

Sharp kernel clustering algorithms and their associated Grothendieck inequalities

Subhash Khot

Assaf Naor

Abstract

In the kernel clustering problem we are given a (large) $n \times n$ symmetric positive semidefinite matrix $A = (a_{ij})$ with $\sum_{i=1}^n \sum_{j=1}^n a_{ij} = 0$ and a (small) $k \times k$ symmetric positive semidefinite matrix $B = (b_{ij})$. The goal is to find a partition $\{S_1, \dots, S_k\}$ of $\{1, \dots, n\}$ which maximizes $\sum_{i=1}^k \sum_{j=1}^k \left(\sum_{(p,q) \in S_i \times S_j} a_{pq} \right) b_{ij}$. We design a polynomial time approximation algorithm that achieves an approximation ratio of $\frac{R(B)^2}{C(B)}$, where $R(B)$ and $C(B)$ are geometric parameters that depend only on the matrix B , defined as follows: if $b_{ij} = \langle v_i, v_j \rangle$ is the Gram matrix representation of B for some $v_1, \dots, v_k \in \mathbb{R}^k$ then $R(B)$ is the minimum radius of a Euclidean ball containing the points $\{v_1, \dots, v_k\}$. The parameter $C(B)$ is defined as the maximum over all measurable partitions $\{A_1, \dots, A_k\}$ of \mathbb{R}^{k-1} of the quantity $\sum_{i=1}^k \sum_{j=1}^k b_{ij} \langle z_i, z_j \rangle$, where for $i \in \{1, \dots, k\}$ the vector $z_i \in \mathbb{R}^{k-1}$ is the Gaussian moment of A_i , i.e., $z_i = \frac{1}{(2\pi)^{(k-1)/2}} \int_{A_i} x e^{-\|x\|_2^2/2} dx$. We also show that for every $\varepsilon > 0$, achieving an approximation guarantee of $(1 - \varepsilon) \frac{R(B)^2}{C(B)}$ is Unique Games hard.

1 Introduction

Kernel Clustering [13] is a combinatorial optimization problem which originates in the theory of machine learning. It is a general framework for clustering massive statistical data so as to uncover a certain hypothesized structure. The problem is defined as follows: let $A = (a_{ij})$ be an $n \times n$ symmetric positive semidefinite matrix which is usually normalized to be centered, i.e., $\sum_{i=1}^n \sum_{j=1}^n a_{ij} = 0$. The matrix A is often thought of as the correlation matrix of random variables (X_1, \dots, X_n) that measure attributes of certain empirical data, i.e., $a_{ij} = \mathbb{E}[X_i X_j]$. We are also given another symmetric positive semidefinite $k \times k$ matrix $B = (b_{ij})$ which functions as a hypothesis, or test matrix. Think of n as huge and k as small. The goal is to cluster A so as to obtain a smaller matrix which most resembles B . Formally, we wish to find a partition $\{S_1, \dots, S_k\}$ of $\{1, \dots, n\}$ so that if we write $c_{ij} := \sum_{(p,q) \in S_i \times S_j} a_{pq}$, i.e., we form a $k \times k$ matrix $C = (c_{ij})$ by clustering A according to the given partition, then the resulting clustered version of A has the maximum correlation $\sum_{i=1}^k \sum_{j=1}^k c_{ij} b_{ij}$ with the hypothesis matrix B . Equivalently, the goal is to evaluate the number:

$$\mathbf{Clust}(A|B) := \max_{\sigma: \{1, \dots, n\} \rightarrow \{1, \dots, k\}} \sum_{i=1}^k \sum_{j=1}^k a_{ij} b_{\sigma(i)\sigma(j)}. \quad (1)$$

The strength of this generic clustering framework is based in part on the flexibility of adapting the matrix B to the problem at hand. Various particular choices of B lead to well studied optimization problems, while other specialized choices of B are based on statistical hypotheses which have been applied with some empirical success. We refer to [13, 7] for additional background and a discussion of specific examples.

In [7] we investigated the computational complexity of the kernel clustering problem. Answering a question posed in [13], we showed that this problem has a constant factor polynomial time approximation

algorithm. We refer to [7] for more information on the best known approximation guarantees. We also obtained hardness results for kernel clustering under various complexity assumptions. For example, we showed in [7] that when $B = I_3$ is the 3×3 identity matrix then a $\frac{16\pi}{27}$ approximation guarantee for $\mathbf{Clust}(A|I_3)$ is achievable, while any approximation guarantee smaller than $\frac{16\pi}{27}$ is Unique Games hard. We will discuss the Unique Games Conjecture (UGC) presently. At this point it suffices to say that the above statement is evidence that the hardness threshold of the problem of approximating $\mathbf{Clust}(A|I_3)$ is $\frac{16\pi}{27}$, or more modestly that obtaining a polynomial time algorithm which approximates $\mathbf{Clust}(A|I_3)$ up to a factor smaller than $\frac{16\pi}{27}$ would require a major breakthrough.

Another result proved in [7] is that when $k \geq 3$ and B is either the $k \times k$ identity matrix or is spherical (i.e., $b_{ii} = 1$ for all $i \in \{1, \dots, k\}$) and centered (i.e., $\sum_{i=1}^k \sum_{j=1}^k b_{ij} = 0$) then there is a polynomial time approximation algorithm which, given A , approximates $\mathbf{Clust}(A|B)$ to within a factor of $\frac{8\pi}{9} \left(1 - \frac{1}{k}\right)$. We also presented in [7] a conjecture (called the Propeller Conjecture) which we proved would imply that $\frac{8\pi}{9} \left(1 - \frac{1}{k}\right)$ is the UGC hardness threshold when $B = I_k$. We refer to [7] for more information on the Propeller Conjecture, which at present remains open.

The above quoted result from [7] settles the problem of evaluating the UGC hardness threshold of the following type of algorithmic task: given A and an hypothesis matrix B which is guaranteed to belong to a certain class of matrices (in our case centered and spherical), approximate efficiently the number $\mathbf{Clust}(A|B)$. Naturally this can be refined to a family of optimization problems which depend on a fixed B : for each B , what is the UGC hardness threshold of the problem of, given A , approximating $\mathbf{Clust}(A|B)$? In [7] we answered this question only when $B = I_3$, and for $B = I_k$ assuming the Propeller Conjecture, and asked about the case of general B (we did give some B -dependent bounds in [7], but they were not sharp for $B \neq I_k$ for reasons that will become clear presently). This is a natural question since it makes sense to use the best possible polynomial time algorithm if we know B in advance.

Here we answer the above question in full generality. To explain our results we need to define two geometric parameters which are associated to B . Since B is symmetric and positive semidefinite we can find vectors $v_1, \dots, v_k \in \mathbb{R}^k$ such that B is their Gram matrix, i.e., $b_{ij} = \langle v_i, v_j \rangle$ for all $i, j \in \{1, \dots, k\}$. Let $R(B)$ be the smallest possible radius of a Euclidean ball in \mathbb{R}^k which contains $\{v_1, \dots, v_k\}$ and let $w(B)$ be the center of this ball. Let $C(B)$ be the maximum over all partitions $\{A_1, \dots, A_k\}$ of \mathbb{R}^{k-1} into measurable sets of the quantity $\sum_{i=1}^k \sum_{j=1}^k b_{ij} \langle z_i, z_j \rangle$, where for $i \in \{1, \dots, k\}$ the vector $z_i \in \mathbb{R}^{k-1}$ is the Gaussian moment of A_i , i.e., $z_i = \frac{1}{(2\pi)^{(k-1)/2}} \int_{A_i} x e^{-\|x\|_2^2/2} dx$ (this maximum exists, as shown in Section 2). Our main result is the following theorem¹:

Theorem 1.1. *For every symmetric positive semidefinite $k \times k$ matrix B there exists a randomized polynomial time algorithm which given an $n \times n$ symmetric positive semidefinite centered matrix A , outputs a number $\text{Alg}(A)$ such that*

$$\mathbf{Clust}(A|B) \leq \mathbb{E} [\text{Alg}(A)] \leq \frac{R(B)^2}{C(B)} \mathbf{Clust}(A|B).$$

On the other hand, assuming the Unique Games Conjecture, no polynomial time algorithm approximates $\mathbf{Clust}(A|B)$ to within a factor strictly smaller than $\frac{R(B)^2}{C(B)}$.

As an example of Theorem 1.1 for a particular hypothesis matrix consider the following perturbation of

¹We refer to the discussion in Question 1 in Section 1.1 below which addresses the issue of computing efficiently good approximate clusterings rather than approximating only the value $\mathbf{Clust}(A|B)$.

the previously studied case $B = I_3$:

$$B_c := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & c \end{pmatrix},$$

where $c > 0$ is a parameter. The problem of approximating efficiently $\mathbf{Clust}(A|B_c)$ corresponds to partitioning the rows of A into 3 sets $S_1, S_2, S_3 \subseteq \{1, \dots, n\}$ and maximizing the sum of the total masses of A on $S_1 \times S_1, S_3 \times S_2, S_3 \times S_3$, where the parameter c can be used to tune the weight of the set S_3 . This problem is not particularly important—we chose it just as a concrete example for the sake of illustration. In Section 6 we compute the parameters $R(B_c), C(B_c)$ and deduce that the UGC hardness threshold of the problem of computing $\mathbf{Clust}(A|B_c)$ equals $\frac{4\pi c(1+c)^2}{(1+2c)^3}$ if $c \geq \frac{1}{2}$ and equals $\frac{\pi(1+c)^2}{2+4c}$ if $c \leq \frac{1}{2}$. The change at $c = \frac{1}{2}$ corresponds in a qualitative change in the best algorithm for computing $\mathbf{Clust}(A|B_c)$ —we refer to Section 6 for an explanation.

In the remainder of this introduction we will explain the various ingredients of Theorem 1.1 (in particular the Unique Games Conjecture), and the new ideas used in its proof.

The main tool in the design of the algorithm in Theorem 1.1 is a natural generalization of the positive semidefinite Grothendieck inequality. In [4] Grothendieck proved that there exists a universal constant $K > 0$ such that for every $n \times n$ symmetric positive semidefinite matrix $A = (a_{ij})$ we have²:

$$\max_{x_1, \dots, x_n \in S^{n-1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle \leq K \max_{\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \varepsilon_i \varepsilon_j. \quad (2)$$

The best constant K in (2) was shown in [11] to be equal to $\frac{\pi}{2}$. A natural variant of (2) is to replace the numbers $-1, 1$ by general $v_1, \dots, v_k \in \mathbb{R}^k$, namely one might ask for the smallest constant $K > 0$ such that for every symmetric positive semidefinite $n \times n$ matrix A we have:

$$\max_{x_1, \dots, x_n \in S^{n-1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle \leq K \max_{u_1, \dots, u_n \in \{v_1, \dots, v_k\}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle u_i, u_j \rangle. \quad (3)$$

In Section 3 we prove that (3) holds with $K = \frac{1}{C(B)}$, where $B = (\langle v_i, v_j \rangle)$ is the Gram matrix of v_1, \dots, v_k , and that this constant is sharp. This inequality is proved along the following lines. Fix n unit vectors $x_1, \dots, x_n \in S^{n-1}$. Let $G = (g_{ij})$ be a $(k-1) \times n$ random matrix whose entries are i.i.d. standard Gaussian random variables. Let $A_1, \dots, A_k \subseteq \mathbb{R}^{k-1}$ be a measurable partition of \mathbb{R}^{k-1} at which $C(B)$ is attained. Define a random choice of $u_i \in \{v_1, \dots, v_k\}$ by setting $u_i = v_\ell$ for the unique $\ell \in \{1, \dots, k\}$ such that $Gx_i \in A_\ell$. The fact that (3) holds with $K = \frac{1}{C(B)}$ is a consequence of the following fact, which we prove in Section 3:

$$\mathbb{E} \left[\sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle u_i, u_j \rangle \right] \geq C(B) \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle. \quad (4)$$

The crucial point in the proof of (4) is the following identity, proved in Lemma 3.2 as a corollary of the closed-form formula for the Poisson kernel of the Hermite polynomials: for every two measurable subsets

²This inequality is sometimes written as $\max_{x_i, y_j \in S^{n-1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, y_j \rangle \leq K \max_{\varepsilon_i, \delta_j \in \{-1, 1\}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \varepsilon_i \delta_j$, but it is easy (and standard) to verify that since A is positive semidefinite this formulation coincides with (2).

$E, F \subseteq \mathbb{R}^{k-1}$ and any two unit vectors $x, y \in \mathbb{R}^n$, we have

$$\begin{aligned} & \Pr[Gx \in E \text{ and } Gy \in F] \\ &= \gamma_{k-1}(E)\gamma_{k-1}(F) + \langle x, y \rangle \left\langle \int_E u d\gamma_{k-1}(u), \int_F u d\gamma_{k-1}(u) \right\rangle + \sum_{\ell=2}^{\infty} \langle x^{\otimes \ell}, y^{\otimes \ell} \rangle \sum_{\substack{s \in (\mathbb{N} \cup \{0\})^{k-1} \\ s_1 + \dots + s_{k-1} = \ell}} \alpha_s(E)\alpha_s(F), \quad (5) \end{aligned}$$

for some real coefficients $\{\alpha_s(E)\}_{s \in (\mathbb{N} \cup \{0\})^{k-1}}, \{\alpha_s(F)\}_{s \in (\mathbb{N} \cup \{0\})^{k-1}} \subseteq \mathbb{R}$. Here γ_{k-1} denotes the standard Gaussian measure on \mathbb{R}^{k-1} . The product structure of the decomposition (5) hints at the role of the fact that A is positive semidefinite in the proof of (4)—the complete details appear in Section 3.

Once the generalized Grothendieck inequality (18) is obtained with $K = \frac{1}{C(B)}$ it is simple to design the algorithm whose existence is claimed in Theorem 1.1, which is based on semidefinite programming—this is done in Section 4.

We shall now pass to an explanation of the hardness result in Theorem 1.1. The Unique Games Conjecture, posed by Khot in [6], is as follows. A Unique Game is an optimization problem with an instance $\mathcal{L} = \mathcal{L}(G(V, W, E), n, \{\pi_{vw}\}_{(v,w) \in E})$. Here $G(V, W, E)$ is a regular bipartite graph with vertex sets V and W and edge set E . Each vertex is supposed to receive a label from the set $\{1, \dots, n\}$. For every edge $(v, w) \in E$ with $v \in V$ and $w \in W$, there is a given permutation $\pi_{vw} : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$. A labeling of the Unique Game instance is an assignment $\rho : V \cup W \rightarrow \{1, \dots, n\}$. An edge (v, w) is satisfied by a labeling ρ if and only if $\rho(v) = \pi_{vw}(\rho(w))$. The goal is to find a labeling that maximizes the fraction of edges satisfied (call this maximum $\text{OPT}(\mathcal{L})$). We think of the number of labels n as a constant and the size of the graph $G(V, W, E)$ as the size of the problem instance. The Unique Games Conjecture (UGC) asserts that for arbitrarily small constants $\varepsilon, \delta > 0$, there exists a constant $n = n(\varepsilon, \delta)$ such that no polynomial time algorithm can distinguish whether a Unique Games instance $\mathcal{L} = \mathcal{L}(G(V, W, E), n, \{\pi_{vw}\}_{(v,w) \in W})$ satisfies $\text{OPT}(\mathcal{L}) \leq \delta$ (soundness) or there exists a labeling such that for $1 - \varepsilon$ fraction of the vertices $v \in V$ all the edges incident with v are satisfied (completeness)³. This conjecture is (by now) a commonly used complexity assumption to prove hardness of approximation results. Despite several recent attempts to get better polynomial time approximation algorithms for the Unique Game problem (see the table in [3] for a description of known results), the unique games conjecture still stands.

Our UGC hardness result follows the standard “dictatorship test” approach which is prevalent in PCP based hardness proofs, with a new twist which seems to be of independent interest. Since the kernel clustering problem is concerned with an assignment of one of k labels to each of the rows of the matrix A , the natural setting of our hardness proof is a dictatorship test for functions on $\{1, \dots, k\}^n$ taking values in $\{1, \dots, k\}$ (this was already the case in [7]). The general “philosophy” of such hardness proofs is to associate to every such function a certain numerical parameter called the “objective value” (which is adapted to the optimization problem at hand). The general scheme is to show that for some numbers $a, b > 0$, if f depends on only one coordinate (i.e., it is a “dictatorship”) then the objective value of f is at least a , while if f does not have any coordinate which is too influential then the objective value of f is at most $b + o(1)$ (the $o(1)$ depends on the notion of having no influential coordinates and its exact form is not important for the purpose of this overview—we refer to Section 5 for details). Once such a result is proved, techniques from the theory of Probabilistically Checkable Proofs can show that under a suitable complexity theoretic assumption (in our case the UGC) no polynomial time algorithm can achieve an approximation factor smaller than $\frac{a}{b}$.

³This version of the UGC is not the standard version as stated in [6], which only requires $\text{OPT}(\mathcal{L}) \geq 1 - \varepsilon$ in the completeness. However, it was shown in [8] that this seemingly stronger version of the UGC actually follows from the original UGC—we will require this stronger statement in our proofs.

Implicit to the above discussion is an underlying product distribution on $\{1, \dots, k\}^n$ with respect to which we measure the influence of variables. In [7] the case of $B = I_k$ was solved using the uniform distribution on $\{1, \dots, k\}$. It turns out that in order to prove the sharp hardness result in Theorem 1.1 we need to use a non-uniform distribution which depends on the geometry of B . Namely, writing B as a Gram matrix $b_{ij} = \langle v_i, v_j \rangle$, recall that $R(B)$ is the radius of the smallest Euclidean ball containing $\{v_1, \dots, v_k\}$ and $w(B)$ is the center of this ball. A simple separation argument shows that $w(B)$ is in the convex hull of the vectors in $\{v_1, \dots, v_k\}$ whose distance from $w(B)$ is exactly $R(B)$. Writing $w(B)$ as a convex combination of these points and considering the coefficients of this convex combination results in a probability distribution on $\{1, \dots, k\}$. In our hardness proof we use the n -fold product of (a small perturbation of) this probability distribution as the underlying distribution on $\{1, \dots, k\}$ for our dictatorship test—see Figure 1 for a schematic description of the situation described above. The full details of this approach, including all the relevant definitions, are presented in Section 5.

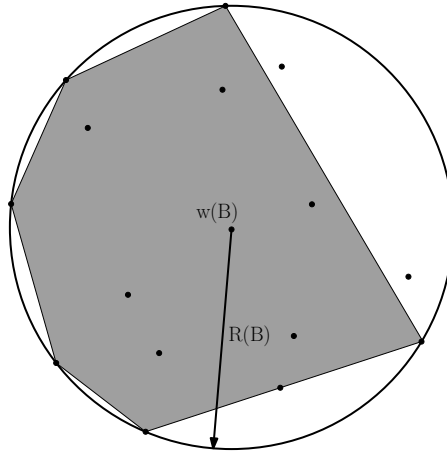


Figure 1: *The geometry of the test matrix B induces a dictatorship test: the points above are the vectors $\{v_1, \dots, v_k\} \subseteq \mathbb{R}^k$ such that B is their Gram matrix. The ball depicted above is the smallest Euclidean ball containing $\{v_1, \dots, v_k\}$, $R(B)$ is its radius and $w(B)$ is its center. Then $w(B)$ is in the convex hull of the points in $\{v_1, \dots, v_k\}$ which are at distance exactly $R(B)$ from $w(B)$. Writing $w(B)$ as a convex combination of these boundary points yields a distribution over the labels $\{1, \dots, k\}$. Our dictatorship test corresponds to selecting a point from the n -fold power of this probability space and comparing the behavior of a certain “objective value” (defined in equation (31) below), which depends only on the singleton Fourier coefficients, for dictatorships and for functions with low influences.*

1.1 Open problems

We end this introduction with a statement of some open problems.

Question 1. Theorem 1.1 shows that the UGC hardness threshold of the problem of computing $\mathbf{Clust}(A|B)$ for a fixed hypothesis matrix B equals $\frac{R(B)^2}{C(B)}$. It is natural to ask if there is also a polynomial time algorithm which outputs a clustering of A whose value is within a factor of $\frac{R(B)^2}{C(B)}$ of the optimal clustering. The issue is that our rounding algorithm uses the partition $\{A_1, \dots, A_k\}$ of \mathbb{R}^{k-1} at which $C(B)$ is attained. In Section 2 we study this optimal partition, and show that it has a relatively simple structure rather than being composed

of general measurable sets: it corresponds to cones which are induced by the faces of a simplex. This information allows us to compute efficiently a partition which comes as close as we wish to the optimal partition when k is fixed, or grows slowly with n (to be safe lets just say for the sake of argument that $k \approx \log \log n$ works). We refer to Remark 2.3 for details. We currently do not know if there is polynomial time rounding algorithm when, say, $k \approx \sqrt{n}$. Given $\varepsilon > 0$, is there an algorithm which, given A and B , computes $\mathbf{Clust}(A|B)$ to within a factor of $(1 + \varepsilon) \frac{R(B)^2}{C(B)}$, and runs in time which is polynomial in both n and k (and maybe even $1/\varepsilon$)?

Question 2. We remind the reader that the Propeller Conjecture remains open. This conjecture is about the value of $C(I_k)$ when $k \geq 4$. It states that the partition at which $C(I_k)$ is attained is actually much simpler than what one might initially expect: only 3 of the sets have positive measure and they form a cylinder over a planar 120° “propeller”. We refer to [7] for a precise formulation and some evidence for the validity of the Propeller Conjecture.

Question 3. The kernel clustering problem was stated in [13] for matrices A which are centered. This makes sense from the perspective of machine learning, but it seems meaningful to also ask for the UGC hardness threshold of the same problem when A is not assumed to be centered. In the present paper we did not investigate this case at all, and it seems that the exact UGC hardness threshold when A is not necessarily centered is not known for any interesting hypothesis matrix B . Note that in [7] we showed that there is a constant factor polynomial time approximation algorithm when A is not necessarily centered: we obtained in [7] an approximation guarantee of $1 + \frac{3\pi}{2}$ in this case, but this is probably suboptimal.

2 Preliminaries on the parameter $C(B)$

Let $B = (b_{ij})_{i,j=1}^k \in M_k(\mathbb{R})$ be a $k \times k$ symmetric positive semidefinite matrix. In what follows we fix $k \geq 2$ and the matrix B . We also fix vectors $v_1, \dots, v_k \in \mathbb{R}^k$ for which $b_{ij} = \langle v_i, v_j \rangle$ for all $i, j \in \{1, \dots, k\}$.

Let γ_n denote the standard Gaussian measure on \mathbb{R}^n , i.e., the density of γ_n is $\frac{1}{(2\pi)^{n/2}} e^{-\|x\|_2^2/2}$. We denote by H_k the Hilbert space $L_2(\gamma_n) \oplus L_2(\gamma_n) \oplus \dots \oplus L_2(\gamma_n)$ (k times) and we consider the convex subset $\Delta_k(\gamma_n) \subseteq H_k$ give by:

$$\Delta_k(\gamma_n) := \left\{ (f_1, \dots, f_k) \in H_k : \forall j \in \{1, \dots, k\} f_j \geq 0 \wedge \sum_{j=1}^k f_j = 1 \right\}. \quad (6)$$

Define:

$$C(n, B) := \sup_{(f_1, \dots, f_k) \in \Delta_k(\gamma_n)} \sum_{i=1}^k \sum_{j=1}^k b_{ij} \cdot \left\langle \int_{\mathbb{R}^n} x f_i(x) d\gamma_n(x), \int_{\mathbb{R}^n} x f_j(x) d\gamma_n(x) \right\rangle. \quad (7)$$

The following lemma is a variant of Lemma 3.1 in [7] (but see Remark 2.1 for an explanation of a subtle difference). It simply states that the supremum in (7) is attained at a k -tuple of functions which correspond to a partition of \mathbb{R}^n .

Lemma 2.1. *There exist disjoint measurable sets $A_1, \dots, A_k \subseteq \mathbb{R}^n$ such that $A_1 \cup A_2 \cup \dots \cup A_k = \mathbb{R}^n$ and*

$$\sum_{i=1}^k \sum_{j=1}^k b_{ij} \cdot \left\langle \int_{A_j} x d\gamma_n(x), \int_{A_i} x d\gamma_n(x) \right\rangle = C(n, B).$$

Proof. Define $\Psi : \Delta_k(\gamma_n) \rightarrow \mathbb{R}$ by

$$\Psi(f_1, \dots, f_k) := \sum_{i=1}^k \sum_{j=1}^k b_{ij} \cdot \left\langle \int_{\mathbb{R}^n} x f_i(x) d\gamma_n(x), \int_{\mathbb{R}^n} x f_j(x) d\gamma_n(x) \right\rangle. \quad (8)$$

We first observe that Ψ is a convex function. Indeed, fix $\lambda \in [0, 1]$ and $(f_1, \dots, f_k), (g_1, \dots, g_k) \in \Delta_k(\gamma_n)$. Denote $z_i := \int_{\mathbb{R}^n} x f_i(x) d\gamma_n(x)$ and $w_i := \int_{\mathbb{R}^n} x g_i(x) d\gamma_n(x)$ for every $i \in \{1, \dots, k\}$. Then:

$$\begin{aligned} & \lambda \Psi(f_1, \dots, f_k) + (1 - \lambda) \Psi(g_1, \dots, g_k) - \Psi(\lambda f_1 + (1 - \lambda)g_1, \dots, \lambda f_k + (1 - \lambda)g_k) \\ &= \sum_{i=1}^k \sum_{j=1}^k \langle v_i, v_j \rangle \left(\lambda \langle z_i, z_j \rangle + (1 - \lambda) \langle w_i, w_j \rangle - \langle \lambda z_i + (1 - \lambda)w_i, \lambda z_j + (1 - \lambda)w_j \rangle \right) \\ &= \lambda(1 - \lambda) \sum_{i=1}^k \sum_{j=1}^k \langle v_i, v_j \rangle \langle z_i - w_i, z_j - w_j \rangle \\ &= \lambda(1 - \lambda) \left\| \sum_{i=1}^k v_i \otimes (z_i - w_i) \right\|_2^2 \geq 0. \end{aligned}$$

Since $\Delta_k(\gamma_n)$ is a weakly compact subset of H_k and Ψ is weakly continuous and convex, Ψ attains its maximum (which equals $C(n, B)$) on $\Delta_k(\gamma_n)$ at an extreme point of $\Delta_k(\gamma_n)$, say at $(f_1^*, \dots, f_k^*) \in \Delta_k(\gamma_n)$. It follows that there exist measurable sets $A_1, \dots, A_k \subseteq \mathbb{R}^n$ which form a partition of \mathbb{R}^n such that $(f_1^*, \dots, f_k^*) = (\mathbf{1}_{A_1}, \dots, \mathbf{1}_{A_k})$ almost everywhere⁴, as required. \square

Remark 2.1. In [7] a stronger result was proved when $B = I_k$ (the $k \times k$ identity matrix). Namely, using the notation of the proof of Lemma 2.1 it was shown that the maximum of Ψ on the larger convex set

$$\widetilde{\Delta_k(\gamma_n)} := \left\{ (f_1, \dots, f_k) \in H_k : \forall j \in \{1, \dots, k\} f_j \geq 0 \wedge \sum_{j=1}^k f_j \leq 1 \right\}$$

is also attained at $(f_1^*, \dots, f_k^*) = (\mathbf{1}_{A_1}, \dots, \mathbf{1}_{A_k})$ for some measurable sets $A_1, \dots, A_k \subseteq \mathbb{R}^n$ which form a partition of \mathbb{R}^n . It turns out that this stronger fact helps to slightly simplify the proof of the corresponding UGC hardness result. However, we do not know how to prove this stronger statement for general B , so we formulated the weaker statement in Lemma 2.1, at the cost of needing to modify our proof of the UGC hardness result for general B in Section 5.

The same extreme point argument as in the proof of Lemma 2.1 shows that the maximum of Ψ on $\widetilde{\Delta_k(\gamma_n)}$ is attained at $(f_1^*, \dots, f_k^*) = (\mathbf{1}_{A_1}, \dots, \mathbf{1}_{A_k})$ for some disjoint measurable sets $A_1, \dots, A_k \subseteq \mathbb{R}^n$, but now it does not follow that they necessarily cover all of \mathbb{R}^n . When $B = I_k$ it can be shown as in [7] that these sets do cover \mathbb{R}^n . The same statement is true when B is diagonal, as we now show by arguing as in the proof in [7], but we do not know if it is true for general B . So, assume that B is diagonal with positive diagonal entries (b_1, \dots, b_k) . Let $A = \mathbb{R}^n \setminus \bigcup_{i=1}^k A_i$. Denote $z_j := \int_{A_j} x d\gamma_n(x)$ and $w = \int_A x d\gamma_n(x)$. Note that

⁴To see this standard fact observe that otherwise there would be some $A \subseteq \mathbb{R}^n$ of positive measure, $\varepsilon \in (0, 1/2)$, and distinct $i, j \in \{1, \dots, k\}$ such that $f_i \mathbf{1}_A, f_j \mathbf{1}_A \in (\varepsilon, 1 - \varepsilon)$. But (f_1^*, \dots, f_k^*) would then not be an extreme point since it is the average of $(g_1, \dots, g_k), (h_1, \dots, h_k) \in \Delta_k(\gamma_n) \setminus \{(f_1^*, \dots, f_k^*)\}$, where $g_\ell = h_\ell = f_\ell^*$ for $\ell \in \{1, \dots, k\} \setminus \{i, j\}$ and $g_i = (f_i^* + \varepsilon) \mathbf{1}_A + f_i^* \mathbf{1}_{\mathbb{R}^n \setminus A}$, $h_i = (f_i^* - \varepsilon) \mathbf{1}_A + f_i^* \mathbf{1}_{\mathbb{R}^n \setminus A}$, $g_j = (f_j^* - \varepsilon) \mathbf{1}_A + f_j^* \mathbf{1}_{\mathbb{R}^n \setminus A}$, $h_j = (f_j^* + \varepsilon) \mathbf{1}_A + f_j^* \mathbf{1}_{\mathbb{R}^n \setminus A}$.

$w + z_1 + \dots + z_k = 0$. If $w = 0$ then Ψ attains its maximum on the partition $\{A \cup A_1, A_2, \dots, A_k\}$, so assume for the sake of contradiction that $w \neq 0$. For every $i \in \{1, \dots, k\}$ we have:

$$\begin{aligned} \sum_{j=1}^n b_j \|z_j\|_2^2 &= \Psi(\mathbf{1}_{A_1}, \dots, \mathbf{1}_{A_k}) \geq \Psi(\mathbf{1}_{A_1}, \dots, \mathbf{1}_{A_{i-1}}, \mathbf{1}_{A \cup A_i}, \mathbf{1}_{A_{i+1}}, \dots, \mathbf{1}_{A_k}) \\ &= \sum_{\substack{1 \leq j \leq k \\ j \neq i}} b_j \|z_j\|_2^2 + b_i \|z_i + w\|_2^2 = \sum_{j=1}^n b_j \|z_j\|_2^2 + 2b_i \langle z_i, w \rangle + b_i \|w\|_2^2. \end{aligned}$$

Thus $2\langle z_i, w \rangle + \|w\|_2^2 \leq 0$, and if we sum this inequality over $i \in \{1, \dots, k\}$ while recalling that $w = -\sum_{i=1}^k z_i$ we see that $(k-2)\|w\|_2^2 \leq 0$, which is a contradiction. Note that for general B the same argument shows that for all $i \in \{1, \dots, k\}$ we have $2\sum_{j=1}^k b_{ij} \langle z_j, w \rangle + b_{ii} \|w\|_2^2 \leq 0$. These inequalities do not seem to lend themselves to the same type of easy contradiction as in the case of diagonal matrices. \triangleleft

The proof of the following lemma is an obvious modification of the proof of Lemma 3.2 in [7].

Lemma 2.2. *If $n \geq k-1$ then $C(n, B) = C(k-1, B)$.*

Proof. The inequality $C(n, B) \geq C(k-1, B)$ is easy since for every $(f_1, \dots, f_k) \in \Delta_k(\gamma_{k-1})$ we can define $(\tilde{f}_1, \dots, \tilde{f}_k) \in \Delta_k(\gamma_n)$ by $\tilde{f}_j(x, y) = f_j(x)$ (thinking here of \mathbb{R}^n as $\mathbb{R}^{k-1} \times \mathbb{R}^{n-k+1}$). Then for all $j \in \{1, \dots, k\}$ we have $\int_{\mathbb{R}^{k-1}} x f_j(x) d\gamma_{k-1}(x) = \int_{\mathbb{R}^n} x \tilde{f}_j(x) d\gamma_n(x)$, implying that $\Psi(\tilde{f}_1, \dots, \tilde{f}_k) = \Psi(f_1, \dots, f_k)$.

In the reverse direction, by Lemma 2.1 there is a measurable partition A_1, \dots, A_k of \mathbb{R}^n such that if we define $z_j := \int_{A_j} x d\gamma_n(x) \in \mathbb{R}^n$ then we have $\sum_{i=1}^k \sum_{j=1}^k b_{ij} \langle z_i, z_j \rangle = C(n, B)$. Note that $\sum_{j=1}^k z_j = 0$. Hence the dimension of the subspace $V := \text{span}\{z_1, \dots, z_k\}$ is $d \leq k-1$. Define $g_1, \dots, g_k : V \rightarrow [0, 1]$ by $g_j(x) = \gamma_{V^\perp}((A_j - x) \cap V^\perp)$. Then $(g_1, \dots, g_k) \in \Delta_k(\gamma_V)$, so that

$$\begin{aligned} C(k-1, B) &\geq C(d, B) \\ &\geq \sum_{i=1}^k \sum_{j=1}^k b_{ij} \left\langle \int_V x g_i(x) d\gamma_V(x), \int_V x g_j(x) d\gamma_V(x) \right\rangle \\ &= \sum_{i=1}^k \sum_{j=1}^k b_{ij} \left\langle \int_V \int_{V^\perp} \mathbf{1}_{A_i}(x+y) x d\gamma_V(x) d\gamma_{V^\perp}(y), \int_V \int_{V^\perp} \mathbf{1}_{A_j}(x+y) x d\gamma_V(x) d\gamma_{V^\perp}(y) \right\rangle \\ &= \sum_{i=1}^k \sum_{j=1}^k b_{ij} \left\langle \int_{A_i} \text{Proj}_V(w) d\gamma_n(w), \int_{A_j} \text{Proj}_V(w) d\gamma_n(w) \right\rangle \\ &= \sum_{i=1}^k \sum_{j=1}^k b_{ij} \langle \text{Proj}_V(z_i), \text{Proj}_V(z_j) \rangle \\ &= \sum_{i=1}^k \sum_{j=1}^k b_{ij} \langle z_i, z_j \rangle = C(n, B), \end{aligned}$$

as required. \square

In light of Lemma 2.2 we define $C(B) := C(k-1, B)$. We shall now prove an analogue of Lemma 3.3 in [7] which gives structural information on the partition $\{A_1, \dots, A_k\}$ of \mathbb{R}^{k-1} at which $C(B)$ is attained. We

first recall some notation and terminology from [7]. Given distinct $z_1, \dots, z_k \in \mathbb{R}^{k-1}$ and $j \in \{1, \dots, k\}$ define a set $P_j(z_1, \dots, z_k) \subseteq \mathbb{R}^k$ by

$$P_j(z_1, \dots, z_k) := \left\{ x \in \mathbb{R}^k : \langle x, z_j \rangle = \max_{i \in \{1, \dots, k\}} \langle x, z_i \rangle \right\}.$$

Thus $\{P_j(z_1, \dots, z_k)\}_{j=1}^k$ is a partition of \mathbb{R}^{k-1} which we call the simplicial partition induced by z_1, \dots, z_k (strictly speaking the elements of this partition are not disjoint, but they intersect at sets of measure 0).

Lemma 2.3. *Let $A_1, \dots, A_k \subseteq \mathbb{R}^{k-1}$ be a partition into measurable sets such that if we set $z_j := \int_{A_j} x d\gamma_{k-1}(x)$ then*

$$C(B) = \sum_{i=1}^k \sum_{j=1}^k b_{ij} \langle z_i, z_j \rangle. \quad (9)$$

Assume also that this partition is minimal in the sense that the number of elements of positive measure in this partition is minimum among all the possible partitions satisfying (9). Define

$$J := \left\{ j \in \{1, \dots, k\} : \gamma_{k-1}(A_j) > 0 \right\}$$

and set $|J| = \ell$. Then up to an orthogonal transformation $\{z_j\}_{j \in J} \subseteq \mathbb{R}^{\ell-1}$ and the vectors $\{z_j\}_{j \in J}$ are non-zero and distinct. Moreover, if we define $\{w_j\}_{j \in J} \subseteq \mathbb{R}^{\ell-1}$ by

$$w_j := \sum_{s \in J} b_{js} z_s, \quad (10)$$

then the vectors $\{w_j\}_{j \in J}$ are distinct and for each $j \in J$ we have

$$A_j = P_j((w_i)_{i \in J}) \times \mathbb{R}^{k-\ell} \quad (11)$$

up to sets of measure zero.

Proof. Since $\sum_{j \in J} \mathbf{1}_{A_j} = 1$ almost everywhere we have $\sum_{j \in J} z_j = 0$. Thus the dimension of the span of $\{z_j\}_{j \in J}$ is at most $|J| - 1 = \ell - 1$, and by applying an orthogonal transformation we may assume that $\{z_j\}_{j \in J} \subseteq \mathbb{R}^{\ell-1}$. Also, for every distinct $i, j \in J$ replace A_i by $A_i \cup A_j$ and A_j by the empty set and obtain a partition of \mathbb{R}^{k-1} which contains exactly $\ell - 1$ elements of positive measure and for which we have (by the minimality of ℓ):

$$\begin{aligned} C(B) &> \sum_{s,t \in J \setminus \{i,j\}} b_{st} \langle z_s, z_t \rangle + 2 \sum_{s \in J \setminus \{i,j\}} b_{is} \langle z_s, z_i + z_j \rangle + b_{ii} \|z_i + z_j\|_2^2 \\ &= \sum_{s,t \in J} b_{st} \langle z_s, z_t \rangle + 2 \sum_{s \in J} (b_{is} - b_{js}) \langle z_s, z_j \rangle + (b_{ii} + b_{jj} - 2b_{ij}) \|z_j\|_2^2 \\ &= C(B) + 2 \langle w_i - w_j, z_j \rangle + \|v_i - v_j\|_2^2 \cdot \|z_j\|_2^2, \end{aligned}$$

where we used the fact that $b_{st} = \langle v_s, v_t \rangle$. Thus

$$2 \langle w_i - w_j, z_j \rangle + \|v_i - v_j\|_2^2 \cdot \|z_j\|_2^2 < 0, \quad (12)$$

and by symmetry we also have the inequality:

$$2 \langle w_j - w_i, z_i \rangle + \|v_i - v_j\|_2^2 \cdot \|z_i\|_2^2 < 0. \quad (13)$$

It follows in particular from (12) and (13) that z_i and z_j are non-zero and that $w_i \neq w_j$. Moreover if we sum (12) and (13) we get that

$$2 \langle w_i - w_j, z_j - z_i \rangle + \|v_i - v_j\|_2^2 (\|z_i\|_2^2 + \|z_j\|_2^2) < 0$$

which implies that $z_i \neq z_j$.

The above reasoning implies in particular that $\{P_j((w_i)_{i \in J}) \times \mathbb{R}^{k-\ell}\}_{j \in J}$ is a partition of \mathbb{R}^{k-1} (up to pairwise intersections at sets of measure 0). Assume for the sake of contradiction that there exist $i \in J$ such that

$$\gamma_{k-1}(A_i \setminus (P_i((w_s)_{s \in J}) \times \mathbb{R}^{k-\ell})) > 0.$$

Arguing as in the proof of Lemma 3.3 in [7] we see that there exists $\varepsilon > 0$ and $j \in J \setminus \{i\}$ such that if we denote $E := \{x \in A_i : \langle x, z_j \rangle \geq \langle x, z_i \rangle + \varepsilon\}$ then $\gamma_{k-1}(E) > 0$.

Define a partition $\tilde{A}_1, \dots, \tilde{A}_k$ of \mathbb{R}^{k-1} by

$$\tilde{A}_r := \begin{cases} A_r & r \notin \{i, j\} \\ A_i \setminus E & r = i \\ A_j \cup E & r = j. \end{cases}$$

Then for $w := \int_E x d\gamma_{k-1}(x)$ we have

$$\begin{aligned} C(B) &\geq \sum_{s,t \in J} b_{st} \left\langle \int_{\tilde{A}_s} x d\gamma_{k-1}(x), \int_{\tilde{A}_t} x d\gamma_{k-1}(x) \right\rangle \\ &= \sum_{s,t \in J \setminus \{i,j\}} b_{st} \langle z_s, z_t \rangle + 2 \sum_{s \in J \setminus \{i,j\}} b_{is} \langle z_s, z_i - w \rangle + 2 \sum_{s \in J \setminus \{i,j\}} b_{js} \langle z_s, z_j + w \rangle \\ &\quad + 2b_{ij} \langle z_i - w, z_j + w \rangle + b_{ii} \|z_i - w\|_2^2 + b_{jj} \|z_j + w\|_2^2 \\ &= C(B) - 2 \sum_{s \in J} b_{is} \langle z_s, w \rangle + 2 \sum_{s \in J} b_{js} \langle z_s, w \rangle + (b_{ii} + b_{jj} - 2b_{ij}) \|w\|_2^2 \\ &\stackrel{(10)}{=} C(B) + 2 \langle w_j - w_i, w \rangle + \|v_i - v_j\|_2^2 \cdot \|w\|_2^2 \\ &\geq C(B) + 2 \int_E (\langle z_j, x \rangle - \langle z_i, x \rangle) d\gamma_{k-1}(x) \\ &\geq C(B) + 2\varepsilon \gamma_{k-1}(E) > C(B), \end{aligned}$$

a contradiction. □

Remark 2.2. Note that we have the following non-trivial identity as a corollary of Lemma 2.3 (and using the same notation): For each $i \in J$,

$$z_j = \int_{P_j((w_i)_{i \in J})} x d\gamma_{\ell-1}(x), \tag{14}$$

where we recall that the w_i are defined in (10). This system of equalities seems to contain non-trivial information on the structure of the partition at which $C(B)$ is attained. In future research it would be of interest to exploit this information, though we have no need for it for our present purposes. ◁

Remark 2.3. Given B and $\varepsilon > 0$ we can estimate $C(B)$ up to an error of at most ε in constant time (which depends only on B, k, ε). Moreover, we can compute in constant time a conical simplicial partition of \mathbb{R}^{k-1}

at which the value of Ψ is at least $C(B) - \varepsilon$. These statements are a simple corollary of Lemma 2.3. Indeed, all we have to do is to run over all choices of $\ell \in \{1, \dots, k\}$ and for each such ℓ construct an appropriate net of $z_1, \dots, z_\ell \in \mathbb{R}^{\ell-1}$ of bounded size, and then check each of the induced simplicial partitions of \mathbb{R}^{k-1} as in (11) for the one which maximizes Ψ . To this end we need some a priori bound on the length of z_i : the crude bound

$$\|z_i\|_2 = \left\| \int_{A_i} x d\gamma_{\ell-1}(x) \right\|_2 \leq \int_{\mathbb{R}^{\ell-1}} \|x\|_2 d\gamma_{\ell-1}(x) \leq \sqrt{\ell}$$

will suffice. Fix $\delta > 0$ which will be determined momentarily. Let \mathcal{N} be a δ -net in the Euclidean ball of radius $\sqrt{\ell}$ in $\mathbb{R}^{\ell-1}$. Then $|\mathcal{N}| \leq \left(\frac{3\sqrt{\ell}}{\delta}\right)^\ell$.

Let A_1, \dots, A_k be as in Lemma 2.3, i.e., the true (minimal) partition at which $C(B)$ is attained. Let J, ℓ, z_i and w_i be as in Lemma 2.3. For each $i \in J$ find $z'_i \in \mathcal{N}$ for which $\|z_i - z'_i\|_2 \leq \delta$. Define $w'_i = \sum_{s \in J} b_{js} z'_s$. Then we have the crude bound $\|w_i - w'_i\|_2 \leq \delta \sum_{s=1}^k \sum_{t=1}^k |b_{st}| := \delta \|B\|_1$. We also have the a priori bounds $\|w_i\|_2, \|w'_i\|_2 \leq \sqrt{\ell} \|B\|_1$. By compactness there exists $\delta = \delta(\varepsilon, \ell, B)$ such that these estimates imply that for all $j \in J$,

$$\left\| z_j - \int_{P_j((w'_i)_{i \in J})} x d\gamma_{\ell-1}(x) \right\|_2 = \left\| \int_{P_j((w_i)_{i \in J})} x d\gamma_{\ell-1}(x) - \int_{P_j((w'_i)_{i \in J})} x d\gamma_{\ell-1}(x) \right\|_2 \leq \frac{\varepsilon}{2\sqrt{\ell}\|B\|_1}. \quad (15)$$

(It is actually easy to give a concrete bound on the required δ if so desired, but this is not important for our purposes.) It follows from (15) that:

$$\begin{aligned} C(B) &\geq \sum_{s,t \in J} b_{st} \left\langle \int_{P_s((w'_i)_{i \in J})} x d\gamma_{\ell-1}(x), \int_{P_t((w'_i)_{i \in J})} x d\gamma_{\ell-1}(x) \right\rangle \\ &\geq \sum_{s,t \in J} b_{st} \langle z_s, z_t \rangle - \sum_{s,t \in J} |b_{st}| \cdot \frac{\varepsilon}{2\sqrt{\ell}\|B\|_1} \cdot 2\sqrt{\ell} = C(B) - \varepsilon. \end{aligned}$$

Note that the above integrals can be estimated efficiently (polynomial time in k) with arbitrarily good precision due to the fact that the simplicial cones $P_j((w'_i)_{i \in J})$ have an efficient membership oracle and the Gaussian measure is log-concave. These are very crude bounds that suffice for our algorithmic purposes when k is fixed, but deteriorate exponentially with k . It would be of interest to understand whether we can estimate $C(B)$ (and more importantly the associated partitions, as they are used in our rounding procedure) in time which is polynomial in k . Perhaps the identities (14) can play a role in the design of such an efficient algorithm, but we did not investigate this issue. \triangleleft

We end this section with a simple analytic interpretation of the parameter $C(B)$. Given a square integrable function $f : \mathbb{R}^n \rightarrow \mathbb{R}^k$ its Rademacher projection $\mathbf{Rad}(f) : \mathbb{R}^n \rightarrow \mathbb{R}^k$ (see [10] for an explanation of this terminology) is defined for $x = (x_1, \dots, x_n) \in \mathbb{R}^n$ as:

$$\mathbf{Rad}(f)(x) = \sum_{i=1}^n \left(\int_{\mathbb{R}^n} y_i f(y) d\gamma_n(y) \right) x_i.$$

Assume that f takes values in $\{v_1, \dots, v_k\} \subseteq \mathbb{R}^k$ and define $A_i = f^{-1}(v_i)$ for $i \in \{1, \dots, k\}$. Then $\{A_1, \dots, A_k\}$ is a measurable partition of \mathbb{R}^n . We also have the identity:

$$\mathbf{Rad}(f)(x) = \sum_{i=1}^n \left(\sum_{j=1}^k v_j \int_{A_j} y_i d\gamma_n(y) \right) x_i.$$

Thus

$$\begin{aligned} \|\mathbf{Rad}(f)\|_{L_2(\gamma_n, \mathbb{R}^k)}^2 &= \int_{\mathbb{R}^n} \|\mathbf{Rad}(f)(x)\|_2^2 d\gamma_n(x) = \sum_{i=1}^n \left\| \sum_{j=1}^k v_j \int_{A_j} y_i d\gamma_n(y) \right\|_2^2 \\ &= \sum_{i=1}^n \sum_{j=1}^k \sum_{\ell=1}^k \langle v_j, v_\ell \rangle \left(\int_{A_j} y_i d\gamma_n(y) \right) \left(\int_{A_\ell} y_i d\gamma_n(y) \right) = \sum_{j=1}^k \sum_{\ell=1}^k b_{j\ell} \left\langle \int_{A_j} y d\gamma_n(y), \int_{A_\ell} y d\gamma_n(y) \right\rangle. \end{aligned} \quad (16)$$

The identity (16) implies the following lemma:

Lemma 2.4. *For every $n \geq k - 1$ we have:*

$$C(B) = \max_{f: \mathbb{R}^n \rightarrow \{v_1, \dots, v_k\}} \|\mathbf{Rad}(f)\|_{L_2(\gamma_n, \mathbb{R}^k)}^2.$$

Recall that $R(B)$ is defined as the radius of the smallest ball in \mathbb{R}^k which contains the set $\{v_1, \dots, v_k\}$ and that $w(B)$ is the center of this ball. Lemma 2.4 implies the following corollary:

Corollary 2.5. $C(B) \leq R(B)^2$.

Proof. Let $\{A_1, \dots, A_k\}$ be a partition of \mathbb{R}^{k-1} into measurable sets such that if we define $z_j = \int_{A_j} x d\gamma_{k-1}(x)$ then

$$\begin{aligned} C(B) &= \sum_{i=1}^k \sum_{j=1}^k \langle v_i, v_j \rangle \langle z_i, z_j \rangle \\ &= \sum_{i=1}^k \sum_{j=1}^k \langle v_i - w(B), v_j - w(B) \rangle \langle z_i, z_j \rangle + 2 \sum_{i=1}^k \langle v_i, w(B) \rangle \left\langle z_i, \sum_{j=1}^k z_j \right\rangle + \|w(B)\|_2^2 \cdot \left\| \sum_{j=1}^k z_j \right\|_2^2. \end{aligned} \quad (17)$$

Since $\sum_{j=1}^k z_j = 0$ it follows from (16) and (17) that for $f: \mathbb{R}^{k-1} \rightarrow \{v_i - w(B)\}_{i=1}^k$ defined by $f|_{A_i} = v_i - w(B)$ we have:

$$C(B) = \|\mathbf{Rad}(f)\|_{L_2(\gamma_n, \mathbb{R}^k)}^2 \stackrel{(\star)}{\leq} \|f\|_{L_2(\gamma_n, \mathbb{R}^k)}^2 \leq \|f\|_{L_\infty(\gamma_n, \mathbb{R}^k)}^2 = \max_{i \in \{1, \dots, k\}} \|v_i - w(B)\|_2^2 = R(B)^2,$$

where in (\star) we used the fact that \mathbf{Rad} is an orthogonal projection on the Hilbert space $L_2(\gamma_n, \mathbb{R}^k)$. \square

3 Generalized positive semidefinite Grothendieck inequalities

The purpose of this section is to prove the following theorem, which as explained in the introduction, is an extension of Grothendieck's inequality for positive semidefinite matrices.

Theorem 3.1. *Let $A = (a_{ij}) \in M_n(\mathbb{R})$ be an $n \times n$ symmetric positive semidefinite matrix. Let $v_1, \dots, v_k \in \mathbb{R}^k$ be $k \geq 2$ vectors and let $B = (b_{ij} = \langle v_i, v_j \rangle)$ be the corresponding Gram matrix. Then*

$$\max_{x_1, \dots, x_n \in S^{n-1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle \leq \frac{1}{C(B)} \max_{\sigma: \{1, \dots, n\} \rightarrow \{1, \dots, k\}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle v_{\sigma(i)}, v_{\sigma(j)} \rangle. \quad (18)$$

We shall prove in Section 3.1 that the factor $\frac{1}{C(B)}$ in (18) cannot be improved, even when in (18) A is restricted to be centered, i.e., $\sum_{i=1}^n \sum_{j=1}^n a_{ij} = 0$.

The key tool in the proof of Theorem 3.1 is the following lemma:

Lemma 3.2. *Let $\{g_{ij} : i \in \{1, \dots, m\}, j \in \{1, \dots, n\}\}$ be i.i.d. standard Gaussian random variables and let $G = (g_{ij})$ be the corresponding $m \times n$ random Gaussian matrix. Fix two unit vectors $x, y \in S^{n-1}$ and two measurable subsets $E, F \subseteq \mathbb{R}^m$. Then:*

$$\begin{aligned} & \Pr[Gx \in E \wedge Gy \in F] \\ &= \gamma_m(E)\gamma_m(F) + \langle x, y \rangle \left\langle \int_E u d\gamma_m(u), \int_F u d\gamma_m(u) \right\rangle + \sum_{\ell=2}^{\infty} \langle x^{\otimes \ell}, y^{\otimes \ell} \rangle \sum_{\substack{s \in (\mathbb{N} \cup \{0\})^m \\ s_1 + \dots + s_m = \ell}} \alpha_s(E)\alpha_s(F), \quad (19) \end{aligned}$$

for some real coefficients $\{\alpha_s(E)\}_{s \in (\mathbb{N} \cup \{0\})^m}, \{\alpha_s(F)\}_{s \in (\mathbb{N} \cup \{0\})^m} \subseteq \mathbb{R}$.

Proof. Denote $r = \langle x, y \rangle$. Let $g, h \in \mathbb{R}$ be independent standard Gaussian random variables and let $g_1, \dots, g_m \in \mathbb{R}^n$ be i.i.d. standard Gaussian random vectors in \mathbb{R}^n (i.e., they are independent and distributed according to γ_n). Then for each $i \in \{1, \dots, m\}$ the planar random vector $(\langle g_i, x \rangle, \langle g_i, y \rangle) \in \mathbb{R}^2$ has the same distribution as $(g, rg + \sqrt{1-r^2}h) \in \mathbb{R}^2$, and hence its density is given for $(u, v) \in \mathbb{R}^2$ by:

$$f_r(u, v) := \frac{1}{2\pi\sqrt{1-r^2}} \cdot \exp\left(-\frac{u^2 - 2ruv + v^2}{2(1-r^2)}\right).$$

The Hermite polynomials $\{H_k\}_{k=0}^{\infty}$ are defined as:

$$H_k(t) := (-1)^k e^{t^2} \frac{d^k}{dt^k} (e^{-t^2}) = \sum_{s=0}^{\lfloor k/2 \rfloor} \frac{(-1)^s k!}{s!(k-2s)!} (2t)^{k-2s}.$$

The formula for the Poisson kernel for Hermite polynomials (see for example equation 6.1.13 in [1] or the discussion in [14]) says that

$$f_r(u, v) = \frac{e^{-(u^2+v^2)/2}}{2\pi} \sum_{k=0}^{\infty} \frac{r^k}{2^k k!} H_k\left(\frac{u}{\sqrt{2}}\right) H_k\left(\frac{v}{\sqrt{2}}\right).$$

Since the vector $(Gx, Gy) \in \mathbb{R}^{2m}$ has the same distribution as the vector $(\langle g_i, x \rangle, \langle g_i, y \rangle)_{i=1}^m$, whose (planar) entries are i.i.d. with density f_r , we see that:

$$\begin{aligned} \Pr[Gx \in E \wedge Gy \in F] &= \int_{E \times F} \left(\prod_{i=1}^m f_r(u_i, v_i) \right) dudv \\ &= \int_{E \times F} \frac{e^{-(\|u\|_2^2 + \|v\|_2^2)/2}}{(2\pi)^m} \left(\prod_{i=1}^m \left(\sum_{k=0}^{\infty} \frac{r^k}{2^k k!} H_k\left(\frac{u_i}{\sqrt{2}}\right) H_k\left(\frac{v_i}{\sqrt{2}}\right) \right) \right) dudv \\ &= \int_{E \times F} \left(\sum_{s \in (\mathbb{N} \cup \{0\})^m} \frac{r^{s_1 + \dots + s_m}}{2^{s_1 + \dots + s_m} \prod_{i=1}^m s_i!} \left(\prod_{i=1}^m H_{s_i}\left(\frac{u_i}{\sqrt{2}}\right) \right) \left(\prod_{i=1}^m H_{s_i}\left(\frac{v_i}{\sqrt{2}}\right) \right) \right) d\gamma_m(u) d\gamma_m(v) \\ &= \gamma_m(E)\gamma_m(F) + \langle x, y \rangle \left\langle \int_E u d\gamma_m(u), \int_F u d\gamma_m(u) \right\rangle + \sum_{\ell=2}^{\infty} \langle x^{\otimes \ell}, y^{\otimes \ell} \rangle \sum_{\substack{s \in (\mathbb{N} \cup \{0\})^m \\ s_1 + \dots + s_m = \ell}} \alpha_s(E)\alpha_s(F), \end{aligned}$$

where we used the fact that $H_0(t) = 1$ and $H_1(t) = 2t$, and for every measurable subset $W \subseteq \mathbb{R}^m$ and $s \in (\mathbb{N} \cup \{0\})^m$ the notation

$$\alpha_s(W) := \frac{1}{2^{(s_1 + \dots + s_m)/2} \prod_{i=1}^m \sqrt{s_i!}} \int_W \left(\prod_{i=1}^m H_{s_i} \left(\frac{u_i}{\sqrt{2}} \right) \right) d\gamma_m(u).$$

The proof of the identity (14) is complete. \square

Proof of Theorem 3.1. Fix n unit vectors $x_1, \dots, x_n \in S^{n-1}$. Let $\{A_1, \dots, A_k\}$ be a partition of \mathbb{R}^{k-1} into measurable subsets. Let G be a random Gaussian matrix as in Lemma 3.2 with $m = k - 1$. Define a random assignment $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$ by setting $\sigma(i)$ to be the unique $p \in \{1, \dots, k\}$ for which $Gx_i \in A_p$. Then for every $i, j \in \{1, \dots, n\}$ we have

$$\mathbb{E} \left[\langle v_{\sigma(i)}, v_{\sigma(j)} \rangle \right] = \sum_{p=1}^k \sum_{q=1}^k \langle v_p, v_q \rangle \Pr \left[Gx_i \in A_p \wedge Gx_j \in A_q \right] = \sum_{p=1}^k \sum_{q=1}^k b_{pq} \Pr \left[Gx_i \in A_p \wedge Gx_j \in A_q \right].$$

We may therefore apply Lemma 3.2 to deduce that:

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle v_{\sigma(i)}, v_{\sigma(j)} \rangle \right] &= \left(\sum_{i=1}^n \sum_{j=1}^n a_{ij} \right) \sum_{p=1}^k \sum_{q=1}^k b_{pq} \gamma_{k-1}(A_p) \gamma_{k-1}(A_q) \\ &\quad + \left(\sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle \right) \sum_{p=1}^k \sum_{q=1}^k b_{pq} \left\langle \int_{A_p} x d\gamma_{k-1}(x), \int_{A_q} x d\gamma_{k-1}(x) \right\rangle \\ &\quad + \sum_{\ell=2}^{\infty} \left(\sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i^{\otimes \ell}, x_j^{\otimes \ell} \rangle \right) \sum_{\substack{s \in (\mathbb{N} \cup \{0\})^m \\ s_1 + \dots + s_m = \ell}} \sum_{p=1}^k \sum_{q=1}^k b_{pq} \alpha_s(A_p) \alpha_s(A_q) \\ &\geq \left(\sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle \right) \sum_{p=1}^k \sum_{q=1}^k b_{pq} \left\langle \int_{A_p} x d\gamma_{k-1}(x), \int_{A_q} x d\gamma_{k-1}(x) \right\rangle, \end{aligned}$$

where we used the fact that both A and B are positive semidefinite. It thus follows that there exists an assignment $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$ for which

$$\sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle v_{\sigma(i)}, v_{\sigma(j)} \rangle \geq \left(\sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle \right) \sum_{p=1}^k \sum_{q=1}^k b_{pq} \left\langle \int_{A_p} x d\gamma_{k-1}(x), \int_{A_q} x d\gamma_{k-1}(x) \right\rangle,$$

and since this is true for all measurable partitions $\{A_1, \dots, A_k\}$ of \mathbb{R}^{k-1} we deduce that there exists an assignment $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$ for which:

$$\sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle v_{\sigma(i)}, v_{\sigma(j)} \rangle \geq C(B) \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle,$$

as required. \square

3.1 Optimality

The purpose of this section is to show that Theorem 3.1 is sharp:

Theorem 3.3. *Let $v_1, \dots, v_k \in \mathbb{R}^k$ be $k \geq 2$ vectors and let $B = (b_{ij} = \langle v_i, v_j \rangle)$ be the corresponding Gram matrix. Assume that $K > 0$ is a constant such that for every $n \in \mathbb{N}$ and every centered symmetric positive semidefinite matrix $A = (a_{ij}) \in M_n(\mathbb{R})$ we have:*

$$\max_{x_1, \dots, x_n \in \mathcal{S}^{n-1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle \leq K \max_{\sigma: \{1, \dots, n\} \rightarrow \{1, \dots, k\}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle v_{\sigma(i)}, v_{\sigma(j)} \rangle. \quad (20)$$

Then $K \geq \frac{1}{C(B)}$.

Proof. The proof consists of a discretization of a continuous example. The discretization step is somewhat tedious, but straightforward. We will start with a presentation of the continuous example. Fix $m \in \mathbb{N}$ and let $g, h \in \mathbb{R}^m$ be independent standard gaussian random vectors. Since $(\|g\|_2, \|h\|_2)$ is independent of $(\frac{g}{\|g\|_2}, \frac{h}{\|h\|_2})$ we have:

$$\begin{aligned} \int_{\mathbb{R}^m \times \mathbb{R}^m} \langle x, y \rangle \cdot \left\langle \frac{x}{\|x\|_2}, \frac{y}{\|y\|_2} \right\rangle d\gamma_m(x) d\gamma_m(y) &= \mathbb{E} \left[\|g\|_2 \cdot \|h\|_2 \left\langle \frac{g}{\|g\|_2}, \frac{h}{\|h\|_2} \right\rangle^2 \right] \\ &= \mathbb{E} [\|g\|_2 \cdot \|h\|_2] \cdot \mathbb{E} \left[\left\langle \frac{g}{\|g\|_2}, \frac{h}{\|h\|_2} \right\rangle^2 \right] = \mathbb{E} [\|g\|_2]^2 \mathbb{E} \left[\frac{g_1^2}{\|g\|_2^2} \right] = \mathbb{E} [\|g\|_2]^2 \frac{1}{m} \sum_{i=1}^m \mathbb{E} \left[\frac{g_i^2}{\|g\|_2^2} \right] = \frac{1}{m} \mathbb{E} [\|g\|_2]^2, \end{aligned} \quad (21)$$

where we used the rotation invariance of the distribution of h .

The distribution of $\|g\|_2^2$ is the χ^2 distribution with m degrees of freedom, and therefore its density at $u > 0$ equals $\frac{1}{2^{m/2} \Gamma(m/2)} u^{\frac{m}{2}-1} e^{-u/2}$. It follows that

$$\mathbb{E} [\|g\|_2] = \frac{1}{2^{m/2} \Gamma(m/2)} \int_0^\infty \sqrt{u} \cdot u^{\frac{m}{2}-1} e^{-u/2} du = \sqrt{2} \cdot \frac{\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m}{2})} \geq \sqrt{m} \left(1 - O\left(\frac{1}{m}\right) \right), \quad (22)$$

where the last step is an application of Stirling's formula. Plugging (22) into (21) we see that:

$$\int_{\mathbb{R}^m \times \mathbb{R}^m} \langle x, y \rangle \cdot \left\langle \frac{x}{\|x\|_2}, \frac{y}{\|y\|_2} \right\rangle d\gamma_m(x) d\gamma_m(y) \geq 1 - O\left(\frac{1}{m}\right). \quad (23)$$

Now, assuming that $m \geq k - 1$, for every $f: \mathbb{R}^m \rightarrow \{v_1, \dots, v_k\}$ we have

$$\begin{aligned} \int_{\mathbb{R}^m \times \mathbb{R}^m} \langle x, y \rangle \cdot \langle f(x), f(y) \rangle d\gamma_m(x) d\gamma_m(y) &= \left\| \int_{\mathbb{R}^m} x \otimes f(x) d\gamma_m(x) \right\|_2^2 \\ &= \left\| \sum_{i=1}^m e_i \otimes \left(\int_{\mathbb{R}^m} x_i f(x) d\gamma_m(x) \right) \right\|_2^2 = \|\mathbf{Rad}(f)\|_{L_2(\gamma_m, \mathbb{R}^k)}^2 \leq C(B), \end{aligned} \quad (24)$$

where we used Lemma 2.4 (and here e_1, \dots, e_m is the standard basis or \mathbb{R}^m).

We shall now perform a simple discretization argument to conclude the proof of Theorem 3.3. Fix $\varepsilon > 0$ and $M \in \mathbb{N}$. Let \mathcal{F} be the set of all axis parallel cubes in $[-\varepsilon M, \varepsilon M]^m$ which are a product of m intervals

whose endpoints are consecutive integer multiples of ε in $[-M, M]$. Thus $|\mathcal{F}| = (2M)^m$ and each $Q \in \mathcal{F}$ has volume ε^m . For $Q \in \mathcal{F}$ let z_Q be the center of Q . For every $P, Q \in \mathcal{F}$ define

$$a_{PQ} := \varepsilon^{2m} e^{-\frac{\|z_P\|_2^2 + \|z_Q\|_2^2}{2}} \langle z_P, z_Q \rangle.$$

By our assumption (20) there is an assignment $\sigma : \mathcal{F} \rightarrow \{1, \dots, k\}$ such that

$$\sum_{P, Q \in \mathcal{F}} a_{PQ} \left\langle \frac{z_P}{\|z_P\|_2}, \frac{z_Q}{\|z_Q\|_2} \right\rangle \leq K \sum_{P, Q \in \mathcal{F}} a_{PQ} \langle v_{\sigma(P)}, v_{\sigma(Q)} \rangle. \quad (25)$$

We shall now use the following straightforward (and crude) estimates:

$$\begin{aligned} & \left| \int_{\mathbb{R}^m \times \mathbb{R}^m} \langle x, y \rangle \left\langle \frac{x}{\|x\|_2}, \frac{y}{\|y\|_2} \right\rangle d\gamma_m(x) d\gamma_m(y) - \sum_{P, Q \in \mathcal{F}} a_{PQ} \left\langle \frac{z_P}{\|z_P\|_2}, \frac{z_Q}{\|z_Q\|_2} \right\rangle \right| \\ & \leq \sum_{P, Q \in \mathcal{F}} \int_{P \times Q} \left| e^{-\frac{\|z_P\|_2^2 + \|z_Q\|_2^2}{2}} \langle z_P, z_Q \rangle \left\langle \frac{z_P}{\|z_P\|_2}, \frac{z_Q}{\|z_Q\|_2} \right\rangle - e^{-\frac{\|x\|_2^2 + \|y\|_2^2}{2}} \langle x, y \rangle \left\langle \frac{x}{\|x\|_2}, \frac{y}{\|y\|_2} \right\rangle \right| dx dy \\ & \quad + \left| \int_{(\mathbb{R}^m \times \mathbb{R}^m) \setminus ([-\varepsilon M, \varepsilon M]^m \times [-\varepsilon M, \varepsilon M]^m)} \langle x, y \rangle \left\langle \frac{x}{\|x\|_2}, \frac{y}{\|y\|_2} \right\rangle d\gamma_m(x) d\gamma_m(y) \right| \\ & \leq O(1) \sqrt{m} \varepsilon (\sqrt{m} M \varepsilon)^3 \sum_{P, Q \in \mathcal{F}} \int_{P \times Q} e^{-\frac{\|x\|_2^2 + \|y\|_2^2}{2}} dx dy + O(1) m^2 e^{-\frac{(\varepsilon M)^2}{4}} \\ & \leq O(1) \sqrt{m} \varepsilon (\sqrt{m} M \varepsilon)^3 + O(1) m^2 e^{-\frac{(\varepsilon M)^2}{4}}. \end{aligned}$$

We shall require in what follows that $\varepsilon M = 2m$. Hence, using (23) we deduce that:

$$\sum_{P, Q \in \mathcal{F}} a_{PQ} \left\langle \frac{z_P}{\|z_P\|_2}, \frac{z_Q}{\|z_Q\|_2} \right\rangle \geq 1 - O\left(m^5 \varepsilon + \frac{1}{m}\right). \quad (26)$$

On the other hand, define $f : \mathbb{R}^m \rightarrow \{v_1, \dots, v_k\}$ by

$$f(x) = \begin{cases} v_{\sigma(Q)} & x \in Q \in \mathcal{F}, \\ v_1 & x \notin [-\varepsilon M, \varepsilon M]^m. \end{cases}$$

Observe that by symmetry

$$\int_{(\mathbb{R}^m \times \mathbb{R}^m) \setminus ([-\varepsilon M, \varepsilon M]^m \times [-\varepsilon M, \varepsilon M]^m)} \langle x, y \rangle \cdot \langle f(x), f(y) \rangle d\gamma_m(x) d\gamma_m(y) = 0,$$

and therefore a similar crude estimate yields:

$$\begin{aligned} & \left| \int_{\mathbb{R}^m \times \mathbb{R}^m} \langle x, y \rangle \cdot \langle f(x), f(y) \rangle d\gamma_m(x) d\gamma_m(y) - \sum_{P, Q \in \mathcal{F}} a_{PQ} \langle v_{\sigma(P)}, v_{\sigma(Q)} \rangle \right| \\ & \leq \sum_{P, Q \in \mathcal{F}} \int_{P \times Q} \left| e^{-\frac{\|x\|_2^2 + \|y\|_2^2}{2}} \langle x, y \rangle - e^{-\frac{\|z_P\|_2^2 + \|z_Q\|_2^2}{2}} \langle z_P, z_Q \rangle \right| |\langle v_{\sigma(P)}, v_{\sigma(Q)} \rangle| dx dy \\ & \leq O(m^5 \varepsilon) \max_{i \in \{1, \dots, k\}} \|v_i\|_2^2. \quad (27) \end{aligned}$$

Choosing $\varepsilon = m^{-6}$ (and thus $M = 2m^7$), and combining (27) with (24) and (26), yields in combination with (25) the bound:

$$1 - O\left(\frac{1}{m}\right) \leq K \left(C(B) + O\left(\frac{1}{m}\right) \max_{i \in \{1, \dots, k\}} \|v_i\|_2^2 \right).$$

Letting $m \rightarrow \infty$ concludes the proof of Theorem 3.3. \square

4 A sharp approximation algorithm for kernel clustering

Let $A = (a_{ij}) \in M_n(\mathbb{R})$ be a centered symmetric positive semidefinite matrix and let $B = (b_{ij}) \in M_k(\mathbb{R})$ be a symmetric positive semidefinite matrix. Our goal is to design a polynomial time algorithm which approximates the value:

$$\mathbf{Clust}(A|B) = \max_{\sigma: \{1, \dots, n\} \rightarrow \{1, \dots, k\}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} b_{\sigma(i)\sigma(j)}.$$

We proceed as follows. We first find vectors $v_1, \dots, v_k \in \mathbb{R}^k$ such that $b_{ij} = \langle v_i, v_j \rangle$ for all $i, j \in \{1, \dots, k\}$. This can be done in polynomial time (Cholesky decomposition). Let $R(B)$ be the minimum radius of the Euclidean ball in \mathbb{R}^k that contains $\{v_1, \dots, v_k\}$ and let $w(B)$ be the center of this ball. Both $R(B)$ and $w(B)$ can be efficiently computed by solving an appropriate semidefinite program.

We now use semidefinite programming to compute the value:

$$\begin{aligned} \text{SDP}(A|B) &:= \max \left\{ \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle : x_1, \dots, x_n \in \mathbb{R}^n \wedge \|x_i\|_2 \leq 1 \forall i \in \{1, \dots, n\} \right\} \\ &= \max \left\{ \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle : x_1, \dots, x_n \in S^{n-1} \right\}, \end{aligned} \quad (28)$$

where the last equality in (28) holds since the function $(x_1, \dots, x_n) \mapsto \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle$ is convex (by virtue of the fact that A is positive semidefinite). We claim that

$$\frac{\mathbf{Clust}(A|B)}{R(B)^2} \leq \text{SDP}(A|B) \leq \frac{\mathbf{Clust}(A|B)}{C(B)}, \quad (29)$$

which implies that if we output the number $R(B)^2 \mathbf{Clust}(A|B)$ we will obtain a polynomial time algorithm which approximates $\mathbf{Clust}(A|B)$ up to a factor of $\frac{R(B)^2}{C(B)}$.

To verify (29) let $x_1^*, \dots, x_n^* \in S^{n-1}$ and $\sigma^* : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$ be such that

$$\text{SDP}(A|B) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i^*, x_j^* \rangle,$$

and

$$\mathbf{Clust}(A|B) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} b_{\sigma^*(i)\sigma^*(j)}.$$

Write $(a_{ij})_{i,j=1}^n = \langle \langle u_i, u_j \rangle \rangle_{i,j=1}^n$ for some $u_1, \dots, u_n \in \mathbb{R}^n$. The assumption that A is centered means that $\sum_{i=1}^n u_i = 0$. The right-hand side of inequality in (29) is simply a restatement of Theorem 3.1. The left-hand

side inequality (29) follows from the fact that $\frac{v_{\sigma^*(i)} - w(B)}{R(B)}$ has norm at most 1 for all $i \in \{1, \dots, n\}$. Indeed, these norm bounds imply that:

$$\begin{aligned}
\text{SDP}(A|B) &\geq \sum_{i=1}^n \sum_{j=1}^n a_{ij} \left\langle \frac{v_{\sigma^*(i)} - w(B)}{R(B)}, \frac{v_{\sigma^*(j)} - w(B)}{R(B)} \right\rangle \\
&= \frac{1}{R(B)^2} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle v_{\sigma^*(i)}, v_{\sigma^*(j)} \rangle - \frac{2}{R(B)^2} \sum_{i=1}^n \langle w(B), v_{\sigma^*(i)} \rangle \left\langle u_i, \sum_{j=1}^n u_j \right\rangle + \frac{\|w(B)\|_2^2}{R(B)^2} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \\
&= \frac{\mathbf{Clust}(A|B)}{R(B)^2}.
\end{aligned}$$

This completes the proof that our algorithm approximates efficiently the number $\mathbf{Clust}(A|B)$, but does not address the issue of how to efficiently compute an assignment $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$ for which the induced clustering of A has the required value. An inspection of the proof of Theorem 3.1 shows that the issue here is to find efficiently a conical simplicial partition A_1, \dots, A_k of \mathbb{R}^{k-1} at which $C(B)$ is almost attained, say

$$\sum_{p=1}^k \sum_{q=1}^k b_{pq} \left\langle \int_{A_p} x d\gamma_{k-1}(x), \int_{A_q} x d\gamma_{k-1}(x) \right\rangle \geq (1 - \varepsilon)C(B).$$

Once this partition is computed, using the notation in the proof of Theorem 3.1 we have a randomized algorithm which outputs an assignment $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$ such that

$$\mathbb{E}_\sigma \left[\sum_{i=1}^n \sum_{j=1}^n a_{ij} b_{\sigma(i)\sigma(j)} \right] \geq \frac{(1 - \varepsilon)C(B)}{R(B)^2} \mathbf{Clust}(A|B).$$

Note that there is no difficulty to compute σ efficiently once the partition $\{A_1, \dots, A_k\}$ is given, since these sets are simplicial cones. The issue with efficiency here is how to compute this partition in polynomial time. As we discussed in Remark 2.3, this can be done when k is fixed (or grows very slowly with n), but we do not know how to do this when, say, $k = \sqrt{n}$.

5 Matching Unique Games hardness

In this section we show that for a fixed positive semi-definite matrix B , approximating $\mathbf{Clust}(A|B)$ within a ratio strictly smaller than $\frac{R(B)^2}{C(B)}$ is Unique Games hard. We will study functions $f : \{1, \dots, k\}^n \rightarrow \mathbb{R}$ and their Fourier spectrum at the first level. A novel feature of our proof is that our Fourier analysis will be carried out with respect to a distribution on $\{1, \dots, k\}$ that is not necessarily uniform. In fact, the choice of the distribution itself is dictated by the matrix B as described in Section 5.1.

5.1 Choosing a special probability distribution on $\{1, \dots, k\}$

Fact 5.1. *Let $B = (b_{ij})$ be a $k \times k$ symmetric positive semi-definite matrix and $b_{ij} = \langle v_i, v_j \rangle$ be its Gram representation, where v_1, \dots, v_k are vectors (w.l.o.g.) in \mathbb{R}^k . Let $R(B)$ be the minimum radius of a Euclidean ball containing all these vectors, and $w(B)$ be the center of this ball. Then $w(B)$ is a convex combination of the v_i 's that are on the boundary of the ball. In other words, there exist non-negative coefficients $p(1), \dots, p(k)$ such that $\sum_{i=1}^k p(i) = 1$, $w(B) = \sum_{i=1}^k p(i)v_i$ and $p(i) \neq 0$ only if $\|v_i - w(B)\|_2 = R(B)$.*

Fact 5.1 is well known (see for example the proof of Proposition 1.13 in [2]). Its proof is a simple separation argument. Indeed, define $J := \{j \in \{1, \dots, k\} : \|v_j - w(B)\|_2 = R(B)\}$ and let K be the convex hull of $\{v_j\}_{j \in J}$. Assume for the sake of contradiction that $w(B) \notin K$. Then there would be a hyperplane H separating $w(B)$ from K . Moving $w(B)$ a little in the direction of H would turn the equalities on J to strict inequalities, while preserving the strict inequalities off J . This contradicts the minimality of $R(B)$.

We intend to use the probability distribution $(p(1), \dots, p(k))$ from fact 5.1. However, for technical reasons, we need the probability mass for each atom to be non-zero, and therefore, we will use a very small perturbation of this distribution. Towards this end we define $\mu(i) = (1 - \beta)p(i) + \frac{\beta}{k}$ for every $i \in \{1, \dots, k\}$. The value of $\beta > 0$ is chosen to be sufficiently small as in the following lemma.

Lemma 5.2. *Fix any $\varepsilon > 0$ and the matrix B . Then for a sufficiently small $\beta = \beta(\varepsilon, B) > 0$,*

$$\sum_{i=1}^k \mu(i) \left\| v_i - \sum_{j=1}^k \mu(j)v_j \right\|_2^2 \geq R(B)^2 - \varepsilon. \quad (30)$$

Proof. Note that if $\beta = 0$, then $\mu(i) = p(i)$ for all $i \in \{1, \dots, k\}$, and

$$\sum_{i=1}^k \mu(i) \left\| v_i - \sum_{j=1}^k \mu(j)v_j \right\|_2^2 = \sum_{i=1}^k p(i) \|v_i - w(B)\|_2^2 = R(B)^2,$$

since $p(i) \neq 0$ only if $\|v_i - w(B)\|_2 = R(B)$. Thus by continuity for sufficiently small β the inequality (30) holds. For concreteness we also give a direct argument which gives a reasonable bound on β . Assume that $\beta < \frac{1}{7}$. Then, using the fact that $\mu \geq (1 - \beta)p$ (point-wise), we see that:

$$\begin{aligned} \left(\sum_{i=1}^k \mu(i) \left\| v_i - \sum_{j=1}^k \mu(j)v_j \right\|_2^2 \right)^{1/2} &\geq \sqrt{1 - \beta} \left(\sum_{i=1}^k p(i) \left\| (1 - \beta) \left(v_i - \sum_{j=1}^k p(j)v_j \right) + \frac{\beta}{k} \sum_{j=1}^k (v_i - v_j) \right\|_2^2 \right)^{1/2} \\ &\geq \sqrt{1 - \beta} \left(\sum_{i=1}^k p(i) \|(1 - \beta)(v_i - w(B))\|_2^2 \right)^{1/2} - \sqrt{1 - \beta} \left(\sum_{i=1}^k p(i) \left\| \frac{\beta}{k} \sum_{j=1}^k (v_i - v_j) \right\|_2^2 \right)^{1/2} \\ &\geq (1 - \beta)^{3/2} R(B) - \beta \sqrt{1 - \beta} \left(\sum_{i=1}^k p(i) \frac{1}{k} \sum_{j=1}^k \|v_i - v_j\|_2^2 \right)^{1/2} \\ &\geq (1 - \beta)^{3/2} R(B) - \beta \sqrt{1 - \beta} \max_{i,j \in \{1, \dots, k\}} \|v_i - v_j\|_2 \\ &\geq \sqrt{1 - \beta} (1 - 3\beta) R(B) \\ &\geq \sqrt{1 - 7\beta} \cdot R(B), \end{aligned}$$

where in the penultimate inequality we used the trivial fact that $\max_{i,j \in \{1, \dots, k\}} \|v_i - v_j\|_2 \leq 2R(B)$. Thus we can take $\beta = \frac{\varepsilon}{7R(B)^2}$ to ensure the validity of (30). \square

Henceforth we fix the probability space $(\Omega = \{1, \dots, k\}, \mu)$. Let $U = (u_{ij})$ be a $k \times k$ orthogonal matrix such that $u_{1j} = \sqrt{\mu(j)}$ for all $j \in \{1, \dots, k\}$ (such an orthogonal matrix exists since this ensures that $\sum_{j=1}^k u_{1j}^2 = 1$). Now define random variables $X_1, \dots, X_k : \{1, \dots, k\} \rightarrow \mathbb{R}$ by $X_i(j) = \frac{u_{ij}}{\sqrt{\mu(j)}}$ (here is one place where we need the atoms of μ to have positive mass. We will also use this fact to allow for the application

of the result of [9] in the proof of Theorem 5.4 below). Then by design X_1 is the constant 1 function, and for all $i, j \in \{1, \dots, k\}$ we have:

$$\sum_{\ell=1}^k \mu(\ell) X_i(\ell) X_j(\ell) = \sum_{\ell=1}^k u_{i\ell} u_{j\ell} = (UU^t)_{ij} = \delta_{ij},$$

where δ_{ij} is the Kronecker delta. Similarly:

$$\sum_{\ell=1}^k X_\ell(i) X_\ell(j) = \frac{1}{\sqrt{\mu(i)\mu(j)}} \sum_{\ell=1}^k u_{\ell i} u_{\ell j} = \frac{(U^t U)_{ij}}{\sqrt{\mu(i)\mu(j)}} = \frac{\delta_{ij}}{\mu(i)}.$$

By relabeling these random variables (for the sake for simplicity of later notation) we thus obtain the following lemma:

Lemma 5.3. *There exist random variables X_0, X_1, \dots, X_{k-1} on Ω such that:*

- $X_0 \equiv 1$.
- For $i, j \in \{0, \dots, k-1\}$ we have

$$\mathbb{E}_\mu[X_i X_j] = \begin{cases} 0 & \text{if } i \neq j, \\ 1 & \text{if } i = j. \end{cases}$$

- For every $\omega, \omega' \in \Omega$ we have

$$\sum_{i=0}^{k-1} X_i(\omega) X_i(\omega') = \begin{cases} 0 & \text{if } \omega \neq \omega', \\ \frac{1}{\mu(\omega)} & \text{if } \omega = \omega'. \end{cases}$$

5.2 Dictatorships vs. functions with small influences

In this section we will associate to every function from $\{1, \dots, k\}^n$ to

$$\Delta_k := \left\{ x \in \mathbb{R}^k : x_i \geq 0 \forall i \in \{1, \dots, k\}, \sum_{i=1}^k x_i = 1 \right\}$$

a numerical parameter, or “objective value”. We will show that the value of this parameter for functions which depend only on a single coordinate (i.e. dictatorships) differs markedly from its value on functions which do not depend significantly on any particular coordinate (i.e. functions with small influences). This step is an analog of the “dictatorship test” which is prevalent in PCP based hardness proofs.

We begin with some notation and preliminaries on Fourier-type expansions. For any function $f : \mathbb{R}^n \rightarrow \Delta_k$ we write $f = (f_1, f_2, \dots, f_k)$ where $f_i : \mathbb{R}^n \rightarrow [0, 1]$ and $\sum_{i=1}^k f_i = 1$. With this notation we have

$$C(B) = \sup_{f: \mathbb{R}^{k-1} \rightarrow \Delta_k} \sum_{i=1}^k \sum_{j=1}^k b_{ij} \left\langle \int_{\mathbb{R}^{k-1}} x f_i(x) d\gamma_{k-1}(x), \int_{\mathbb{R}^{k-1}} x f_j(x) d\gamma_{k-1}(x) \right\rangle$$

where $C(B)$ is as in Section 2. We have already seen that the supremum above is actually attained. Also $C(B)$ remains the same if the supremum is taken over functions over \mathbb{R}^n with $n \geq k-1$, i.e. for every $n \geq k-1$,

$$C(B) = \sup_{f: \mathbb{R}^n \rightarrow \Delta_k} \sum_{i=1}^k \sum_{j=1}^k b_{ij} \left\langle \int_{\mathbb{R}^n} x f_i(x) d\gamma_n(x), \int_{\mathbb{R}^n} x f_j(x) d\gamma_n(x) \right\rangle.$$

Let $(\Omega = \{1, \dots, k\}, \mu)$ be the probability space as chosen in Section 5.1. Let (Ω^n, μ^n) be the associated product space. We will be analyzing functions $f : \Omega^n \rightarrow \Delta_k$ (and more generally into \mathbb{R}^k). As in Lemma 5.3, fix a basis of orthonormal random variables on Ω where one of them is the constant 1 function, that is $\{X_0 \equiv 1, X_1, \dots, X_{k-1}\}$. Then any function $f : \Omega \rightarrow \mathbb{R}$ can be written as a linear combination of the X_i 's.

In order to analyze functions $f : \Omega^n \rightarrow \mathbb{R}$, we let $\mathcal{X} = (\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_n)$ be an “ensemble” of random variables where for $i \in \{1, \dots, n\}$ we write $\mathcal{X}_i = \{X_{i,0}, X_{i,1}, \dots, X_{i,k-1}\}$, and for every i , $\{X_{i,j}\}_{j=0}^{k-1}$ are independent copies of the $\{X_j\}_{j=0}^{k-1}$. Any $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n) \in \{0, 1, 2, \dots, k-1\}^n$ will be called a multi-index. We shall denote by $|\sigma|$ the number on non-zero entries in σ . Each multi-index defines a monomial

$$x_\sigma := \prod_{\substack{i \in \{1, \dots, n\} \\ \sigma_i \neq 0}} x_{i, \sigma_i}$$

on a set of $n(k-1)$ indeterminates $\{x_{ij} \mid i \in \{1, \dots, n\}, j \in \{1, 2, \dots, k-1\}\}$, and also a random variable $X_\sigma : \Omega^n \rightarrow \mathbb{R}$ as

$$X_\sigma(\omega) := \prod_{i=1}^n X_{i, \sigma_i}(\omega_i).$$

The random variables $\{X_\sigma\}_\sigma$ form an orthonormal basis for the space of functions $f : \Omega^n \rightarrow \mathbb{R}$. Thus, every such f can be written uniquely as (the “Fourier expansion”)

$$f = \sum_{\sigma} \widehat{f}(\sigma) X_\sigma, \quad \widehat{f}(\sigma) \in \mathbb{R}.$$

We denote the corresponding multi-linear polynomial as $Q_f = \sum_{\sigma} \widehat{f}(\sigma) x_\sigma$. One can think of f as the polynomial Q_f applied to the ensemble \mathcal{X} , i.e. $f = Q_f(\mathcal{X})$. Of course, one can also apply Q_f to any other ensemble, and specifically to the Gaussian ensemble $\mathcal{G} = (\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_n)$ where $\mathcal{G}_i = \{G_{i,0} \equiv 1, G_{i,1}, \dots, G_{i,k-1}\}$ and $G_{i,j}, i \in \{1, \dots, n\}, j \in \{1, \dots, k-1\}$ are i.i.d. standard Gaussians. Define the influence of the i 'th variable on f as

$$\text{Inf}_i(f) := \sum_{\sigma_i \neq 0} \widehat{f}(\sigma)^2.$$

Roughly speaking, the results of [12, 9] say that if $f : \Omega^n \rightarrow [0, 1]$ is a function all of whose influences are small, then $f = Q_f(\mathcal{X})$ and $Q_f(\mathcal{G})$ are almost identically distributed, and in particular, the values of $Q_f(\mathcal{G})$ are essentially contained in $[0, 1]$. Note that $Q_f(\mathcal{G})$ is a random variable on the probability space $(\mathbb{R}^{n(k-1)}, \gamma_{n(k-1)})$.

Consider functions $f : \Omega^n \rightarrow \Delta_k$. We write $f = (f_1, f_2, \dots, f_k)$ where $f_i : \Omega^n \rightarrow [0, 1]$ with $\sum_{i=1}^k f_i = 1$. Each f_i has a unique representation (along with the corresponding multi-linear polynomial)

$$f_i = \sum_{\sigma} \widehat{f}_i(\sigma) X_\sigma, \quad Q_i := Q_{f_i} = \sum_{\sigma} \widehat{f}_i(\sigma) x_\sigma.$$

We shall define an objective function $\text{OBJ}(f)$ that is a positive semidefinite quadratic form on the table of values of f which corresponds to a centered symmetric positive semidefinite bilinear form. Then we analyze the value of this objective function when f is a “dictatorship” versus when f has all low influences.

The objective value

For a function $f : \Omega^n \rightarrow \Delta_k$ (or more generally, $f : \Omega^n \rightarrow \mathbb{R}^k$) define

$$\text{OBJ}(f) := \sum_{i=1}^k \sum_{j=1}^k b_{ij} \left(\sum_{\sigma: |\sigma|=1} \widehat{f}_i(\sigma) \widehat{f}_j(\sigma) \right). \quad (31)$$

Note that there are $n(k-1)$ multi-indices σ such that $|\sigma| = 1$.

The objective value for dictatorships

For $\ell \in \{1, \dots, n\}$ we define a dictatorship function $f^{dict, \ell} : \Omega^n \rightarrow \Delta_k$ as follows. The range of the function is limited to only k points in Δ_k , namely the points $\{e_1, e_2, \dots, e_k\}$ where e_i is a vector with i^{th} coordinate 1 and all other coordinates zero.

$$f^{dict, \ell}(\omega) := e_i \quad \text{if } \omega_\ell = i. \quad (32)$$

In other words, when one writes $f^{dict, \ell} = (f_1, f_2, \dots, f_k)$, for $i \in \{1, \dots, k\}$, f_i is $\{0, 1\}$ -valued and $f_i(\omega) = 1$ if and only if $\omega_\ell = i$. The Fourier expansion of f_i is

$$f_i(\omega) = \mu(i) \sum_{\sigma: \sigma_j=0 \forall j \neq \ell} X_{\sigma_\ell}(i) X_\sigma(\omega). \quad (33)$$

Indeed, the right hand side of (33) equals

$$\mu(i) \sum_{0 \leq \sigma_\ell \leq k-1} X_{\sigma_\ell}(i) X_{\sigma_\ell}(\omega_\ell) = \begin{cases} 1 & \text{if } \omega_\ell = i, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{see Lemma 5.3})$$

Thus,

$$\begin{aligned} \text{OBJ}(f^{dict, \ell}) &= \sum_{i=1}^k \sum_{j=1}^k b_{ij} \left(\sum_{\sigma: |\sigma|=1} \widehat{f}_i(\sigma) \widehat{f}_j(\sigma) \right) \\ &= \sum_{i=1}^k \sum_{j=1}^k b_{ij} \left(\sum_{r=1}^{k-1} \mu(i) X_r(i) \mu(j) X_r(j) \right) \\ &= \sum_{i=1}^k \sum_{j=1}^k b_{ij} \cdot \mu(i) \mu(j) \left(\sum_{r=0}^{k-1} X_r(i) X_r(j) - 1 \right) \\ &= \sum_{\substack{i, j \in \{1, \dots, k\} \\ i \neq j}} \langle v_i, v_j \rangle \cdot \mu(i) \mu(j) (-1) + \sum_{i=1}^k \langle v_i, v_i \rangle \cdot \mu(i)^2 \left(\frac{1}{\mu(i)} - 1 \right) \\ &= \sum_{i=1}^k \mu(i) \left\| v_i - \sum_{j=1}^k \mu(j) v_j \right\|_2^2 \\ &\geq R(B)^2 - \varepsilon, \end{aligned} \quad (34)$$

using Lemma 5.2.

The objective value for functions with low influences

For $f : \Omega^n \rightarrow \mathbb{R}$, $j \in \{1, \dots, n\}$ and $m \in \mathbb{N}$ denote (the “degree m -influence” of f):

$$\text{Inf}_j^{\leq m}(f) := \sum_{\substack{|\sigma| \leq m \\ \sigma_j \neq 0}} \widehat{f}(\sigma)^2.$$

For every $0 \leq \rho \leq 1$ we will use the smoothing operator:

$$T_\rho f = \sum_{\sigma} \rho^{|\sigma|} \widehat{f}(\sigma) X_\sigma.$$

Equivalently,

$$T_\rho f(\omega_1, \dots, \omega_n) = \mathbb{E}[f(\omega'_1, \dots, \omega'_n)],$$

where independently for each i , ω'_i is chosen to be ω_i with probability ρ and a random (with respect to the underlying distribution μ) element in Ω with probability $1 - \rho$.

The following theorem is the key analytic fact used in our UGC hardness result:

Theorem 5.4. *For every $\varepsilon > 0$, there exists $\tau > 0$ so that the following holds: for any function $f : \Omega^n \rightarrow \Delta_k$ which satisfies*

$$\forall i \in \{1, \dots, k\}, \forall j \in \{1, \dots, n\}, \quad \text{Inf}_j^{\leq \log(1/\tau)}(f_i) \leq \tau$$

we have,

$$\text{OBJ}(f) \leq C(B) + \varepsilon.$$

Proof. Let $\delta, \eta > 0$ be sufficiently small constants to be chosen later. Let $Q_i = Q_{f_i}$ be the multi-linear polynomial associated with f_i . Recall that Q_i is a multi-linear polynomial in the $n(k-1)$ indeterminates $\{x_{jp} \mid j \in \{1, \dots, n\}, p \in \{1, \dots, k-1\}\}$. Moreover $f_i = Q_i(\mathcal{X})$ has range $[0, 1]$ and $\sum_{i=1}^k f_i = 1$.

Let $R_i = (T_{1-\delta} Q_i)(\mathcal{X})$ and $S_i = (T_{1-\delta} Q_i)(\mathcal{G})$ (the smoothening operator $T_{1-\delta}$ helps us meet some technical pre-conditions before applying the invariance principle of [9]). Note that R_i has range $[0, 1]$ and S_i has range \mathbb{R} . It will follow however from [9] that S_i is *essentially* in $[0, 1]$. First we relate $\text{OBJ}(f)$ to the functions S_i which will, up to truncation, induce a partition of $\mathbb{R}^{n(k-1)}$, which in turn will give the bound in terms of $C(B)$.

$$\begin{aligned} (1-\delta)^2 \cdot \text{OBJ}(f) &= (1-\delta)^2 \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \sum_{\sigma:|\sigma|=1} \widehat{f}_i(\sigma) \widehat{f}_\ell(\sigma) \\ &= (1-\delta)^2 \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \sum_{j=1}^n \sum_{p=1}^{k-1} \left(\int_{\mathbb{R}^{n(k-1)}} x_{jp} Q_i(x) d\gamma_{n(k-1)}(x) \right) \cdot \left(\int_{\mathbb{R}^{n(k-1)}} x_{jp} Q_\ell(x) d\gamma_{n(k-1)}(x) \right) \\ &= (1-\delta)^2 \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \left\langle \int_{\mathbb{R}^{n(k-1)}} x Q_i(x) d\gamma_{n(k-1)}(x), \int_{\mathbb{R}^{n(k-1)}} x Q_\ell(x) d\gamma_{n(k-1)}(x) \right\rangle \\ &= \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \left\langle \int_{\mathbb{R}^{n(k-1)}} x (T_{1-\delta} Q_i)(x) d\gamma_{n(k-1)}(x), \int_{\mathbb{R}^{n(k-1)}} x (T_{1-\delta} Q_\ell)(x) d\gamma_{n(k-1)}(x) \right\rangle \\ &= \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \left\langle \int_{\mathbb{R}^{n(k-1)}} x S_i(x) d\gamma_{n(k-1)}(x), \int_{\mathbb{R}^{n(k-1)}} x S_\ell(x) d\gamma_{n(k-1)}(x) \right\rangle. \end{aligned} \tag{35}$$

We shall now bound the last term above by $C(B) + o(1)$. For any real-valued function h on $\mathbb{R}^{n(k-1)}$, let

$$\text{chop}(h)(x) := \begin{cases} 0 & \text{if } h(x) < 0, \\ h(x) & \text{if } h(x) \in [0, 1], \\ 1 & \text{if } h(x) > 1. \end{cases}$$

Applying Theorem 3.20 in [9] to the polynomial Q_i , it follows that (provided τ is sufficiently small compared to δ and η),

$$\|S_i - \text{chop}(S_i)\|_{L_2(\gamma_{n(k-1)})}^2 = \int_{\mathbb{R}^{n(k-1)}} |S_i(x) - \text{chop}(S_i)(x)|^2 d\gamma_{n(k-1)}(x) \leq \eta. \quad (36)$$

The functions $\text{chop}(S_i)$ are almost what we want except that they might not sum up to 1. So further define

$$S_i^*(x) := \frac{\text{chop}(S_i)(x)}{\sum_{i=1}^k \text{chop}(S_i)(x)}.$$

Clearly, $\{S_i^*\}_{i=1}^k$ have range $[0, 1]$ and $\sum_{i=1}^k S_i^* \equiv 1$. Observe that the following holds point-wise:

$$\sum_{j=1}^k |\text{chop}(S_j) - S_j^*| = \left| \sum_{j=1}^k \text{chop}(S_j) - 1 \right| = \left| \sum_{j=1}^k \text{chop}(S_j) - \sum_{j=1}^k S_j \right| \leq \sum_{j=1}^k |S_j - \text{chop}(S_j)|,$$

where we used that $\sum_{j=1}^k S_j = T_{1-\delta} \sum_{j=1}^k Q_j = T_{1-\delta} 1 = 1$. It follows that for all $i \in \{1, \dots, k\}$ we have:

$$\|\text{chop}(S_i) - S_i^*\|_{L_2(\gamma_{n(k-1)})} \leq \sum_{j=1}^k \|\text{chop}(S_j) - S_j^*\|_{L_2(\gamma_{n(k-1)})} \leq \sum_{j=1}^k \|S_j - \text{chop}(S_j)\|_{L_2(\gamma_{n(k-1)})} \leq k\sqrt{\eta},$$

where we used (36). Finally,

$$\|S_i - S_i^*\|_{L_2(\gamma_{n(k-1)})} \leq \|S_i - \text{chop}(S_i)\|_{L_2(\gamma_{n(k-1)})} + \|\text{chop}(S_i) - S_i^*\|_{L_2(\gamma_{n(k-1)})} \leq (k+1)\sqrt{\eta}. \quad (37)$$

Now write

$$u_i = \int_{\mathbb{R}^{n(k-1)}} x S_i(x) d\gamma_{n(k-1)}(x), \quad w_i = \int_{\mathbb{R}^{n(k-1)}} x S_i^*(x) d\gamma_{n(k-1)}(x). \quad (38)$$

The norm of $u_i - w_i$ is bounded by $(k+1)\sqrt{\eta}$ using (37) and Lemma 5.5 below. Since $|S_i^*| \leq 1$, the norm of w_i is bounded by 1. Returning to the estimation in Equation (35) and applying Lemma 5.6 below, we see that:

$$(1-\delta)^2 \cdot \text{OBJ}(f) = \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \langle u_i, u_\ell \rangle \leq \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \langle w_i, w_\ell \rangle + O(k\sqrt{\eta}) \left(\sum_{i=1}^k \sum_{\ell=1}^k |b_{i\ell}| \right).$$

Since $\sum_{i=1}^k S_i^* \equiv 1$ we have

$$\begin{aligned} \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \langle w_i, w_\ell \rangle &= \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \left\langle \int_{\mathbb{R}^{n(k-1)}} x S_i^*(x) d\gamma_{n(k-1)}(x), \int_{\mathbb{R}^{n(k-1)}} x S_\ell^*(x) d\gamma_{n(k-1)}(x) \right\rangle \\ &\leq \sup_{f: \mathbb{R}^{n(k-1)} \rightarrow \Delta_k} \left(\sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \left\langle \int_{\mathbb{R}^{n(k-1)}} x f_i(x) d\gamma_{n(k-1)}(x), \int_{\mathbb{R}^{n(k-1)}} x f_\ell(x) d\gamma_{n(k-1)}(x) \right\rangle \right) = C(B). \end{aligned}$$

It follows that $\text{OBJ}(f) \leq C(B) + \varepsilon$, provided that η and δ are small enough. \square

Lemma 5.5. Let $g \in L_2(\mathbb{R}^n, \gamma_n)$. Then

$$\left\| \int_{\mathbb{R}^n} x g(x) d\gamma_n(x) \right\|_2 \leq \|g\|_{L_2(\mathbb{R}^n, \gamma_n)}.$$

Proof. Note that the square of the left hand side equals

$$\sum_{i=1}^n \left| \int_{\mathbb{R}^n} x_i g(x) d\gamma_n(x) \right|^2 = \sum_{i=1}^n \langle x_i, g \rangle^2.$$

Since $x_i \in L_2(\mathbb{R}^n, \gamma_n)$ are an orthonormal set of functions, the sum of squares of projections of g onto them is at most the squared norm of g . \square

Lemma 5.6. Suppose $\{u_i\}_{i=1}^k$ and $\{w_i\}_{i=1}^k$ are vectors in \mathbb{R}^n such that $\|u_i - w_i\|_2 \leq d$ for every $i \in \{1, \dots, k\}$ and $\|w_i\|_2 \leq 1$. Let $B = (b_{ij})$ be a $k \times k$ matrix. Then

$$\left| \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \langle u_i, u_\ell \rangle - \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \langle w_i, w_\ell \rangle \right| \leq (2d + d^2) \sum_{i=1}^k \sum_{\ell=1}^k |b_{i\ell}|.$$

Proof. From the given conditions on the norms of $a_i = u_i - w_i$ and w_i , it follows that for any $i, \ell \in \{1, \dots, k\}$,

$$|\langle u_i, u_\ell \rangle - \langle w_i, w_\ell \rangle| \leq |\langle a_i, w_\ell \rangle| + |\langle a_\ell, w_i \rangle| + |\langle a_i, a_\ell \rangle| \leq 2d + d^2.$$

Hence,

$$\left| \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \langle u_i, u_\ell \rangle - \sum_{i=1}^k \sum_{\ell=1}^k b_{i\ell} \langle w_i, w_\ell \rangle \right| \leq \sum_{i=1}^k \sum_{\ell=1}^k |b_{i\ell}| |\langle u_i, u_\ell \rangle - \langle w_i, w_\ell \rangle| \leq (2d + d^2) \sum_{i=1}^k \sum_{\ell=1}^k |b_{i\ell}|,$$

as required. \square

The intended hardness factor

As we show next, the dictatorship test can be translated (in a more or less standard way by now) into a Unique Games hardness result. The hardness factor (as usual) turns out to be the ratio of the objective value when the function is a dictatorship versus when the function has all low influences, i.e.

$$\frac{R(B)^2 - \varepsilon}{C(B) + \varepsilon} = \frac{R(B)^2}{C(B)} - o(1).$$

5.3 The reduction from unique games to kernel clustering

Given a Unique Games Instance $\mathcal{L}(G(V, W, E), n, \{\pi_{vw}\}_{(v,w) \in E})$, we construct an instance of the clustering problem.

Reformulation of the clustering problem

As in our earlier paper [7], we first reformulate the kernel clustering problem for the ease of presentation. As observed there, we can reformulate it as (the matrix A in the problem **Clust**($A|B$) is captured by the quadratic form Q below):

Kernel Clustering Problem: Given a $k \times k$ symmetric positive semidefinite matrix B , and a symmetric positive semidefinite quadratic form $Q(\cdot, \cdot)$ on $\mathbb{R}^N \times \mathbb{R}^N$, find $F : \{1, \dots, N\} \rightarrow \Delta_k$, $F = (F_1, F_2, \dots, F_k)$, so as to maximize $\sum_{i=1}^k \sum_{j=1}^k b_{ij} Q(F_i, F_j)$.

The clustering problem instance

Given a Unique Games instance $\mathcal{L}(G(V, W, E), n, \{\pi_{vw}\}_{(v,w) \in E})$, the clustering problem is to find a function $F : W \times \Omega^n \rightarrow \Delta_k$ so as to maximize $\sum_{i=1}^k \sum_{j=1}^k b_{ij} Q(F_i, F_j)$ where Q is a suitably defined symmetric positive semidefinite quadratic form. For notational convenience, we write:

$$F_w := F(w, \cdot), \quad F_w : \Omega^n \rightarrow \Delta_k.$$

Also, for every $v \in V$, we write:

$$F_v := \mathbb{E}_{(v,w) \in E} [F_w \circ \pi_{vw}], \quad F_v : \Omega^n \rightarrow \Delta_k.$$

We used the following notation: for any function $g : \Omega^n \rightarrow \Delta_k$ and $\pi : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ we write $g \circ \pi : \Omega^n \rightarrow \Delta_k$ for the function $(g \circ \pi)(\omega) := g(\omega_{\pi(1)}, \omega_{\pi(2)}, \dots, \omega_{\pi(n)})$. As usual, we denote $F_w = (F_{w,1}, F_{w,2}, \dots, F_{w,k})$ where each $F_{w,i}$ has range $[0, 1]$ and $\sum_{i=1}^k F_{w,i} = 1$. Similarly, $F_v = (F_{v,1}, F_{v,2}, \dots, F_{v,k})$ and $\sum_{i=1}^k F_{v,i} = 1$. Now we are ready to define the clustering problem instance.

Clustering instance: The goal is to find $F : W \times \Omega^n \rightarrow \Delta_k$ so as to maximize:

$$\max_{F: W \times \Omega^n \rightarrow \Delta_k} \mathbb{E}_{v \in V} [\text{OBJ}(F_v)] = \max_{F: W \times \Omega^n \rightarrow \Delta_k} \mathbb{E}_{v \in V} \left[\sum_{i=1}^k \sum_{j=1}^k b_{ij} \sum_{\sigma: |\sigma|=1} \widehat{F}_{v,i}(\sigma) \cdot \widehat{F}_{v,j}(\sigma) \right]. \quad (39)$$

Completeness

We will show that if the Unique Games instance has an almost satisfying labeling, then the objective value of the clustering problem is at least $R(B)^2 - o(1)$. So, let $\rho : V \cup W \rightarrow \{1, \dots, n\}$ be the labeling, such that for at least $1 - \varepsilon$ fraction of the vertices $v \in V$ (call such v good) we have

$$\pi_{vw}(\rho(w)) = \rho(v) \quad \forall (v, w) \in E.$$

Define $F : W \times \Omega^n \rightarrow \Delta_k$ as follows: for every $w \in W$, $F_w : \Omega^n \rightarrow \Delta_k$ equals the dictatorship corresponding to $\rho(w) \in \{1, \dots, n\}$, i.e.,

$$F_w := f^{\text{dict}, \rho(w)}.$$

Lemma 5.7 ([7]). *For a good $v \in V$ we have $F_v = f^{\text{dict}, \rho(v)}$.*

Thus the contribution of v in (39) is $\text{OBJ}(f^{\text{dict}, \rho(v)}) \geq R(B)^2 - \varepsilon$ as observed in Equation (34). Since $1 - \varepsilon$ fraction of $v \in V$ are good, (39) is at least $(1 - \varepsilon) \cdot (R(B)^2 - \varepsilon) = R(B)^2 - o(1)$.

Soundness

Suppose for the sake of contradiction that the value of (39) is at least $C(B) + 2\varepsilon$. As in [7], it can be proved that the Unique Games instance must have a labeling that satisfies at least a constant fraction of its edges, the constant depending on the parameter τ used in Theorem 5.4. This is a contradiction, provided the soundness of the Unique Games instance is chosen to be even lower to begin with. The proof is the same as in [7], by replacing the $C(k)$ therein by $C(B)$ ([7] focused on the case when B is the $k \times k$ identity matrix. The constant $C(k)$ therein is same as our constant $C(B)$ when B is the $k \times k$ identity matrix).

6 A concrete example

In this section we will use our results to evaluate the UGC hardness threshold of the problem of computing

$$\mathbf{Clust} \left(A \left| \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & c \end{pmatrix} \right. \right), \quad (40)$$

where $A \in M_n(\mathbb{R})$ is centered, symmetric and positive semidefinite and $c \in (0, \infty)$ is a parameter. The case $c = 1$, corresponding to $B = I_3$ (the 3×3 identity matrix) was evaluated in [7], where it was shown that the UGC hardness threshold in this case equals $\frac{16\pi}{27}$.

For general $c > 0$ the optimization problem in (40) corresponds to the following question: given n random variables X_1, \dots, X_n the goal is to partition them into three sets $S_1, S_2, S_3 \subseteq \{1, \dots, n\}$ such that

$$\sum_{i,j \in S_1} \mathbb{E}[X_i X_j] + \sum_{i,j \in S_2} \mathbb{E}[X_i X_j] + c \sum_{i,j \in S_3} \mathbb{E}[X_i X_j] \quad (41)$$

is maximized. Thus we wish to cluster the variables into three clusters so as to maximize the intra-cluster correlations, while the parameter c allows us to tune the relative importance of one of the clusters. We stress that we do not claim that this optimization problem is of particular intrinsic importance. We chose it as a way to concretely demonstrate our results for the simplest possible perturbation of the case of $B = I_3$. We remark that it is also possible to explicitly solve the case of general 3×3 diagonal matrices B , i.e., the case of a general weighting of the clusters in (41). The formula for the UGC hardness threshold for general 3×3 diagonal matrices turns out to be quite complicated, so we chose to deal only with (40) as a simple example for the sake of illustration. Note that for 3×3 matrices the characterization of $C(B)$ in terms of planar conical partitions is particularly simple, and allows for explicit computations of the UGC hardness threshold in additional cases.

Denote $B := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & c \end{pmatrix} = (\langle v_i, v_j \rangle)_{i,j=1}^3$, where $v_1 = (1, 0, 0)$, $v_2 = (0, 1, 0)$, $v_3 = (0, 0, \sqrt{c}) \in \mathbb{R}^3$. The side

lengths of the triangle whose vertices are v_1, v_2, v_3 are $\{\ell_1 = \sqrt{1+c}, \ell_2 = \sqrt{1+c}, \ell_3 = \sqrt{2}\}$. Note that this is an acute triangle, so its smallest bounding circle coincides with its circumcircle, and therefore its radius is given by [5]:

$$R(B)^2 = \frac{\ell_1^2 \ell_2^2 \ell_3^2}{(\ell_1 + \ell_2 + \ell_3)(-\ell_1 + \ell_2 + \ell_3)(\ell_1 - \ell_2 + \ell_3)(\ell_1 + \ell_2 - \ell_3)} = \frac{(1+c)^2}{2+4c}. \quad (42)$$

We shall now compute $C(B)$. By Lemma 2.3 the partition $\{A_1, A_2, A_3\}$ of \mathbb{R}^2 at which $C(B)$ is attained consists of disjoint cones of angles $\alpha_1, \alpha_2, \alpha_3 \in [0, 2\pi]$ where $\alpha_1 + \alpha_2 + \alpha_3 = 2\pi$. A direct computation shows that for $j \in \{1, 2, 3\}$ we have:

$$\left\| \int_{A_j} x d\gamma_2(x) \right\|_2^2 = \frac{1}{2\pi} \sin^2\left(\frac{\alpha_j}{2}\right).$$

Hence

$$C(B) = \frac{1}{2\pi} \max_{\substack{\alpha_1, \alpha_2, \alpha_3 \in [0, 2\pi] \\ \alpha_1 + \alpha_2 + \alpha_3 = 2\pi}} \left(\sin^2\left(\frac{\alpha_1}{2}\right) + \sin^2\left(\frac{\alpha_2}{2}\right) + c \sin^2\left(\frac{\alpha_3}{2}\right) \right). \quad (43)$$

Assume for the moment that the maximum in (43) is attained when $\alpha_1, \alpha_2, \alpha_3 \in (0, 2\pi)$. Then using Lagrange multipliers we see that $\sin \alpha_1 = \sin \alpha_2 = c \sin \alpha_3$. This implies in particular that either $\alpha_1 = \alpha_2$ or (since $\alpha_1, \alpha_2, \alpha_3 \in (0, 2\pi)$ and $\alpha_1 + \alpha_2 + \alpha_3 = 2\pi$) $\alpha_1 + \alpha_2 = \pi$. In the latter case $\alpha_3 = \pi$, and it follows from the Lagrange multiplier equations that $\sin \alpha_1 = \sin \alpha_2 = 0$, which forces one of $\{\alpha_1, \alpha_2\}$ to vanish, contrary to our assumption. Hence we know that $\alpha_1 = \alpha_2 := \alpha$. Then $\alpha_3 = 2\pi - 2\alpha$, and since $\alpha_3 \in (0, 2\pi)$ we also know that $\alpha \in (0, \pi)$. The Lagrange multiplier equations imply that $\sin \alpha = c \sin(2\pi - 2\alpha) = -2c \sin \alpha \cos \alpha$. Thus $\cos \alpha = -\frac{1}{2c}$, and in particular we see that necessarily $c \geq \frac{1}{2}$. It follows that

$$\sin^2\left(\frac{\alpha}{2}\right) = \frac{1 - \cos \alpha}{2} = \frac{2c + 1}{4c},$$

and

$$\sin^2\left(\frac{\alpha_3}{2}\right) = \sin^2(\pi - \alpha) = 1 - \cos^2 \alpha = 1 - \frac{1}{4c^2}.$$

Hence in this case:

$$\sin^2\left(\frac{\alpha_1}{2}\right) + \sin^2\left(\frac{\alpha_2}{2}\right) + c \sin^2\left(\frac{\alpha_3}{2}\right) = 2\frac{2c+1}{4c} + c\frac{4c^2-1}{4c^2} = \frac{(2c+1)^2}{4c}. \quad (44)$$

It remains to deal with the boundary case $\{\alpha_1, \alpha_2, \alpha_3\} \cap \{0, 2\pi\} \neq \emptyset$, which as we have seen above is where the maximum in (43) is necessarily attained if $c < \frac{1}{2}$. If one of $\{\alpha_1, \alpha_2, \alpha_3\}$ equals 2π then the others must vanish, in which case $\sin^2\left(\frac{\alpha_1}{2}\right) + \sin^2\left(\frac{\alpha_2}{2}\right) + c \sin^2\left(\frac{\alpha_3}{2}\right) = 0$. If one of $\{\alpha_1, \alpha_2, \alpha_3\}$ vanishes then in order to maximize $\sin^2\left(\frac{\alpha_1}{2}\right) + \sin^2\left(\frac{\alpha_2}{2}\right) + c \sin^2\left(\frac{\alpha_3}{2}\right)$ the other two must equal π , in which case the maximum value of this quantity is $\max\{2, 1+c\}$. Since $\max\{2, 1+c\}$ never exceeds the quantity $\frac{(2c+1)^2}{4c}$ from (44) it follows that the maximum of $\sin^2\left(\frac{\alpha_1}{2}\right) + \sin^2\left(\frac{\alpha_2}{2}\right) + c \sin^2\left(\frac{\alpha_3}{2}\right)$ over $\{\alpha_1 + \alpha_2 + \alpha_3 = 2\pi \wedge \alpha_1, \alpha_2, \alpha_3 \in [0, 2\pi]\}$ equals $\frac{(2c+1)^2}{4c}$ when $c \geq \frac{1}{2}$ and equals 2 when $c \leq \frac{1}{2}$. We therefore proved that

$$C(B) = \begin{cases} \frac{(2c+1)^2}{8\pi c} & \text{if } c \geq \frac{1}{2}, \\ \frac{1}{\pi} & \text{if } c \leq \frac{1}{2}. \end{cases} \quad (45)$$

By combining (42) with (45) we conclude that the UGC hardness threshold for computing (40) is:

$$\frac{R(B)^2}{C(B)} = \begin{cases} \frac{4\pi c(1+c)^2}{(1+2c)^3} & \text{if } c \geq \frac{1}{2}, \\ \frac{\pi(1+c)^2}{2+4c} & \text{if } c \leq \frac{1}{2}. \end{cases} \quad (46)$$

Remark 6.1. An inspection of the above argument, in combination with our algorithm that was presented in Section 4, shows that the phase transition in (46) at $c = \frac{1}{2}$ corresponds to a qualitative change in the optimal algorithm: after shifting the vectors $\{v_1, \dots, v_k\}$ so that $w(B) = 0$ and renormalizing by $R(B)$, for $c > \frac{1}{2}$ the algorithm projects the points obtained from the SDP to \mathbb{R}^2 and classifies them according to a partition of \mathbb{R}^2 into three cones of positive measure, while for $c < \frac{1}{2}$ the partitioning is into two half-planes and the third set (the one weighted by c) is empty.

References

- [1] G. E. Andrews, R. Askey, and R. Roy. *Special functions*, volume 71 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1999.

- [2] Y. Benyamini and J. Lindenstrauss. *Geometric nonlinear functional analysis. Vol. 1*, volume 48 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2000.
- [3] M. Charikar, K. Makarychev, and Y. Makarychev. Near-optimal algorithms for unique games (extended abstract). In *STOC'06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 205–214, New York, 2006. ACM.
- [4] A. Grothendieck. Résumé de la théorie métrique des produits tensoriels topologiques. *Bol. Soc. Mat. São Paulo*, 8:1–79, 1953.
- [5] R. A. Johnson. *Advanced Euclidean geometry: An elementary treatise on the geometry of the triangle and the circle*. Under the editorship of John Wesley Young. Dover Publications Inc., New York, 1960.
- [6] S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, pages 767–775 (electronic), New York, 2002. ACM.
- [7] S. Khot and A. Naor. Approximate kernel clustering. In *49th Annual IEEE Symposium on Foundations of Computer Science*, pages 561–570. IEEE Computer Society, 2008.
- [8] S. Khot and O. Regev. Vertex cover might be hard to approximate to within $2 - \epsilon$. *J. Comput. System Sci.*, 74(3):335–349, 2008.
- [9] E. Mossel, R. O’Donnell, and K. Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. In *46th Annual Symposium on Foundations of Computer Science*, pages 21–30. IEEE Computer Society, 2005.
- [10] G. Pisier. *The volume of convex bodies and Banach space geometry*, volume 94 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1989.
- [11] R. E. Rietz. A proof of the Grothendieck inequality. *Israel J. Math.*, 19:271–276, 1974.
- [12] V. I. Rotar’. Limit theorems for polylinear forms. *J. Multivariate Anal.*, 9(4):511–530, 1979.
- [13] L. Song, A. Smola, A. Gretton, and K. A. Borgwardt. A dependence maximization view of clustering. In *Proceedings of the 24th international conference on Machine learning*, pages 815 – 822, 2007.
- [14] E. M. Stein. *Harmonic analysis: real-variable methods, orthogonality, and oscillatory integrals*, volume 43 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1993. With the assistance of Timothy S. Murphy, Monographs in Harmonic Analysis, III.