization* T_L as follows:

$$T_L(X) = T(I)^{-1/2} \cdot T(X) \cdot T(I)^{-1/2}$$

Note that $T_L(I) = I$.

Algorithm 2 was suggested by Gurvits [Gur04] to find a doubly stochastic scaling of an operator T .

Input: Completely positive operator T .
 Perform right and left normalizations starting from $T_0 = T$ alternately for t steps. Let T_j be the operator after j steps.
Output: T_t

Algorithm 2: Algorithm G

Theorem 2.12 ([Gur04]). *Let T be a completely positive operator. Then T is rank non-decreasing iff $\text{cap}(T) > 0$.*

Theorem 2.13 ([Gur04]). *Let T be a completely positive operator s.t. $\text{cap}(T) > 0$. Then for every $\epsilon > 0$, there exists an operator scaling T' of T s.t. $\text{ds}(T') \leq \epsilon$. Furthermore this T' can be found by running Algorithm G for an appropriate number of steps $t(T, \epsilon)$.*

The first polynomial time algorithms for checking whether T is rank non-decreasing and also computing a multiplicative approximation to capacity were given by [GGOW16]. In the algorithms below, the input will be a completely positive operator T which is given in terms of its Kraus operators A_1, \dots, A_m . We will assume that the entries of the matrices A_i can be described using b bits.

Theorem 2.14 ([GGOW16]). *Let $T : M_n(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ be a completely positive operator s.t. the bit-complexity of the description of T (in terms of its Kraus operators) is b . Then there is a $\text{poly}(n, b)$ time (deterministic) algorithm to test whether T is rank non-decreasing. There is also a $\text{poly}(n, b, 1/\epsilon)$ time algorithm to compute a $(1 + \epsilon)$ -multiplicative approximation to $\text{cap}(T)$. Furthermore, the algorithm outputs a scaling T' of T which is almost doubly stochastic i.e. $\text{cap}(T') \geq 1 - \epsilon$.*

The above mentioned algorithm does not find a witness to the rank decreasing property. Ivanyos, Qiao and Subrahmanyam [IQS15] were the first ones to give an efficient algorithm that computes such a witness. Their algorithm, however, is quite different from the one in [GGOW16] and does not compute capacity.

Theorem 2.15 ([IQS15]). *Let $T : M_n(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ be a completely positive operator s.t. the bit-complexity of the description of T (in terms of its Kraus operators) is b . Then there is a $\text{poly}(n, b)$ time (deterministic) algorithm to test whether T is rank non-decreasing. Furthermore, if T is rank decreasing, then the algorithm also outputs an $X \succeq 0$ s.t. $\text{Rank}(T(X)) < \text{Rank}(X)$.*

3 Rectangular operators

In this section we extend the definitions and main results of the theory of square operators, presented in [Gur04, GGOW16], to rectangular operators, which we will define below. In particular, at the end of the section (Corollaries 3.17 and 3.18), we make explicit what the operator scaling algorithms for square operators imply for rectangular operators.

We define a rectangular operator as any linear map $T : \mathcal{M}_{n_1}(\mathbb{C}) \rightarrow \mathcal{M}_{n_2}(\mathbb{C})$. Choi's characterization [Cho75] extends to rectangular operators as well. And this is how we define completely positive rectangular operators.

Definition 3.1 (Rectangular Completely Positive Operators). An operator (or map) $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$ is called completely positive if there are $n_2 \times n_1$ complex matrices A_1, \dots, A_m s.t. $T(X) = \sum_{i=1}^m A_i X A_i^\dagger$. The matrices A_1, \dots, A_m are called Kraus operators of T (and they are not unique). T is called completely positive trace preserving (cptp) if in addition, $\text{tr}(T(X)) = \text{tr}(X)$ for all X . This is equivalent to the condition $T^*(I) = \sum_{i=1}^m A_i^\dagger A_i = I$.

The following is a natural extension of the definition of doubly stochastic operators to the rectangular case.

Definition 3.2. If $T(X) = \sum_{i=1}^m A_i X A_i^\dagger$ is a completely positive operator, we define its dual T^* by $T^*(X) = \sum_{i=1}^m A_i^\dagger X A_i$. We say T is doubly stochastic if $T\left(\frac{n_2}{n_1} I_{n_1}\right) = I_{n_2}$ and T is trace preserving. This is the same as

$$T\left(\frac{n_2}{n_1} I_{n_1}\right) = I_{n_2} \quad \text{and} \quad T^*(I_{n_2}) = I_{n_1}.$$

Remark 3.3. Note that there is an asymmetry in the above definition w.r.t to T and T^* . This is an arbitrary convention on our part just to ensure that doubly stochastic operators are trace preserving. Not much in the theory would change if one defined doubly stochastic in a different way.

The following is a natural extension of the rank decreasing property to the rectangular case. For a matrix X of dimension n we defined its *fractional-rank* as $\text{Rank}(X)/n$.

Definition 3.4 (Fractional-Rank Decreasing Operators). A completely positive operator T is said to be fractional-rank decreasing if there exists an $X \succeq 0$ s.t. $\frac{\text{rank}(T(X))}{n_2} < \frac{\text{rank}(X)}{n_1}$.

The following is a natural extension of capacity to the rectangular case.

Definition 3.5 (Capacity). The capacity of a completely positive operator $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$, denoted by $\text{cap}(T)$, is defined as

$$\text{cap}(T) = \inf \left\{ \frac{\text{Det} \left(\frac{n_2}{n_1} T(X) \right)}{\text{Det}(X)^{\frac{n_2}{n_1}}} : X \succ 0 \right\} = \inf \left\{ \text{Det} \left(\frac{n_2}{n_1} T(X) \right) : X \succ 0, \text{Det}(X) = 1 \right\}$$

Definition 3.6 (Operator Scaling). An operator T' is called an operator scaling of T if there exist invertible matrices B, C (which are $n_2 \times n_2$ and $n_1 \times n_1$ respectively) s.t.

$$T'(X) = B \cdot T \left(C \cdot X \cdot C^\dagger \right) \cdot B^\dagger$$

Alternatively, if A_1, \dots, A_m are Kraus operators for T , then BA_1C, \dots, BA_mC are Kraus operators for T' .

Definition 3.7 (Distance to double stochasticity).

$$\text{ds}(T) = \text{tr} \left[\left(T \left(\frac{n_2}{n_1} I_{n_1} \right) - I_{n_2} \right)^2 \right] + \text{tr} \left[\left(T^*(I_{n_2}) - I_{n_1} \right)^2 \right]$$

The next two operations are natural extensions of the right and left normalizations to the case of rectangular operators.

Definition 3.8 (Right Normalization). Given a completely positive operator $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$, define its *right normalization* T_R as follows:

$$T_R(X) = T \left(T^*(I_{n_2})^{-1/2} \cdot X \cdot T^*(I_{n_2})^{-1/2} \right)$$

Note that $T_R^*(I_{n_2}) = I_{n_1}$.

Definition 3.9 (Left Normalization). Given a completely positive operator $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$, define its *left normalization* T_L as follows:

$$T_L(X) = \frac{n_1}{n_2} \cdot T(I_{n_1})^{-1/2} \cdot T(X) \cdot T(I_{n_1})^{-1/2}$$

Input: Completely positive operator $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$.

Perform right and left normalizations starting from $T_0 = T$ alternately for t steps. Let T_j be the operator after j steps.

Output: T_t

Algorithm 3: Algorithm G for rectangular operators

Note that $T_L \left(\frac{n_2}{n_1} I_{n_1} \right) = I_{n_2}$.

Algorithm 3 is a natural extension of Algorithm G to the case of rectangular operators.

We will now see how from a rectangular operator $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$, we can define a square operator $\tilde{T} : M_{n_1 n_2}(\mathbb{C}) \rightarrow M_{n_1 n_2}(\mathbb{C})$ that captures most of the properties of T (we can also embed in dimension $\text{lcm}(n_1, n_2)$ but since it only saves a quadratic factor, we ignore this to simplify notation). \tilde{T} is intended to act on block diagonal matrices (the off-diagonal blocks are ignored).

Construction 3.10. Let $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$ be an operator. Given an $n_1 n_2 \times n_1 n_2$ psd matrix X , view it as a $n_2 \times n_2$ block matrix, where each block $X_{i,j}$ is $n_1 \times n_1$. Define $\tilde{T}(X)$ as $n_1 \times n_1$ block diagonal matrix where each diagonal block is the same $n_2 \times n_2$ matrix: $\frac{1}{n_1} \sum_{i=1}^{n_2} T(X_{i,i})$. In other words, we have that

$$\tilde{T}(X) = I_{n_1} \otimes \left(\frac{1}{n_1} \sum_{i=1}^{n_2} T(X_{i,i}) \right)$$

In terms of Kraus operators, if A_1, \dots, A_m are Kraus operators for T , then a set of Kraus operators for \tilde{T} are given by the set $\left\{ \frac{1}{n_1} E_{i,j} \otimes A_k \right\}_{i=1, j=1, k=1}^{n_1, n_2, m}$. Here $E_{i,j}$ is the elementary matrix with a 1 in position (i, j) and 0's everywhere else.

The next lemma shows some useful relations between the operators T and \tilde{T} .

Lemma 3.11. If $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$ is an operator and \tilde{T} is the operator obtained from Construction 3.10, the following hold:

1. $\text{cap}(\tilde{T}) = \text{cap}(T)^{n_1}$.
2. T is fractional-rank non-decreasing iff \tilde{T} is rank non-decreasing.

Proof.

$$\begin{aligned} \text{cap}(\tilde{T}) &= \inf \left\{ \text{Det} \left(\tilde{T}(X) \right) : X \succ 0, \text{Det}(X) = 1 \right\} \\ &= \inf \left\{ \text{Det} \left(\frac{1}{n_1} \sum_{i=1}^{n_2} T(X_{i,i}) \right)^{n_1} : X \succ 0, \text{Det}(X) = 1 \right\} \\ &= \inf \left\{ \text{Det} \left(\frac{n_2}{n_1} \cdot \frac{1}{n_2} \sum_{i=1}^{n_2} T(X_{i,i}) \right)^{n_1} : X_{i,i} \succ 0 \text{ for all } i, \prod_{i=1}^{n_2} \text{Det}(X_{i,i}) = 1 \right\} \\ &= \inf \left\{ \text{Det} \left(\frac{n_2}{n_1} T(Y) \right)^{n_1} : Y \succ 0, \text{Det}(Y) = 1 \right\} \\ &= \text{cap}(T)^{n_1} \end{aligned}$$

The second equality follows from a generalization of Hadamard's inequality, determinant of a block diagonal psd matrix is smaller than the product of determinants of its diagonal blocks:

$$\text{Det}(X) \leq \prod_{i=1}^{n_2} \text{Det}(X_{i,i})$$

The third equality follows from replacing $\frac{1}{n_2} \sum_{i=1}^{n_2} X_{i,i}$ by Y and using the log-concavity of determinant:

$$\text{Det}(Y) = \text{Det} \left(\frac{1}{n_2} \sum_{i=1}^{n_2} X_{i,i} \right) \geq \prod_{i=1}^{n_2} \text{Det}(X_{i,i})^{1/n_2} = 1$$

Now let us move on to the second part of the lemma. Suppose T is fractional-rank decreasing. So there exists $Y \succeq 0$ s.t.

$$\frac{\text{Rank}(T(Y))}{n_2} < \frac{\text{Rank}(Y)}{n_1}$$

Let X be the $n_2 \times n_2$ block diagonal matrix, all of whose diagonal blocks are Y . Then $\tilde{T}(X)$ is a $n_1 \times n_1$ block diagonal matrix all of whose diagonal blocks are $\frac{n_2}{n_1} T(Y)$. Hence

$$\text{Rank}(\tilde{T}(X)) = n_1 \text{Rank}(T(Y)) < n_2 \text{Rank}(Y) = \text{Rank}(X)$$

and thus \tilde{T} is rank decreasing. The other direction is more interesting. Suppose \tilde{T} is rank decreasing and it decreases the rank of X . Then we will prove that T is fractional-rank decreasing and $\sum_{i=1}^{n_2} X_{i,i}$ is a witness i.e.

$$\frac{\text{Rank}(T(\sum_{i=1}^{n_2} X_{i,i}))}{n_2} < \frac{\text{Rank}(\sum_{i=1}^{n_2} X_{i,i})}{n_1}$$

Now

$$\begin{aligned} n_1 \cdot \text{Rank} \left(T \left(\sum_{i=1}^{n_2} X_{i,i} \right) \right) &= \text{Rank}(\tilde{T}(X)) \\ &< \text{Rank}(X) \end{aligned}$$

So we would be done if we can prove that

$$\text{Rank}(X) \leq n_2 \cdot \text{Rank} \left(\sum_{i=1}^{n_2} X_{i,i} \right)$$

This follows from the following chain of inequalities:

$$\text{Rank}(X) \leq \sum_{i=1}^{n_2} \text{Rank}(X_{i,i}) \leq n_2 \cdot \max_{i=1 \text{ to } n_2} \text{Rank}(X_{i,i}) \leq n_2 \cdot \text{Rank} \left(\sum_{i=1}^{n_2} X_{i,i} \right)$$

The last inequality is true because for psd matrices Y_1, Y_2 , $\text{Rank}(Y_1 + Y_2) \geq \text{Rank}(Y_1)$ and also the fact that $X_{i,i}$ are psd matrices since X is psd. The second inequality is obvious. The first inequality follows from the following claim:

Claim 3.12. *Let*

$$X = \begin{bmatrix} X_{1,1} & X_{1,2} \\ X_{2,1} & X_{2,2} \end{bmatrix}$$

be a psd matrix s.t. $X_{1,1}$ and $X_{2,2}$ are square matrices of possibly different dimension. Then $\text{Rank}(X) \leq \text{Rank}(X_{1,1}) + \text{Rank}(X_{2,2})$.

Proof. We will prove that $\dim(\ker(X_{1,1})) + \dim(\ker(X_{2,2})) \leq \dim(\ker(X))$ from which the claim easily follows. Let $v \in \ker(X_{1,1})$. Then look at the vector

$$v' = \begin{bmatrix} v \\ 0 \end{bmatrix}$$

Since $X_{1,1}v = 0$, we get that $v'^{\dagger}Xv' = 0$. Since X is psd, this implies that $v' \in \ker(X)$. Similarly, if $w \in \ker(X_{2,2})$, the vector

$$w' = \begin{bmatrix} 0 \\ w \end{bmatrix}$$

lies in $\ker(X)$. The proof is finished by observing that v' and w' lie in orthogonal subspaces. \square

\square

Now we will see that with Construction 3.10, most of the results about square operators carry over to the rectangular case.

Corollary 3.13. *Let $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$ be a completely positive operator. Then T is fractional-rank non-decreasing iff $\text{cap}(T) > 0$.*

Proof. Immediate from Theorem 2.12 and Lemma 3.11. \square

For any completely positive operator, we have that $\text{cap}(T) > 0$ iff $\text{cap}(T^*) > 0$. This is proved below:

Proposition 3.14. *Let $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$ be a completely positive operator. Then $\text{cap}(T)^{1/n_2} = \frac{n_2}{n_1} \text{cap}(T^*)^{1/n_1}$.*

Proof.

$$\text{cap}(T)^{1/n_2} = \inf_{X \succ 0} \left\{ \frac{\text{Det} \left(\frac{n_2}{n_1} T(X) \right)^{1/n_2}}{\text{Det}(X)^{1/n_1}} \right\}$$

and

$$\frac{1}{\text{cap}(T^*)^{1/n_1}} = \sup_{Y \succ 0} \left\{ \frac{\text{Det}(Y)^{1/n_2}}{\text{Det} \left(\frac{n_1}{n_2} T^*(Y) \right)^{1/n_1}} \right\}$$

Thus

$$\begin{aligned}
\frac{\text{cap}(T)^{1/n_2}}{\text{cap}(T^*)^{1/n_1}} &= \inf_{X \succ 0} \sup_{Y \succ 0} \left\{ \frac{\text{Det} \left(\frac{n_2}{n_1} T(X) \cdot Y \right)^{1/n_2}}{\text{Det} \left(\frac{n_1}{n_2} X \cdot T^*(Y) \right)^{1/n_1}} \right\} \\
&\geq \inf_{X \succ 0} \left\{ \frac{\text{Det} \left(\frac{n_2}{n_1} I_{n_2} \right)^{1/n_2}}{\text{Det} \left(\frac{n_1}{n_2} X \cdot T^*(T(X)^{-1}) \right)^{1/n_1}} \right\} \\
&\geq \frac{n_2}{n_1} \cdot \inf_{X \succ 0} \left\{ \frac{n_1}{\text{tr} \left[\frac{n_1}{n_2} X \cdot T^*(T(X)^{-1}) \right]} \right\} \\
&= \frac{n_2}{n_1} \cdot \inf_{X \succ 0} \left\{ \frac{n_1}{\text{tr} \left[\frac{n_1}{n_2} T(X) \cdot T(X)^{-1} \right]} \right\} \\
&= \frac{n_2}{n_1}
\end{aligned}$$

The first inequality follows by taking $Y = T(X)^{-1}$. Second inequality follows from AM-GM. Second equality follows from the fact that $\text{tr}[A \cdot T^*(B)] = \text{tr}[T(A) \cdot B]$ for all matrices A, B of appropriate sizes. We have proved that

$$\frac{\text{cap}(T)^{1/n_2}}{\text{cap}(T^*)^{1/n_1}} \geq \frac{n_2}{n_1}$$

By applying the same argument to T^* instead of T , we get that

$$\frac{\text{cap}(T^*)^{1/n_1}}{\text{cap}(T)^{1/n_2}} \geq \frac{n_1}{n_2}$$

Combining these completes the proof. \square

Thus, we can deduce the following corollary:

Corollary 3.15. *For any completely positive operator $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$, $\text{cap}(T) > 0$ iff T^* is fractional-rank non-decreasing.*

Similar to the square case, whenever a rectangular operator T is fractional-rank non-decreasing, there exists an operator scaling which is almost doubly stochastic.

Corollary 3.16. *Let $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$ be a completely positive operator s.t. $\text{cap}(T) > 0$. Then for every $\epsilon > 0$, there exists an operator scaling T' of T s.t. $\text{ds}(T') \leq \epsilon$. Furthermore this T' can be found by running Algorithm 3 for an appropriate number of steps $t(T, \epsilon)$.*

Proof. Given a $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$, let us construct a $\tilde{T} : M_{n_1 n_2}(\mathbb{C}) \rightarrow M_{n_1 n_2}(\mathbb{C})$ via Construction 3.10. Since $\text{cap}(T) > 0$, by Lemma 3.11, $\text{cap}(\tilde{T}) > 0$. Then by Theorem 2.13, if we run Algorithm G on \tilde{T} for an appropriate number of steps $t(\tilde{T}, \epsilon')$, the output $\tilde{Z} = \tilde{T}_{t(\tilde{T}, \epsilon')}$ would satisfy $\text{ds}(\tilde{Z}) \leq \epsilon'$. It is not hard to see that \tilde{Z} is of the form

$$\tilde{Z}(X) = I_{n_1} \otimes \left(\frac{1}{n_1} \sum_{i=1}^{n_2} Z(X_{i,i}) \right)$$

where Z is the result of running Algorithm 3 on T for $t(\tilde{T}, \epsilon')$ steps. Also $\text{ds}(\tilde{Z}) = n_1 \text{ds}(Z)$, so choosing $\epsilon' = \epsilon/n_1$ proves that running Algorithm 3 on T for $t(\tilde{T}, \epsilon'/n_1)$ steps outputs T' s.t. $\text{ds}(T') \leq \epsilon$. \square

Algorithms for computing capacity of square operators can be used in a black box way to compute the capacity of rectangular operators because of Construction 3.10 and Lemma 3.11.

Corollary 3.17. *Let $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$ be a completely positive operator s.t. the bit-complexity of the description of T (in terms of its Kraus operators) is b . Then there is a $\text{poly}(n_1, n_2, b)$ time (deterministic) algorithm to test whether T is fractional-rank non-decreasing. There is also a $\text{poly}(n_1, n_2, b, 1/\epsilon)$ time algorithm to compute a $(1 + \epsilon)$ -multiplicative approximation to $\text{cap}(T)$. Furthermore, the algorithm outputs a scaling T' of T which is almost doubly stochastic i.e. $\text{cap}(T') \geq 1 - \epsilon$.*

Similarly, the algorithm in Theorem 2.15 [IQS15] for computing witnesses to rank decreasing property of square operators can be used in a black box way to compute witnesses to fractional-rank decreasing property of rectangular operators.

Corollary 3.18. *Let $T : M_{n_1}(\mathbb{C}) \rightarrow M_{n_2}(\mathbb{C})$ be a completely positive operator s.t. the bit-complexity of the description of T (in terms of its Kraus operators) is b . Then there is a $\text{poly}(n_1, n_2, b)$ time (deterministic) algorithm to test whether T is fractional-rank non-decreasing. Furthermore, if T is fractional-rank decreasing, then the algorithm also outputs an $X \succeq 0$ s.t. $\frac{\text{Rank}(T(X))}{n_2} < \frac{\text{Rank}(X)}{n_1}$.*

4 Reduction from Brascamp-Lieb to completely positive operators

In this section we will show how each Brascamp-Lieb datum (\mathbf{B}, \mathbf{p}) corresponds to an operator $T_{(\mathbf{B}, \mathbf{p})}$ for which $\text{cap}(T_{(\mathbf{B}, \mathbf{p})}) = 1/\text{BL}(\mathbf{B}, \mathbf{p})^2$. From this correspondence, and from the connection between capacity and rank expansion of an operator, we will show how to derive the [BCCT08a] conditions for finiteness of $\text{BL}(\mathbf{B}, \mathbf{p})$. Recall that in the Brascamp-Lieb setting, we are given a datum (\mathbf{B}, \mathbf{p}) consisting of m matrices B_1, \dots, B_m , where each B_i is a $n_i \times n$ matrix. We also have m non-negative rational numbers p_1, \dots, p_m s.t. $p_i = c_i/d$ for some $c_i, d \in \mathbb{N}$. Since one necessary condition for finiteness of the Brascamp-Lieb constant is that $n = \sum_{i=1}^m p_i n_i$, we will assume this equality throughout the section.

We begin by recalling the characterization of the Brascamp-Lieb constant due to Lieb [Lie90]:

Theorem 4.1 (Brascamp-Lieb constant). *Given a Brascamp-Lieb datum (\mathbf{B}, \mathbf{p}) , the Brascamp-Lieb constant $\text{BL}(\mathbf{B}, \mathbf{p})$ equals:*

$$\text{BL}(\mathbf{B}, \mathbf{p}) = \sup \left\{ \left(\frac{\prod_{i=1}^m \text{Det}(X_i)^{p_i}}{\text{Det} \left(\sum_{i=1}^m p_i B_i^\dagger X_i B_i \right)} \right)^{1/2} : X_i \succ 0 \text{ is a } n_i \times n_i \text{ matrix} \right\}$$

We now describe our main reduction from the BL datum to (rectangular) completely positive operators.

Construction 4.2. Let $m' = \sum_{i=1}^m c_i$. We will construct m' $nd \times n$ matrices $A_1, \dots, A_{m'}$ from B_1, \dots, B_m . Essentially $A_1, \dots, A_{m'}$ will consist of c_i copies of B_i acting on different parts of the space. Consider the following matrix A :

$$A = \left(\begin{array}{c} B_1 \\ \vdots \\ B_1 \\ B_2 \\ \vdots \\ B_2 \\ \vdots \\ B_m \\ \vdots \\ B_m \end{array} \right) \left. \begin{array}{l} \left. \begin{array}{c} \\ \\ \\ \end{array} \right\} c_1 \text{ times} \\ \left. \begin{array}{c} \\ \\ \\ \end{array} \right\} c_2 \text{ times} \\ \vdots \\ \left. \begin{array}{c} \\ \\ \\ \end{array} \right\} c_m \text{ times} \end{array} \right\}$$

A has a natural block structure consisting of a total of m' blocks. The matrices A_j are obtained by keeping one block intact and zeroing out all other blocks. More formally, suppose j is such that $\sum_{i=1}^k c_i < j \leq \sum_{i=1}^{k+1} c_i$. Then all entries of A_j will be zero except rows

$$\sum_{i=1}^k c_i n_i + \left(j - \sum_{i=1}^k c_i - 1 \right) \cdot n_{k+1} + 1 \text{ to } \sum_{i=1}^k c_i n_i + \left(j - \sum_{i=1}^k c_i - 1 \right) \cdot n_{k+1} + n_{k+1}$$

which are occupied by B_{k+1} . Define the operator $T_{(\mathbf{B}, \mathbf{p})} : M_{nd}(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ as

$$T_{(\mathbf{B}, \mathbf{p})}(X) = \sum_{j=1}^{m'} A_j^\dagger X A_j$$

Definition 4.3 (Brascamp-Lieb operator). The completely positive operator $T_{(\mathbf{B}, \mathbf{p})}$ constructed in Construction 4.2 is called the *Brascamp-Lieb operator* for datum (\mathbf{B}, \mathbf{p}) .

The next lemma relates important properties of a BL datum and its Brascamp-Lieb operator.

Lemma 4.4. Let (\mathbf{B}, \mathbf{p}) be a Brascamp-Lieb datum, where $p_i = c_i/d$ for $c_i, d \in \mathbb{N}$ and $B_i : \mathbb{R}^n \rightarrow \mathbb{R}^{n_i}$. Moreover, assume that $n = \sum_{i=1}^m c_i n_i/d$. Let $T_{(\mathbf{B}, \mathbf{p})}$ be the Brascamp-Lieb operator corresponding to (\mathbf{B}, \mathbf{p}) . Then

1. If $\text{BL}(\mathbf{B}, \mathbf{p}) < \infty$, then

$$\text{cap}(T_{(\mathbf{B}, \mathbf{p})}) = 1/\text{BL}(\mathbf{B}, \mathbf{p})^2$$

2. The inequality

$$\dim(V) \leq \sum_j p_j \dim(B_j V)$$

holds for all subspaces $V \subseteq \mathbb{R}^n$ iff $T_{(\mathbf{B}, \mathbf{p})}^*$ is fractional-rank non-decreasing.

Moreover, if B_i 's are integer matrices such that $\|B_i\|_\infty \leq M$, then the same holds for the matrices $A_1, \dots, A_{m'}$.

Proof. Let us start by proving the first part of the lemma. For j s.t. $\sum_{i=1}^k c_i < j \leq \sum_{i=1}^{k+1} c_i$ (let us denote this by $i_j = k+1$), let us denote by X_j the $n_{k+1} \times n_{k+1}$ submatrix of X indexed by rows and columns from

$$\sum_{i=1}^k c_i n_i + \left(j - \sum_{i=1}^k c_i - 1 \right) \cdot n_{k+1} + 1 \text{ to } \sum_{i=1}^k c_i n_i + \left(j - \sum_{i=1}^k c_i - 1 \right) \cdot n_{k+1} + n_{k+1}$$

Then it can be seen that

$$T_{(\mathbf{B}, \mathbf{p})}(X) = \sum_{j=1}^{m'} B_{i_j}^\dagger X_j B_{i_j}$$

Now let us calculate the capacity of $T_{(\mathbf{B}, \mathbf{p})}$:

$$\begin{aligned} \text{cap}(T_{(\mathbf{B}, \mathbf{p})}) &= \inf \left\{ \text{Det} \left(\frac{T_{(\mathbf{B}, \mathbf{p})}(X)}{d} \right) : X \succ 0, \text{Det}(X) = 1 \right\} \\ &= \inf \left\{ \text{Det} \left(\frac{\sum_{j=1}^{m'} B_{i_j}^\dagger X_j B_{i_j}}{d} \right) : X \succ 0, \text{Det}(X) = 1 \right\} \\ &= \inf \left\{ \text{Det} \left(\frac{\sum_{j=1}^{m'} B_{i_j}^\dagger X_j B_{i_j}}{d} \right) : X_j \succ 0, \prod_{j=1}^{m'} \text{Det}(X_j) = 1, X_j \text{ is } n_{i_j} \times n_{i_j} \right\} \\ &= \inf \left\{ \text{Det} \left(\frac{\sum_{i=1}^m c_i B_i^\dagger Y_i B_i}{d} \right) : Y_i \succ 0, \prod_{i=1}^m \text{Det}(Y_i)^{c_i} = 1, Y_i \text{ is } n_i \times n_i \right\} \\ &= \inf \left\{ \frac{\text{Det} \left(\frac{\sum_{i=1}^m c_i B_i^\dagger Y_i B_i}{d} \right)}{\prod_{i=1}^m \text{Det}(Y_i)^{c_i/d}} : Y_i \succ 0, Y_i \text{ is } n_i \times n_i \right\} \\ &= \inf \left\{ \frac{\text{Det} \left(\sum_{i=1}^m p_i B_i^\dagger Y_i B_i \right)}{\prod_{i=1}^m \text{Det}(Y_i)^{p_i}} : Y_i \succ 0, Y_i \text{ is } n_i \times n_i \right\} \\ &= 1/\text{BL}(\mathbf{B}, \mathbf{p})^2 \end{aligned}$$

For third equality, we use the generalization of Hadamard's inequality. For fourth equality, we use log-concavity of determinant after replacing $\sum_{j:i_j=k} X_j$ by $c_k Y_k$.

Now we move on to prove the second part of the lemma. The condition that the dual operator $T_{(\mathbf{B}, \mathbf{p})}^* : \mathcal{M}_n(\mathbb{C}) \rightarrow \mathcal{M}_{nd}(\mathbb{C})$ is fractional-rank non-decreasing is the following:

$$\text{Rank}(T_{(\mathbf{B}, \mathbf{p})}^*(X))/d \geq \text{Rank}(X), \text{ for all } X \succeq 0$$

As the matrices A_j defined in Construction 4.2 correspond to distinct copies of B_j acting on X , the dual operator corresponds to the following transformation:

$$T_{(\mathbf{B}, \mathbf{p})}^*(X) = \bigoplus_{j=1}^m \left(\bigoplus_{i=1}^{c_j} B_i X B_i^\dagger \right) \quad (1)$$

Hence, the condition $\text{Rank}(T^*(X))/d \geq \text{Rank}(X)$ becomes:

$$\text{Rank}(X) \leq \text{Rank}(T^*(X))/d = \frac{1}{d} \cdot \sum_{j=1}^m c_j \text{Rank}(B_j X B_j^\dagger) = \sum_{j=1}^m p_j \text{Rank}(B_j X B_j^\dagger).$$

Let $V = \langle \mathbf{v}_1, \dots, \mathbf{v}_k \rangle$ be an orthonormal eigenvector basis for X , assuming $\text{Rank}(X) = k$. Hence, we have $\lambda_i > 0$ such that $X = \sum_{i=1}^k \lambda_i \mathbf{v}_i \mathbf{v}_i^\dagger$ and the inequality above becomes:

$$\begin{aligned} \dim(V) = \text{Rank}(X) &\leq \sum_{j=1}^m p_j \text{Rank}(B_j X B_j^\dagger) \\ &= \sum_{j=1}^m p_j \text{Rank} \left(B_j \left(\sum_{i=1}^k \lambda_i \mathbf{v}_i \mathbf{v}_i^\dagger \right) B_j^\dagger \right) \\ &= \sum_{j=1}^m p_j \dim(B_j V). \end{aligned}$$

As this inequality must hold for any $X \succeq 0$, it must be the case that it holds for all spaces $V \subseteq \mathbb{R}^n$. This is precisely the condition of [BCCT08a].

Notice that by our construction of the matrices A_j , if B_i 's are integer matrices with $\|B_i\|_\infty \leq M$, then matrices A_j will also be integer matrices whose entries are bounded by M . \square

Now we will see that with Construction 4.2 and Lemma 4.4, we can deduce the following corollaries for the BL datum. The next corollary is Theorem 1.1 from the introduction proved in [BCCT08a].

Corollary 4.5. *Let (\mathbf{B}, \mathbf{p}) be a Brascamp-Lieb datum, where $p_i = c_i/d$ for $c_i, d \in \mathbb{N}$ and $B_i : \mathbb{R}^n \rightarrow \mathbb{R}^{n_i}$. Moreover, assume that $n = \sum_{i=1}^m c_i n_i/d$. Then $\text{BL}(\mathbf{B}, \mathbf{p})$ is finite iff for every subspace $V \subseteq \mathbb{R}^n$, it holds that*

$$\dim(V) \leq \sum_{i=1}^m p_i \dim(B_i(V))$$

Proof. Follows immediately from Corollary 3.15 and Lemma 4.4. \square

Next we will derive an algorithm for checking finiteness of the Brascamp-Lieb constant from Construction 4.2 and Corollary 3.18.

Corollary 4.6. *Let (\mathbf{B}, \mathbf{p}) be a Brascamp-Lieb datum, where $p_i = c_i/d$ for $c_i, d \in \mathbb{N}$ and $B_i : \mathbb{R}^n \rightarrow \mathbb{R}^{n_i}$. Moreover, assume that $n = \sum_{i=1}^m c_i n_i/d$. Let b be the total bit size of \mathbf{B} . Then there is a $\text{poly}(b, d)$ time (deterministic) algorithm to check if $\text{BL}(\mathbf{B}, \mathbf{p})$ is finite. Furthermore, if $\text{BL}(\mathbf{B}, \mathbf{p})$ is not finite, then the algorithm outputs a subspace $V \subseteq \mathbb{R}^n$ s.t.*

$$\dim(V) > \sum_{i=1}^m p_i \dim(B_i(V))$$

Proof. Let $T_{(\mathbf{B}, \mathbf{p})} : \mathcal{M}_{nd}(\mathbb{C}) \rightarrow \mathcal{M}_n(\mathbb{C})$ be the operator defined in Construction 4.2 and let $T_{(\mathbf{B}, \mathbf{p})}^* : \mathcal{M}_n(\mathbb{C}) \rightarrow \mathcal{M}_{nd}(\mathbb{C})$ be its dual. By Lemma 4.4 and Proposition 3.14, we have that $\text{BL}(\mathbf{B}, \mathbf{p})$ is finite

iff $T_{(\mathbf{B}, \mathbf{p})}^*$ is fractional-rank non-decreasing. Corollary 3.18 gives us a deterministic $\text{poly}(b, d)$ -time algorithm to test whether the operator $T_{(\mathbf{B}, \mathbf{p})}^*$ is fractional-rank non-decreasing, and if it is not, the algorithm outputs $X \succeq 0$ such that

$$\frac{\text{rank}(T(X))}{nd} < \frac{\text{rank}(X)}{n} \quad (2)$$

In the latter case, to find the subspace V which violates the dimension inequality, let $V = \langle \mathbf{v}_1, \dots, \mathbf{v}_k \rangle$ be an orthonormal eigenvector basis for X . Hence, we have $X = \sum_{i=1}^k \lambda_i \mathbf{v}_i \mathbf{v}_i^\dagger$. According to equation (1),

$$T_{(\mathbf{B}, \mathbf{p})}^*(X) = \bigoplus_{j=1}^m \left(\bigoplus_{i=1}^{c_j} B_i X B_i^\dagger \right)$$

Therefore, from (2) we obtain:

$$\begin{aligned} \frac{\text{rank}(X)}{n} > \frac{\text{rank}(T(X))}{nd} &\Rightarrow \dim(V) = \text{rank}(X) > \frac{\text{rank}(T(X))}{d} \\ &= \sum_{j=1}^m p_j \text{Rank}(B_j X B_j^\dagger) \\ &= \sum_{j=1}^m p_j \text{Rank} \left(B_j \left(\sum_{i=1}^k \lambda_i \mathbf{v}_i \mathbf{v}_i^\dagger \right) B_j^\dagger \right) \\ &= \sum_{j=1}^m p_j \dim(B_j V). \end{aligned}$$

The above proves that $\text{Im}(X)$ is a vector space which violates the dimension inequalities. \square

4.1 Connection to quiver semi-invariants

In this short subsection, we provide an invariant theoretic view of our reduction¹⁵. It is far from being self-contained and we refer the reader to the survey [DW00] about semi-invariants of quiver representations. In this theory, there is a general way of generating semi-invariants for general quivers by a reduction to the Kronecker quiver (left-right action) ([DW00, DZ01, SdB01], also see Section 5 in [DM15]). As it happens, the data to both BL inequalities and to completely positive operators fits the descriptions of the “star quiver” and the “Kronecker quiver” respectively, and with hindsight, the reduction described in Construction 4.2 (when composed with the reduction from rectangular to square operators) can be viewed as a special case of the reduction above. More specifically, semi-invariants for the star quiver with dimensions $(n, n_1, n_2, \dots, n_m)$ and weight vector $(d, -c_1, -c_2, \dots, -c_m)$ can be generated by reducing to a Kronecker quiver of appropriate dimensions (see Figure 1). It is worth emphasizing here the intriguing connection (that this paper exploits) between analysis (BL inequalities) and algebra (semi-invariants of the star quiver): a BL datum $(B_1, \dots, B_m, c_1/d, \dots, c_m/d)$ is not feasible iff (B_1, \dots, B_m) is the common

¹⁵We remark that other connections between BL inequalities and quiver representations are mentioned in [BCCT08a].

zero of the semi-invariants of the star quiver with dimensions $(n, n_1, n_2, \dots, n_m)$ and weight vectors $\{(l \cdot d, -l \cdot c_1, -l \cdot c_2, \dots, -l \cdot c_m) : l \in \mathbb{N}\}$ ¹⁶.

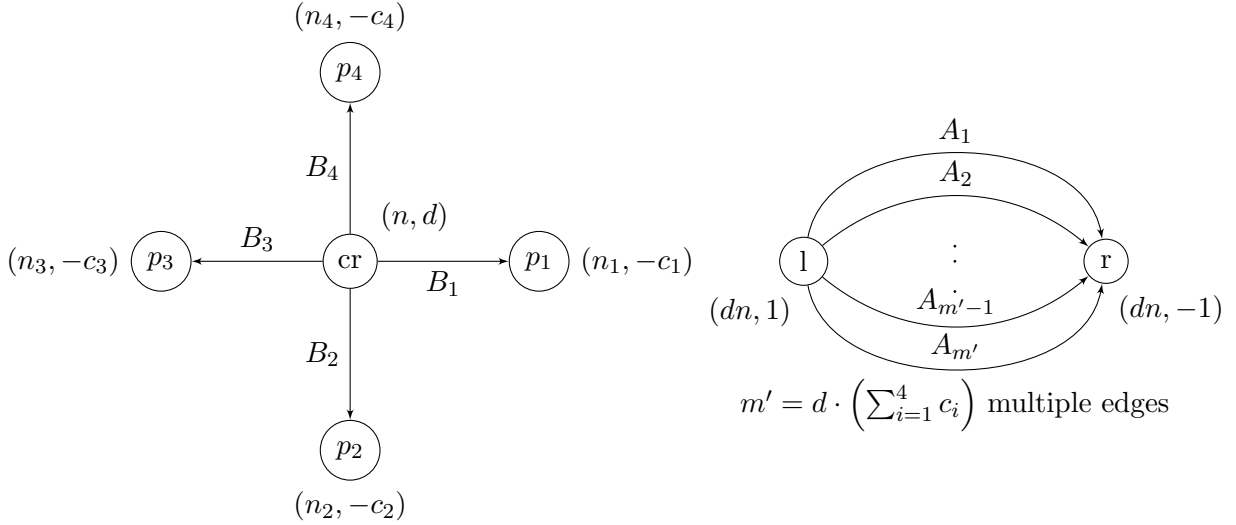


Figure 1: Reduction from star quiver to the Kronecker quiver. Nodes are labelled with (d, w) : d is the dimension of the vector space at this node and w is the weight corresponding to this node. Edges are labelled with maps from the vector space at the head of the edge to the vector space at the tail of the edge.

5 Upper bound on the Brascamp-Lieb constant

In this section, we will prove an exponential upper bound on the BL constant $\text{BL}(\mathbf{B}, \mathbf{p})$ (Theorem 1.4). The main ingredient is a lower bound on the capacity of square completely positive operators with integer entries from [GGOW16].

Theorem 5.1 (Capacity of Square Operators, Theorem 2.18 in [GGOW16]). *Suppose T_A is a completely positive operator which has positive capacity and has Kraus operators $A_1, \dots, A_m \in \mathcal{M}_n(\mathbb{Z})$. In this case:*

$$\text{cap}(T_A) \geq \exp(-2n \log(n))$$

We now derive the upper bound on the BL-constant from the lower bound on capacity above via our main reduction. We do so first for BL-data with integer entries (which is independent of their magnitude!), and then for general ones with rational entries.

Theorem 5.2 (Lower Bound on Brascamp-Lieb Capacity). *Let (\mathbf{B}, \mathbf{p}) be a Brascamp-Lieb datum, with $B_i \in \mathcal{M}_{n_i \times n}(\mathbb{Z})$ and $p_i = c_i/d$, where $c_i, d \in \mathbb{N}$. Let $T_{(\mathbf{B}, \mathbf{p})}$ be the Brascamp-Lieb operator, as defined in Construction 4.2. If $T_{(\mathbf{B}, \mathbf{p})}$ has positive capacity, then the following bound holds:*

$$\text{cap}(T_{(\mathbf{B}, \mathbf{p})}) \geq \exp(-2n \log(n^2 d)).$$

¹⁶One can even define such a connection for the finiteness of a specific BL constant: $\text{BL}(\mathbf{B}, \mathbf{p}) = \infty$ iff \mathbf{B} is in the null cone of the subgroup of $\text{GL}_n(\mathbb{C}) \times \text{GL}_{n_1}(\mathbb{C}) \times \dots \times \text{GL}_{n_m}(\mathbb{C})$ satisfying $\text{Det}(C) = \prod_{i=1}^m \text{Det}(C_i)^{p_i}$.

Proof. According to Construction 4.2, we construct an operator $T_{(\mathbf{B}, \mathbf{p})} : \mathcal{M}_{nd}(\mathbb{C}) \rightarrow \mathcal{M}_n(\mathbb{C})$ such that $\text{cap}(T_{(\mathbf{B}, \mathbf{p})}) = 1/\text{BL}(\mathbf{B}, \mathbf{p})^2$ and $T_{(\mathbf{B}, \mathbf{p})}(X) = \sum_{j=1}^{m'} A_j^\dagger X A_j$, where $m' \leq nd$, each A_j has integer coordinates. From operator $T_{(\mathbf{B}, \mathbf{p})}$, construct operator $\tilde{T}_{(\mathbf{B}, \mathbf{p})} : \mathcal{M}_{n^2d} \rightarrow \mathcal{M}_{n^2d}$ according to Construction 3.10. Item 1 of Lemma 3.11 implies that

$$\text{cap}(T_{(\mathbf{B}, \mathbf{p})}) = \text{cap}(\tilde{T}_{(\mathbf{B}, \mathbf{p})})^{1/nd}.$$

Moreover, item 1 of Lemma 3.11, together with the assumption that $T_{(\mathbf{B}, \mathbf{p})}$ has positive capacity, implies that $\tilde{T}_{(\mathbf{B}, \mathbf{p})}$ has positive capacity. Hence, Theorem 5.1 implies that

$$\text{cap}(\tilde{T}_{(\mathbf{B}, \mathbf{p})}) \geq \exp(-2n^2d \log(n^2d)).$$

Thus, we have:

$$\text{cap}(T_{(\mathbf{B}, \mathbf{p})}) = \text{cap}(\tilde{T}_{(\mathbf{B}, \mathbf{p})})^{1/nd} \geq \exp(-2n \log(n^2d)).$$

□

From the theorem above we obtain the following upper bound on the Brascamp-Lieb constant. Note that we get a bound which is independent of the magnitude of the entries of A_i 's. This is possible because we restrict A_i 's to be integer matrices.

Corollary 5.3 (Upper Bound on Brascamp-Lieb Constant). *Let (\mathbf{B}, \mathbf{p}) be a Brascamp-Lieb datum, with $B_i \in \mathcal{M}_{n_i \times n}(\mathbb{Z})$ and $p_i = c_i/d$, where $c_i, d \in \mathbb{N}$. If $\text{BL}(\mathbf{B}, \mathbf{p}) < \infty$, then the following bound holds:*

$$\text{BL}(\mathbf{B}, \mathbf{p}) \leq \exp(4n \log(n^2d)).$$

As a corollary, we can also get an exponential bound that is independent of d .

Corollary 5.4. *Let (\mathbf{B}, \mathbf{p}) be a Brascamp-Lieb datum, with $B_i \in \mathcal{M}_{n_i \times n}(\mathbb{Z})$ and p_1, \dots, p_m arbitrary positive reals. If $\text{BL}(\mathbf{B}, \mathbf{p}) < \infty$, then the following bound holds:*

$$\text{BL}(\mathbf{B}, \mathbf{p}) \leq \exp(12nm \log(mn)).$$

Proof. It is proven in [CDK⁺15] that $\log(\text{BL}(\mathbf{B}, \mathbf{p}))$ is convex in \mathbf{p} . This can also be seen easily from the entropy formulation of Brascamp-Lieb inequalities [CCE09, LCCV16]. Thus to upper bound, $\text{BL}(\mathbf{B}, \mathbf{p})$, it is enough to look at the vertices of the BL polytope $P_{\mathbf{B}}$ (from the log-convexity, it follows that if $\text{BL}(\mathbf{B}, \mathbf{p}) < \infty$ for some \mathbf{p} , then the maximum finite value will be achieved at the vertices of $P_{\mathbf{B}}$). It is not hard to see that the vertices of the BL polytope $P_{\mathbf{B}}$ have common denominator $d \leq m! \cdot n^m$. Plugging this into Corollary 5.3 gives us the desired bound. □

From the bound on the BL constant for integer datum, we can also get a bound for arbitrary rational datum in terms of the bit sizes (which is a restatement of Theorem 1.4 in the introduction.).

Corollary 5.5. *Let (\mathbf{B}, \mathbf{p}) be a Brascamp-Lieb datum, with $B_i \in \mathcal{M}_{n_i \times n}(\mathbb{Q})$ and $p_i = c_i/d$, where each entry of B_i has bit-size at most \tilde{b} and $c_i, d \in \mathbb{N}$. If $\text{BL}(\mathbf{B}, \mathbf{p}) < \infty$, then the following bound holds:*

$$\text{BL}(\mathbf{B}, \mathbf{p}) \leq \exp(4n \log(n^2d) + n\tilde{b}).$$

Proof. Note that $\tilde{B}_i = 2^{\tilde{b}} B_i$ are integer matrices. Also it is easy to see that $\text{BL}(\mathbf{B}, \mathbf{p}) = \text{BL}(\tilde{\mathbf{B}}, \mathbf{p}) \cdot 2^{n\tilde{b}}$. Then the desired bound follows from Corollary 5.3. \square

Remark 5.6. *As before, d can be eliminated from the above corollary if needed but the bound in the exponential would become quadratic.*

6 Computing the Brascamp-Lieb constant

In this section we show how Algorithm G with truncation, described in [Gur04, GGOW16], computes the Brascamp-Lieb constant up to a multiplicative factor of $(1 \pm \epsilon)$. We begin by restating Theorem 3.5 from [GGOW16], which states that we can approximate the capacity of a completely positive operator.

Theorem 6.1 (Theorem 3.5 from [GGOW16]). *Let T be a completely positive operator, whose Kraus operators are given by $n \times n$ matrices $A_1, \dots, A_m \in \mathcal{M}_n(\mathbb{Z})$, such that each entry of A_i has absolute value at most M . Algorithm G, with truncation parameter*

$$P(n, 1/\epsilon, \log(M)) = \frac{1}{\epsilon} \cdot (n^{12} \log^4(Mn)) \cdot \log(n^4/\epsilon^2)$$

when applied for $t = \frac{4n^3}{\epsilon^2} \cdot (1 + 10n^2 \log(Mn))$ steps approximates $\text{cap}(T)$ within a multiplicative factor of $1 \pm \epsilon$.

Theorem 6.1 tells us how to approximate the capacity of an operator. Construction 4.2 gives us an operator whose capacity is the inverse square of the Brascamp-Lieb constant. Hence, to approximate the Brascamp-Lieb constant corresponding to datum (\mathbf{B}, \mathbf{p}) , we need only obtain operator $T_{(\mathbf{B}, \mathbf{p})}$ and use algorithm G to compute an approximation to $\text{cap}(T_{(\mathbf{B}, \mathbf{p})})$. Then, $\text{BL}(\mathbf{B}, \mathbf{p}) = \left(\frac{1}{\text{cap}(T_{(\mathbf{B}, \mathbf{p})})} \right)^{1/2}$. This yields the following algorithm, and theorem (from which Theorem 1.5 in the introduction follows by a simple scaling argument):

Theorem 6.2 (Approximating the Brascamp-Lieb Constant). *Given a Brascamp-Lieb datum (\mathbf{B}, \mathbf{p}) , where $p_i = c_i/d$, with $c_i, d \in \mathbb{N}$, $d \neq 0$, $B_i \in \mathcal{M}_{n_i \times n}(\mathbb{Z})$ and approximation parameter $\epsilon > 0$. Assume that each entry of B_i has absolute value at most M . Then Algorithm \tilde{G} runs in time $\text{poly}(n, d, \log M, 1/\epsilon)$ and computes $\text{BL}(\mathbf{B}, \mathbf{p})$ to a multiplicative approximation factor of $(1 \pm \epsilon)$.*

Proof. If $n \neq \sum_{i=1}^m p_i n_i$ then we know that $\text{BL}(\mathbf{B}, \mathbf{p}) = \infty$. Otherwise we correctly proceed. In step 2, we construct the operator $\tilde{T}_{(\mathbf{B}, \mathbf{p})}$ according to Construction 3.10. By item 1 of Lemma 3.11 we have that $\text{cap}(T_{(\mathbf{B}, \mathbf{p})}) = \text{cap}(\tilde{T}_{(\mathbf{B}, \mathbf{p})})^{1/nd}$. By Lemma 4.4, we have $\text{cap}(T_{(\mathbf{B}, \mathbf{p})}) = 1/\text{BL}(\mathbf{B}, \mathbf{p})^2$. Thus,

$$\text{BL}(\mathbf{B}, \mathbf{p}) = (1/\text{cap}(\tilde{T}_{(\mathbf{B}, \mathbf{p})}))^{1/2nd}.$$

As step 3 returns α in the interval $[(1 - \epsilon) \cdot \text{cap}(\tilde{T}_{(\mathbf{B}, \mathbf{p})}), \text{cap}(\tilde{T}_{(\mathbf{B}, \mathbf{p})})]$ we have that

$$(1/\alpha)^{1/2nd} \in [\text{cap}(\tilde{T}_{(\mathbf{B}, \mathbf{p})})^{1/2nd}, (1 + \epsilon)\text{cap}(\tilde{T}_{(\mathbf{B}, \mathbf{p})})^{1/2nd}] = [\text{BL}(\mathbf{B}, \mathbf{p}), (1 + \epsilon)\text{BL}(\mathbf{B}, \mathbf{p})].$$

Since algorithm \tilde{G} only uses once algorithm G on an operator of dimension $nd \times nd$ with entries bounded by M , the running time of algorithm \tilde{G} is given by $\text{poly}(n, d, \log(M), 1/\epsilon)$. This concludes the analysis of algorithm \tilde{G} . \square

Input: Brascamp-Lieb datum (\mathbf{B}, \mathbf{p}) , where each $p_i = c_i/d$, with $c_i, d \in \mathbb{N}$, $d \neq 0$, $B_i \in \mathcal{M}_{n_i \times n}(\mathbb{Z})$ and approximation parameter $\epsilon > 0$. Assume that each entry of B_i has absolute value at most M .
Output: $\text{BL}(\mathbf{B}, \mathbf{p})$ with multiplicative error of $(1 \pm \epsilon)$.

1. Check if $n = \sum_{i=1}^m p_i n_i$. If yes, go to step 2. If not, return ∞ .
2. Let $T_{(\mathbf{B}, \mathbf{p})}$ be the operator constructed in Construction 4.2. Let $\tilde{T}_{(\mathbf{B}, \mathbf{p})}$ be the operator obtained from $T_{(\mathbf{B}, \mathbf{p})}$ via Construction 3.10. Go to step 3.
3. Apply Algorithm G on $\tilde{T}_{(\mathbf{B}, \mathbf{p})}$, with approximation parameter ϵ . Let α be the output of Algorithm G .
4. Return $(1/\alpha)^{1/2nd}$.

Algorithm 4: Algorithm \tilde{G} computing the BL constant

7 Continuity of the Brascamp-Lieb constant

In this section, we note that the explicit bounds on the continuity of capacity of completely positive operators (at rational points) in [GGOW16] directly translate to bounds on the continuity of Brascamp-Lieb constants (at rational points). We will then discuss how this can be extended to irrational points.

The following theorem is proved in [GGOW16].

Theorem 7.1 (Theorem 4.5 in [GGOW16]). *Suppose A_1, \dots, A_m and B_1, \dots, B_m are two tuples of $n \times n$ matrices such that the bit-complexity of elements of A_i 's is b and $\|A_i - B_i\| \leq \delta$ for all i . Let T_A be the operator defined by A_1, \dots, A_m and T_B be the operator defined by B_1, \dots, B_m . Then there exists a polynomial $P(n, b, \log(m))$ s.t. if $\delta \leq \exp(-P(n, b, \log(m)))$ and $\text{cap}(T_A) > 0$, then $\text{cap}(T_B) > 0$. Furthermore*

$$\left(1 - \frac{P(n, b, \log(m))}{\log(1/\delta)^{1/3}}\right) \leq \frac{\text{cap}(T_A)}{\text{cap}(T_B)} \leq \left(1 + \frac{P(n, b, \log(m))}{\log(1/\delta)^{1/3}}\right)$$

From this we get the following theorem (continuity of BL constant at rational points) as an immediate corollary, due to Construction 4.2.

Theorem 7.2 (Theorem 1.12 restated). *Suppose (\mathbf{B}, \mathbf{p}) and $(\tilde{\mathbf{B}}, \mathbf{p})$ are two tuples of Brascamp-Lieb datum. Let (n, n_1, \dots, n_m) be the dimension vector and $p_1, \dots, p_m = c_1/d, \dots, c_m/d$. Also suppose the entries of B_i 's can be described using b bits and $\|B_i - \tilde{B}_i\| \leq \delta$ for all i . Then there exists a polynomial $P(n, d, b, \log(m))$ s.t. if $\delta \leq \exp(-P(n, d, b, \log(m)))$ and $\text{BL}(\mathbf{B}, \mathbf{p})$ is finite, then $\text{BL}(\tilde{\mathbf{B}}, \mathbf{p})$ is finite. Furthermore*

$$\left(1 - \frac{P(n, d, b, \log(m))}{\log(1/\delta)^{1/3}}\right) \leq \frac{\text{BL}(\mathbf{B}, \mathbf{p})}{\text{BL}(\tilde{\mathbf{B}}, \mathbf{p})} \leq \left(1 + \frac{P(n, d, b, \log(m))}{\log(1/\delta)^{1/3}}\right)$$

Continuity at irrational points We note that Theorem 7.1 can be extended to continuity of capacity at irrational A_1, \dots, A_m (without any explicit bounds). We describe the changes that need to be made to the proof of Theorem 4.5 in [GGOW16].

1. First prove that for every $\eta > 0$, there is an η -approximate fixed point C of the operator T_A defined by A_1, \dots, A_m (if $\text{cap}(T_A) > 0$).
2. Using the fact that $\|A_i - B_i\| \leq \delta$, prove that C is also an η' -approximate fixed point of the operator T_B defined by B_1, \dots, B_m . δ can be chosen to be sufficiently small based on the lowest and highest eigenvalues of C . Now an application of Lemma 3.8 in [GGOW16] can finish the proof.

When \mathbf{p} is not rational, our reduction from BL datum to operators doesn't work anymore. Nonetheless, Algorithm 1 described in the introduction can be applied directly to the BL datum. Since this algorithm is continuous in \mathbf{B} , this should imply continuity of $\text{BL}(\mathbf{B}, \mathbf{p})$ in \mathbf{B} . However, it is not clear to us at the moment how long this algorithm will have to be run before the BL constant gets sufficiently close to 1. This is because the progress per step in Theorem 1.9 depends on the common denominator d of \mathbf{p} , which doesn't make sense in the case when \mathbf{p} is not rational. But perhaps the fact that the vertices of the BL polytope $P_{\mathbf{B}}$ are rational can come in handy.

8 Interesting special cases of BL polytopes

In this section, we explore special cases of BL polytopes that give rise to natural polytopes which are well studied in the combinatorial optimization literature. As mentioned before, Barthe [Bar98] proved that the BL polytopes corresponding to rank-1 BL datum are exactly the basis polytopes.

Now we show that some special cases of BL polytopes corresponding to rank-2 BL datum give rise to linear matroid intersection polytopes. Given a collection of vectors in \mathbb{R}^n , $\mathbf{v} = (v_1, \dots, v_m)$, the linear matroid $\mathcal{M}_{\mathbf{v}}$ associated with this collection (over the ground set $[m]$) is $\{I \subseteq [m] : (v_i)_{i \in I} \text{ are linearly independent}\}$ ¹⁷.

Given two linear matroids $\mathcal{M}_{\mathbf{v}}$ and $\mathcal{M}_{\mathbf{w}}$, their intersection polytope $P_{\mathcal{M}_{\mathbf{v}}, \mathcal{M}_{\mathbf{w}}}$ is the convex hull of the characteristic vectors of their common bases (sometimes this is defined as the convex hull of the common independent sets) i.e.

$$P_{\mathcal{M}_{\mathbf{v}}, \mathcal{M}_{\mathbf{w}}} = \text{conv} \{1_I, I \subseteq [m] : (v_i)_{i \in I}, (w_i)_{i \in I} \text{ both form a basis for } \mathbb{R}^n\}$$

For a linear matroid $\mathcal{M}_{\mathbf{v}}$ and a set $J \subseteq [m]$, we define the subspace

$$V_J = \text{span}((v_j)_{j \in J})$$

It is well known (e.g. see Chapter 41 in [Sch03]) $P_{\mathcal{M}_{\mathbf{v}}, \mathcal{M}_{\mathbf{w}}}$ can be described by the following linear

¹⁷Linear matroids can be defined w.r.t. any field \mathbb{F} but we will restrict ourselves to \mathbb{R} as this is the case which embeds into BL.

program (with exponentially many inequalities!) on a variable vector \mathbf{p} :

$$\begin{aligned} \sum_{j=1}^m p_j &= n \\ \sum_{j \in J} p_j &\leq \dim(V_J) \quad \forall J \subseteq [m] \end{aligned} \tag{3}$$

$$\sum_{j \in J'} p_j \leq \dim(W_{J'}) \quad \forall J' \subseteq [m] \tag{4}$$

$$\mathbf{p} \geq 0$$

Now consider the following BL datum: for every $j \in [m]$, $B_j : \mathbb{R}^{2n} \rightarrow \mathbb{R}^2$ is given by

$$B_j(x, y) = (\langle v_j, x \rangle, \langle w_j, y \rangle)$$

We will show that the BL polytope $P_{\mathbf{B}}$ corresponding to this BL datum is the same as the matroid intersection polytope $P_{\mathcal{M}_v, \mathcal{M}_w}$. Our feasibility algorithm for BL-polytopes thus automatically gives a polynomial time algorithm for the linear matroid intersection problem (here over \mathbb{R}). Of course it was already known that this problem is in P , and moreover Gurvits [Gur04] already noticed that operator scaling efficiently solves this problem; our hope is that such encodings are possible for other optimization problems with exponential size linear programs.

Lemma 8.1. *The BL polytope $P_{\mathbf{B}}$ corresponding to the above BL datum is the same as $P_{\mathcal{M}_v, \mathcal{M}_w}$.*

Proof. Recall that by Theorem 1.1, a point $\mathbf{p} \in P_{\mathbf{B}}$ is given by the following constraints:

$$\begin{aligned} \sum_{j=1}^m p_j \cdot 2 &= 2n \\ \sum_{j=1}^m p_j \dim(B_j(U)) &\geq \dim(U) \quad \forall \text{ subspaces } U \text{ of } \mathbb{R}^{2n} \\ \mathbf{p} &\geq 0 \end{aligned} \tag{5}$$

Given a subspace $U \subseteq \mathbb{R}^{2n}$, let $U_1, U_2 \subseteq \mathbb{R}^n$ denote the projections of U onto the first and second halves of \mathbb{R}^{2n} . More formally,

$$\begin{aligned} U_1 &= \{u_1 \in \mathbb{R}^n : (u_1, u_2) \in U \text{ for some } u_2 \in \mathbb{R}^n\} \\ U_2 &= \{u_2 \in \mathbb{R}^n : (u_1, u_2) \in U \text{ for some } u_1 \in \mathbb{R}^n\} \end{aligned}$$

Also let U_1^\perp, U_2^\perp denote the orthogonal complements of U_1 and U_2 respectively (inside \mathbb{R}^n). It is not hard to see that

$$\dim(B_j(U)) = 2 - 1\{v_j \in U_1^\perp\} - 1\{w_j \in U_2^\perp\}$$

Then inequality (5) can be written as:

$$\sum_{j=1}^m p_j \cdot \left(2 - 1\{v_j \in U_1^\perp\} - 1\{w_j \in U_2^\perp\}\right) \geq \dim(U)$$

This inequality is the strongest when $U = U_1 \oplus U_2$. So we only need to worry about the inequalities of the following form:

$$\sum_{j=1}^m p_j \cdot \left(2 - 1\{v_j \in U_1^\perp\} - 1\{w_j \in U_2^\perp\} \right) \geq \dim(U_1) + \dim(U_2)$$

Here U_1 and U_2 arbitrary subspaces of \mathbb{R}^n . Since $\sum_{j=1}^m p_j = n$ for all $\mathbf{p} \in P_{\mathbf{B}}$, the above inequality can be rearranged to give the following:

$$\sum_{j=1}^m p_j \cdot 1\{v_j \in U_1^\perp\} + \sum_{j=1}^m p_j \cdot 1\{w_j \in U_2^\perp\} \leq \dim(U_1^\perp) + \dim(U_2^\perp)$$

By taking $U_1^\perp = \{\mathbf{0}\}$ or $U_2^\perp = \{\mathbf{0}\}$, the above set of inequalities are equivalent to the following:

$$\sum_{j=1}^m p_j \cdot 1\{v_j \in Z\} \leq \dim(Z) \quad \forall \text{ subspaces } Z \text{ of } \mathbb{R}^n \quad (6)$$

$$\sum_{j=1}^m p_j \cdot 1\{w_j \in Z'\} \leq \dim(Z') \quad \forall \text{ subspaces } Z' \text{ of } \mathbb{R}^n \quad (7)$$

For subsets $J, J' \subseteq [m]$, setting $Z = V_J$ and $Z' = W_{J'}$, we see that the inequalities (6) and (7) imply the inequalities (3) and (4). In the other direction, for subspaces $Z, Z' \subseteq \mathbb{R}^n$, setting $J = \{j \in [m] : v_j \in Z\}$ and $J' = \{j \in [m] : w_j \in Z'\}$ recovers the inequalities (6) and (7) from the inequalities (3) and (4). This completes the proof. \square

Since the perfect bipartite matching polytopes are a special case of linear matroid intersection polytopes, we see that the bipartite matching polytopes are a special case of BL polytopes for rank-2 BL datum. As mentioned before, it is an interesting question if the general perfect matching polytopes are a special case of BL polytopes or not.

9 Additional properties of the Brascamp-Lieb constant

In this section, we prove that under appropriate normalization conditions on the BL datum, the BL constant is at least 1, with equality iff the datum is geometric. This is a strengthening of the following theorem proved in [Val11]. We will assume that exponent vectors \mathbf{p} satisfy $p_i > 0$ for all i .

Theorem 9.1 (Theorem 1,5 of [Val11]). *Let (\mathbf{B}, \mathbf{p}) be a projection-normalized and feasible BL datum. Then $\text{BL}(\mathbf{B}, \mathbf{p}) \geq 1$ and equality holds iff the datum is geometric.*

We prove the following strengthening of the above theorem.

Theorem 9.2. *Let (\mathbf{B}, \mathbf{p}) be a feasible BL datum that satisfies*

$$\sum_{i=1}^m p_i \text{tr} \left[B_i^\dagger B_i \right] = n$$

Then $\text{BL}(\mathbf{B}, \mathbf{p}) \geq 1$ and equality holds iff the datum is geometric.

It is easy to see that we get the following corollary from the above theorem.

Corollary 9.3. *Let (\mathbf{B}, \mathbf{p}) be a feasible BL datum that is either projection-normalized or isotropy normalized. Then $\text{BL}(\mathbf{B}, \mathbf{p}) \geq 1$ and equality holds iff the datum is geometric.*

Proof. We will apply Theorem 9.2. The only thing that needs to be verified is that when (\mathbf{B}, \mathbf{p}) is projection-normalized or isotropy normalized, then

$$\sum_{i=1}^m p_i \text{tr} [B_i^\dagger B_i] = n$$

When (\mathbf{B}, \mathbf{p}) is projection-normalized, then $\text{tr} [B_i^\dagger B_i] = \text{tr} [B_i B_i^\dagger] = \text{tr}[I_{n_i}] = n_i$ and hence

$$\sum_{i=1}^m p_i \text{tr} [B_i^\dagger B_i] = \sum_{i=1}^m p_i n_i = n$$

Note that $\sum_{i=1}^m p_i n_i = n$ is satisfied since (\mathbf{B}, \mathbf{p}) is feasible. When (\mathbf{B}, \mathbf{p}) is isotropy-normalized, then

$$\sum_{i=1}^m p_i \text{tr} [B_i^\dagger B_i] = \text{tr} \left[\sum_{i=1}^m p_i B_i^\dagger B_i \right] = \text{tr}[I_n] = n$$

This completes the proof. □

For rational exponent vectors \mathbf{p} , Theorem 9.2 would follow from the reduction to operator scaling and invoking appropriate theorems from [GGOW16]. However we don't assume rationality of \mathbf{p} in this section. It turns out that one can imitate the proofs of corresponding statements in the operator scaling world¹⁸ and the proofs work even in the irrational case. For the proof, we would need the following easy proposition, which is a consequence of the AM-GM inequality applied to the eigenvalues of A .

Proposition 9.4. *Let $A \in \mathcal{M}_n(\mathbb{C})$ be an $n \times n$ Hermitian positive semidefinite matrix s.t. $\text{tr}[A] = n$. Then $\text{Det}(A) \leq 1$ and equality holds iff $A = I_n$.*

Now we are ready to prove our theorem.

Proof. (Of Theorem 9.2) Recall that

$$1/\text{BL}(\mathbf{B}, \mathbf{p})^2 = \inf \frac{\text{Det} \left(\sum_{i=1}^m p_i B_i^\dagger X_i B_i \right)}{\prod_{i=1}^m \text{Det}(X_i)^{p_i}}$$

where the infimum is taken over all choices of positive definite matrices X_i in dimension n_i . By plugging in $X_i = I_{n_i}$, we can see that

$$1/\text{BL}(\mathbf{B}, \mathbf{p})^2 \leq \text{Det} \left(\sum_{i=1}^m p_i B_i^\dagger B_i \right) \leq 1$$

¹⁸Let T be a positive operator acting on $n \times n$ matrices s.t. $\text{tr}[T(I)] = n$. Then $\text{cap}(T) \leq 1$ and equality holds iff T is doubly stochastic.

Here the second inequality follows from Proposition 9.4 and the fact that $\text{tr} \left[\sum_{i=1}^m p_i B_i^\dagger B_i \right] = n$, which follows from the assumption in the theorem statement. We also get from Proposition 9.4 that equality holds only if $\sum_{i=1}^m p_i B_i^\dagger B_i = I_n$ i.e. (\mathbf{B}, \mathbf{p}) is isotropy-normalized.

We also know that for a geometric datum the BL constant is 1. So the only thing that is left is to show that if the BL constant is 1, then (\mathbf{B}, \mathbf{p}) is projection-normalized as well. For this, we plug in $X_i = (B_i B_i^\dagger)^{-1}$ (note that $B_i B_i^\dagger$ is invertible since (\mathbf{B}, \mathbf{p}) is feasible) to get:

$$1/\text{BL}(\mathbf{B}, \mathbf{p})^2 \leq \text{Det} \left(\sum_{i=1}^m p_i B_i^\dagger (B_i B_i^\dagger)^{-1} B_i \right) \cdot \prod_{i=1}^m \text{Det} (B_i B_i^\dagger)^{p_i}$$

First note that

$$\begin{aligned} \text{tr} \left[\sum_{i=1}^m p_i B_i^\dagger (B_i B_i^\dagger)^{-1} B_i \right] &= \sum_{i=1}^m p_i \text{tr} \left[B_i^\dagger (B_i B_i^\dagger)^{-1} B_i \right] \\ &= \sum_{i=1}^m p_i \text{tr} \left[(B_i B_i^\dagger)^{-1} B_i B_i^\dagger \right] \\ &= \sum_{i=1}^m p_i \text{tr} [I_{n_i}] = \sum_{i=1}^m p_i n_i = n \end{aligned}$$

Hence, by Proposition 9.4, we get that

$$\text{Det} \left(\sum_{i=1}^m p_i B_i^\dagger (B_i B_i^\dagger)^{-1} B_i \right) \leq 1$$

and thus

$$1/\text{BL}(\mathbf{B}, \mathbf{p})^2 \leq \prod_{i=1}^m \text{Det} (B_i B_i^\dagger)^{p_i}$$

Let $\lambda_{i,1}, \dots, \lambda_{i,n_i}$ be the eigenvalues of $B_i B_i^\dagger$. Then

$$\begin{aligned} \log \left(\prod_{i=1}^m \text{Det} (B_i B_i^\dagger)^{p_i} \right) &= \sum_{i=1}^m p_i \sum_{j=1}^{n_i} \log (\lambda_{i,j}) \\ &= \sum_{i=1}^m p_i n_i \cdot \frac{1}{n_i} \sum_{j=1}^{n_i} \log (\lambda_{i,j}) \\ &\leq \sum_{i=1}^m p_i n_i \log \left(\frac{1}{n_i} \sum_{j=1}^{n_i} \lambda_{i,j} \right) \\ &= n \cdot \sum_{i=1}^m \frac{p_i n_i}{n} \log \left(\text{tr} [B_i B_i^\dagger] / n_i \right) \\ &\leq n \cdot \log \left(\sum_{i=1}^m \frac{p_i}{n} \cdot \text{tr} [B_i B_i^\dagger] \right) \\ &= 0 \end{aligned}$$

The two inequalities follow from concavity of log. The last equality follows from the assumption that

$$\sum_{i=1}^m p_i \operatorname{tr} [B_i^\dagger B_i] = n$$

In the first inequality, equality holds iff for all i , $\lambda_{i,j}$'s are equal for all j i.e. for all i , $B_i B_i^\dagger = c_i I_{n_i}$ for some $c_i > 0$. In the second inequality, equality holds iff $\operatorname{tr} [B_i B_i^\dagger] / n_i$ is equal for all i i.e. $c_i = c$ for some $c > 0$. By the normalization condition

$$\sum_{i=1}^m p_i \operatorname{tr} [B_i^\dagger B_i] = n$$

we get that $c = 1$ and hence $B_i B_i^\dagger = I_{n_i}$ for all i i.e. (\mathbf{B}, \mathbf{p}) is isotropy normalized. This completes the proof. \square

10 Iteration analysis for Algorithm 1

In this section, we give an analysis for the number of iterations required in Algorithm 1 to get close to the geometric position. While for the rational exponent case, such an analysis follows from the reduction to operators, here we give an analysis that works even in the irrational case. The proof closely mimicks the argument of [GGOW16] in the operator world.

Let us first define the analogue of the ds measure for BL datum which measures how close the BL datum is to geometric position.

Definition 10.1 (Distance to geometric position).

$$g(\mathbf{B}, \mathbf{p}) = \operatorname{tr} \left[\left(\sum_{i=1}^m p_i B_i^\dagger B_i - I_n \right)^2 \right] + \sum_{i=1}^n \operatorname{tr} \left[\left(B_i B_i^\dagger - I_{n_i} \right)^2 \right]$$

Note that (\mathbf{B}, \mathbf{p}) is geometric iff $g(\mathbf{B}, \mathbf{p}) = 0$. We will need the following lemma which is a robust version of the AM-GM inequality and is a consequence of Lemma 3.10 in [LSW98].

Lemma 10.2. *Let $A \in \mathcal{M}_n(\mathbb{C})$ be a Hermitian positive semidefinite matrix s.t. $\operatorname{tr}[A] = n$ and $\operatorname{tr} [(A - I_n)^2] \geq \epsilon$, where $\epsilon \leq 1$. Then $\operatorname{Det}(A) \leq \exp(-\epsilon/6)$.*

We are now ready to state the bound on the number of iterations.

Theorem 10.3. *Let (\mathbf{B}, \mathbf{p}) is a feasible datum that is isotropy or projection normalized s.t. $B_i \in \mathcal{M}_{n_i \times n}(\mathbb{Q})$ and each entry of B_i has bit-size atmost b . Then for any $\epsilon \leq 1$ and $t = O\left(\frac{nm \log(nm) + nb}{\epsilon}\right)$, the sequence $\mathbf{B}^0 = \mathbf{B}, \mathbf{B}^1, \dots, \mathbf{B}^t$ in Algorithm 1 satisfies $\min_{i \in \{0, 1, \dots, t\}} g(\mathbf{B}^i, \mathbf{p}) \leq \epsilon$.*

We remark that the assumption on the datum being normalized is without loss of generality as it can be achieved in one extra iteration. The proof will be along the lines described in the introduction and will proceed by analyzing the effect of normalization operations on the BL constant.

Proof. By Corollaries 5.4 and 5.5, we have that $\text{BL}(\mathbf{B}, \mathbf{p}) \leq \exp(12nm \log(nm) + nb)$. Assume on the contrary that $\min_{i \in \{0, 1, \dots, t\}} g(\mathbf{B}^i, \mathbf{p}) > \epsilon$. We will prove that an isotropy normalization step decreases the BL constant by a multiplicative factor of $\exp(-\epsilon/12)$ and a projection normalization step decreases the BL constant¹⁹. We also know by Corollary 9.3 that the BL constant is at least 1 in the isotropy or projection normalized case. This will complete the contradiction and the proof.

Let us first analyze the isotropy normalization step. Let $(\mathbf{B}', \mathbf{p})$ be the datum before the normalization and $(\mathbf{B}'', \mathbf{p})$ after. Since $(\mathbf{B}', \mathbf{p})$ is projection normalized, we have that

$$\text{tr} \left[\left(\sum_{i=1}^m p_i B_i'^{\dagger} B_i' - I_n \right)^2 \right] = g(\mathbf{B}', \mathbf{p}) > \epsilon$$

We also know that

$$\text{BL}(\mathbf{B}'', \mathbf{p}) = \text{Det} \left(\sum_{i=1}^m p_i B_i'^{\dagger} B_i' \right)^{1/2} \cdot \text{BL}(\mathbf{B}', \mathbf{p})$$

Hence we can use Lemma 10.2 to conclude that the BL constant decreases by a multiplicative factor of $\exp(-\epsilon/12)$.

Now let us look at the projection normalization step. Let $(\mathbf{B}', \mathbf{p})$ be the datum before the normalization and $(\mathbf{B}'', \mathbf{p})$ after. We know that

$$\text{BL}(\mathbf{B}'', \mathbf{p}) = \prod_{i=1}^m \text{Det} \left(B_i' B_i'^{\dagger} \right)^{p_i/2} \cdot \text{BL}(\mathbf{B}', \mathbf{p})$$

Since $(\mathbf{B}', \mathbf{p})$ is isotropy normalized, we can conclude that $\prod_{i=1}^m \text{Det} \left(B_i' B_i'^{\dagger} \right)^{p_i/2} \leq 1$ by the same log-concavity argument as in the proof of Theorem 9.2. This concludes the proof. \square

11 Conclusion and open questions

In this paper we address some of the central algorithmic questions regarding Brascamp-Lieb inequalities. In particular, we gave polynomial time algorithms to test membership in the BL polytopes, compute a violating inequality (separation oracle) and approximate the BL constants. These algorithms were obtained by a simple reduction from the BL data to completely positive operators and then using recent polynomial time algorithms that test whether such operators are rank decreasing and compute the capacity of such operators [GGOW16, IQS15]. The reduction itself provides in some cases different proofs of some of the known structural results regarding BL inequalities, in particular the characterization of feasibility.

On the algorithmic side, our work leaves some concrete open questions.

1. While our reduction is polynomial time in the binary description length of the matrices (B_1, \dots, B_m) in the BL datum, it is not polynomial time in the binary description length of the exponent vector p_1, \dots, p_m . When $p_1, \dots, p_m = \frac{c_1}{d}, \dots, \frac{c_m}{d}$, the running times of our algorithms scale as $\text{poly}(d)$. Can one improve them to get (the optimal) $\text{poly}(\log(d))$?
2. Our algorithm for approximating the BL constants up to an $(1+\epsilon)$ factor run in time $\text{poly}(1/\epsilon)$. Can one improve it to a running time that scales as $\text{poly}(\log(1/\epsilon))$?

¹⁹It is also possible to get a quantitative bound on the decrease in the projection step but it is not needed.

3. We provide a weak separation oracle for the BL polytope. Sufficiently strong separation oracles for convex programs yield optimization algorithms via the general reduction of [GLS88]. However, our current separation oracle does not seem strong enough. Can one efficiently optimize linear functions over BL polytopes?

On the structural side, our reduction may provide a higher level view of BL inequalities, simply since BL data is just a proper subset of completely positive operators, and viewing them as such may lead to better understanding. One intriguing question is whether this viewpoint leads to an even larger family of inequalities. It is obvious that for any doubly stochastic operator T , the fact that $\text{cap}(T) = 1$ implies the following determinantal inequality (that holds for all psd matrices X),

$$\text{Det} \left(\sum_{i=1}^m A_i X A_i^\dagger \right) = \text{Det}(T(X)) \geq \text{Det}(X)$$

which can be translated to an inequality on Gaussian densities. One can wonder for which operators T the Gaussian densities can be replaced by arbitrary non-negative densities, as in BL-inequalities. There are simple counterexamples showing this does not hold in full generality. On the other hand, our reduction and the connection to quiver representation (Section 4.1) tells us that for operators T arising from the star quiver, these inequalities are in 1-1 correspondence to the BL-inequalities. It would be interesting to better understand what is special about star quivers in this context. In particular, do other quivers lead to new inequalities?

Acknowledgments

We would like to thank Zeev Dvir for pointing us to the paper [BCCT08a]. We would also like to thank Jonathan Bennett, Larry Guth and anonymous reviewers for helpful comments. We thank Rohit Gurjar for suggesting that BL polytopes can encode the linear matroid intersection polytopes.

References

- [Bal89] Keith Ball. Volumes of sections of cubes and related problems. *Geometric Aspects of Functional Analysis*, pages 251–260, 1989.
- [Bar98] Franck Barthe. On a reverse form of the Brascamp–Lieb inequality. *Inventiones Mathematicae*, 134:335–361, 1998.
- [BB10] Jonathan Bennett and Neal Bez. Some nonlinear Brascamp–Lieb inequalities and applications to harmonic analysis. *Journal of Functional Analysis*, 259(10):2520–2556, 2010.
- [BBCF17] Jonathan Bennett, Neal Bez, Michael Cowling, and Taryn Flock. Behaviour of the Brascamp–Lieb constant. *Bulletin of the London Mathematical Society*, 49(3):512–518, 2017.
- [BBFL17] Jonathan Bennett, Neal Bez, Taryn Flock, and Sanghyuk Lee. Stability of the Brascamp–Lieb constant and applications. *American Journal of Mathematics*, 2017.

- [BCCT08a] Jonathan Bennett, Anthony Carbery, Michael Christ, and Terence Tao. The Brascamp-Lieb inequalities: finiteness, structure, and extremals. *Geometric and Functional Analysis*, 17(5):1343–1415, 2008.
- [BCCT08b] Jonathan Bennett, Anthony Carbery, Michael Christ, and Terence Tao. Finite bounds for Holder-Brascamp-Lieb multilinear inequalities. *Mathematical Research Letters*, 2008.
- [BD16] Jean Bourgain and Ciprian Demeter. Mean value estimates for Weyl sums in two dimensions. *Journal of the London Mathematical Society*, 94(3):814–838, 2016.
- [BDG16] Jean Bourgain, Ciprian Demeter, and Larry Guth. Proof of the main conjecture in Vinogradov’s mean value theorem for degrees higher than three. *Annals of Mathematics*, 184(2):633–682, 2016.
- [BDWY12] Boaz Barak, Zeev Dvir, Avi Wigderson, and Amir Yehudayoff. Fractional Sylvester-Gallai theorems. *Proceedings of the National Academy of Sciences*, 110(48):19213–19219, 2012.
- [BKS14] Boaz Barak, Jonathan A Kelner, and David Steurer. Rounding sum-of-squares relaxations. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 31–40. ACM, 2014.
- [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 Symposium on Theory of Computing*, pages 143–151. ACM, 2015.
- [BL76] Herm Brascamp and Elliot Lieb. Best constants in Young’s inequality, its converse and its generalization to more than three functions. *Advances in Mathematics*, 20:151–172, 1976.
- [CCE09] Eric Carlen and Dario Cordero-Erausquin. Subadditivity of the entropy and its relation to Brascamp-Lieb type inequalities. *Geometric and Functional Analysis*, 19:373–405, 2009.
- [CDK⁺13] Michael Christ, James Demmel, Nicholas Knight, Thomas Scanlon, and Katherine Yelick. Communication lower bounds and optimal algorithms for programs that reference arrays - Part 1. arXiv preprint arXiv:1308.0068, 2013.
- [CDK⁺15] Michael Christ, James Demmel, Nicholas Knight, Thomas Scanlon, and Katherine Yelick. On Holder-Brascamp-Lieb inequalities for torsion-free discrete Abelian groups. arXiv preprint arXiv:1510.04190, 2015.
- [Cho75] Man-Duen Choi. Completely positive linear maps on complex matrices. *Linear Algebra and Its Applications*, pages 285–290, 1975.
- [CT84] Imre Csiszár and Gábor Tusnády. Information geometry and alternating minimization procedures. *Stat. Decis. Suppl.*, 1:205–237, 1984.

- [DGOS16] Zeev Dvir, Ankit Garg, Rafael Oliveira, and József Solymosi. Rank bounds for design matrices with block entries and geometric applications. *arXiv preprint arXiv:1610.08923*, 2016.
- [DM15] Harm Derksen and Visu Makam. Polynomial degree bounds for matrix semi-invariants. arXiv preprint arXiv:1512.03393, 2015.
- [DSW14a] Zeev Dvir, Shubhangi Saraf, and Avi Wigderson. Breaking the quadratic barrier for 3-LCC’s over the reals. *STOC*, 2014.
- [DSW14b] Zeev Dvir, Shubhangi Saraf, and Avi Wigderson. Improved rank bounds for design matrices and a new proof of KELLY’s theorem. *Forum of Mathematics, Sigma*, 2:24, 2014.
- [DW00] Harm Derksen and Jerzy Weyman. Semi-invariants of quivers and saturation for littlewood-richardson coefficients. *Journal of the American Mathematical Society*, 13(3):467–479, 2000.
- [DW07] Harm Derksen and Jerzy Weyman. The combinatorics of quiver representations. arXiv preprint arXiv:0608288, 2007.
- [DZ01] Mátyás Domokos and A. N. Zubkov. Semi-invariants of quivers as determinants. *Transformation Groups*, 6(1):9–24, 2001.
- [Edm65] Jack Edmonds. Maximum matching and a polyhedron with 0,1-vertices. *Journal of research of the National Bureau of Standards*, 69B:125–130, 1965.
- [For02] Jürgen Forster. A linear lower bound on the unbounded error probabilistic communication complexity. *Journal of Computer and System Sciences*, 65(4):612–625, 2002.
- [GGOW16] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. A deterministic polynomial time algorithm for non-commutative rational identity testing. *FOCS*, arXiv:1511.03730, 2016.
- [GLS88] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer, 2 edition, 1988.
- [GS00] Leonid Gurvits and Alex Samorodnitsky. A deterministic polynomial-time algorithm for approximating mixed discriminant and mixed volume. *STOC*, pages 48–57, 2000.
- [Gur04] Leonid Gurvits. Classical complexity and quantum entanglement. *Journal of Computer and System Sciences*, 69(3):448–484, 2004.
- [Har14] Moritz Hardt. Understanding alternating minimization for matrix completion. *FOCS*, pages 651–660, 2014.
- [HM13] Moritz Hardt and Ankur Moitra. Algorithms and hardness for robust subspace recovery. *JMLR: Workshop and Conference Proceedings*, 30:1–22, 2013.

- [HSS15] Samuel B Hopkins, Tselil Schramm, Jonathan Shi, and David Steurer. Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors. *arXiv preprint arXiv:1512.02337*, 2015.
- [IQS15] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Constructive noncommutative rank computation in deterministic polynomial time over fields of arbitrary characteristics. *arXiv preprint arXiv:1512.03531*, December 2015.
- [JNS13] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank matrix completion using alternating minimization. *STOC*, pages 665–674, 2013.
- [LCCV16] Jingbo Liu, Thomas Courtarde, Paul Cuff, and Sergio Verdu. Smoothing Brascamp-Lieb inequalities and strong converses for common randomness generation. *ISIT*, pages 1043–1047, 2016.
- [LCCV17] Jingbo Liu, Thomas Courtarde, Paul Cuff, and Sergio Verdu. Information-theoretic perspectives on Brascamp-Lieb inequality and its reverse. *arXiv preprint arXiv:1702.06260*, 2017.
- [LH64] Carlton Lemke and J. T. Howson. Equilibrium points of bimatrix games. *SIAM Journal on Applied Mathematics*, 12(2):413–423, 1964.
- [Lie90] Elliot Lieb. Gaussian kernels have only Gaussian maximizers. *Inventiones Mathematicae*, 102:179–208, 1990.
- [LS17] Zhixian Lei and Yueqi Sheng. Sum of square proof for Brascamp-Lieb type inequality. *arXiv preprint arXiv:1710.01458v2*, 2017.
- [LSW98] Nati Linial, Alex Samorodnitsky, and Avi Wigderson. A deterministic strongly polynomial algorithm for matrix scaling and approximate permanents. *STOC*, pages 644–652, 1998.
- [MS06] Ketan Mulmuley and Milind Sohoni. Geometric complexity theory I: An approach to the P. vs. NP and related problems. *SIAM J. Comput*, 31(2):496–526, 2006.
- [MSS16] Tengyu Ma, Jonathan Shi, and David Steurer. Polynomial-time tensor decompositions with sum-of-squares. *arXiv preprint arXiv:1610.01980*, 2016.
- [Mul12] Ketan Mulmuley. Geometric Complexity Theory V: Equivalence between blackbox derandomization of polynomial identity testing and derandomization of Noether’s normalization lemma. *FOCS*, pages 629–638, 2012.
- [NS16] Aleksandar Nikolov and Mohit Singh. Maximizing determinants under partition constraints. *STOC*, pages 192–201, 2016.
- [Sch03] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency, Volume B*, volume 24 of *Algorithms and Combinatorics*. Springer, 1 edition, 2003.
- [SdB01] Aidan Schofield and Michel Van den Bergh. Semi-invariants of quivers for arbitrary dimension vectors. *Indagationes Mathematicae*, 12(1):125–138, 2001.

- [Sin64] R. Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The Annals of Mathematical Statistics*, 35:876–879, 1964.
- [Val10] Stefán Valdimarsson. The Brascamp-Lieb polyhedron. *Canadian Journal of Mathematics*, 62:870–888, 2010.
- [Val11] Stefán Valdimarsson. Geometric Brascamp-Lieb has the optimal best constant. *Journal of Geometric Analysis*, 21(4):1036–1043, 2011.
- [ZP11] Stefanos Zafeiriou and Maria Petrou. Nonnegative tensor factorization as an alternative Csiszar–Tusnady procedure: algorithms, convergence, probabilistic interpretations and novel probabilistic tensor latent variable analysis algorithms. *Data Mining and Knowledge Discovery*, 22(3):419–466, 2011.