

# Quadratic Memory is Necessary for Optimal Query Complexity in Convex Optimization: Center-of-Mass is Pareto-Optimal

Moïse Blanchard  
MIT  
moiseb@mit.edu

Junhui Zhang  
MIT  
junhuiz@mit.edu

Patrick Jaillet  
MIT  
jaillet@mit.edu

## Abstract

We give query complexity lower bounds for convex optimization and the related feasibility problem. We show that quadratic memory is necessary to achieve the optimal oracle complexity for first-order convex optimization. In particular, this shows that center-of-mass cutting-planes algorithms in dimension  $d$  which use  $\tilde{O}(d^2)$  memory and  $\tilde{O}(d)$  queries are Pareto-optimal for both convex optimization and the feasibility problem, up to logarithmic factors. Precisely, building upon techniques introduced in [1], we prove that to minimize 1-Lipschitz convex functions over the unit ball to  $1/d^4$  accuracy, any deterministic first-order algorithms using at most  $d^{2-\delta}$  bits of memory must make  $\tilde{\Omega}(d^{1+\delta/3})$  queries, for any  $\delta \in [0, 1]$ . For the feasibility problem, in which an algorithm only has access to a separation oracle, we show a stronger trade-off: for at most  $d^{2-\delta}$  memory, the number of queries required is  $\tilde{\Omega}(d^{1+\delta})$ . This resolves a COLT 2019 open problem of Woodworth and Srebro.

**Keywords.** Convex optimization, feasibility problem, first-order methods, cutting-planes, center-of-mass, memory lower bounds, query complexity

## 1 Introduction

We consider the canonical problem of first-order convex optimization in which one aims to minimize a convex function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  with access to an oracle that for any query  $\mathbf{x}$  returns  $(f(\mathbf{x}), \nabla f(\mathbf{x}))$  the value of the function and a subgradient of  $f$  at  $\mathbf{x}$ . Arguably, this is one of the most fundamental problems in optimization, mathematical programming and machine learning.

A classical question is how many oracle queries are required to guarantee finding an  $\epsilon$ -approximate minimizer for any 1-Lipschitz convex functions  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  over the unit ball. We denote by  $B_d(\mathbf{x}, r) = \{\mathbf{x}' \in \mathbb{R}^d : \|\mathbf{x} - \mathbf{x}'\|_2 \leq r\}$  the ball centered in  $\mathbf{x}$  of radius  $r$ . There exist methods that given first-order oracle access only need  $\mathcal{O}(d \log 1/\epsilon)$  queries and this query complexity is worst-case optimal [2] when  $\epsilon \ll 1/\sqrt{d}$ . Known methods achieving the optimal  $\mathcal{O}(d \log 1/\epsilon)$  query complexity fall in the broad class of cutting plane methods, that build upon the well-known ellipsoid method [3, 4] which uses  $\mathcal{O}(d^2 \log 1/\epsilon)$  queries. These include the inscribed ellipsoid [5, 6], volumetric center or Vaidya’s method [7, 8], approximate center-of-mass via sampling techniques [9, 10] and recent improvements [11, 12]. Unfortunately, all these methods suffer from at least  $\Omega(d^3 \log 1/\epsilon)$  time complexity and further require storing all subgradients, or at least an ellipsoid in  $\mathbb{R}^d$ , therefore at least  $\Omega(d^2 \log 1/\epsilon)$  bits of memory. These limitations are prohibitive for large-scale optimization, hence cutting plane methods are viewed as rather impractical and less frequently used for high-dimensional applications. On the other hand, the simplest, perhaps most commonly used and practical gradient descent requires  $\mathcal{O}(1/\epsilon^2)$  queries, which is not optimal for  $\epsilon \ll 1/\sqrt{d}$ , but only needs  $\mathcal{O}(d)$  time per query and  $\mathcal{O}(d \log 1/\epsilon)$  memory.

A natural question is whether one can preserve the optimal query lower bounds from cutting-planes methods with simpler methods, for instance, inspired by gradient descent techniques. Such hope is largely motivated by the fact that in many different theoretical settings, cutting plane methods have achieved state-of-the-art runtimes including semidefinite programming [11, 13], submodular optimization [11, 14–16] or equilibrium computation [17, 18]. Towards this goal, [19] first posed this question in terms of query complexity / memory trade-off: given a certain number of bits of memory, which query complexity is achievable? While cutting planes methods require  $\Omega(d^2 \log 1/\epsilon)$  memory, gradient descent only requires storing one vector and as a result, uses  $\mathcal{O}(d \log 1/\epsilon)$  memory, which is information-theoretically optimal [19]<sup>1</sup>. Understanding this trade-off could pave the way for the design of more efficient methods in convex optimization.

The first result in this direction was provided in [1], where they showed that it is impossible to be both optimal in query complexity and in memory. Specifically, they proved that any potentially randomized algorithm that uses at most  $d^{1.25-\delta}$  memory must make at least  $\tilde{\Omega}(d^{1+4/3\delta})$  queries. This implies that a super-linear amount of memory  $d^{1.25}$  is required to achieve the optimal rate of convergence (that is achieved by algorithms using more than quadratic memory). However, this leaves open the fundamental question of whether one can improve over the memory of cutting-plane methods while keeping optimal query complexity.

**Question (COLT 2019 [19]).** Is it possible for a first-order algorithm that uses at most  $\mathcal{O}(d^{2-\delta})$  bits of memory to achieve query complexity  $\tilde{\mathcal{O}}(d \text{polylog } 1/\epsilon)$  when  $d = \Omega(\log^c 1/\epsilon)$  but  $d = o(1/\epsilon^c)$  for all  $c > 0$ ?

In this paper, building upon the techniques introduced in [1], we provide a negative answer to this question: quadratic memory is necessary to achieve the optimal query complexity with deterministic algorithms. As a result, cutting plane methods including the standard center-of-mass algorithm are Pareto-optimal up to logarithmic factors within the query complexity / memory trade-off. Our main result for convex optimization is the following.

**Theorem 1.** *For  $\epsilon = 1/d^4$  and any  $\delta \in [0, 1]$ , a deterministic first-order algorithm guaranteed to minimize 1-Lipschitz convex functions over the unit ball with  $\epsilon$  accuracy uses at least  $d^{2-\delta}$  bits or makes  $\tilde{\Omega}(d^{1+\delta/3})$  queries.*

A key component of cutting plane methods is that they merely rely on the subgradient information at each query to restrict the search space. As a result, these can be used to solve the larger class of feasibility problems that are essential in mathematical programming and optimization. In a feasibility problem, one aims to find an  $\epsilon$ -approximation of an unknown vector  $\mathbf{x}^*$ , and has access to a separation oracle. For any query  $\mathbf{x}$ , the separation oracle either returns a separating hyperplane  $\mathbf{g}$  from  $\mathbf{x}$  to  $B_d(\mathbf{x}^*, \epsilon)$ —such that  $\langle \mathbf{g}, \mathbf{x} - \mathbf{z} \rangle > 0$  for any  $\mathbf{z} \in B_d(\mathbf{x}^*, \epsilon)$ —or signals that  $\|\mathbf{x} - \mathbf{x}^*\| \leq \epsilon$ . This class of problems is broader than convex optimization since the negative subgradient always provides a separating hyperplane from a suboptimal query to the optimal set. Hence, feasibility and convex minimization problem are closely related and it is often the case that obtaining query lower bounds for the feasibility problem simplifies the analysis while still providing key insights for the more restrictive convex optimization problem [2, 20].

As a result, a similar fundamental question is to understand the query complexity / memory trade-off for the feasibility problem. As noted above, any lower bound for convex optimization yields the same lower bound for the feasibility problem. Here, we can significantly improve over the previous trade-off.

---

<sup>1</sup> $\Omega(d \log 1/\epsilon)$  bits of memory are already required just to represent the answer to the optimization problem.

**Theorem 2.** For  $\epsilon = 1/(48d^2\sqrt{d})$  and any  $\delta \in [0, 1]$ , a deterministic algorithm guaranteed to solve the feasibility problem over the unit ball with  $\epsilon$  accuracy uses at least  $d^{2-\delta}$  bits of memory or makes at least  $\tilde{\Omega}(d^{1+\delta})$  queries.

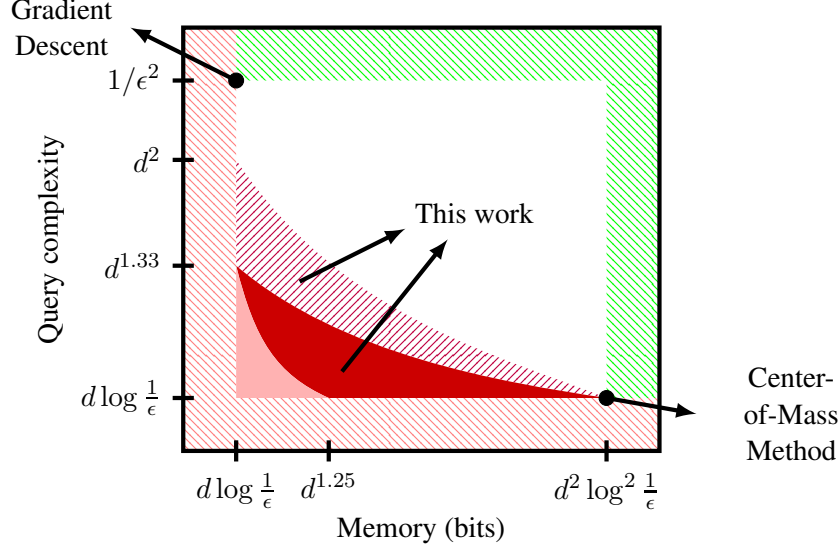


Figure 1: Trade-offs between available memory and first-order oracle complexity for minimizing 1-Lipschitz convex functions over the unit ball (adapted from [1, 19]). The dashed pink “L” (resp. green inverted “L”) shaped region corresponds to historical information-theoretic lower bounds (resp. upper bounds) on the memory and query-complexity. The solid pink region corresponds to the recent lower bound trade-off from [1], which holds for randomized algorithms. In our work, we show that the solid red region is not achievable for any deterministic algorithms. For the feasibility problem, we also show that the dashed red region is not achievable either for any deterministic algorithms.

## 1.1 Literature review

Recently, there has been a series of studies exploring the trade-offs between sample complexity and memory constraints for learning problems, such as linear regression [21, 22], principal component analysis (PCA) [23], learning under the statistical query model [24] and other general learning problems [25–31].

For parity problems that meet certain spectral (mixing) requirements, [32] first proved by a computation tree argument that an exponential number of random samples is needed if the memory is sub-quadratic. Similar trade-offs have been obtained when the learning problem satisfies other types of properties [27–31]. It should be noted that all the above-mentioned results hold for learning problems over finite fields, i.e. the concept classes are finite. For continuous problems, [22] was the first to apply [32]’s framework and showed a sample-complexity lower bound for memory-constrained linear regression.

In contrast to learning with random samples, there is limited understanding of the memory-constrained optimization and feasibility problem. [33] demonstrated that, in the absence of memory constraints, finding an  $\epsilon$ -approximate solution for Lipschitz convex functions requires  $\Omega(d \log 1/\epsilon)$  queries, which can be achieved by the center-of-mass method using  $O(d^2 \log^2 1/\epsilon)$  bits of memory. At the other extreme, gradient descent needs  $\Omega(1/\epsilon^2)$  queries but only  $O(d \log 1/\epsilon)$  bits of memory, the minimum memory needed to represent a solution. These two extreme cases are represented by dashed pink “impossible

region” and dashed green “achievable region” in Figure 1. Since then, [1] showed that there is a trade-off between memory and query for convex optimization: it is impossible to be both optimal in query complexity and memory. Their lower bound is represented by the solid pink “impossible region” in Figure 1. In this paper, we significantly improve these results to match the quadratic upper bound of cutting plane methods. Additionally, there has been recent progress in the study of query complexity for randomized algorithms [34, 35].

On the algorithmic side, the afore-mentioned methods that achieve  $O(\text{poly}(d))$  query complexity [3–12] all require at least  $\Omega(d^2 \log 1/\epsilon)$  bits of memory. There is also significant literature on memory-efficient optimization algorithms, such as the Limited-memory-BFGS [36, 37]. However, the convergence behavior for even the original BFGS on non-smooth convex objectives is still a challenging, open question [38].

**Comparison with [1]** Our proof techniques build upon those introduced in [1]. We follow the proof strategy that they introduced to derive lower bounds for the memory/query complexity. Below, we delineate which ideas and techniques are borrowed from [1] and which are the novel elements that we introduce. Details on these proof elements are given in Section 2.1.

First, [1] define a class of difficult functions for convex optimization of the following form

$$\max \left\{ \|\mathbf{A}\mathbf{x}\|_\infty - \eta_0, \eta_1 \left( \max_{i \leq N} \mathbf{v}_i^\top \mathbf{x} - i\gamma \right) \right\}, \quad (1)$$

where  $\mathbf{A} \sim \mathcal{U}(\{\pm 1\}^{d/2 \times d})$  is a matrix with  $\pm 1$  entries sampled uniformly, and  $\mathbf{v}_i \sim \mathcal{U}(d^{-1/2} \{\pm 1\}^d)$  are sampled independently, uniformly within the rescaled hypercube. To give intuition on this class, the term  $\|\mathbf{A}\mathbf{x}\|_\infty - \eta_0$  acts as barrier : in order to observe subgradients from the other term, one needs to use queries  $\mathbf{x}$  that are approximately within the nullspace of  $\mathbf{A}$ . The second term  $\max_{i \leq N} \mathbf{v}_i^\top \mathbf{x} - i\gamma$  is the “Nemirovski” function, which was used in previous works [39–41] to obtain lower bounds in parallel convex optimization. At a high level, the limitation in the lower bounds from [1] comes from the fact that one is limited in the number  $N$  of vectors  $\mathbf{v}_1, \dots, \mathbf{v}_N$  that can be used in the Nemirovski function. To resolve this issue, we introduce adaptivity within the choice of a modified Nemirovski function. At a high level, we choose the vectors  $\mathbf{v}_1, \dots, \mathbf{v}_N$  depending on the queries of the algorithm which allows to fit in more terms. In turn, this allows to improve the lower bounds.

As a second step, [1] relate the optimization problem on the defined class of functions to an Orthogonal Vector Game. In this game, the goal is to find vectors that are approximately orthogonal to a matrix  $\mathbf{A}$  with access to row queries of  $\mathbf{A}$ . The argument is as follows: because of the barrier term  $\|\mathbf{A}\mathbf{x}\|_\infty - \eta_0$ , optimizing the Nemirovski function requires exploring independent directions of the nullspace of  $\mathbf{A}$ , which is performed at *informative queries*. With our new class of functions, we can adapt this logic. However, the adaptivity in the vectors  $\mathbf{v}_i$  provides information to the learner on  $\mathbf{A}$  in addition to the queried rows of  $\mathbf{A}$ . We therefore need to modify the game by introducing an Orthogonal Vector Game with Hints, where hints encapsulate this extra information.

For the last step, [1] give an information-theoretic argument to provide a query complexity lower bound on the defined Orthogonal Vector Game. Following the same structure, we show that a similar argument holds for our modified game. The main added difficulty resides in bounding the information leakage from the hints, and we show that these provide no more information than the memory itself.

As a last remark, the lower bounds provided in [1] hold for randomized algorithms, while the adaptivity of our procedure only applies to deterministic algorithms.

## 1.2 Outline of paper

Our main results for the trade-off between memory and query complexity for optimization and feasibility problem have been presented in Section 1 (Theorem 1, 2). In Section 2, we formally define memory-constrained algorithms and provide a brief overview of our proof techniques and contributions. Our proofs for convex optimization are given in Section 3. We introduce the *optimization procedure* which adaptively constructs a hard family of functions, provide a reduction from this hard family to an *orthogonal vector game with hints*, and show a memory-sample trade-off (Proposition 14) for the game, which completes the proof of the Theorem 1. Last, in Section 4, we consider the feasibility problem and, with a similar methodology, prove Theorem 2.

## 2 Formal setup and overview of techniques

Standard results in oracle complexity give the minimal number of queries for algorithms to solve a given problem. However, this does not account for possible restrictions on the memory available to the algorithm. In this paper, we are interested in the trade-off between memory and query complexity for both convex optimization and the feasibility problem. Our results apply to a large class of *memory-constrained* algorithms. We give below a general definition of the memory constraint for algorithms with access to an oracle  $\mathcal{O} : \mathcal{S} \rightarrow \mathcal{R}$  taking as input a query  $q \in \mathcal{S}$  and returning as response  $\mathcal{O}(q) \in \mathcal{R}$ .

**Definition 3** (*M-bit memory-constrained deterministic algorithm*). *Let  $\mathcal{O} : \mathcal{S} \rightarrow \mathcal{R}$  be an oracle. An  $M$ -bit memory-constrained deterministic algorithm is specified by a query function  $\psi_{\text{query}} : \{0, 1\}^M \rightarrow \mathcal{S}$  and an update function  $\psi_{\text{update}} : \{0, 1\}^M \times \mathcal{S} \times \mathcal{R} \rightarrow \{0, 1\}^M$ . The algorithm starts with the memory state  $\text{Memory}_0 = 0^M$  and iteratively makes queries to the oracle. At iteration  $t$ , it makes the query  $q_t = \psi_{\text{query}}(\text{Memory}_{t-1})$  to the oracle, receives the response  $r_t = \mathcal{O}(q_t)$  then updates its memory  $\text{Memory}_t = \psi_{\text{update}}(\text{Memory}_{t-1}, q_t, r_t)$ .*

The algorithm can stop making queries at any iteration and the last query is its final output. Notice that the memory constraint applies only between each query but not for internal computations, i.e. the computation of the update  $\psi_{\text{update}}$  and the query  $\psi_{\text{query}}$  can potentially use unlimited memory. This is a rather weak memory constraint on the algorithm; a fortiori, our negative results also apply to stronger notions of memory-constrained algorithms. In Definition 3, we ask the query and update functions to be time-invariant. In our context, this is without loss of generality: any  $M$ -bit algorithm using  $T$  queries with time-dependent query and update functions [1, 19] can be turned into an  $(M + \lceil \log T \rceil)$ -bit time-invariant algorithm by storing the iteration number  $t$  as part of the memory. The query lower bounds we provide are at most  $T \leq \text{poly}(d)$ . Hence, an additional  $\log T = O(\log d)$  bits to the memory size  $M$ , does not affect our main results, Theorems 1 and 2.

In this paper, we use the above described framework to study the interplay between query complexity and memory for two fundamental problems in optimization and machine learning.

**Convex optimization.** We first consider convex optimization in which one aims to minimize a 1-Lipschitz convex function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  over the unit ball  $B_d(0, 1) \subset \mathbb{R}^d$ . The goal is to output a point  $\tilde{\mathbf{x}} \in B_d(0, 1)$  such that  $f(\tilde{\mathbf{x}}) \leq \min_{\mathbf{x} \in B_d(0, 1)} f(\mathbf{x}) + \epsilon$ , referred to as  $\epsilon$ -approximate points. The optimization algorithm has access to a first order oracle  $\mathcal{O}_{CO} : \mathbb{R}^d \rightarrow \mathbb{R} \times \mathbb{R}^d$ , which for any query  $\mathbf{x}$  returns the couple  $(f(\mathbf{x}), \partial f(\mathbf{x}))$  where  $\partial f(\mathbf{x})$  is a subgradient of  $f$  at the query point  $\mathbf{x}$ .

**Remark 4.** *The above requirement for  $\epsilon$ -approximate optimality is weaker than asking to find a point that is at distance  $\epsilon$  from  $\arg \min_{\mathbf{x} \in B_d(0, 1)} f(\mathbf{x})$  (for 1-Lipschitz convex functions). As a result, our lower*

bounds for  $\epsilon$ -approximate optimality hold a fortiori for the problem where one aims to find a point at distance at most  $\epsilon$  from the solution set.

**Feasibility problem.** Second, we consider the trade-off between memory and query complexity for the feasibility problem, where the goal is to find an element  $\tilde{\mathbf{x}} \in Q$  for a convex set  $Q \subset B_d(0, 1)$ . Instead of a first-order oracle, the algorithm has access to a separation oracle  $\mathcal{O}_F : \mathbb{R}^d \rightarrow \{\text{Success}\} \cup \mathbb{R}^d$ . For any query  $\mathbf{x} \in \mathbb{R}^d$ , the separation oracle either returns Success reporting that  $\mathbf{x} \in Q$ , or provides a separating vector  $\mathbf{g} \in \mathbb{R}^d$ , i.e., such that for all  $\mathbf{x}' \in Q$ ,

$$\langle \mathbf{g}, \mathbf{x} - \mathbf{x}' \rangle > 0.$$

We say that an algorithm solves the feasibility problem with accuracy  $\epsilon > 0$  if it can solve any feasibility problem for which the successful set contains a ball of radius  $\epsilon$ , i.e., such that there exists  $\mathbf{x}^* \in B_d(0, 1)$  satisfying  $B_d(\mathbf{x}^*, \epsilon) \subset Q$ .

The feasibility problem is at least as hard as convex optimization in the following sense: an algorithm that solves the feasibility problem with accuracy  $\epsilon/L$  can be used to solve  $L$ -Lipschitz convex optimization problems by feeding the subgradients from first-order queries to the algorithm as separating hyperplanes. Alternatively, from any 1-Lipschitz function  $f$  one can derive a feasibility problem, where the feasibility set is  $Q = \{\mathbf{x} \in B_d(0, 1), f(\mathbf{x}) \leq f^* + \epsilon\}$  and the separating oracle at  $\mathbf{x} \notin Q$  is a subgradient  $\partial f(\mathbf{x})$  at  $\mathbf{x}$ .

## 2.1 Overview of proof techniques and innovations

We prove the two main Theorems 1 and 2 with similar techniques, hence for conciseness, we only give here the main ideas used to derive lower bounds for convex optimization. Although our proof borrows its structure and techniques from [1], we introduce key innovations involving adaptivity to improve the lower bounds up to the maximum quadratic memory for deterministic algorithms—up to logarithmic factors. We recall, however, that the bounds in [1] hold for randomized algorithms as well. In the proofs, we aim to optimize the dependence of the parameters in  $d$ . Constants, however, are not necessarily optimized.

**An adaptive optimization procedure.** At the high level, we design an *optimization procedure* which for any algorithm constructs a hard family of convex functions adaptively on its queries. To be precise, the procedure constructs functions from the following family of convex functions with appropriately chosen parameters  $\eta, \gamma_1, \gamma_2, p_{max}, l_p, \delta$ :

$$F_{\mathbf{A}, \mathbf{v}}(\mathbf{x}) = \max \left\{ \|\mathbf{A}\mathbf{x}\|_\infty - \eta, \eta \mathbf{v}_0^\top \mathbf{x}, \eta \left( \max_{p \leq p_{max}, l \leq l_p} \mathbf{v}_{p,l}^\top \mathbf{x} - p\gamma_1 - l\gamma_2 \right) \right\}. \quad (2)$$

We take  $\mathbf{A} \sim \mathcal{U}(\{\pm 1\}^{n \times d})$  and  $\mathbf{v}_0 \sim \mathcal{U}(\mathcal{D}_\delta)$  uniformly sampled in the beginning, where  $\mathcal{D}_\delta \subset \mathcal{S}^{d-1}$  is a (finite) discretization of the sphere. The first term  $\|\mathbf{A}\mathbf{x}\|_\infty - \eta$  acts as a barrier term: in order to observe subgradients from the other terms, one needs the query  $\mathbf{x}$  to satisfy  $\|\mathbf{A}\mathbf{x}\|_\infty \leq 2\eta$ . These are called *informative queries* as introduced in [1]. Hence, informative queries must lie approximately in the orthogonal space to the lines of  $\mathbf{A}$ . The second term  $\eta \mathbf{v}_0^\top \mathbf{x}$  is used to ensure that solutions with low objective (in particular with objective at most  $\eta\gamma_1/2$ ) have norm bounded away from 0. As a result, these informative queries, once renormalized, will still belong approximately to the nullspace of  $\mathbf{A}$  denoted  $\text{Ker}(\mathbf{A})$ .

The adaptivity to the algorithm is captured in the third term, which is constructed along the optimization process. This construction proceeds by periods  $p = 1, 2, \dots, p_{max}$  designed so that during each period  $p$ , the algorithm is forced to visit a subspace of  $Ker(\mathbf{A})$  of dimension  $k$ . To do so, we iteratively construct vectors  $\mathbf{v}_{p,1}, \dots, \mathbf{v}_{p,l_p}$  as follows. Suppose that at the beginning of step  $t$  of period  $p$ , one has defined vectors  $\mathbf{v}_{p,1}, \dots, \mathbf{v}_{p,l}$ .

- The procedure first evaluates the explored subspace of the algorithm during this period. In practice, the procedure keeps in memory *exploratory* queries  $\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,r}}$  during period  $p$  up to time  $t$ . The exploratory subspace is then  $Span(\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,r}})$ .
- If a query with a sufficiently low objective is queried, we sample a new vector  $\mathbf{v}_{p,l+1}$  which is approximately orthogonal to the exploratory subspace. The corresponding new term in the objective is  $\mathbf{v}_{p,l+1}^\top \mathbf{x} - p\gamma_1 - (l+1)\gamma_2$ .

Once this new term is added to the objective, the algorithm is constrained to make queries with an additional component along the direction  $-\mathbf{v}_{p,l+1}$ . Since this vector is approximately orthogonal to all previous queries, this forces the algorithm to query vectors linearly independent from all previous queries in period  $p$ . The period then ends once the dimension of the exploratory subspace reaches  $k$ , having defined  $l_p$  vectors  $\mathbf{v}_{p,1}, \dots, \mathbf{v}_{p,l_p}$ . As discussed above, the exploratory subspace must increase dimension for any additional such vector. Thus, after  $l_p \leq k$  vectors, period  $p$  ends.

The constructed family of convex functions in Eq (2) is similar to the family described in Eq (1) that were considered in [1]. However, by sampling the vectors  $\mathbf{v}_{p,l}$  adaptively, the *optimization procedure* is able to fit in more terms, thereby providing a significant improvement in the lower bounds.

**Benefits of adaptivity.** We now expand on how the adaptive terms allow improving the lower bound of [1] to match the quadratic upper bound of cutting plane methods. The limitation in the functions of the form Eq (1) comes from the fact that the offset in the Nemirovski function is  $\gamma = \Omega(\sqrt{k \log d/d})$ . This offset is necessary to ensure that with high probability, 1. subgradients  $\mathbf{v}_1, \dots, \mathbf{v}_N$  are discovered exactly in this order and 2. that any query which visits a new vector  $\mathbf{v}_i$  must not lie in the subspace formed by the last  $k$  last informative vectors. Indeed, for the last claim, from high-dimensional concentration, for a random unit vector  $\mathbf{v}$  and a  $k$  dimensional subspace  $E$ ,  $\|P_E(\mathbf{v})\| = \Theta(\sqrt{k \log d/d})$ . This offset is not necessary for our procedure, since by construction, at each period, a  $k$ -dimensional subspace of  $Ker(\mathbf{A})$  is forced to be explored. As a result, we can take  $\gamma_1 = \Theta(\sqrt{\log d/d})$ . This offset is still necessary to ensure that vectors  $\mathbf{v}_{p,l}$  are discovered in their order of construction (lexicographic order on  $(p, l)$ ) with high probability.

**An Orthogonal Vector Game with Hints.** The next step of the proof involves linking the optimization of the above-mentioned constructed functions with an Orthogonal Vector Game with Hints. Similarly to the game introduced by [1], the goal for the player is to find  $k$  linearly-independent vectors approximately in  $Ker(\mathbf{A})$ . To do so, the player can access an  $M$ -bit message Message and make  $m$  queries, where  $M = ckd$  for a small constant  $c > 0$ . In the game introduced by [1], the queries are lines of the matrix  $\mathbf{A}$ . They then show that to find  $k$  dimensions of  $\mathbf{A}$ , where  $\mathbf{A}$  is taken uniformly at random  $\mathbf{A} \sim \{\pm 1\}^{d/2 \times d}$ , (nearly) all the lines of  $\mathbf{A}$  must be queried. The argument is information-theoretic: each new dimension of  $Ker(\mathbf{A})$  must be (approximately) orthogonal to all lines of  $\mathbf{A}$ . Hence, this provides additional mutual information  $O(k)$  for every line of  $\mathbf{A}$ , including the  $d/2 - m$  lines that were not observed through queries. This extra information on  $\mathbf{A}$  can only be explained by the message, which has  $M$  bits. Hence,  $M \geq O(k)(d/2 - m)$ . Setting the constant  $c > 0$  appropriately, this shows that  $m = \Omega(d)$ .

In our case, the optimization procedure ensures that the algorithm needs to explore  $k$  dimensions of  $\text{Ker}(\mathbf{A})$  in each period. However, each query yields a response from the optimization oracle that can either be a line of  $\mathbf{A}$  (corresponding to the term  $\|\mathbf{A}\mathbf{x}\|_\infty - \eta$  of Eq (2)) or  $\mathbf{v}_0$  (term  $\eta\mathbf{v}_0^\top\mathbf{x}$  of Eq (2)), or previously defined vectors  $\mathbf{v}_{p',l'}$ . Now since the vectors  $\mathbf{v}_{p',l'}$  have been constructed adaptively on the queries of the algorithm, which themselves may depend on lines of  $\mathbf{A}$ , during a period  $p$ , responses  $\mathbf{v}_{p',l'}$  for  $p' < p$  are a source of information leakage for  $\mathbf{A}$  from previous periods. As a result, the query lower bound on the game introduced by [1] is not sufficient for our purposes. Instead, we introduce an Orthogonal Vector Game with Hints, where hints correspond exactly to these vectors  $\mathbf{v}_{p',l'}$  from previous periods. Informally, the game corresponds to a simulation of one of the periods of the optimization procedure: for each query  $\mathbf{x}$ , the oracle returns the subgradient that would have been returned in the optimization procedure, up to minor details.

**Bounding the information leakage.** Once the link is settled, the goal is to prove lower bounds on the number of queries needed to solve the Orthogonal Vector Game with Hints. The main difficulty is to bound the information leakage from these hints. We recall that hints are of the form  $\mathbf{v}_{p',l'}$ , which have been constructed adaptively on the queries of the algorithm during period  $p'$ . In particular, these contain information on the lines of  $\mathbf{A}$  queried during period  $p' < p$ , which may be complementary with those queried during period  $p$ . If this total information leakage through the hints yields a mutual information with  $\text{Ker}(\mathbf{A})$  significantly higher than that of the  $M$  bits of Message, obtained lower bounds cannot possibly reflect any trade-off with memory constraints. It is therefore essential to obtain information leakage at most  $\mathcal{O}(M) = \tilde{\mathcal{O}}(dk)$ .

To solve this issue, we introduce a discretization  $\mathcal{D}_\delta$  of the unit sphere where the vectors  $\mathbf{v}_{p,l}$  take value. Next, we show that each individual vector  $\mathbf{v}_{p',l'}$  from previous periods can only provide information  $\tilde{\mathcal{O}}(k)$  on the matrix  $\mathbf{A}$ . To have an intuition on this, note that for any (at most)  $k$  vectors  $\mathbf{x}_1, \dots, \mathbf{x}_k$ , the volume of the subset of the unit sphere  $S^{d-1}$  of vectors approximately orthogonal to  $\mathbf{x}_1, \dots, \mathbf{x}_k$ , say  $S(\mathbf{x}_1, \dots, \mathbf{x}_k) = \{\mathbf{y} \in S^{d-1} : |\mathbf{y}^\top \mathbf{x}_i| \leq d^{-3}, i \leq k\}$  is  $q_k = \mathcal{O}(1/d^{3k})$ . Hence, since the vector  $\mathbf{v}$  is roughly taken uniformly at random within  $\mathcal{D}_\delta \cap S(\mathbf{x}_1, \dots, \mathbf{x}_k)$ , we can show that the mutual information of  $\mathbf{v}$  with the initial vectors  $\mathbf{x}_1, \dots, \mathbf{x}_k$  is at most  $\mathcal{O}(-\log q_k) = \mathcal{O}(k \log d)$ . As a result, even if  $m = d$ , the total information leakage through the vectors  $\mathbf{v}_{p',l'}$  from previous periods, is at most  $\mathcal{O}(kd \log d)$ . The formal proof involves an anti-concentration bounds on the distance of a random unit vector to a linear subspace of dimension  $k$ , as well as a more involved discretization procedure than the one presented above. In summary, by introducing adaptive functions through the optimization procedure, we show that the same memory-sample trade-off holds for the Orthogonal Vector Game with Hints and the game without hints introduced in [1], up to logarithmic factors.

### 3 Memory-constrained convex optimization

To prove our results we need to use discretizations of the unit sphere  $S^{d-1}$ . It will be convenient to ensure that the partitions induced by these discretizations have equal area, which can be done with the following lemma.

**Lemma 5** ([42] Lemma 21). *For any  $0 < \delta < \pi/2$ , the sphere  $S^{d-1}$  can be partitioned into  $N(\delta) = (\mathcal{O}(1)/\delta)^d$  equal volume cells, each of diameter at most  $\delta$ .*

We denote by  $\mathcal{V}_\delta = \{V_i(\delta), i \in [N(\delta)]\}$  the corresponding partition, and consider a set of representatives  $\mathcal{D}_\delta = \{\mathbf{b}_i(\delta), i \in [N(\delta)]\} \subset S^{d-1}$  such that for all  $i \in [N(\delta)]$ ,  $\mathbf{b}_i(\delta) \in V_i(\delta)$ . With these



notations we can define the discretization function  $\phi_\delta$  as follows

$$\phi_\delta(\mathbf{x}) = \mathbf{b}_i(\delta), \quad \mathbf{x} \in V_i(\delta).$$

### 3.1 Definition of the difficult class of optimization problems

In this section we present the class of functions that we use to prove our lower bounds. Throughout the paper, we pose  $n = \lceil d/4 \rceil$ . We first define some useful functions. For any  $\mathbf{A} \in \mathbb{R}^{n \times d}$ , we define  $\mathbf{g}_\mathbf{A}$  as follows

$$\mathbf{g}_\mathbf{A}(\mathbf{x}) = \mathbf{a}_{i_{\min}}, \quad i_{\min} = \min\{i \in [n], |\mathbf{a}_i^\top \mathbf{x}| = \|\mathbf{A}\mathbf{x}\|_\infty\}.$$

With this function we can define a subgradient function for  $\mathbf{x} \mapsto \|\mathbf{A}\mathbf{x}\|_\infty$ ,

$$\tilde{\mathbf{g}}_\mathbf{A}(\mathbf{x}) = \epsilon \mathbf{g}_\mathbf{A}(\mathbf{x}), \quad \epsilon = \text{sign}(\mathbf{g}_\mathbf{A}(\mathbf{x})^\top \mathbf{x}).$$

We are now ready to introduce the class of functions which we use for our lower bounds. These are of the following form.

$$F_{\mathbf{A},\mathbf{v}}(\mathbf{x}) = \max \left\{ \|\mathbf{A}\mathbf{x}\|_\infty - \eta, \eta \mathbf{v}_0^\top \mathbf{x}, \eta \left( \max_{p \leq p_{\max}} \max_{l \leq l_p} \mathbf{v}_{p,l}^\top \mathbf{x} - p\gamma_1 - l\gamma_2 \right) \right\}.$$

Here,  $\mathbf{A} \in \{\pm 1\}^{n \times d}$  is a matrix. Also,  $\mathbf{v}_0$  and the terms  $\mathbf{v}_{p,l}$  are vectors in  $\mathbb{R}^d$ . More precisely, these vectors will lie in the discretization  $\mathcal{D}_\delta$  for  $\delta = 1/d^3$ . We postpone the definition of  $p_{\max}$  and  $l_p$  for  $p \leq p_{\max}$ . Last, we use the following choice for the remaining parameters:  $\eta = 2/d^3$ ,  $\gamma_1 = 12\sqrt{\frac{\log d}{d}}$  and  $\gamma_2 = \frac{71}{4d}$ . For convenience, we also define the functions

$$F_\mathbf{A}(\mathbf{x}) = \max\{\|\mathbf{A}\mathbf{x}\|_\infty - \eta, \eta \mathbf{v}_0^\top \mathbf{x}\}$$

$$F_{\mathbf{A},\mathbf{v},p,l}(\mathbf{x}) = \max \left\{ \|\mathbf{A}\mathbf{x}\|_\infty - \eta, \eta \mathbf{v}_0^\top \mathbf{x}, \eta \left( \max_{(p',l') \leq_{\text{lex}} (p,l), l' \leq l_{p'}} \mathbf{v}_{p',l'}^\top \mathbf{x} - p'\gamma_1 - l'\gamma_2 \right) \right\},$$

with the convention  $F_{\mathbf{A},\mathbf{v},1,0} = F_\mathbf{A}$ . The functions  $F_{\mathbf{A},\mathbf{v},p,l}$  will encapsulate the current state of the function to be minimized: it will be updated adaptively on the queries of the algorithm. We also define a subgradient function for  $F_{\mathbf{A},\mathbf{v},p,l}$  by first favoring lines of  $\mathbf{A}$ , then vectors from  $\mathbf{v}$  in case of ties, as follows,

$$\partial F_{\mathbf{A},\mathbf{v},p,l}(\mathbf{x}) = \begin{cases} \tilde{\mathbf{g}}_\mathbf{A}(\mathbf{x}_t) & \text{if } F_{\mathbf{A},\mathbf{v},p,l}(\mathbf{x}) = \|\mathbf{A}\mathbf{x}\|_\infty - \eta, \\ \eta \mathbf{v}_0 & \text{otherwise and if } F_{\mathbf{A},\mathbf{v},p,l}(\mathbf{x}) = \eta \mathbf{v}_0^\top \mathbf{x}, \\ \eta \mathbf{v}_{p,l} & \text{otherwise and if } (p,l) = \arg \max_{(p',l') \leq_{\text{lex}} (p,l)} \mathbf{v}_{p',l'}^\top \mathbf{x} - p'\gamma_1 - l'\gamma_2. \end{cases}$$

In the last case, ties are broken by lexicographic order. We define  $\partial F_{\mathbf{A},\mathbf{v}} = \partial F_{\mathbf{A},\mathbf{v},p_{\max},l_{p_{\max}}}$  similarly.

We consider a so-called *optimization procedure*, which will construct the sequence of vectors  $\mathbf{v} = (\mathbf{v}_{p,l})$  adaptively on the responses of the considered algorithm. Throughout this section, we use a parameter  $1 \leq k \leq d/3 - 1$  — which will be taken as  $k = \tilde{\Theta}(M/d)$  where  $M$  is the memory of the algorithm — and let  $p_{\max}$  be the largest number which satisfies the following constraint.

$$p_{\max} \leq \min\{(c_{d,1}d - 1)/k, c_{d,2}(d/k)^{1/3} - 1\}, \quad (3)$$

where  $c_{d,1} = 1/(90^2 \log^2 d)$  and  $c_{d,2} = 1/(81 \log^{2/3} d)$ .

---

**Input:**  $d, k, p_{max}$ , algorithm  $alg$

**Part 1:** Procedure to adaptively construct  $\mathbf{v}$ ;

```

1 Sample  $\mathbf{A} \sim \mathcal{U}(\{\pm 1\}^{n \times d})$  and  $\mathbf{v}_0 \sim \mathcal{U}(\mathcal{D}_\delta)$ .;
2 Initialize the memory of  $alg$  to  $\mathbf{0}$  and let  $p = 1, r = l = 0$ .;
3 for  $t \geq 1$  do
4   if  $t > d^2$  then Set  $(P, L) = (p, l)$  and break the for loop ;
5   Run  $alg$  with current memory to obtain a query  $\mathbf{x}_t$ ;
6   if  $F_{\mathbf{A}}(\mathbf{x}) > \eta$  then // Non-informative query
7     | return  $(\|\mathbf{A}\mathbf{x}_t\|_\infty - \eta, \tilde{\mathbf{g}}_{\mathbf{A}}(\mathbf{x}_t))$  as response to  $alg$ .
8   else // Informative query
9     | if  $r \leq k - 1$  and  $F_{\mathbf{A}, \mathbf{v}, p, l}(\mathbf{x}_t) \leq -\eta\gamma_1/2$  and  $\|P_{S_{\text{Span}(\mathbf{x}_{i_{p,r'}, r' \leq r})^\perp}}(\mathbf{x}_t)\|/\|\mathbf{x}_t\| \geq \frac{\gamma_2}{4}$  then
10    |   Set  $i_{p,r+1} = t$  and increment  $r \leftarrow r + 1$ .
11    | if  $F_{\mathbf{A}, \mathbf{v}, p, l}(\mathbf{x}_t) < -\eta(p\gamma_1 + l\gamma_2 + \gamma_2/2)$  and  $r < k$  then
12    |   Compute Gram-Schmidt decomposition  $\mathbf{b}_{p,1}, \dots, \mathbf{b}_{p,r}$  of  $\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,r}}$ .;
13    |   Sample  $\mathbf{y}_{p,l+1}$  uniformly on  $\mathcal{S}^{d-1} \cap \{\mathbf{z} \in \mathbb{R}^d : |\mathbf{b}_{p,r'}^\top \mathbf{z}| \leq d^{-3}, \forall r' \leq r\}$ .;
14    |   Define  $\mathbf{v}_{p,l+1} = \phi_\delta(\mathbf{y}_{p,l+1})$  and increment  $l \leftarrow l + 1$ .
15    | else if  $F_{\mathbf{A}, \mathbf{v}, p, l}(\mathbf{x}_t) < -\eta(p\gamma_1 + l\gamma_2 + \gamma_2/2)$  and  $p + 1 \leq p_{max}$  then
16    |   Set  $l_p = l$  and  $i_{p+1,1} = t$ .;
17    |   Compute the Gram-Schmidt decomposition  $\mathbf{b}_{p+1,1}$  of  $\mathbf{x}_{i_{p+1,1}}$ .;
18    |   Sample  $\mathbf{y}_{p+1,1}$  uniformly on  $\mathcal{S}^{d-1} \cap \{\mathbf{z} \in \mathbb{R}^d : |\mathbf{b}_{p+1,1}^\top \mathbf{z}| \leq d^{-3}\}$ .;
19    |   Define  $\mathbf{v}_{p+1,1} = \phi_\delta(\mathbf{y}_{p+1,1})$ , increment  $p \leftarrow p + 1$  and reset  $l = r = 1$ .
20    | else if  $F_{\mathbf{A}, \mathbf{v}, p, l}(\mathbf{x}_t) < -\eta(p\gamma_1 + l\gamma_2 + \gamma_2/2)$  then // End of the construction
21    |   Set  $l_{p_{max}} = l, i_{p_{max}+1,1} = t$ .;
22    |   Set  $(P, L) = (p_{max}, l)$  and break the for loop.
23    | return  $(F_{\mathbf{A}, \mathbf{v}, p, l}(\mathbf{x}_t), \partial F_{\mathbf{A}, \mathbf{v}, p, l}(\mathbf{x}_t))$  as response to  $alg$ .
24 end

```

**Part 2:** Procedure once  $\mathbf{v}, P, L$  are constructed;

```

25 for  $t' \geq t$  do return  $(F_{\mathbf{A}, \mathbf{v}, P, L}(\mathbf{x}_{t'}), \partial F_{\mathbf{A}, \mathbf{v}, P, L}(\mathbf{x}_{t'}))$  as response to the query  $\mathbf{x}_{t'}$  ;

```

---

**Procedure 1:** The optimization procedure for algorithm  $alg$

The optimization procedure is described in Procedure 1. First, we sample independently  $\mathbf{A} \sim \mathcal{U}(\{\pm 1\}^{n \times d})$  and  $\mathbf{v}_0 \sim \mathcal{U}(\mathcal{D}_\delta)$ . The matrix  $\mathbf{A}$  and vector  $\mathbf{v}_0$  are then fixed for the rest of the learning procedure. Next, we describe the adaptive procedure to return subgradients. It proceeds by periods, until  $p_{max}$  periods are completed, unless the total number of iterations reaches  $d^2$ , in which case the construction procedure ends as well. First, we say that a query is informative if  $F_{\mathbf{A}}(\mathbf{x}) \leq \eta$ . The procedure proceeds by periods  $p \in [p_{max}]$  and in each period constructs the vectors  $\mathbf{v}_{p,1}, \dots, \mathbf{v}_{p,k}$  iteratively. We are now ready to describe the procedure at time  $t$  when the new query  $\mathbf{x}_t$  is queried. Let  $p \geq 1$  be the index of the current period and  $\mathbf{v}_{p,1}, \dots, \mathbf{v}_{p,l}$  be the vectors of this period constructed so far: the first period is  $p = 1$  and we allow  $l = 0$  here. As will be seen in the construction, we always have  $l \geq 1$  except at the very beginning for which we use the notation  $F_{\mathbf{A}, \mathbf{v}, 1, 0} = F_{\mathbf{A}}$ . Together with these vectors, the oracle keeps in memory indices  $i_{p,1}, \dots, i_{p,r}$  with  $r \leq k$  of *exploratory* queries. The constructed vectors from previous periods are  $\mathbf{v}_{p',l'}$  for  $p' < p$  and  $l' \leq l_{p'}$ .

1. If  $\mathbf{x}_t$  is not informative, i.e.  $F_{\mathbf{A}}(\mathbf{x}) > \eta$ , then procedure returns  $(\|\mathbf{A}\mathbf{x}_t\|_\infty - \eta, \tilde{\mathbf{g}}_{\mathbf{A}}(\mathbf{x}_t))$ .

2. Otherwise, we follow the next steps. If  $r \leq k - 1$  and

$$F_{\mathbf{A}, \mathbf{v}, p, l}(\mathbf{x}_t) \leq -\frac{\eta\gamma_1}{2} \quad \text{and} \quad \frac{\|P_{\text{Span}(\mathbf{x}_{i_{p,r'}, r' \leq r})^\perp}(\mathbf{x}_t)\|}{\|\mathbf{x}_t\|} \geq \frac{\gamma_2}{4},$$

we set  $i_{p, r+1} = t$  and increment  $r$ . In this case, we say that  $\mathbf{x}_t$  is *exploratory*. Next,

- (a) Recalling that  $F_{\mathbf{A}, \mathbf{v}, p, l}$  is constructed so far, if  $F_{\mathbf{A}, \mathbf{v}, p, l}(\mathbf{x}_t) \geq \eta(-p\gamma_1 - l\gamma_2 - \gamma_2/2)$ , we do not do anything.
- (b) Otherwise, and if  $r < k$ , let  $\mathbf{b}_{p,1}, \dots, \mathbf{b}_{p,r}$  be the result from the Gram-Schmidt decomposition of  $\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,r}}$ . Then, let  $\mathbf{y}_{p, l+1}$  be a sample of the distribution obtained by the uniform distribution  $\mathbf{y}_{p, l+1} \sim \mathcal{U}(S^{d-1} \cap \{\mathbf{z} \in \mathbb{R}^d : |\mathbf{b}_{p, r'}^\top \mathbf{z}| \leq \frac{1}{d^{\frac{1}{\beta}}}, \forall r' \leq r\})$ . We then pose  $\mathbf{v}_{p, l+1} = \phi_\delta(\mathbf{y}_{p, l+1})$ . Having defined this new vector, we increment  $l$ .
- (c) Otherwise, if  $r = k$ , this ends period  $p$ . We write the total number of vectors defined during period  $p$  as  $l_p := l$ . If  $p + 1 \leq p_{max}$ , period  $p + 1$  starts from  $t = i_{p+1,1}$ . Similarly to above, let  $\mathbf{b}_{p+1,1}$  be the result of the Gram-Schmidt procedure on  $\mathbf{x}_{p+1,1}$ , and we sample  $\mathbf{y}_{p+1,1}$  according to a uniform distribution  $\mathbf{y}_{p+1,1} \sim \mathcal{U}(S^{d-1} \cap \{\mathbf{z} \in \mathbb{R}^d : |\mathbf{b}_{p+1,1}^\top \mathbf{z}| \leq \frac{1}{d^{\frac{1}{\beta}}}\})$ . Then, we pose  $\mathbf{v}_{p+1,1} = \phi_\delta(\mathbf{y}_{p+1,1})$ . We can then increment  $p$  and reset  $l = r = 1$ .

After these steps, with the current values of  $p$  and  $l$ , we return  $(F_{\mathbf{A}, \mathbf{v}, p, l}(\mathbf{x}_t), \partial F_{\mathbf{A}, \mathbf{v}, p, l}(\mathbf{x}_t))$ .

If we finished the last period  $p = p_{max}$ , or if we reached a total number of iterations  $d^2$ , the construction phase of the function ends. In both cases, let us denote by  $P, L$  the last defined period and vector  $\mathbf{v}_{P, L}$ . In particular, we have  $p \leq p_{max}$ . From now on, the final function to optimize is  $F_{\mathbf{A}, \mathbf{v}, P, L}$  and the oracle is a standard first-order oracle for this function, using the subgradient function  $\partial F_{\mathbf{A}, \mathbf{v}, P, L}$ .

We will relate this procedure to the standard convex optimization problem and prove query lower bounds under memory constraints for this procedure. Before doing so, we formally define what we mean by solving this optimization procedure.

**Definition 6.** *Let  $alg$  be an algorithm for convex optimization. We say that an algorithm  $alg$  is successful for the optimization procedure with probability  $q \in [0, 1]$  and accuracy  $\epsilon > 0$ , if taking  $\mathbf{A} \sim \mathcal{U}(\{\pm 1\}^{n \times d})$ , running  $alg$  with the responses given by the procedure, and denoting by  $\mathbf{x}^*(alg)$  the final answer returned by  $alg$ , with probability at least  $q$  over the randomness of  $\mathbf{A}$  and of the procedure, one has*

$$F_{\mathbf{A}, \mathbf{v}, P, L}(\mathbf{x}^*(alg)) \leq \min_{\mathbf{x} \in B_d(0, 1)} F_{\mathbf{A}, \mathbf{v}, P, L}(\mathbf{x}) + \epsilon.$$

### 3.2 Properties and validity of the optimization procedure

We begin this section with a simple lemma showing that during each period  $p$  at most  $l_p \leq k$  vectors  $\mathbf{v}_{p,1}, \dots, \mathbf{v}_{p, l_p}$  are constructed.

**Lemma 7.** *At any time of the construction procedure,  $l \leq r$ . In particular, since  $r \leq k$ , we have  $l_p \leq k$  for all periods  $p \leq p_{max}$ .*

**Proof** Fix a period  $p$ . We prove this by induction. The claim is satisfied for any  $l = 1$  when  $p \geq 2$  since in this case, at the first time  $t = i_{p,1}$  of the period  $p$  we also construct the first vector  $\mathbf{v}_{p,1}$ . For  $p = 1$ , note that the first informative query  $t$  that falls in scenarios (2b) or (2c) is exploratory. Indeed, in these cases we have  $F_{\mathbf{A}, \mathbf{v}, 1, 0}(\mathbf{x}_t) < \eta(-\gamma_1 - \gamma_2/2) \leq -\eta\gamma_1/2$ , and the second criterion for an exploratory query is immediate  $\|P_{\text{Span}(\mathbf{x}_{i_{1,r'}, r' \leq 0})}(\mathbf{x}_t)\| = 0$  since no indices  $i_{1, r'}$  have been defined yet.

We now suppose that the claim holds for  $l - 1 \geq 1$ . Let  $t_{p,l}$  be the time when  $\mathbf{v}_{p,l}$  is constructed and  $i_{p,1}, \dots, i_{p,r}$  the indices constructed until the beginning of iteration  $t_{p,l}$ . If a new index  $i_{p,r'}$  was constructed in times  $(t_{p,l-1}, t_{p,l})$  then the claim holds immediately. Suppose that this is not the case. Note that  $t_{p,l}$  falls in scenario (2b) which means in particular that

$$\eta(\mathbf{v}_{p,l-1}^\top \mathbf{x}_{t_{p,l}} - p\gamma_1 - (l-1)\gamma_2) \leq F_{\mathbf{A}, \mathbf{v}, p, l-1}(\mathbf{x}_{t_{p,l}}) < \eta(-p\gamma_1 - (l-1)\gamma_2 - \gamma_2/2).$$

As a result,

$$|\mathbf{y}_{p,l-1}^\top \mathbf{x}_{t_{p,l}}| \geq |\mathbf{v}_{p,l-1}^\top \mathbf{x}_{t_{p,l}}| - \delta > \frac{\gamma_2}{2} - \delta.$$

Next, when  $r \geq l - 1$  is the number of indices constructed so far, we decompose  $\mathbf{y}_{p,l-1} = \alpha_1 \mathbf{b}_{p,1} + \dots + \alpha_r \mathbf{b}_{p,r} + \tilde{\mathbf{y}}_{p,l-1}$  where  $\tilde{\mathbf{y}}_{p,l-1} \in \text{Span}(\mathbf{x}_{i_{p,r'}}, r' \leq r)^\perp$ . Now by construction of  $\mathbf{y}_{p,l-1}$  one has  $|\alpha_{r'}| \leq d^{-3}$  for all  $r' \leq r$ . Thus,

$$\|\tilde{\mathbf{y}}_{p,l-1} - \mathbf{y}_{p,l-1}\| \leq \frac{\sqrt{r}}{d^3} \leq \frac{1}{d^2 \sqrt{d}}.$$

Therefore,

$$\|P_{\text{Span}(\mathbf{x}_{i_{p,r'}}, r' \leq r)^\perp}(\mathbf{x}_{t_{p,l}})\| \geq |\tilde{\mathbf{y}}_{p,l-1}^\top \mathbf{x}_{t_{p,l}}| \geq |\mathbf{y}_{p,l-1}^\top \mathbf{x}_{t_{p,l}}| - \frac{1}{d^2 \sqrt{d}} > \frac{\gamma_2}{2} - \frac{1}{d^2 \sqrt{d}} - \delta \geq \frac{\gamma_2}{4}.$$

As a result,  $t_{p,l}$  is exploratory, hence  $i_{p,r+1} = t_{p,l}$ . This ends the proof of the recursion and the lemma.  $\blacksquare$

We recall that  $P$  and  $L$  denote the last defined period and vector  $\mathbf{v}_{P,L}$ . From Lemma 7, we have in particular  $P \leq p_{\max}$  and  $L \leq k$ . In the next result, we show that with high probability, the returned values and vectors returned by the above procedure are consistent with a first-order oracle for minimizing the function  $F_{\mathbf{A}, \mathbf{v}, P, L}$ .

**Proposition 8.** *Let  $\mathbf{A} \in \{\pm 1\}^{n \times d}$  and  $\mathbf{v}_0 \in \mathcal{D}_\delta$ . On an event  $\mathcal{E}$  of probability at least  $1 - C\sqrt{\log d}/d^2$  on the randomness of the procedure for some universal constant  $C > 0$ , all responses of the optimization procedure are consistent with a first-order oracle for the function  $F_{\mathbf{A}, \mathbf{v}, P, L}$ : for any  $t \geq 1$ , if  $(f_t, \mathbf{g}_t)$  is the response of the procedure at time  $t$  for query  $\mathbf{x}_t$ , then  $f_t = F_{\mathbf{A}, \mathbf{v}, P, L}(\mathbf{x}_t)$  and  $\mathbf{g}_t = \partial F_{\mathbf{A}, \mathbf{v}, P, L}(\mathbf{x}_t)$ .*

**Proof** Consider a given iteration  $t$ . We aim to show that  $(f_t, \mathbf{g}_t) = (F_{\mathbf{A}, \mathbf{v}, P, L}(\mathbf{x}_t), \partial F_{\mathbf{A}, \mathbf{v}, P, L}(\mathbf{x}_t))$ . By construction, if  $t \geq d^2$ , the result is immediate. Now suppose  $t \leq d^2$ . We first consider the case when  $\mathbf{x}_t$  is non-informative (1). By definition,  $F_{\mathbf{A}}(\mathbf{x}_t) > \eta$ . Since for any  $(p, l) \leq_{\text{lex}} (P, L)$  one has  $|\mathbf{v}_{p,l}^\top \mathbf{x}_t| \leq \|\mathbf{v}_{p,l}\| \|\mathbf{x}_t\| \leq 1$ , we have

$$F_{\mathbf{A}, \mathbf{v}, P, L}(\mathbf{x}_t) = \max \left\{ F_{\mathbf{A}}(\mathbf{x}_t), \eta \left( \max_{(p,l) \leq_{\text{lex}} (P,L)} \mathbf{v}_{p,l}^\top \mathbf{x}_t - p\gamma_1 - l\gamma_2 \right) \right\} = F_{\mathbf{A}}(\mathbf{x}_t).$$

As a result, the response of the procedure for  $\mathbf{x}_t$  is consistent with  $F_{\mathbf{A}, \mathbf{v}, P, L}$  and the returned subgradient is  $\tilde{\mathbf{g}}_{\mathbf{A}}(\mathbf{x}_t) = \partial F_{\mathbf{A}, \mathbf{v}, P, L}(\mathbf{x}_t)$ . Therefore, it suffices to focus on informative queries (2). We will denote by  $t_{p,l}$  the index of the iteration when  $\mathbf{v}_{p,l}$  has been defined, for  $(p, l) \leq_{\text{lex}} (P, L)$ . Consider a specific couple  $(p, l) \leq_{\text{lex}} (P, L)$ , and let  $r$  denote the number of constructed indices on or before  $t_{p,l}$ . Let  $\mathbf{b}_{p,1}, \dots, \mathbf{b}_{p,r}$  the corresponding vectors resulting from the Gram-Schmidt procedure on  $\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,r}}$ . Then, conditionally on the history until time  $t_{p,l}$ , the vector  $\mathbf{v}_{p,l}$  was defined as  $\mathbf{v}_{p,l} = \phi_\delta(\mathbf{y}_{p,l})$ , where

$\mathbf{y}_{p,l}$  is sampled as  $\sim \mathcal{U}(S^{d-1} \cap \{\mathbf{z} \in \mathbb{R}^d : |\mathbf{b}_{p,r'}^\top \mathbf{z}| \leq d^{-3}, \forall r' \leq r\})$ . As a result, from Lemma 21, for any  $t \leq t_{p,l}$ , we have

$$\mathbb{P} \left( |\mathbf{x}_t^\top \mathbf{v}_{p,l}| \geq 3\sqrt{\frac{2 \log d}{d}} + \frac{2}{d^2} \right) \leq \frac{6\sqrt{2 \log d}}{d^6}.$$

We then define the following event

$$\mathcal{E} = \bigcap_{(p,l) \leq_{lex} (P,L)} \bigcap_{t \leq t_{p,l}} \left\{ |\mathbf{x}_t^\top \mathbf{v}_{p,l}| < 3\sqrt{\frac{2 \log d}{d}} + \frac{2}{d^2} \right\},$$

which by the union bound has probability  $\mathbb{P}(\mathcal{E}) \geq 1 - 3\sqrt{2 \log d}/d^2$ . We are now ready to show that the construction procedure is consistent with optimizing  $F_{\mathbf{A},v,P,L}$  on the event  $\mathcal{E}$ . As seen before, we can suppose that  $\mathbf{x}_t$  is informative (2). Using the same notations as before, because  $\mathcal{E}$  is met, for any  $p < p' \leq P$  and  $l' \leq l_{p'}$ , we have for  $d \geq 2$ ,

$$\mathbf{v}_{p',l'}^\top \mathbf{x}_t - p'\gamma_1 - l'\gamma_2 < 3\sqrt{\frac{2 \log d}{d}} + \frac{1}{d} - p\gamma_1 - \gamma_1 \leq -p\gamma_1 - \frac{\gamma_1}{2} \leq -p\gamma_1 - d\gamma_2 - \frac{\gamma_2}{2},$$

where we used  $3\sqrt{2} + 1 \leq 6$  and  $2d\gamma_2 \leq \gamma_1/2$ . As a result, we obtain that

$$\max_{(p',l') \leq_{lex} (P,L), p' > p} \mathbf{v}_{p',l'}^\top \mathbf{x}_t - p'\gamma_1 - l'\gamma_2 < -p\gamma_1 - l\gamma_2 - \frac{\gamma_2}{2}.$$

Next, we consider the case of vectors  $\mathbf{v}_{p,l'}$  where  $l \leq l' \leq l_p$  and  $t_{p,l'} \geq t$  (this also includes the case when we defined  $\mathbf{v}_{p,l}$  at time  $t = t_{p,l}$ ). We write  $\tilde{l}$  for the smallest such index  $l$ . As a remark,  $\tilde{l} \in \{l, l+1\}$ . Note that if such indices exist, this means that before starting iteration  $t$ , the procedure has not yet reached  $r = k$ . There are two cases. If  $\mathbf{x}_t$  was exploratory, we have  $t = i_{p,r}$  hence  $\|P_{Span(\mathbf{b}_{p,r'}, r' \leq r)}^\top(\mathbf{x}_t)\| = 0$ . If  $\mathbf{x}_t$  is not exploratory, either

$$\|P_{Span(\mathbf{b}_{p,r'}, r' \leq r)}^\top(\mathbf{x}_t)\| < \frac{\gamma_2}{4} \|\mathbf{x}_t\| \leq \frac{\gamma_2}{4}, \quad (4)$$

or we have  $F_{\mathbf{A},v,p,l}(\mathbf{x}_t) > -\eta\gamma_1/2$ . We start with the last scenario when  $F_{\mathbf{A},v,p,l}(\mathbf{x}_t) > -\eta\gamma_1/2$ . Then, on  $\mathcal{E}$ , one has

$$\max_{(p,l) <_{lex} (p',l') \leq_{lex} (P,L)} \mathbf{v}_{p',l'}^\top \mathbf{x}_t - p'\gamma_1 - l'\gamma_2 \leq -\gamma_1 + 3\sqrt{\frac{2 \log d}{d}} + \frac{1}{d} \leq -\frac{\gamma_1}{2}$$

As a result, this shows that  $F_{\mathbf{A},v,P,L}(\mathbf{x}_t) = F_{\mathbf{A},v,p,l}(\mathbf{x}_t)$ . Hence using a first-order oracle from  $F_{\mathbf{A},v,l,p}$  at  $\mathbf{x}_t$  is already consistent with  $F_{\mathbf{A},v,P,L}$ . Thus, for whichever step (2a), (2b) or (2c) is performed, since these can only increase the knowledge on  $\mathbf{v}$ , the response given by the construction procedure is consistent with minimizing  $F_{\mathbf{A},v}$ .

It remains to treat the first two scenarios in which we always have Eq (4). In particular, when writing  $\mathbf{x}_t = \alpha_1 \mathbf{b}_{p,1} + \dots + \alpha_r \mathbf{b}_{p,r} + \tilde{\mathbf{x}}_t$  where  $\tilde{\mathbf{x}}_t = P_{Span(\mathbf{b}_{p,r'}, r' \leq r)}^\perp(\mathbf{x}_t)$ , we have  $\|\tilde{\mathbf{x}}_t\| < \frac{\gamma_2}{4}$ . As a result, for  $\tilde{l} \leq l' \leq l_p$ , one has for

$$\begin{aligned} |\mathbf{v}_{p,l'}^\top \mathbf{x}_t| &\leq |\mathbf{y}_{p,l'}^\top \mathbf{x}_t| + \delta \leq |\alpha_1| |\mathbf{y}_{p,l'}^\top \mathbf{b}_{p,1}| + \dots + |\alpha_r| |\mathbf{y}_{p,l'}^\top \mathbf{b}_{p,r}| + \|\tilde{\mathbf{x}}_t\| + \delta \\ &< \|\boldsymbol{\alpha}\|_1 \frac{1}{d^3} + \frac{\gamma_2}{4} + \delta \\ &\leq \frac{\gamma_2}{4} + \frac{1}{d^2 \sqrt{d}} + \frac{1}{d^3} \leq \frac{\gamma_2}{2}, \end{aligned}$$

where in the last inequality we used  $d \geq 3$ . As a result, provided that  $\tilde{l}$  exists, this shows that

$$\max_{\tilde{l} \leq l' \leq l_p} \mathbf{v}_{p,l'}^\top \mathbf{x}_t - p\gamma_1 - l'\gamma_2 = \mathbf{v}_{p,\tilde{l}}^\top \mathbf{x}_t - p\gamma_1 - \tilde{l}\gamma_2 < -p\gamma_1 - \tilde{l}\gamma_2 + \frac{\gamma_2}{2}. \quad (5)$$

On the other hand, if  $t = i_{p+1,1}$ , the same reasoning works for  $t$  viewing it as in period  $p + 1$ , which shows for this case that

$$\max_{l' \leq l_{p+1}} \mathbf{v}_{p+1,l'}^\top \mathbf{x}_t - (p+1)\gamma_1 - l'\gamma_2 = \mathbf{v}_{p+1,1}^\top \mathbf{x}_t - (p+1)\gamma_1 - \gamma_2 < -(p+1)\gamma_1 - \frac{\gamma_2}{2}. \quad (6)$$

As a conclusion of these estimates, we showed that on  $\mathcal{E}$ , we have

$$F_{\mathbf{A},\mathbf{v},P,L}(\mathbf{x}_t) = \max \left\{ F_{\mathbf{A},\mathbf{v},p,l}(\mathbf{x}_t), \eta(\mathbf{v}_{p',l'}^\top \mathbf{x}_t - p'\gamma_1 - l'\gamma_2) \right\} := \tilde{F}_{\mathbf{A},\mathbf{v},t}(\mathbf{x}_t)$$

where  $(p', l')$  is the very next vector that is defined after starting iteration  $t$  (potentially, it has  $t_{p',l'} = t$  if we defined a vector at this time). It now suffices to check that the value and vector returned by the procedure are consistent with the right-hand side. By construction, if we constructed  $\mathbf{v}_{p',l'}$  at step  $t$ : case (2b) or (2c), then the procedure directly uses a first-order oracle for  $\tilde{F}_{\mathbf{A},\mathbf{v},t}$ . Further, by construction of the subgradients since they break ties lexicographically in  $(p, l)$ , the returned subgradient is exactly  $\partial F_{\mathbf{A},\mathbf{v},P,L}(\mathbf{x}_t)$ . It remains to check that this is the case when no vector  $\mathbf{v}_{p',l'}$  is defined at step  $t$ : case (2a). This corresponds to the case when  $F_{\mathbf{A},\mathbf{v},p,l}(\mathbf{x}_t) \geq \eta(-p\gamma_1 - l\gamma_2 - \gamma/2)$ . Now in this case, the upper bound estimates from Eq (5) and Eq (6) imply that

$$\mathbf{v}_{p',l'}^\top \mathbf{x}_t - p'\gamma_1 - l'\gamma_2 < -p\gamma_1 - l\gamma_2 - \gamma/2,$$

and as a result,  $F_{\mathbf{A},\mathbf{v},P,L}(\mathbf{x}_t) = F_{\mathbf{A},\mathbf{v},p,l}(\mathbf{x}_t)$ . Therefore, using a first-order oracle of  $F_{\mathbf{A},\mathbf{v},p,l}$  at  $\mathbf{x}_t$  is valid, and the break of ties of the subgradient of  $\tilde{F}_{\mathbf{A},\mathbf{v},t}$  is the same as the break of ties of  $\partial F_{\mathbf{A},\mathbf{v},P,L}(\mathbf{x}_t)$ . This ends the proof that on  $\mathcal{E}$  the procedure gives responses consistent with an optimization oracle for  $F_{\mathbf{A},\mathbf{v},P,L}$  with subgradient function  $\partial F_{\mathbf{A},\mathbf{v},P,L}$ . Because  $\mathbb{P}(\mathcal{E}) \geq 1 - C\sqrt{\log d}/d^2$  for some constant  $C > 0$ , this ends the proof of the proposition.  $\blacksquare$

Last, we provide an upper bound on the optimal value of  $F_{\mathbf{A},\mathbf{v},P,L}$ .

**Proposition 9.** *Let  $\mathbf{A} \sim \mathcal{U}(\{\pm 1\}^{n \times d})$  and  $\mathbf{v}_0 \sim \mathcal{U}(\mathcal{D}_\delta)$ . For any algorithm  $\text{alg}$  for convex optimization, let  $\mathbf{v}$  be the resulting set of vectors constructed by the randomized procedure. With probability at least  $1 - C\sqrt{\log d}/d$  over the randomness of  $\mathbf{A}$ ,  $\mathbf{v}_0$  and  $\mathbf{v}$ , we have*

$$\min_{\mathbf{x} \in B_d(0,1)} F_{\mathbf{A},\mathbf{v}}(\mathbf{x}) \leq -\frac{\eta}{40\sqrt{(kp_{\max} + 1)\log d}},$$

for some universal constant  $C > 0$ .

**Proof** For simplicity, let us enumerate all the constructed vectors  $\mathbf{v}_1, \dots, \mathbf{v}_{l_{\max}}$  by order of construction. Hence,  $l_{\max} \leq p_{\max}k$ . We use the same numerotation for  $\mathbf{y}_1, \dots, \mathbf{y}_{l_{\max}}$ . Now let  $C_d = \sqrt{40(l_{\max} + 1)\log d}$  and consider the following vector.

$$\bar{\mathbf{x}} = -\frac{1}{C_d} \sum_{l=0}^{l_{\max}} P_{\text{Span}(\mathbf{a}_i, i \leq n)^\perp}(\mathbf{v}_l).$$

In particular, note that we included  $\mathbf{v}_0$  in the sum. For convenience, we write  $P_{\mathbf{A}^\perp}$  instead of  $P_{\text{Span}(\mathbf{a}_i, i \leq n)^\perp}$ . Also, for convenience let us define  $\mathbf{z}_l = \sum_{l' \leq l} P_{\mathbf{A}^\perp}(\mathbf{v}_{l'})$ . Fix an index  $1 \leq l \leq l_{\max}$ . Then, by Lemma 21, with  $t_0 := \sqrt{\frac{6 \log d}{d}} + \frac{2}{d^2}$ , we have

$$\begin{aligned} \mathbb{P}\left(|P_{\mathbf{A}^\perp}(\mathbf{v}_{l+1})^\top \mathbf{z}_l| > t_0 \|\mathbf{z}_l\|\right) &= \mathbb{P}\left(|\mathbf{v}_{l+1}^\top P_{\mathbf{A}^\perp}(\mathbf{z}_l)| > t_0 \|\mathbf{z}_l\|\right) \\ &\leq \mathbb{P}\left(|\mathbf{v}_{l+1}^\top P_{\mathbf{A}^\perp}(\mathbf{z}_l)| > t_0 \|P_{\mathbf{A}^\perp}(\mathbf{z}_l)\|\right) \\ &\leq \frac{2\sqrt{6 \log d}}{d^2}. \end{aligned}$$

Similarly, we have that

$$\mathbb{P}\left(|\mathbf{v}_{l+1}^\top \mathbf{z}_l| > t_0 \|\mathbf{z}_l\|\right) \leq \frac{2\sqrt{6 \log d}}{d^2}.$$

Now consider the event  $\mathcal{E} = \bigcap_{l \leq l_{\max}} \{|\mathbf{v}_l^\top \mathbf{z}_{l-1}|, |P_{\mathbf{A}^\perp}(\mathbf{v}_l)^\top \mathbf{z}_{l-1}| \leq t_0 \|\mathbf{z}_l\|\}$ , which since  $l_{\max} \leq d$ , by the union bound has probability at least  $1 - 4\sqrt{6 \log d}/d$ . Then, on  $\mathcal{E}$ , for any  $l < l_{\max}$ ,

$$\|\mathbf{z}_{l+1}\|^2 \leq \|\mathbf{z}_l\|^2 + \|P_{\mathbf{A}^\perp}(\mathbf{v}_{l+1})\|^2 + 2|P_{\mathbf{A}^\perp}(\mathbf{v}_{l+1})^\top \mathbf{z}_l| \leq \|\mathbf{z}_l\|^2 + 1 + 2t_0 \|\mathbf{z}_l\|.$$

We now prove by induction that  $\|\mathbf{z}_l\|^2 \leq 40 \log d \cdot (l+1)$ , which is clearly true for  $\mathbf{z}_0$  since  $\|\mathbf{z}_0\| = \|P_{\mathbf{A}^\perp}(\mathbf{v}_0)\| \leq \|\mathbf{v}_0\| \leq 1$ . Suppose this is true for  $l < l_{\max}$ . Then, using the above equation and the fact that  $t_0 \leq 3\sqrt{\frac{\log d}{d}}$  for  $d \geq 4$ ,

$$\|\mathbf{z}_{l+1}\|^2 \leq 40 \log d \cdot (l+1) + 1 + 6\sqrt{40} \log d \sqrt{\frac{l+1}{d}} \leq 40 \log d \cdot (l+2),$$

where we used  $l_{\max} + 1 \leq d$ , which completes the induction. In particular, on  $\mathcal{E}$ , we have that  $\|\bar{\mathbf{x}}\| \leq 1$ . Now observe that by construction  $\bar{\mathbf{x}} \in \text{Span}(\mathbf{a}_i, i \leq n)^\perp$  so that  $\|\mathbf{A}\bar{\mathbf{x}}\|_\infty = 0$ . Next, for any  $0 \leq l \leq l_{\max}$ , we have

$$\mathbf{v}_l^\top \bar{\mathbf{x}} = -\frac{\mathbf{v}_l^\top \mathbf{z}_{l_{\max}}}{C_d} = -\frac{1}{C_d} \left( \|P_{\mathbf{A}^\perp}(\mathbf{v}_l)\|^2 + \mathbf{v}_l^\top \mathbf{z}_{l-1} + \sum_{l' < l \leq l_{\max}} \mathbf{v}_l^\top P_{\mathbf{A}^\perp}(\mathbf{v}_{l'}) \right).$$

We will give estimates on each term of the above equation. First, if the indices  $i_{p,1}, \dots, i_{p,r}$  were defined before defining  $\mathbf{v}_l$ , we denote  $\tilde{\mathbf{y}} = P_{\text{Span}(\mathbf{x}_{i_{p,r'}}, r' \leq r)^\perp}(\mathbf{y}_l)$ , the component of  $\mathbf{y}_l$  which is perpendicular to the explored space at that time. Then, we can write  $\mathbf{y}_l = \alpha_1^l \mathbf{b}_{p,1} + \dots + \alpha_r^l \mathbf{b}_{p,r} + \tilde{\mathbf{y}}_l$ , and note that

$$\|\tilde{\mathbf{y}}_l\| = \sqrt{\|\mathbf{y}_l\|^2 - (\alpha_1^l)^2 - \dots - (\alpha_r^l)^2} \geq \sqrt{1 - \frac{k}{d^6}} \geq 1 - \frac{1}{d^5}.$$

Then, we have

$$\begin{aligned} \|P_{\mathbf{A}^\perp}(\mathbf{v}_l)\| &\geq \|P_{\mathbf{A}^\perp}(\mathbf{y}_l)\| - \delta \\ &\geq \|P_{\text{Span}(\mathbf{a}_i, i \leq n, \mathbf{b}_{p,r'}, r' \leq r)^\perp}(\mathbf{y}_l)\| - \delta \\ &= \|P_{\text{Span}(\mathbf{a}_i, i \leq n, \mathbf{b}_{p,r'}, r' \leq r)^\perp}(\tilde{\mathbf{y}}_l)\| - \delta \\ &\geq \left\| P_{\text{Span}(\mathbf{a}_i, i \leq n, \mathbf{b}_{p,r'}, r' \leq r)^\perp} \left( \frac{\tilde{\mathbf{y}}_l}{\|\tilde{\mathbf{y}}_l\|} \right) \right\| - \frac{1}{d^5} - \delta. \end{aligned}$$

As a result, since  $\delta = d^{-3}$ , this shows that

$$\|P_{\mathbf{A}^\perp}(\mathbf{v}_l)\|^2 \geq \left\| P_{\text{Span}(\mathbf{a}_i, i \leq n, \mathbf{b}_{p,r'}, r' \leq r)^\perp} \left( \frac{\tilde{\mathbf{y}}_l}{\|\tilde{\mathbf{y}}_l\|} \right) \right\|^2 - 2\delta.$$

Now observe that  $\dim(\text{Span}(\mathbf{a}_i, i \leq n, \mathbf{b}_{p,r'}, r' \leq r)^\perp) \geq d - n - k$ , while  $\frac{\tilde{\mathbf{y}}_l}{\|\tilde{\mathbf{y}}_l\|}$  is a uniformly random unit vector in  $\text{Span}(\mathbf{b}_{p,r'}, r \leq r')^\perp$ . Therefore, using Proposition 20 we obtain for  $t < 1$ ,

$$\begin{aligned} & \mathbb{P} \left( \|P_{\mathbf{A}^\perp}(\mathbf{v}_l)\|^2 + 2\delta - \frac{d - n - k}{d} \leq -t \right) \\ & \leq \mathbb{P} \left( \left\| P_{\text{Span}(\mathbf{a}_i, i \leq n, \mathbf{b}_{p,r'}, r' \leq r)^\perp} \left( \frac{\tilde{\mathbf{y}}_l}{\|\tilde{\mathbf{y}}_l\|} \right) \right\|^2 - \frac{d - n - k}{d} \leq -t \right) \\ & \leq e^{-(d-k)t^2}. \end{aligned}$$

As a result since  $d - n - k \geq d/2$ , we obtain

$$\mathbb{P} \left( \|P_{\mathbf{A}^\perp}(\mathbf{v}_l)\|^2 \leq \frac{1}{2} - 2\sqrt{\frac{\log d}{d}} - 2\delta \right) \leq \frac{1}{d^2}.$$

Now define  $\mathcal{F} = \bigcap_{l \leq l_{max}} \{ \|P_{\mathbf{A}^\perp}(\mathbf{v}_l)\|^2 \geq \frac{1}{2} - 2\sqrt{\frac{\log d}{d}} - 2\delta \}$ , which since  $l_{max} + 1 \leq d$  and by the union bound has probability at least  $\mathbb{P}(\mathcal{F}) \geq 1 - 1/d$ . Next, we turn to the last term. For any  $0 \leq l < l_{max}$ , we now focus on the sequence  $(\sum_{l'=l+1}^{l+u} \mathbf{v}_l^\top P_{\mathbf{A}^\top}(\mathbf{y}_{l'}))_{1 \leq u \leq l_{max}-l}$  and first note that this is a martingale. These increments are symmetric (because  $\mathbf{y}_{l'}$  is symmetric) even conditionally on  $\mathbf{A}$  and  $\mathbf{v}_l, \mathbf{y}_l, \dots, \mathbf{y}_{l'-1}$ . Next, let  $t_1 = 2\sqrt{\frac{3 \log d}{d}} + \frac{2}{d^2}$ . Note that for  $d \geq 4$ , we have  $t_1 \leq 4\sqrt{\frac{\log d}{d}}$ . Further, by Lemma 21,

$$\mathbb{P}(|\mathbf{v}_l^\top P_{\mathbf{A}^\top}(\mathbf{y}_{l'})| > t_1) = \mathbb{P}(|P_{\mathbf{A}^\top}(\mathbf{v}_l)^\top \mathbf{y}_{l'}| > t_1) \leq \frac{4\sqrt{3 \log d}}{d^4},$$

where we used the fact that  $P_{\mathbf{A}^\perp}$  is a projection. Let  $\mathcal{G}_l = \bigcap_{l < l' \leq l_{max}} \{ |\mathbf{v}_l^\top P_{\mathbf{A}^\top}(\mathbf{v}_{l'})| \leq t_1 \}$ , which by the union bound has probability  $\mathbb{P}(\mathcal{G}_l) \geq 1 - 4\sqrt{3 \log d}/d^3$ . Next, we define  $I_{l,u} = (\mathbf{v}_l^\top P_{\mathbf{A}^\top}(\mathbf{y}_{l+u}) \wedge t_1) \vee (-t_1)$ , the increments capped at absolute value  $t_1$ . Because  $\mathbf{v}_l^\top P_{\mathbf{A}^\top}(\mathbf{y}_{l+u})$  is symmetric, so is  $I_{l,u}$ . As a result, these are bounded increments of a martingale, to which we can apply the Azuma-Hoeffding inequality.

$$\mathbb{P} \left( \left| \sum_{u=1}^{l_{max}-l} I_{l,u} \right| \leq 2t_1 \sqrt{(l_{max}-l) \log d} \right) \geq 1 - \frac{2}{d^2}.$$

We denote by  $\mathcal{H}_l$  this event. Now observe that on  $\mathcal{G}_l$ , the increments  $I_{l,u}$  and  $\mathbf{v}_l^\top P_{\mathbf{A}^\top}(\mathbf{y}_{l+u})$  coincide for all  $1 \leq u \leq l_{max} - l$ . As a result, on  $\mathcal{G}_l \cap \mathcal{H}_l$  we obtain

$$\begin{aligned} \left| \sum_{l < l' \leq l_{max}} \mathbf{v}_l^\top P_{\mathbf{A}^\perp}(\mathbf{v}_{l'}) \right| & \leq \left| \sum_{l < l' \leq l_{max}} \mathbf{v}_l^\top P_{\mathbf{A}^\perp}(\mathbf{y}_{l'}) \right| + (l_{max} - 1)\delta \\ & \leq \left| \sum_{u=1}^{l_{max}-l} I_{l,u} \right| + (d-2)\delta \\ & \leq 2t_1 \sqrt{l_{max} \log d} + (d-2)\delta. \end{aligned}$$



Then, on the event  $\mathcal{E} \cap \mathcal{F} \cap \bigcap_{l \leq l_{max}} \mathcal{G}_l \cap \mathcal{H}_l$ , for any  $1 \leq l \leq l_{max}$  one has

$$\begin{aligned} \mathbf{v}_l^\top \mathbf{z}_{l_{max}} &\geq \frac{1}{2} - 2\sqrt{\frac{\log d}{d}} - t_0 \|\mathbf{z}_l\| - 2t_1 \sqrt{l_{max} \log d} - \frac{1}{d^2} \\ &\geq \frac{1}{2} - 2\sqrt{\frac{\log d}{d}} - 3 \log d \sqrt{40 \frac{l_{max} + 1}{d}} - 8 \log d \sqrt{\frac{l_{max}}{d}} - \frac{1}{d^2} \\ &\geq \frac{1}{2} - 30 \log d \sqrt{\frac{l_{max} + 1}{d}} \\ &\geq \frac{1}{6}, \end{aligned}$$

where in the last inequalities we used the fact that  $l_{max} \leq kp_{max} \leq c_{d,1}d - 1$  where  $c_{d,1} = \frac{1}{90^2 \log^2 d}$  as per Eq (3). As a result, we obtain that on  $\mathcal{E} \cap \mathcal{F} \cap \bigcap_{l \leq l_{max}} \mathcal{G}_l \cap \mathcal{H}_l$ , which has probability at most  $1 - C\sqrt{\log d}/d$  for some constant  $C > 0$ ,

$$\max_{p \leq p_{max}, l \leq k} \mathbf{v}_{p,l}^\top \bar{\mathbf{x}} \leq -\frac{1}{6C_d} \leq -\frac{1}{40\sqrt{(kp_{max} + 1) \log d}}.$$

Since  $\|\mathbf{A}\bar{\mathbf{x}}\|_\infty = 0$ , and  $\eta \geq \frac{\eta}{40\sqrt{(kp_{max} + 1) \log d}}$ , this shows that

$$F_{\mathbf{A},\mathbf{v}}(\bar{\mathbf{x}}) \leq -\frac{\eta}{40\sqrt{(kp_{max} + 1) \log d}}.$$

This ends the proof of the proposition. ■

### 3.3 Reduction from convex optimization to the optimization procedure

According to Proposition 8, with probability at least  $1 - C\sqrt{\log d}/d^2$ , the procedure returns responses that are consistent with a first-order oracle of the function  $F_{\mathbf{A},\mathbf{v},P,L}$  where  $\mathbf{v}_{P,L}$  is the last vector to have been defined. Now observe that for any constructed vectors  $\mathbf{v}$ , the function  $F_{\mathbf{A},\mathbf{v},P,L}$  is  $\sqrt{d}$ -Lipschitz. As a result, if there exists an algorithm for convex optimization that guarantees  $\epsilon$  accuracy for 1-Lipschitz functions, by rescaling, there exists an algorithm *alg* which is successful for the optimization procedure with probability  $1 - C\sqrt{\log d}/d^2$  and  $\epsilon\sqrt{d}$  accuracy. In the next proposition, we show that to be successful, such an algorithm needs to properly define the complete function  $F_{\mathbf{A},\mathbf{v}}$ , i.e., to complete all periods until  $p_{max}$ .

**Proposition 10.** *Let  $alg$  be a successful algorithm for the optimization procedure with probability  $q \in [0, 1]$  and precision  $\eta/(2\sqrt{d})$ . Suppose that  $alg$  performs at most  $d^2$  queries during the optimization procedure. Then when running  $alg$  with the responses of the optimization procedure,  $alg$  succeeds and ends the period  $p_{max}$  with probability at least  $q - C\sqrt{\log d}/d$  for some universal constant  $C > 0$ .*

**Proof** Let  $\mathbf{x}^*(alg) = \mathbf{x}_T$  denote the final answer of  $alg$  when run with the optimization procedure. By hypothesis, we have  $T \leq d^2$ . As before, let  $P \leq p_{max}$  and  $L \leq k$  be the indices such that the last vector constructed by the optimization procedure is  $\mathbf{v}_{P,L}$ . Let  $\mathcal{E}$  be the event when  $alg$  run on the optimization procedure does not end period  $p_{max}$ . We focus on  $\mathcal{E}$  and consider two cases.

First, suppose that  $T > t_{P,L}$ , i.e., the last vector was not constructed at time  $T$ . As a result, this means that  $\mathbf{x}_T$  corresponds either to a non-informative query—scenario (1)—in which case  $F_{\mathbf{A},\mathbf{v},P,L}(\mathbf{x}_T) \geq F_{\mathbf{A}}(\mathbf{x}_T) \geq \eta$ , or this means that  $F_{\mathbf{A},\mathbf{v},P,L}(\mathbf{x}_T) \geq \eta(-P\gamma_1 - L\gamma_2 - \gamma/2)$ —scenario (2a).

Second, we now suppose that  $T = t_{P,L}$ , i.e., the last vector was constructed at time  $T$ . Then, by construction of  $\mathbf{v}_{P,L}$  and  $\mathbf{y}_{P,L}$ , we have indices  $i_{P,1}, \dots, i_{P,r} \leq T$  such that with the Gram-Schmidt decomposition  $\mathbf{b}_{P,1}, \dots, \mathbf{b}_{P,r}$  of  $\mathbf{x}_{i_{P,1}}, \dots, \mathbf{x}_{i_{P,r}}$ , we have  $|\mathbf{b}_{P,r'}^\top \mathbf{y}_{P,L}| \leq d^{-3}$  for all  $r' \leq r$ . In particular, writing  $\mathbf{x}_T = \alpha_1 \mathbf{b}_{P,1} + \dots + \alpha_r \mathbf{b}_{P,r} + \tilde{\mathbf{x}}_T$ , where  $\tilde{\mathbf{x}}_T \in \text{Span}(\mathbf{x}_{i_{P,r'}}, r' \leq r)^\perp$ , either we have  $i_{P,r} = T$ , in which case  $\tilde{\mathbf{x}}_T = \mathbf{0}$ , or  $\mathbf{x}_T$  was not exploratory in which case we directly have  $F_{\mathbf{A},\mathbf{v},P,L}(\mathbf{x}_T) \geq F_{\mathbf{A},\mathbf{v},P,L-1}(\mathbf{x}_T) > -\eta\gamma_1/2$ , or we have  $\|\tilde{\mathbf{x}}_T\| < \|\mathbf{x}_T\|\gamma_2/4 \leq \gamma_2/4$ . For all remaining cases to consider, we obtain

$$|\mathbf{v}_{P,L}^\top \mathbf{x}_T| \leq |\mathbf{y}_{P,L}^\top \mathbf{x}_T| + \delta \leq \frac{\|\boldsymbol{\alpha}\|_1}{d^3} + \|\tilde{\mathbf{x}}_T\| + \delta \leq \frac{1}{d^3} + \frac{1}{d^2\sqrt{d}} + \frac{\gamma_2}{4} < \frac{\gamma_2}{2}.$$

In the last inequality, we used  $d \geq 4$ . This shows that  $F_{\mathbf{A},\mathbf{v},P,L}(\mathbf{x}_T) \geq \eta(-P\gamma_1 - L\gamma_2 - \gamma_2/2)$ . As a result, in all cases this shows that  $F_{\mathbf{A},\mathbf{v},P,L}(\mathbf{x}^*(\text{alg})) \geq \eta(-P\gamma_1 - L\gamma_2 - \gamma_2/2) \geq -\eta(p_{\max} + 1)\gamma_1$ . Now define the event

$$\mathcal{F} = \left\{ \min_{\mathbf{x} \in B_d(0,1)} F_{\mathbf{A},\mathbf{v}}(\mathbf{x}) \leq -\frac{\eta}{40\sqrt{(kp_{\max} + 1)\log d}} \right\}.$$

By Proposition 9 we have  $\mathcal{P}(\mathcal{F}) \geq 1 - C\sqrt{\log d}/d$ . Now from Eq (3),

$$(p_{\max} + 1)^{3/2} \leq \frac{1}{60\gamma_1\sqrt{k\log d}}.$$

Thus,

$$(p_{\max} + 1)\gamma_1 \leq \frac{1}{60\sqrt{k(p_{\max} + 1)\log d}} \leq \frac{1}{60\sqrt{(kp_{\max} + 1)\log d}}$$

Then, since  $F_{\mathbf{A},\mathbf{v},P,L} \leq F_{\mathbf{A},\mathbf{v}}$ , this shows that on  $\mathcal{E} \cap \mathcal{F}$ ,

$$\begin{aligned} F_{\mathbf{A},\mathbf{v},P,L}(\mathbf{x}^*(\text{alg})) &\geq -\eta(p_{\max} + 1)\gamma_1 \geq \min_{\mathbf{x} \in B_d(0,1)} F_{\mathbf{A},\mathbf{v}}(\mathbf{x}) + \frac{\eta}{120\sqrt{(kp_{\max} + 1)\log d}} \\ &> \min_{\mathbf{x} \in B_d(0,1)} F_{\mathbf{A},\mathbf{v},P,L}(\mathbf{x}) + \frac{\eta}{2\sqrt{d}} \end{aligned}$$

where in the last inequality, we used  $kp_{\max} \leq c_{d,1}d - 1$ . As a result, letting  $\mathcal{G}$  be the event when  $\text{alg}$  succeeds for precision  $\epsilon = \eta/(2\sqrt{d})$ . By hypothesis,  $\mathcal{P}(\mathcal{G}) \geq q$ . Now from the above equations, one has  $\mathcal{E} \cap \mathcal{F} \cap \mathcal{G} = \emptyset$ . Therefore,  $\mathbb{P}(\mathcal{G} \cap \mathcal{E}^c) \geq \mathcal{P}(\mathcal{G}) - \mathbb{P}(\mathcal{G} \cap \mathcal{E} \cap \mathcal{F}) - \mathbb{P}(\mathcal{F}^c) \geq q - C\sqrt{\log d}/d$ . This ends the proof of the proposition.  $\blacksquare$

### 3.4 Reduction of the optimization procedure to an Orthogonal Vector Game with Hints

We are now ready to introduce an orthogonal vector game where the main difference with the game introduced in [1] is that the player can provide additional hints.

We first prove that solving the optimization procedure implies solving the Orthogonal Vector Game with Hints.

**Proposition 11.** *Let  $m \leq d$ . Suppose that there is an  $M$ -bit algorithm that is successful for the optimization procedure with probability  $q$  for accuracy  $\epsilon = \eta/(2\sqrt{d})$  and uses at most  $mp_{\max}$  queries. Then, there is an algorithm for Game 2 for parameters  $(d, k, m, M, \alpha = \frac{2\eta}{\gamma_1}, \beta = \frac{\gamma_2}{4})$ , for which the Player wins with probability at least  $q - C\sqrt{\log d}/d$  for some universal constant  $C > 0$ .*

---

**Input:**  $d, k, m, M, \alpha, \beta$

- 1 *Oracle:* Set  $n \leftarrow \lfloor d/4 \rfloor$ , sample  $\mathbf{A} \sim \mathcal{U}(\{\pm 1\}^{n \times d})$ ;
- 2 *Player:* Observe  $\mathbf{A}$ ;
- 3 **for**  $l \in [d]$  **do**
- 4     *Player:* Based on  $\mathbf{A}$  and any previous queries and responses, submit at most  $k$  vectors  $\mathbf{x}_{l,1}, \dots, \mathbf{x}_{l,r_l}$ ;
- 5     *Oracle:* Perform the Gram-Schmidt decomposition  $\mathbf{b}_{l,1}, \dots, \mathbf{b}_{l,r_l}$  of  $\mathbf{x}_{l,1}, \dots, \mathbf{x}_{l,r_l}$ . Then, sample a vector  $\mathbf{y}_l \in S^{d-1}$  according to a uniform distribution  $\mathcal{U}(S^{d-1} \cap \{\mathbf{z} \in \mathbb{R}^d : \forall r \leq r_l, |\mathbf{b}_{l,r}^\top \mathbf{z}| \leq d^{-3}\})$ . As response to the query, return  $\mathbf{v}_l = \phi_\delta(\mathbf{y}_l)$  to the player.
- 6 **end**
- 7 *Player:* Based on  $\mathbf{A}$ , all previous queries and responses, store an  $M$ -bit message Message.;
- 8 *Player:* Based on  $\mathbf{A}$ , all previous queries and responses, submit a function  $\mathbf{g} : B_d(0, 1) \rightarrow (\{\mathbf{a}_j, j \leq n\} \cup \{\mathbf{v}_l, l \leq d\}) \times [d^2]$  to the Oracle.
- 9 **for**  $i \in [m]$  **do**
- 10     *Player:* Based on Message, any previous queries  $\mathbf{x}_1, \dots, \mathbf{x}_{i-1}$  and responses  $\mathbf{g}_1, \dots, \mathbf{g}_{i-1}$  from this loop phase, submit a query  $\mathbf{x}_i \in \mathbb{R}^d$ ;
- 11     *Oracle:* As the response to query  $\mathbf{z}_i$ , return  $\mathbf{g}_i = \mathbf{g}(\mathbf{z}_i)$ .
- 12 **end**
- 13 *Player:* Based on all queries and responses from this phase  $\{\mathbf{z}_i, \mathbf{g}_i, i \in [m]\}$ , and on Message, return some vectors  $\mathbf{y}_1, \dots, \mathbf{y}_k$  to the oracle.;
- 14 The player wins if the returned vectors have unit norm and satisfy for all  $i \in [k]$ 
  1.  $\|\mathbf{A}\mathbf{y}_i\|_\infty \leq \alpha$
  2.  $\|P_{\text{Span}(\mathbf{y}_1, \dots, \mathbf{y}_{i-1})^\perp}(\mathbf{y}_i)\|_2 \geq \beta$ .

---

### Game 2: Orthogonal Vector Game with Hints

**Proof** Let  $alg$  be an  $M$ -bit algorithm solving the feasibility problem with  $mp_{max}$  queries with probability at least  $q$ . We now describe the strategy for Game 2.

In the first part of the strategy, the player observes  $\mathbf{A}$ . First, submit an empty query to the Oracle to obtain a vector  $\mathbf{v}_0$ , which as a result is uniformly distributed among  $\mathcal{D}_\delta$ . We then proceed to simulate the optimization procedure for  $alg$  using parameters  $\mathbf{A}$  and  $\mathbf{v}_0$  (lines 3-6 of Game 2). Precisely, whenever a new vector  $\mathbf{v}_{p,l}$  needs to be defined according to the optimization procedure, the player submits the corresponding vectors  $\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,r}}$  to the oracle and receives in return a vector which defines  $\mathbf{v}_{p,l}$ . In this manner, the player simulates exactly the optimization procedure. In all cases, the number of queries in this first phase is at most  $1 + kp_{max} \leq d$ . For the remaining queries to perform, the player can query whichever vectors, these will not be used in the rest of the strategy. If the simulation did not end period  $p_{max}$ , the complete procedure fails. We now describe the rest of the procedure when period  $p_{max}$  was ended. During the simulation, the algorithm records the time  $i_{p,1}$  when period  $p$  started for all  $p \leq p_{max} + 1$ . Recall that for  $p_{max} + 1$ , we only define  $i_{p_{max}+1,1}$ , this is the time that ends period  $p_{max}$ . Now by hypothesis,  $i_{p_{max}+1,1} \leq mp_{max}$ . As a result, there must be a period  $p \leq p_{max}$  which uses at most  $m$  queries:  $i_{p+1,1} - i_{p,1} \leq m$ . We define the memory Message to be the memory of  $alg$  just before starting iteration  $i_{p,1}$ , at the beginning of period  $p$  (line 7 of Game 2). Next, since the period  $p_{max}$  was ended, the vectors  $\mathbf{v}_{p,l}$  for  $p \leq p_{max}, l \leq l_p$  were all defined. The player can therefore submit

the function  $\mathbf{g}_{\mathbf{A},\mathbf{v}}$  to the Oracle (line 8 of Game 2) as follows,

$$\mathbf{g}_{\mathbf{A},\mathbf{v}} : \mathbf{x} \mapsto \begin{cases} (\mathbf{g}_{\mathbf{A}}(\mathbf{x}), 1) & \text{if } F_{\mathbf{A},\mathbf{v}}(\mathbf{x}) = \|\mathbf{A}\mathbf{x}\|_{\infty} - \eta, \\ (\mathbf{v}_0, 2) & \text{otherwise and if } F_{\mathbf{A},\mathbf{v}}(\mathbf{x}) = \eta\mathbf{v}_0^{\top}\mathbf{x}, \\ (\mathbf{v}_{p,l}, 2 + (p-1)k + l) & \text{otherwise and if} \\ & (p,l) = \arg \max_{(p',l') \leq \text{lex}(p_{\max}, l_{\max})} \mathbf{v}_{p',l'}^{\top}\mathbf{x} - p\gamma_1 - l\gamma_2. \end{cases} \quad (7)$$

Intuitively, the first component of  $\mathbf{g}_{\mathbf{A},\mathbf{v}}$  gives the subgradient  $\partial F_{\mathbf{A},\mathbf{v}}$  to the following two exceptions: we always return  $\mathbf{a}_i$  instead of  $\pm\mathbf{a}_i$  and we return  $\mathbf{v}_0$  (resp.  $\mathbf{v}_{p,l}$ ) instead of  $\eta\mathbf{v}_0$  (resp.  $\eta\mathbf{v}_{p,l}$ ). The second term of  $\mathbf{g}_{\mathbf{A},\mathbf{v}}$  has values in  $[2 + p_{\max}k]$ . Hence, since  $2 + p_{\max}k \leq d^2$ , the function  $\mathbf{g}_{\mathbf{A},\mathbf{v}}$  takes values in  $(\{\mathbf{a}_j, j \leq n\} \cup \{\mathbf{v}_l, l \leq d\}) \times [d^2]$ .

The strategy then proceeds to play the Orthogonal Vector Game in a second part (lines 9-12 of Game 2) and uses the responses of the Oracle to simulate the run of *alg* for the optimization procedure in period  $p$ . To do so, we set the memory state of the algorithm *alg* to be Message. Then, for the next  $m$  iterations we proceed as follows. At iteration  $i$  of the process, we run *alg* with its current state to obtain a new query  $\mathbf{z}_i$  which is then submitted to the oracle of the Orthogonal Vector Game, to get a response  $(\mathbf{g}_i, s_i)$ . We then use this response to simulate the response that was given by the optimization procedure in the first phase, computing  $(v_i, \tilde{\mathbf{g}}_i)$  as follows

$$(v_i, \tilde{\mathbf{g}}_i) = \begin{cases} (|\mathbf{g}_i^{\top}\mathbf{z}_i| - \eta, \text{sign}(\mathbf{g}_i^{\top}\mathbf{z}_i)\mathbf{g}_i) & s_i = 1, \\ (\eta\mathbf{g}_i^{\top}\mathbf{z}_i, \eta\mathbf{g}_i) & s_i = 2, \\ (\eta(\mathbf{g}_i^{\top}\mathbf{z}_i - p\gamma_1 - l\gamma_2), \eta\mathbf{g}_i) & s_i = 2 + (p-1)k + l, p \leq p_{\max}, 1 \leq l \leq k. \end{cases} \quad (8)$$

We can easily check that in all cases,  $v_i = F_{\mathbf{A},\mathbf{v}}(\mathbf{z}_i)$  and that  $\tilde{\mathbf{g}}_i = \partial F_{\mathbf{A},\mathbf{v}}(\mathbf{z}_i)$ . We then pass  $(v_i, \tilde{\mathbf{g}}_i)$  as response to *alg* for the query  $\mathbf{z}_i$  so it can update its state. Further, having defined  $i_1 = 1$ , the player can keep track of exploratory queries by checking whether

$$v_i \leq -\frac{\eta\gamma_1}{2} \quad \text{and} \quad \frac{\|P_{\text{Span}(\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_r})^{\perp}}(\mathbf{z}_i)\|}{\|\mathbf{z}_i\|} \geq \frac{\gamma_2}{4},$$

where  $i_1, \dots, i_r$  are the indices defined so far. We perform  $m$  such iterations unless *alg* stops and use the last remaining queries arbitrarily. Next, we check if the last index  $i_k$  was defined. If not, we pose  $i_k = m + 1$  and let  $\mathbf{z}_{m+1}$  be the next query of *alg*. The final returned vectors are  $\frac{\mathbf{z}_{i_1}}{\|\mathbf{z}_{i_1}\|}, \dots, \frac{\mathbf{z}_{i_k}}{\|\mathbf{z}_{i_k}\|}$ . This ends the description of the player's strategy.

We now show that the player wins with good probability. First, since *alg* makes at most  $mp_{\max} \leq d^2$  queries, by Proposition 10, on an event  $\mathcal{E}$  of probability at least  $q - C\sqrt{\log d}/d$ , *alg* succeeds and ends the period  $p_{\max}$ . On  $\mathcal{E}$ , by construction, the first phase of the strategy does not fail. Now we show that in the second phase (lines 9-12 of Game 2), the queried vectors coincide exactly with the queried vectors from the corresponding period  $p$  in the first phase (lines 3-6 of Game 2). To do so, we only need to check that the responses provided to *alg* coincide with the response given by the optimization procedure. First, recall that on  $\mathcal{E}$ , all periods are completed, hence  $F_{\mathbf{A},\mathbf{v},P,L} = F_{\mathbf{A},\mathbf{v}}$ . Next, by Proposition 8, the responses of the procedure are consistent with optimizing  $F_{\mathbf{A},\mathbf{v},P,L}$  and subgradients  $\partial F_{\mathbf{A},\mathbf{v},P,L}$  on an event  $\mathcal{F}$  of probability at least  $1 - C'\sqrt{\log d}/d^2$ . Therefore, on  $\mathcal{E} \cap \mathcal{F}$ , it suffices to check that the responses provided to *alg* are consistent with  $F_{\mathbf{A},\mathbf{v}}$ , which we already noted: at every step  $i$ ,  $(v_i, \tilde{\mathbf{g}}_i) = (F_{\mathbf{A},\mathbf{v}}(\mathbf{z}_i), \partial F_{\mathbf{A},\mathbf{v}}(\mathbf{z}_i))$ . This proves that the responses and queries coincide exactly with those given by the optimization procedure on  $\mathcal{E} \cap \mathcal{F}$ .

---

**Input:**  $d, k, p_{max}, m$ , algorithm  $alg$

**Part 1:** Strategy to store Message knowing  $A$ ;

- 1 Initialize the memory of  $alg$  to be  $\mathbf{0}$ ;
- 2 Submit  $\emptyset$  to the Oracle and use the response as  $\mathbf{v}_0$ ;
- 3 Run  $alg$  with the optimization procedure knowing  $A$  and  $\mathbf{v}_0$  until the first exploratory query  $\mathbf{x}_{i_1,1}$ .
- 4 **for**  $p \in [p_{max}]$  **do**
- 5     Let  $\text{Memory}_p$  be the current memory state of  $alg$  and  $i_{p,1}$  the current iteration step. ;
- 6     Run  $alg$  with the feasibility procedure until period  $p$  ends at iteration step  $i_{p+1,1}$ . If  $alg$  stopped before, **return** the strategy fails. When needed to sample a unit vector  $\mathbf{v}_{p',l'}$ , submit vectors  $\mathbf{x}_{i_{p',1}}, \dots, \mathbf{x}_{i_{p',r'}}$  to the Oracle where  $i_{p',1}, \dots, i_{p',r'}$  are the exploratory queries defined at that stage. We use the corresponding response of the Oracle as  $\mathbf{v}_{p',l'}$ ;
- 7     **if**  $i_{p+1,1} - i_{p,1} \leq m$  **then**
- 8         | Set Message =  $\text{Memory}_p$
- 9 **end**
- 10 **for** Remaining queries to perform to Oracle **do** Submit arbitrary query, e.g.  $\emptyset$  ;
- 11 **if** Message has not been defined yet **then return** The strategy fails;
- 12 Submit  $\mathbf{g}_{A,v}$  to the Oracle as defined in Eq (7).;

**Part 2:** Strategy to make queries;

- 13 Set the memory state of  $alg$  to be Message and define  $i_1 = 1, r = 1$ ;
- 14 **for**  $i \in [m]$  **do**
- 15     Run  $alg$  with current memory to obtain a query  $\mathbf{z}_i$ ;
- 16     Submit  $\mathbf{z}_i$  to the Oracle from Game 2, to get response  $(\mathbf{g}_i, s_i)$ ;
- 17     Compute  $(v_i, \tilde{\mathbf{g}}_i)$  using  $\mathbf{z}_i, \mathbf{g}_i$  and  $s_i$  as defined in Eq (8) and pass  $(v_i, \tilde{\mathbf{g}}_i)$  as response to  $alg$ ;
- 18     **if**  $v_i \leq -\eta\gamma_1/2$  and  $\|P_{\text{Span}(\mathbf{z}_{i_r}, r' \leq r)}^\perp(\mathbf{z}_i)\|/\|\mathbf{z}_i\| \geq \frac{\gamma_2}{4}$  **then**
- 19         | Set  $i_{r+1} = i$  and increment  $r \leftarrow r + 1$ .
- 20 **end**

**Part 3:** Strategy to return vectors;

- 21 **if** index  $i_k$  has not been defined yet **then**
  - 22     | With the current memory of  $alg$  find a new query  $\mathbf{z}_{m+1}$  and set  $i_k = m + 1$ ;
  - 23 **return**  $\left\{ \frac{\mathbf{z}_{i_1}}{\|\mathbf{z}_{i_1}\|}, \dots, \frac{\mathbf{z}_{i_k}}{\|\mathbf{z}_{i_k}\|} \right\}$  to the Oracle.
- 

**Algorithm 3:** Strategy of the Player for the Orthogonal Vector Game with Hints

Next, by construction, the chosen phase  $p$  had at most  $m$  iterations. Thus, on  $\mathcal{E} \cap \mathcal{F}$ , among  $\mathbf{z}_1, \dots, \mathbf{z}_{m+1}$ , we have the vectors  $\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,k}}$ . Further, if  $i_k$  was not defined during part 2 of the strategy, this means that  $i_k = m + 1$ , as defined in the player's strategy (line 21-22 of Algorithm 3). As a result, for all  $u \leq k$ , we have  $\mathbf{z}_{i_u} = \mathbf{x}_{i_{p,u}}$ . We now show that the returned vectors  $\frac{\mathbf{x}_{i_{p,1}}}{\|\mathbf{x}_{i_{p,1}}\|}, \dots, \frac{\mathbf{x}_{i_{p,k}}}{\|\mathbf{x}_{i_{p,k}}\|}$  are successful for Game 2. First, because  $i_{p,1}, \dots, i_{p,k}$  are exploratory queries, we have directly for  $u \leq k$ ,

$$\frac{\|P_{\text{Span}(\mathbf{x}_{i_{p,v}, v < u})^\perp}(\mathbf{x}_{i_{p,u}})\|}{\|\mathbf{x}_{i_{p,u}}\|} \geq \frac{\gamma_2}{4}.$$

Next, if  $l$  is the index of the last constructed vector  $\mathbf{v}_{p,l}$  before  $i_{p,u}$  in the optimization procedure, one has  $F_{\mathbf{A}, \mathbf{v}_{p,l}}(\mathbf{x}_{i_{p,u}}) \leq -\eta\gamma_1/2$ . Therefore,  $\|\mathbf{A}\mathbf{x}_{i_{p,u}}\|_\infty \leq F_{\mathbf{A}, \mathbf{v}_{p,l}}(\mathbf{x}_{i_{p,u}}) + \eta \leq \eta$ . Further,  $\eta\mathbf{v}_0^\top \mathbf{x}_{i_{p,u}} \leq F_{\mathbf{A}, \mathbf{v}_{p,l}}(\mathbf{x}_{i_{p,u}}) \leq -\eta\gamma_1/2$ . This proves that  $\|\mathbf{x}_{i_{p,u}}\| \geq \gamma_1/2$ . Putting the previous two inequalities together yields

$$\frac{\|\mathbf{A}\mathbf{x}_{i_{p,u}}\|_\infty}{\|\mathbf{x}_{i_{p,u}}\|} \leq \frac{2\eta}{\gamma_1}.$$

As a result, this shows that the returned vectors are successful for Game 2 for the desired parameters  $\alpha = 2\eta/\gamma_1$  and  $\beta = \gamma_2/4$ . Thus, the player wins on  $\mathcal{E} \cap \mathcal{F}$ , which has probability at least  $q - (C + C')\sqrt{\log d}/d^2$  by the union bound. This ends the proof of the proposition.  $\blacksquare$

### 3.5 Query lower bound for the Orthogonal Vector Game with Hints

Before proving a lower bound on the necessary number of queries for Game 2, we need to introduce two results. The first one is a known concentration result for vectors in the hypercube. It shows that for a uniform vector in the hypercube, being approximately orthogonal to  $k$  orthonormal vectors has exponentially small probability in  $k$ .

**Lemma 12 ([1]).** *Let  $\mathbf{h} \sim \mathcal{U}(\{\pm 1\}^d)$ . Then, for any  $t \in (0, 1/2]$  and any matrix  $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_k] \in \mathbb{R}^{d \times k}$  with orthonormal columns,*

$$\mathbb{P}(\|\mathbf{Z}^\top \mathbf{h}\|_\infty \leq t) \leq 2^{-c_H k}.$$

We will also need an anti-concentration bound for random vectors, which intuitively provides a lower bound for the previous concentration result. The following lemma shows that for a uniformly random unit vector, being orthogonal to  $k$  orthonormal vectors is still achievable with exponentially small probability in  $k$ .

**Lemma 13.** *Let  $k < d$  and  $\mathbf{x}_1, \dots, \mathbf{x}_k$  be  $k$  orthonormal vectors. Then,*

$$\mathbb{P}_{\mathbf{y} \sim \mathcal{U}(S^{d-1})} \left( |\mathbf{x}_i^\top \mathbf{y}| \leq \frac{1}{d^3}, \forall i \leq k \right) \geq \frac{1}{e^{d-4} d^{3k}}.$$

**Proof** Let  $\mathbf{y} \sim \mathcal{U}(S^{d-1})$  be a uniformly random unit vector. Then, for  $i < k$  and any  $y_1, \dots, y_{i-1}$  such

that  $|y_1|, \dots, |y_{i-1}| \leq \frac{1}{d^3}$ , we have

$$\begin{aligned} \mathbb{P}\left(|y_i| \leq \frac{1}{d^3} \mid y_1, \dots, y_{i-1}\right) &= \mathbb{P}_{\mathbf{u} \sim \mathcal{U}(S^{d-i})}\left(|u_1| \leq \frac{1}{d^3 \sqrt{1 - (y_1^2 + \dots + y_{i-1}^2)}}\right) \\ &\geq \frac{\int_0^{1/d^3} (1 - y^2)^{(d-i-1)/2} dy}{\int_0^1 (1 - y^2)^{(d-i-1)/2} dy} \\ &\geq \frac{(1 - d^{-6})^{d/2}}{d^3} \geq \frac{e^{-d^{-5}}}{d^3}, \end{aligned}$$

where in the last equation we used  $d \geq 2$ . Therefore, we can show by induction that  $\mathbb{P}(|y_i| \leq 1/d^3, \forall i \leq k) \geq \frac{e^{-kd^{-5}}}{d^{3k}}$ . Thus, by isometry this shows that

$$\mathbb{P}\left(|\mathbf{x}_i^\top \mathbf{y}| \leq \frac{1}{d^3}, \forall i \leq k\right) \geq \frac{1}{e^{d^{-4}} d^{3k}}.$$

This ends the proof of the lemma.  $\blacksquare$

We are now ready to prove a query lower bound for Game 2. Precisely, we show that for appropriate choices of parameters, one needs  $m = \tilde{\Omega}(d)$  queries. The proof is closely inspired from the arguments given in [1]. The main added difficulty arises from bounding the information leakage of the provided hints. As such, our goal is to show that these do not provide more information than the message itself.

**Proposition 14.** *Let  $k \geq 20 \frac{M+3d \log(2d)+1}{c_H n}$ . And let  $0 < \alpha, \beta \leq 1$  such that  $\alpha(\sqrt{d}/\beta)^{5/4} \leq \frac{1}{2}$ . If the Player wins the Orthogonal Vector Game with Hints (Game 2) with probability at least  $1/2$ , then  $m \geq \frac{c_H}{8(30 \log d + c_H)} d$ .*

**Proof** We first define some notations. Let  $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_k]$  be the matrix storing the final outputs from the algorithm. Next, for the responses of the oracle  $(\mathbf{g}_1, s_1), \dots, (\mathbf{g}_m, s_m)$ , we first store all the scalar responses in a vector  $\mathbf{c} = [s_1, \dots, s_m]$ . We now focus on the responses  $\mathbf{g}_1, \dots, \mathbf{g}_m$ . Next, let  $\tilde{\mathbf{G}}$  denote the matrix containing these responses of the oracle which are lines of  $\mathbf{A}$ . Let  $\mathbf{G}$  be the matrix containing unique columns from  $\tilde{\mathbf{G}}$ , augmented with rows of  $\mathbf{A}$  so that it has exactly  $m$  columns which are all different rows of  $\mathbf{A}$ . Last, let  $\mathbf{A}'$  be the matrix  $\mathbf{A}$  once the rows from  $\mathbf{G}$  are removed. Next, let  $\tilde{\mathbf{V}}$  be a matrix containing the responses of the oracle which are vectors  $\mathbf{v}_l$ , ordered by increasing index  $l$ . As before, let  $\mathbf{V}$  be the matrix  $\tilde{\mathbf{V}}$  where we only conserve unique columns and append it with additional vectors  $\mathbf{v}_l$  so that  $\mathbf{V}$  has exactly  $m$  columns. We denote by  $\mathbf{w}_1, \dots, \mathbf{w}_m$  these vectors, and recall that they are vectors  $\mathbf{v}_l$  ordered by increasing order of index  $l$ . Last, we define a vector  $\mathbf{j}$  of indices such that  $j(i)$  contains the information of which column of the matrices  $\mathbf{G}$  or  $\mathbf{V}$  corresponds  $\mathbf{g}_i$ . Precisely, if  $\mathbf{g}_i$  is a line  $\mathbf{a}$  from  $\mathbf{A}$ , we set  $j(i) = j$  where  $j$  is the index of the column from  $\mathbf{G}$  corresponding to  $\mathbf{a}$ . Otherwise, if  $j$  is the index of the column from  $\mathbf{V}$  corresponding to  $\mathbf{g}_i$ , we set  $j(i) = m + j$ .

Next, we argue that  $\mathbf{Y}$  is a deterministic function of Message, the matrices  $\mathbf{G}$ ,  $\mathbf{V}$  and the vector of indices  $\mathbf{j}$  and  $\mathbf{c}$ . First,  $\mathbf{c}$  provides the scalar responses directly. For the  $d$ -dimensional component of the responses, first, note that from  $\mathbf{G}$ ,  $\mathbf{V}$  and  $\mathbf{j}$  one can easily recover the vectors  $\mathbf{g}_1, \dots, \mathbf{g}_m$ . Next, using the algorithm for the second section of the Orthogonal Vector Game with Hints set with initial memory Message and the vectors  $\mathbf{g}_1, \dots, \mathbf{g}_m$  as responses of the oracle, one can inductively compute the queries  $\mathbf{x}_1, \dots, \mathbf{x}_m$ . Last,  $\mathbf{Y}$  is a deterministic function of  $\mathbf{x}_i, \mathbf{g}_i, i \in [m]$  and Message. This ends the claim that there is a function  $\phi$  such that  $\mathbf{Y} = \phi(\text{Message}, \mathbf{G}, \mathbf{V}, \mathbf{j}, \mathbf{c})$ . Now by the data processing inequality,

$$I(\mathbf{A}'; \mathbf{Y} \mid \mathbf{G}, \mathbf{V}, \mathbf{j}, \mathbf{c}) \leq I(\mathbf{A}'; \text{Message} \mid \mathbf{G}, \mathbf{V}, \mathbf{j}, \mathbf{c}) \leq H(\text{Message} \mid \mathbf{G}, \mathbf{V}, \mathbf{j}, \mathbf{c}) \leq M. \quad (9)$$

In the last inequality we used the fact that Message uses at most  $M$  bits. Now, we have that

$$I(\mathbf{A}'; \mathbf{Y} \mid \mathbf{G}, \mathbf{V}, \mathbf{j}, \mathbf{c}) = H(\mathbf{A}' \mid \mathbf{G}, \mathbf{V}, \mathbf{j}, \mathbf{c}) - H(\mathbf{A}' \mid \mathbf{Y}, \mathbf{G}, \mathbf{V}, \mathbf{j}, \mathbf{c}). \quad (10)$$

In the next steps we bound the two terms. We start with the second term of the right hand side of Eq (10) using similar arguments to the proof given in [1]. Let  $\mathcal{E}$  be the event when the Player succeeds at Game 2. Now consider the case when  $\mathbf{Y}$  is a winning matrix. Then we have  $\|\mathbf{A}\mathbf{y}_i\|_\infty \leq \alpha$  for all  $i \leq k$ . As a result, any line  $\mathbf{a}$  of  $\mathbf{A}'$  satisfies  $\|\mathbf{Y}^\top \mathbf{a}\|_\infty \leq \alpha$ . Further, we have that  $\|P_{\text{Span}(\mathbf{y}_j, j < i)^\perp}(\mathbf{y}_i)\| \leq \beta$  for all  $i \leq k$ . By Lemma 22, there exist  $\lceil k/5 \rceil$  orthonormal vectors  $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_{\lceil k/5 \rceil}]$  such that for any  $\mathbf{x} \in \mathbb{R}^d$  one has  $\|\mathbf{Z}^\top \mathbf{x}\|_\infty \leq \left(\frac{\sqrt{d}}{\beta}\right)^{5/4} \|\mathbf{Y}^\top \mathbf{x}\|_\infty$ . In particular, all lines  $\mathbf{a}$  of  $\mathbf{A}'$  satisfy

$$\|\mathbf{Z}^\top \mathbf{a}\|_\infty \leq \left(\frac{\sqrt{d}}{\beta}\right)^{5/4} \alpha \leq \frac{1}{2},$$

where we used the hypothesis in the parameters  $\alpha$  and  $\beta$ . Now by Lemma 12, one has

$$\left| \left\{ \mathbf{a} \in \{\pm 1\}^d : \|\mathbf{Z}^\top \mathbf{a}\|_\infty \leq \frac{1}{2} \right\} \right| \leq 2^d \mathbb{P}_{\mathbf{h} \sim \mathcal{U}(\{\pm 1\}^d)} \left( \|\mathbf{Z}^\top \mathbf{h}\|_\infty \leq \frac{1}{2} \right) \leq 2^{d - c_H \lceil k/5 \rceil}.$$

Therefore, we proved that if  $\mathbf{Y}'$  is a winning vector,  $H(\mathbf{A}' \mid \mathbf{Y} = \mathbf{Y}') \leq (n - m)(d - c_H k/5)$ . Otherwise, if  $\mathbf{Y}'$  loses, we can directly use  $H(\mathbf{A}' \mid \mathbf{Y} = \mathbf{Y}') \leq (n - m)d$ . Combining these equations gives

$$\begin{aligned} H(\mathbf{A}' \mid \mathbf{Y}, \mathbf{G}, \mathbf{V}, \mathbf{j}, \mathbf{c}) &\leq H(\mathbf{A}' \mid \mathbf{Y}) \\ &\leq \mathbb{P}(\mathcal{E}^c)(n - m)d + \mathbb{P}(\mathcal{E})(n - m)(d - c_H k/5) \\ &\leq (n - m)(d - \mathbb{P}(\mathcal{E})c_H k/5). \end{aligned}$$

Next, we turn to the first term of the right-hand side of Eq (10).

$$\begin{aligned} H(\mathbf{A}' \mid \mathbf{G}, \mathbf{V}, \mathbf{j}, \mathbf{c}) &= H(\mathbf{A} \mid \mathbf{G}, \mathbf{V}, \mathbf{j}, \mathbf{c}) = H(\mathbf{A} \mid \mathbf{V}) - I(\mathbf{A}; \mathbf{G}, \mathbf{j}, \mathbf{c} \mid \mathbf{V}) \\ &\geq H(\mathbf{A} \mid \mathbf{V}) - H(\mathbf{G}, \mathbf{j}, \mathbf{c}) \\ &\geq H(\mathbf{A} \mid \mathbf{V}) - md - m \log(2m) - m \log(d^2) \\ &= H(\mathbf{A}) - I(\mathbf{A}; \mathbf{V}) - md - 3m \log(2d) \\ &= (n - m)d - 3m \log(2d) - I(\mathbf{A}; \mathbf{V}). \end{aligned}$$

In the second inequality, we use the fact that  $\mathbf{G}$  uses  $md$  bits and  $\mathbf{j}$  can be stored with  $m \log(2m)$  bits. Now by the chain rule,

$$I(\mathbf{A}; \mathbf{V}) = \sum_{i \leq m} I(\mathbf{A}; \mathbf{w}_i \mid \mathbf{w}_1, \dots, \mathbf{w}_{i-1}).$$

Now if  $\mathbf{w}_i = \mathbf{v}_l$ , recalling that the vectors  $\mathbf{w}_{i'}$  are ordered by increasing index of  $l'$ , we have

$$\begin{aligned} I(\mathbf{A}; \mathbf{w}_i \mid \mathbf{w}_1, \dots, \mathbf{w}_{i-1}) &= H(\mathbf{w}_i \mid \mathbf{w}_1, \dots, \mathbf{w}_{i-1}) - H(\mathbf{w}_i \mid \mathbf{A}, \mathbf{w}_1, \dots, \mathbf{w}_i) \\ &\leq H(\mathbf{w}_i) - H(\mathbf{w}_i \mid \mathbf{A}, \mathbf{w}_1, \dots, \mathbf{w}_i, \mathbf{x}_{l,1}, \dots, \mathbf{x}_{l,r_l}) \\ &= \log |\mathcal{D}_\delta| - H(\mathbf{w}_i \mid \mathbf{x}_{l,1}, \dots, \mathbf{x}_{l,r_l}). \end{aligned}$$



In the last equality, we used the fact that if  $\mathbf{b}_{l,1}, \dots, \mathbf{b}_{l,r_l}$  are the resulting vectors from the Gram-Schmidt decomposition of  $\mathbf{x}_{l,1}, \dots, \mathbf{x}_{l,r_l}$ ,  $\mathbf{y}_l$  is generated uniformly in  $S^{d-1} \cap \{\mathbf{y} : \forall r \leq r_l, |\mathbf{b}_{l,r}^\top \mathbf{y}| \leq d^{-3}\}$  independently from the past history, and  $\mathbf{v}_l = \phi_\delta(\mathbf{y}_l)$ . Now by Lemma 13, we know that

$$\mathbb{P}_{\mathbf{z} \sim \mathcal{U}(S^{d-1})} \left( \forall r \leq r_l, |\mathbf{b}_{l,r}^\top \mathbf{z}| \leq d^{-3} \right) \geq \frac{1}{e^{d-4} d^{3k}}.$$

As a result, for any  $\mathbf{b}_j(\delta) \in \mathcal{D}_\delta$ , one has

$$\mathbb{P}(\mathbf{w}_i = \mathbf{b}_j(\delta) \mid \mathbf{x}_{l,1}, \dots, \mathbf{x}_{l,r_l}) \leq \frac{\mathbb{P}_{\mathbf{z} \sim \mathcal{U}(S^{d-1})}(\mathbf{z} \in V_j(\delta))}{\mathbb{P}_{\mathbf{z} \sim \mathcal{U}(S^{d-1})}(\forall r \leq r_l, |\mathbf{b}_{l,r}^\top \mathbf{z}| \leq d^{-3})} \leq \frac{e^{d-4} d^{3k}}{|\mathcal{D}_\delta|},$$

where we used the fact that each cell has the same area. In particular, this shows that

$$H(\mathbf{w}_i \mid \mathbf{x}_{l,1}, \dots, \mathbf{x}_{l,r_l}) = \mathbb{E}_{\mathbf{b} \sim \mathbf{w}_i \mid \mathbf{x}_{l,1}, \dots, \mathbf{x}_{l,r_l}} [-\log p_{\mathbf{w}_i \mid \mathbf{x}_{l,1}, \dots, \mathbf{x}_{l,r_l}}(\mathbf{b})] \geq \log \left( \frac{|\mathcal{D}_\delta|}{e^{d-4} d^{3k}} \right).$$

Hence,

$$I(\mathbf{A}; \mathbf{w}_i \mid \mathbf{w}_1, \dots, \mathbf{w}_{i-1}) \leq 3k \log d + d^{-4} \log e.$$

Putting everything together gives

$$\begin{aligned} I(\mathbf{A}'; \mathbf{Y} \mid \mathbf{G}, \mathbf{V}, \mathbf{j}) &\geq (n-m)d - 3m \log(2d) - 3km \log d - 2md^{-4} - (n-m)(d - \mathbb{P}(\mathcal{E})c_H k/5) \\ &\geq \frac{c_H}{10} k(n-m) - 3km \log d - 1 - 3d \log(2d), \end{aligned}$$

where in the last equation we used  $d \geq 2$ . Together with Eq (9), this implies

$$m \geq \frac{c_H kn/10 - M - 1 - 3d \log(2d)}{k(3 \log d + c_H/10)}.$$

As a result, since  $k \geq 20 \frac{M+3d \log(2d)+1}{c_H n}$  and  $n \geq d/4$ , we obtain

$$m \geq \frac{c_H n}{60 \log d + 2c_H} \geq \frac{c_H}{8(30 \log d + c_H)} d.$$

This ends the proof of the proposition. ■

We are now ready to prove the main result.

**Proof of Theorem 1** We set  $n = \lceil d/4 \rceil$  and  $k = \lceil 20 \frac{M+3d \log(2d)+1}{c_H n} \rceil$ . By Proposition 8, with probability at least  $1 - C\sqrt{\log d}/d^2$ , the procedure is consistent with a first-order oracle for convex optimization. Hence, since the functions  $F_{\mathbf{A},v,P,L}$  are  $\sqrt{d}$ -Lipschitz, any  $M$ -bit algorithm guaranteed to solve convex optimization within accuracy  $\epsilon = \eta/(2d) = 1/d^4$  for 1-Lipschitz functions, yields an algorithm that is successful for the optimization procedure with probability at least  $1 - C\sqrt{\log d}/d^2$  and precision  $\epsilon\sqrt{d} = \eta/(2\sqrt{d})$ . Suppose that it uses at most  $Q$  queries. Then, by Proposition 11, there is a strategy for Game 2 for parameters  $(d, k, \lceil Q/p_{max} \rceil + 1, M, \alpha = \frac{2\eta}{\gamma_1}, \beta = \frac{\gamma_2}{4})$  in which the Player wins with probability at least  $1 - C'\sqrt{\log d}/d$ . Now for  $d$  large enough, this probability is at least  $1/2$ . Further,

$$\frac{2\eta}{\gamma_1} \left( \frac{4\sqrt{d}}{\gamma_2} \right)^{5/4} \leq \frac{(4/3)^{5/4}}{6} \eta d^3 \leq \frac{1}{2}.$$

Hence, by Proposition 14, one has

$$\lceil Q/p_{max} \rceil + 1 \geq \frac{c_H}{8(30 \log d + c_H)} d.$$

Because  $p_{max} = \Theta((d/k)^{1/3} \log^{-2/3} d)$ , this implies

$$Q = \Omega\left(\frac{(d/k)^{1/3} d}{\log^{5/3} d}\right) = \Omega\left(\frac{d^{5/3}}{(M + \log d)^{1/3} \log^{5/3} d}\right).$$

In particular, if  $M = d^{1+\delta}$  for  $\delta \in [0, 1]$ , the number of queries is  $Q = \tilde{\Omega}(d^{1+(1-\delta)/3})$ . ■

## 4 Memory-constrained feasibility problem

### 4.1 Defining the feasibility procedure

Similarly to Section 3, we pose  $n = \lceil d/4 \rceil$ . Also, for any matrix  $\mathbf{A} \in \{\pm 1\}^{n \times d}$ , we use the same functions  $\mathbf{g}_\mathbf{A}$  and  $\tilde{\mathbf{g}}_\mathbf{A}$ . We use similar techniques as those we introduced for the optimization problem. However, since in this case, the separation oracle only returns a separating hyperplane, without any value considerations of an underlying function, Procedure 1 can be drastically simplified, which leads to improved lower bounds.

Let  $\eta_0 = 1/(24d^2)$ ,  $\eta_1 = \frac{1}{2\sqrt{d}}$ ,  $\delta = 1/d^3$ , and  $k \leq d/3 - n$  be a parameter. Last, let  $p_{max} = \lfloor (c_{d,1}d - 1)/(k - 1) \rfloor$ , where  $c_{d,1}$  is the same quantity as in Eq (3). The feasibility procedure is defined in Procedure 4. The oracle first randomly samples  $\mathbf{A} \sim \mathcal{U}(\{\pm 1\}^{n \times d})$  and  $\mathbf{v}_0 \sim \mathcal{U}(\mathcal{D}_\delta)$ . This matrix and vector are then fixed in the rest of the procedure. Whenever the player queries a point  $\mathbf{x}$  such that  $\|\mathbf{A}\mathbf{x}\|_\infty > \eta_0$  (resp.  $\mathbf{v}_0^\top \mathbf{x} > -\eta_1$ ), the oracle returns  $\tilde{\mathbf{g}}_\mathbf{A}(\mathbf{x})$  (resp.  $\mathbf{v}_0$ ). All other queries are called *informative* queries. With this definition, it now remains to define the separation oracle on informative queries. The oracle proceeds by periods in which the behavior is different. In each period  $p$ , the oracle constructs vectors  $\mathbf{v}_{p,1}, \dots, \mathbf{v}_{p,k-1}$  inductively and keeps in memory some queries  $i_{p,1}, \dots, i_{p,k}$  that will be called *exploratory*. The first informative query  $t$  will be the first exploratory query and starts period 1.

Given a new query  $\mathbf{x}_t$ ,

1. If  $\|\mathbf{A}\mathbf{x}_t\|_\infty > \eta_0$ , the oracle returns  $\tilde{\mathbf{g}}_\mathbf{A}(\mathbf{x}_t)$ .
2. If  $\mathbf{v}_0^\top \mathbf{x}_t > -\eta_1$ , the oracle returns  $\mathbf{v}_0$ .
3. If  $\mathbf{x}_t$  was queried in the past sequence, the oracle returns the same vector that was returned previously.
4. Otherwise, let  $p$  be the index of the current period and let  $\mathbf{v}_{p,1}, \dots, \mathbf{v}_{p,l}$  be the vectors from the current period constructed so far, together with their corresponding exploratory queries  $i_{p,1}, \dots, i_{p,l} < t$ . Potentially, if  $p = 1$  one may not have defined any such vectors at the beginning of time  $t$ . In this case, let  $l = 0$ .
  - (a) If  $\max_{1 \leq l' \leq l} \mathbf{v}_{p,l'}^\top \mathbf{x}_t > -\eta_1$  (with the convention  $\max_\emptyset = -\infty$ ), the oracle returns  $\mathbf{v}_{p,l'}$  where  $l' = \arg \max_{l' \leq l} \mathbf{v}_{p,l'}^\top \mathbf{x}_t$ . Ties are broken alphabetically.

- (b) Otherwise, if  $l < k - 1$ , we first define  $i_{p,l+1} = t$ . Then, let  $\mathbf{b}_{p,1}, \dots, \mathbf{b}_{p,l+1}$  be the result from the Gram-Schmidt decomposition of  $\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,l+1}}$  and let  $\mathbf{y}_{p,l+1}$  be a sample of the distribution obtained by the uniform distribution  $\mathbf{y}_{p,l+1} \sim \mathcal{U}(S^{d-1} \cap \{\mathbf{z} \in \mathbb{R}^d : |\mathbf{b}_{p,r}^\top \mathbf{z}| \leq \frac{1}{d^3}, \forall r \leq l+1\})$ . We then pose  $\mathbf{v}_{p,l+1} = \phi_\delta(\mathbf{y}_{p,l+1})$ . Having defined this new vector, the oracle returns  $\mathbf{v}_{p,l}$ . We then increment  $l$ .
- (c) Otherwise, if  $r = k$ , we define  $i_{p,k} = i_{p+1,1} = t$ . If  $p+1 \leq p_{max}$ , this starts the next period  $p+1$ . As above, let  $\mathbf{b}_{p+1,1}$  be the result of the Gram-Schmidt decomposition of  $\mathbf{x}_{i_{p+1,1}}$  and sample  $\mathbf{y}_{p+1,1}$  according to a uniform  $\mathbf{y}_{p+1,1} \sim \mathcal{U}(S^{d-1} \cap \{\mathbf{z} \in \mathbb{R}^d : |\mathbf{b}_{p+1,1}^\top \mathbf{z}| \leq \frac{1}{d^3}\})$ . We then pose  $\mathbf{v}_{p+1,1} = \phi_\delta(\mathbf{y}_{p+1,1})$  and the oracle returns  $\mathbf{v}_{p+1,1}$ . We can then increment  $p$  and reset  $l = 1$ .

The above construction ends when the period  $p_{max}$  is finished. At this point, the oracle has defined the vectors  $\mathbf{v}_{p,l}$  for all  $p \leq p_{max}$  and  $l \leq k$ . We then define the successful set as

$$Q_{\mathbf{A},\mathbf{v}} = \left\{ \mathbf{x} \in B_d(0,1) : \|\mathbf{A}\mathbf{x}\|_\infty \leq \eta_0, \mathbf{v}_0^\top \mathbf{x} \leq -\eta_1, \max_{p \leq p_{max}, l \leq k-1} \mathbf{v}_{p,l}^\top \mathbf{x} \leq -\eta_1 \right\}.$$

From now on, the procedure uses any separation oracle for  $Q_{\mathbf{A},\mathbf{v}}$  as responses to the algorithm, while making sure to be consistent with previous oracle responses if a query is exactly duplicated. We now define what we mean by solving the above feasibility procedure.

**Definition 15.** *Let  $alg$  be an algorithm for the feasibility problem. When running  $alg$  with the responses of the feasibility procedure, we denote by  $\mathbf{v}$  the set of constructed vectors and  $\mathbf{x}^*(alg)$  the final answer returned by  $alg$ . We say that an algorithm  $alg$  is successful for the feasibility procedure with probability  $q \in [0,1]$ , if taking  $\mathbf{A} \sim \mathcal{U}(\{\pm 1\}^{n \times d})$ , with probability at least  $q$  over the randomness of  $\mathbf{A}$  and of the procedure,  $\mathbf{x}^*(alg) \in Q_{\mathbf{A},\mathbf{v}}$ .*

In the rest of this section, we first relate this feasibility procedure to the standard feasibility problem, then prove query lower bounds to solve the feasibility procedure.

## 4.2 Reduction from the feasibility problem to the feasibility procedure

In the next proposition, we check that the above procedure indeed corresponds to a valid feasibility problem.

**Proposition 16.** *On an event of probability at least  $1 - C\sqrt{\log d}/d$ , the procedure described above is a valid feasibility problem. More precisely, the following hold.*

- There exists  $\bar{\mathbf{x}} \in B_d(0,1)$  such that  $\|\mathbf{A}\bar{\mathbf{x}}\|_\infty = 0$ ,  $\mathbf{v}_0^\top \bar{\mathbf{x}} \leq -4\eta_1$ , and

$$\max_{p \leq p_{max}, l \leq k-1} \mathbf{v}_{p,l}^\top \bar{\mathbf{x}} \leq -4\eta_1.$$

- Let  $\epsilon = \min\{\eta_0/\sqrt{d}, \eta_1\}/2$ . Then,  $B_d\left(\bar{\mathbf{x}} - \epsilon \frac{\bar{\mathbf{x}}}{\|\bar{\mathbf{x}}\|}, \epsilon\right) \subseteq B_d(0,1) \cap B_d(\bar{\mathbf{x}}, 2\epsilon) \subseteq Q_{\mathbf{A},\mathbf{v}}$ .
- Throughout the run of the feasibility problem, the separation oracle always returned a valid cut, i.e., for any iteration  $t$ , if  $\mathbf{x}_t$  denotes the query and  $\mathbf{g}_t$  is the returned vector from the oracle, one has

$$\forall \mathbf{x} \in Q_{\mathbf{A},\mathbf{v}}, \quad \langle \mathbf{g}_t, \mathbf{x}_t - \mathbf{x} \rangle > 0.$$

Further, responses are consistent: if  $\mathbf{x}_t = \mathbf{x}_{t'}$ , the responses of the procedure at times  $t$  and  $t'$  coincide.

---

**Input:**  $d, k, p_{max}$ , algorithm  $alg$

- 1 Sample  $\mathbf{A} \sim \mathcal{U}(\{\pm 1\}^{n \times d})$  and  $\mathbf{v}_0 \sim \mathcal{U}(\mathcal{D}_\delta)$ ;
- 2 Initialize the memory of  $alg$  to  $\mathbf{0}$  and let  $p = 1, l = 0$ ;
- 3 **for**  $t \geq 1$  **do**
- 4 Run  $alg$  with current memory to obtain a query  $\mathbf{x}_t$ ;
- 5 **if**  $\|\mathbf{A}\mathbf{x}_t\| > \eta_0$  **then return**  $\tilde{\mathbf{g}}_{\mathbf{A}}(\mathbf{x}_t)$  as response to  $alg$  ;
- 6 **else if**  $\mathbf{v}_0^\top \mathbf{x}_t > -\eta_1$  **then return**  $\mathbf{v}_0$  as response to  $alg$  ;
- 7 **else if** *Query  $\mathbf{x}_t$  was made in the past* **then return** the same vector that was returned for  $\mathbf{x}_t$  ;
- 8 **else**
- 9 **if**  $\max_{1 \leq l' \leq l} \mathbf{v}_{p,l'}^\top \mathbf{x}_t > -\eta_1$  **then**
- 10 **return**  $\mathbf{v}_{p,l'}$  where  $l' = \arg \max_{1 \leq l' \leq l} \mathbf{v}_{p,l'}^\top \mathbf{x}_t$ .
- 11 **else if**  $l < k - 1$  **then**
- 12 Let  $i_{p,l+1} = t$  and compute Gram-Schmidt decomposition  $\mathbf{b}_{p,1}, \dots, \mathbf{b}_{p,l+1}$  of  $\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,l+1}}$ ;
- 13 Sample  $\mathbf{y}_{p,l+1}$  uniformly on  $\mathcal{S}^{d-1} \cap \{\mathbf{z} \in \mathbb{R}^d : |\mathbf{b}_{p,l'}^\top \mathbf{z}| \leq d^{-3}, \forall l' \leq l+1\}$  and define  $\mathbf{v}_{p,l+1} = \phi_\delta(\mathbf{y}_{p,l+1})$ ;
- 14 **return**  $\mathbf{v}_{p,l+1}$  as response to  $alg$  and increment  $l \leftarrow l + 1$ .
- 15 **else if**  $p + 1 \leq p_{max}$  **then**
- 16 Set  $i_{p,k} = i_{p+1,1} = t$  and compute the Gram-Schmidt decomposition  $\mathbf{b}_{p+1,1}$  of  $\mathbf{x}_{i_{p+1,1}}$ ;
- 17 Sample  $\mathbf{y}_{p+1,1}$  uniformly on  $\mathcal{S}^{d-1} \cap \{\mathbf{z} \in \mathbb{R}^d : |\mathbf{b}_{p+1,1}^\top \mathbf{z}| \leq d^{-3}\}$  and define  $\mathbf{v}_{p+1,1} = \phi_\delta(\mathbf{y}_{p+1,1})$ .
- 18 **return**  $\mathbf{v}_{p+1,1}$  as response to  $alg$ , increment  $p \leftarrow p + 1$  and reset  $l = 1$ .
- 19 **else** Set  $i_{p_{max},k} = t$  and break the **for** loop;
- 20 **end**
- 21 **for**  $t' \geq t$  **do** Use any separation oracle for  $Q_{\mathbf{A},\mathbf{v}}$  consistent with previous responses ;

---

**Procedure 4:** The feasibility procedure for algorithm  $alg$

We use a similar proof to that of Proposition 9.

**Proof** For convenience, we rename  $\mathbf{v}_{p,l} = \mathbf{v}_{(p-1)(k-1)+l}$ . Also, let  $l_{max} = p_{max}(k-1) \leq c_{d,1}d - 1$ . Next, let  $C_d = \sqrt{40l_{max} \log d}$ . We define the vector

$$\bar{\mathbf{x}} = -\frac{1}{C_d} \sum_{l=0}^{l_{max}} P_{\text{Span}(\mathbf{a}_i, i \leq n)^\perp}(\mathbf{v}_l).$$

Since  $l_{max} \leq p_{max}(k-1) \leq c_{d,1}d - 1$ , the same arguments as in the proof of Proposition 9 show that on an event  $\mathcal{E}$  of probability at least  $1 - C\sqrt{\log d}/d$ , we have  $\|\bar{\mathbf{x}}\| \leq 1$  and

$$\max_{0 \leq l \leq l_{max}} \mathbf{v}_l^\top \bar{\mathbf{x}} \leq -\frac{1}{40\sqrt{(l_{max}+1)\log d}} \leq -\frac{2}{\sqrt{d}} = -4\eta_1,$$

where in the second inequality we used  $l_{max} \leq c_{d,1}d - 1$ . Now by construction, one has  $\|\mathbf{A}\bar{\mathbf{x}}\|_\infty = 0$ . This ends the proof of the first claim of the proposition. We now turn to the second claim, which is immediate from the fact that  $\mathbf{x} \mapsto \|\mathbf{A}\mathbf{x}\|_\infty$  is  $\sqrt{d}$ -Lipschitz and both  $\mathbf{x} \mapsto \mathbf{v}_0^\top \mathbf{x}$  and  $\mathbf{x} \mapsto \max_{p \leq p_{max}, l \leq k} \mathbf{v}_{p,l}^\top \mathbf{x}$  are 1-Lipschitz. Therefore,  $B_d(\bar{\mathbf{x}} - \epsilon \bar{\mathbf{x}} / \|\bar{\mathbf{x}}\|, \epsilon) \subseteq B_d(0, 1) \cap B_d(\bar{\mathbf{x}}, 2\epsilon) \subset Q_{\mathbf{A},\mathbf{v}}$ . It now remains to check

that the third claim is satisfied. It suffices to check that this is the case during the construction phase of the feasibility procedure. By construction of  $Q_{A,v} \subset \{\mathbf{x} : \|\mathbf{A}\mathbf{x}\|_\infty \leq \eta_0\}$ .

Hence, it suffices to check that for informative queries  $\mathbf{x}_t$ , the returned vectors  $\mathbf{g}_t$  are valid separation hyperplanes. By construction, these can only be either  $\mathbf{v}_0$  or  $\mathbf{v}_{p,l}$  for  $p \leq p_{max}$ ,  $l \leq k-1$ . We denote by  $\mathbf{w}$  this vector. Let  $t'$  be the first time  $\mathbf{x}_t$  was queried. There are two cases. Either  $\mathbf{w}$  was not constructed at time  $t'$ , in which case, by construction this means that we are in scenario (2) or (4a). Both cases imply  $\mathbf{w}^\top \mathbf{x}_t > -\eta_1$ . Hence,  $\mathbf{w}$  which is returned by the procedure is a valid separation hyperplane. Now suppose that  $\mathbf{w} = \mathbf{v}_{p,l}$  was constructed at time  $t'$ —scenarios (4b) or (4c). By construction, one has  $|\mathbf{b}_{p,r}^\top \mathbf{y}_{p,l}| \leq d^{-3}$  for all  $r \leq l$ . Decomposing  $\mathbf{x}_t = \mathbf{x}_{i_{p,l}} = \alpha_1 \mathbf{b}_{p,1} + \dots + \alpha_l \mathbf{b}_{p,l}$ , we obtain

$$|\mathbf{x}_t^\top \mathbf{y}_{p,l}| \leq \frac{\|\boldsymbol{\alpha}\|_1}{d^3} \leq \frac{1}{d^2 \sqrt{d}}.$$

As a result,  $\mathbf{y}_{p,l}^\top \mathbf{x}_t \geq -1/(d^2 \sqrt{d})$ . Now because  $\mathbf{v}_{p,l} = \phi_\delta(\mathbf{y}_{p,l})$ , we have  $\|\mathbf{v}_{p,l} - \mathbf{y}_{p,l}\| \leq \delta$ . Hence, for any  $d \geq 2$ ,

$$\mathbf{w}^\top \mathbf{x}_t \geq -1/(d^2 \sqrt{d}) - \delta > -\eta_1.$$

Hence,  $\mathbf{w}$  was a valid separation hyperplane. The last claim that the responses of the procedure are consistent over time is a direct consequence from its construction. This ends the proof of the proposition. ■

As a simple consequence of this result, solving the feasibility problem is harder than solving the feasibility procedure with high probability.

**Proposition 17.** *Let  $alg$  be an algorithm that solves the feasibility problem with accuracy  $\epsilon = 1/(48d^2 \sqrt{d})$ . Then, it solves the feasibility procedure with probability at least  $1 - C\sqrt{\log d}/d$ .*

**Proof** Let  $\mathcal{E}$  be the event of probability at least  $1 - C\sqrt{\log d}/d$  defined in Proposition 16. We show that on  $\mathcal{E}$ ,  $alg$  solves the feasibility procedure. On  $\mathcal{E}$ , the feasibility procedure emulates a valid feasibility oracle. Further, on  $\mathcal{E}$ , the successful set contains a closed ball of radius  $\epsilon$ . As a result, on  $\mathcal{E}$ ,  $alg$  finds a solution to the feasibility problem emulated by the procedure. ■

Next, we show that it is necessary to finish the  $p_{max}$  periods to solve the feasibility procedure.

**Proposition 18.** *Fix an algorithm  $alg$ . Then, if  $\mathcal{A}$  denotes the event when  $alg$  succeeds and  $\mathcal{B}$  denotes the event when the procedure ends period  $p_{max}$  with  $alg$ , then  $\mathcal{E} \subseteq \mathcal{B}$ .*

**Proof** Consider the case when the period  $p_{max}$  was not ended. Let  $\mathbf{x}^*$  denote the last query performed by  $alg$ . We consider the scenario in which  $\mathbf{x}^*$  fell. Let  $t$  be the first time when  $alg$  submitted query  $\mathbf{x}^*$ . For any of the scenarios (1), (2), or (4a), by construction of  $Q_{A,v}$ , we already have  $\mathbf{x}_t \notin Q_{A,v}$ . It remains to check scenarios (4b) and (4c) for which the procedure constructs a new vector  $\mathbf{v}_{p,l}$ , where  $p$  is the index of the period of  $t$  and  $i_{p,1}, \dots, i_{p,l} = t$  are the previous exploratory queries in period  $p$ . We decompose  $\mathbf{x}_t = \mathbf{x}_{i_{p,l}} = \alpha_1 \mathbf{b}_{p,1} + \alpha_l \mathbf{b}_{p,l}$ . Now by construction,

$$|\mathbf{x}_t^\top \mathbf{y}_{p,l}| = |\mathbf{x}_{i_{p,l}}^\top \mathbf{y}_{p,l}| \leq \frac{\|\boldsymbol{\alpha}\|_1}{d^3} \leq \frac{1}{d^2 \sqrt{d}}.$$

As a result,  $\mathbf{x}_t^\top \mathbf{v}_{p,l} \geq -|\mathbf{x}_t^\top \mathbf{y}_{p,l}| - \delta \geq -d^{-2.5} - d^{-3} > -\eta_1$ , for any  $d \geq 2$ . Thus,  $\mathbf{x}_t = \mathbf{x}^* \notin Q_{A,v}$ . This shows that in order to succeed at the feasibility procedure, an algorithm needs to end all  $p_{max}$  periods. ■

### 4.3 Reduction to the Orthogonal Vector Game with Hints.

The remaining piece of our argument is to show that solving the feasibility procedure is harder than solving the Orthogonal Vector Game with Hints, Game 2.

**Proposition 19.** *Let  $\mathbf{A} \sim \mathcal{U}(\{\pm 1\}^{n \times d})$ . If there exists an  $M$ -bit algorithm that solves the feasibility problem described above using  $mp_{max}$  queries with probability at least  $q$  over the randomness of the algorithm, choice of  $\mathbf{A}$  and the randomness of the separation oracle, then there is an algorithm for Game 2 for parameters  $(d, k, m, M, \alpha = \frac{\eta_0}{\eta_1}, \beta = \frac{\eta_1}{2})$ , for which the Player wins with probability at least  $q$  over the randomness of the player's strategy and  $\mathbf{A}$ .*

**Proof** Let  $alg$  be an  $M$ -bit algorithm solving the feasibility problem with  $mp_{max}$  queries with probability at least  $q$ . In Algorithm 5, we describe the strategy of the player in Game 2.

---

**Input:**  $d, k, p_{max}, m$ , algorithm  $alg$

**Part 1:** Strategy to store Message knowing  $\mathbf{A}$ ;

- 1 Initialize the memory of  $alg$  to be  $\mathbf{0}$ ;
- 2 Submit  $\emptyset$  to the Oracle and use the response as  $\mathbf{v}_0$ ;
- 3 Run  $alg$  with the optimization procedure knowing  $\mathbf{A}$  and  $\mathbf{v}_0$  until the first exploratory query  $\mathbf{x}_{i_{1,1}}$ .
- 4 **for**  $p \in [p_{max}]$  **do**
- 5     Let Memory $_p$  be the current memory state of  $alg$  and  $i_{p,1}$  the current iteration step. ;
- 6     Run  $alg$  with the feasibility procedure until period  $p$  ends at iteration step  $i_{p+1,1}$ . If  $alg$  stopped before, **return** the strategy fails. When needed to sample a unit vector  $\mathbf{v}_{p',l'}$ , submit vectors  $\mathbf{x}_{i_{p',1}}, \dots, \mathbf{x}_{i_{p',l'}}$  to the Oracle. We use the corresponding response of the Oracle as  $\mathbf{v}_{p',l'}$ ;
- 7     **if**  $i_{p+1,1} - i_{p,1} \leq m$  **then**
- 8         | Set Message = Memory $_p$
- 9 **end**
- 10 **for** Remaining queries to perform to Oracle **do** Submit arbitrary query, e.g.  $\emptyset$  ;
- 11 **if** Message has not been defined yet **then return** The strategy fails;
- 12 Submit  $\tilde{\mathbf{g}}_{\mathbf{A},\mathbf{v}}$  to the Oracle as defined in Eq (11).;

**Part 2:** Strategy to make queries;

- 13 Set the memory state of  $alg$  to be Message.;
- 14 **for**  $i \in [m]$  **do**
- 15     Run  $alg$  with current memory to obtain a query  $\mathbf{z}_i$ ;
- 16     Submit  $\mathbf{z}_i$  to the Oracle from Game 2, to get response  $(\mathbf{g}_i, s_i)$ ;
- 17     Compute  $\tilde{\mathbf{g}}_i$  using  $\mathbf{z}_i, \mathbf{g}_i$  and  $s_i$  as defined in Eq (12) and pass  $\tilde{\mathbf{g}}_i$  as response to  $alg$ ;
- 18 **end**

**Part 3:** Strategy to return vectors;

- 19 **for**  $l \in [k]$  **do** Set  $i_l$  to be the index  $i$  of the first query  $\mathbf{z}_i$  for which  $s_i = l$ , if it exists ;
  - 20 **if** index  $i_k$  has not been defined yet **then**
  - 21     | With the current memory of  $alg$  find a new query  $\mathbf{z}_{m+1}$  and set  $i_k = m + 1$ ;
  - 22 **return**  $\left\{ \frac{\mathbf{z}_{i_1}}{\|\mathbf{z}_{i_1}\|}, \dots, \frac{\mathbf{z}_{i_k}}{\|\mathbf{z}_{i_k}\|} \right\}$  to the Oracle.
- 

**Algorithm 5:** Strategy of the Player for the Orthogonal Vector Game with Hints

In the first part of the strategy, the player observes  $\mathbf{A}$ . Then they proceed to simulate the feasibility problem with  $alg$  using parameters  $\mathbf{A}$ . When needed to sample a vector  $\mathbf{v}_{p,l}$  (resp.  $\mathbf{v}_0$ ), the player

submits the corresponding queries  $\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,l}}$  (resp.  $\emptyset$ ) useful to define  $\mathbf{v}_{p,l}$ . The player then takes the response given by the Oracle as that vector  $\mathbf{v}_{p,l}$  (resp.  $\mathbf{v}_0$ ), which simulates exactly a run of the feasibility procedure. Further, since  $1 + p_{max}(k-1) \leq d$ , the player does not run out of queries. Importantly, during the run, the player keeps track of the length  $i_{p,k} - i_{p,1}$  of period  $p$ . The first time we encounter a period  $p$  with length at most  $m$ , we set  $\text{Message} = \text{Memory}_p$ , the memory state of  $alg$  at the beginning of period  $p$ . If there is no such period, the strategy fails. Also, if  $alg$  stopped before ending period  $p_{max}$ , the strategy fails. Next, the algorithm submits the following function  $\tilde{\mathbf{g}}_{A,v}$  to the Oracle. Since the responses of the feasibility procedure are consistent over time, we adopt the following notation. For a previously queried vector  $\mathbf{x}$  of  $alg$ , we denote  $\mathbf{g}(\mathbf{x})$  the vector which was returned to  $alg$  during the first part (lines 3-9 of Algorithm 5).

$$\tilde{\mathbf{g}}_{A,v} : \mathbf{x} \mapsto \begin{cases} (\mathbf{0}, 1) & \text{if } \mathbf{x} \text{ was never queried in the first part,} \\ (\mathbf{a}_i, 1) & \text{ow. and if } \mathbf{g}(\mathbf{x}) \in \{\pm \mathbf{a}_i\}, i \leq n, \\ (\mathbf{v}_0, 2) & \text{ow. and if } \mathbf{g}(\mathbf{x}) = \mathbf{v}_0, \\ (\mathbf{v}_{p',l'}, 2 + l' \mathbb{1}_{p'=p} + k \mathbb{1}_{p'=p+1, l'=1}) & \text{ow. and if } \mathbf{g}(\mathbf{x}) = \mathbf{v}_{p',l'}, p' \leq p_{max}, l \leq k-1. \end{cases} \quad (11)$$

Intuitively, the first component of  $\tilde{\mathbf{g}}$  gives the returned vector in the first period, at the exception that we always return  $\mathbf{a}_i$  instead of  $\{\pm \mathbf{a}_i\}$ . The second term has values in  $[2 + k \leq d^2]$ . Hence, the submitted function is valid.

Next, in the second part of the algorithm, the player proceeds to simulate a run the feasibility procedure with  $alg$  on period  $p$ . To do so, we first set the memory state of  $alg$  to  $\text{Message}$ . Each new query  $\mathbf{z}_i$  is submitted to the Oracle of Game 2 to get a response  $(\mathbf{g}_i, s_i)$ . Then, we compute  $\tilde{\mathbf{g}}_i$  as follows

$$\tilde{\mathbf{g}}_i = \begin{cases} \mathbf{g}_i & \text{if } s_i \geq 2, \\ \text{sign}(\mathbf{g}_i^\top \mathbf{z}_i) \mathbf{g}_i & \text{if } s_i = 1. \end{cases} \quad (12)$$

One can easily check that  $\tilde{\mathbf{g}}_i$  corresponds exactly to the response that was passed to  $alg$  in the first part of the strategy. The player then passes  $\tilde{\mathbf{g}}_i$  to  $alg$  so that it can update its state. We repeat this process for  $m$  steps. Further, the player can also keep track of the exploratory queries: the index  $i_l$  of the first response satisfying  $s_i = 2 + l$  for  $l \leq k-1$  (resp.  $s_i = 2 + k$ ) is the exploratory query which led to the construction of  $\mathbf{v}_{p,l}$  (resp.  $\mathbf{v}_{p+1,1}$ ) in the first part. Last, we check if the last index  $i_k$  was defined. If not, we pose  $i_k = m+1$  and let  $\mathbf{z}_{m+1}$  be the next query of  $alg$  with the current memory. The player then returns the vectors  $\frac{\mathbf{z}_{i_1}}{\|\mathbf{z}_{i_1}\|}, \dots, \frac{\mathbf{z}_{i_k}}{\|\mathbf{z}_{i_k}\|}$ . This ends the description of the player's strategy.

By Proposition 18, on an event  $\mathcal{E}$  of probability at least  $q$ , the algorithm  $alg$  succeeds and ends period  $p_{max}$ . As a result, similarly as in the proof of Proposition 11, since  $alg$  makes at most  $mp_{max}$  queries, and there are  $p_{max}$  periods, there must be a period of length at most  $m$ . Hence the strategy never fails at this phase of the player's strategy on the event  $\mathcal{E}$ . Further, we already checked that in the second phase, the vectors  $\tilde{\mathbf{g}}_i$  passed to  $alg$  coincide exactly with the responses passed to  $alg$  in the first part. Thus, this shows that during the second part, the player simulates exactly the run of the feasibility problem on period  $p$ . More precisely, the queries coincide with the queries in the feasibility problem at times  $i_{p,1}, \dots, \min\{i_{p,k}, i_{p,1} + m - 1\}$ . Now because the first part succeeded on  $\mathcal{E}$ , we have  $i_{p,k} \leq i_{p,0} + m$ . Therefore, if  $i_k$  has not yet been defined, this means that we had  $i_{p,k} = i_{p,1} + m$ . Hence, the next query with the current memory  $\mathbf{z}_{m+1}$  is exactly the query  $\mathbf{x}_{i_{p,k}}$  for the feasibility problem. This shows that the vectors  $\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_k}$  coincide exactly with the vectors  $\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,k}}$  when running  $alg$  on the feasibility problem in the first part.

We now show that the returned vectors are successful for Game 2. By construction,  $\mathbf{x}_{i_{p,1}}, \dots, \mathbf{x}_{i_{p,k}}$  are all informative. In particular,  $\|\mathbf{A}\mathbf{x}_{i_{p,l}}\|_\infty \leq \eta_0$  for all  $1 \leq l \leq k$ . Further, these queries did not fall

in scenario (2), hence  $\mathbf{v}_0^\top \mathbf{x}_{i_p,l} < -\eta_1$ , which implies  $\|\mathbf{x}_{i_p,l}\| > \eta_1$  for all  $l \leq k$ . As a result,

$$\frac{\|\mathbf{A}\mathbf{x}_{i_p,l}\|_\infty}{\|\mathbf{x}_{i_p,l}\|} \leq \frac{\eta_0}{\eta_1}.$$

Next fix  $l \leq k - 1$ . By construction of  $\mathbf{y}_{p,l}$ ,

$$\|P_{\text{Span}(\mathbf{x}_{i_p,l'}, l' \leq l)}(\mathbf{y}_{p,l})\|^2 = \sum_{l' \leq l} |\mathbf{b}_{p,l'}^\top \mathbf{y}_{p,l}|^2 \leq \frac{k}{d^6} \leq \frac{1}{d^5}.$$

Hence,

$$\|\mathbf{v}_{p,l} - P_{\text{Span}(\mathbf{x}_{i_p,l'}, l' \leq l)^\perp}(\mathbf{y}_{p,l})\| \leq \|P_{\text{Span}(\mathbf{x}_{i_p,l'}, l' \leq l)}(\mathbf{y}_{p,l})\| + \delta \leq \frac{1}{d^5} + \delta.$$

As a result, since  $\mathbf{x}_{p,l+1}^\top \mathbf{v}_{p,l} < -\eta_1$ , we have

$$\|P_{\text{Span}(\mathbf{x}_{i_p,l'}, l' \leq l)^\perp}(\mathbf{x}_{p,l+1})\| \geq |\mathbf{x}_{p,l+1}^\top P_{\text{Span}(\mathbf{x}_{i_p,l'}, l' \leq l)^\perp}(\mathbf{y}_{p,l})| > \eta_1 - \frac{1}{d^5} - \delta \geq \frac{\eta_1}{2}.$$

This shows that the returned vectors  $\frac{\mathbf{x}_{i_p,1}}{\|\mathbf{x}_{i_p,1}\|}, \dots, \frac{\mathbf{x}_{i_p,k}}{\|\mathbf{x}_{i_p,k}\|}$  are successful for Game 2 with parameters  $\alpha = \frac{\eta_0}{\eta_1}$  and  $\beta = \frac{\eta_1}{2}$ . This ends the proof that strategy succeeds on  $\mathcal{E}$  for these parameters, which ends the proof of the proposition.  $\blacksquare$

We are now ready to prove the main result.

**Proof of Theorem 2** Suppose that there is an algorithm *alg* for solving the feasibility problem to optimality  $\epsilon = 1/(48d^2\sqrt{d})$  with memory  $M$  and at most  $Q$  queries. Let  $k = \lceil 20 \frac{M+3d \log(2d)+1}{c_H n} \rceil$ . By Proposition 17, it solves the feasibility procedure with parameter  $k$  with probability at least  $1 - C\sqrt{\log d}/d$ . By Proposition 19 there is an algorithm for Game 2 that wins with probability  $1/3$  with  $m = \lceil Q/p_{max} \rceil$  and parameters  $\alpha = \eta_0/\eta_1$  and  $\beta = \eta_1/2$ . Now we check that

$$\alpha \left( \frac{\sqrt{d}}{\beta} \right)^{5/4} \leq 12d^2\eta_0 = \frac{1}{2}.$$

Hence, by Proposition 14, we have

$$m \geq \frac{c_H}{8(30 \log d + c_H)} d.$$

This shows that

$$Q \geq \Omega \left( p_{max} \frac{d}{\log d} \right) = \Omega \left( \frac{d^2}{k \log^3 d} \right) = \Omega \left( \frac{d^3}{(M + \log d) \log^3 d} \right).$$

This implies that for a memory  $M = d^{2-\delta}$  with  $0 \leq \delta \leq 1$  the number of queries is  $Q = \tilde{\Omega}(d^{1+\delta})$ .  $\blacksquare$

## Acknowledgments

This work was partly funded by ONR grant N00014-18-1-2122 and AFOSR grant FA9550-19-1-0263.



## References

- [1] Annie Marsden, Vatsal Sharan, Aaron Sidford, and Gregory Valiant. Efficient convex optimization requires superlinear memory. In *Conference on Learning Theory*, pages 2390–2430. PMLR, 2022.
- [2] Arkadij Semenovič Nemirovskij and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.
- [3] David B Yudin and Arkadi S Nemirovskii. Informational complexity and efficient methods for the solution of convex extremal problems. *Matekon*, 13(2):22–45, 1976.
- [4] Naum Z Shor. Cut-off method with space extension in convex programming problems. *Cybernetics*, 13(1):94–96, 1977.
- [5] Sergei Pavlovich Tarasov. The method of inscribed ellipsoids. In *Soviet Mathematics-Doklady*, volume 37, pages 226–230, 1988.
- [6] Ju E Nesterov. Self-concordant functions and polynomial-time methods in convex programming. *Report, Central Economic and Mathematic Institute, USSR Acad. Sci*, 1989.
- [7] David S Atkinson and Pravin M Vaidya. A cutting plane algorithm for convex programming that uses analytic centers. *Mathematical programming*, 69(1-3):1–43, 1995.
- [8] Pravin M Vaidya. A new algorithm for minimizing convex functions over convex sets. *Mathematical programming*, 73(3):291–341, 1996.
- [9] Anatoly Yur’evich Levin. An algorithm for minimizing convex functions. In *Doklady Akademii Nauk*, volume 160, pages 1244–1247. Russian Academy of Sciences, 1965.
- [10] Dimitris Bertsimas and Santosh Vempala. Solving convex programs by random walks. *Journal of the ACM (JACM)*, 51(4):540–556, 2004.
- [11] Yin Tat Lee, Aaron Sidford, and Sam Chiu-wai Wong. A faster cutting plane method and its implications for combinatorial and convex optimization. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1049–1065. IEEE, 2015.
- [12] Haotian Jiang, Yin Tat Lee, Zhao Song, and Sam Chiu-wai Wong. An improved cutting plane method for convex optimization, convex-concave games, and its applications. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 944–953, 2020.
- [13] Kurt M Anstreicher. The volumetric barrier for semidefinite programming. *Mathematics of Operations Research*, 25(3):365–380, 2000.
- [14] S Thomas McCormick. Submodular function minimization. *Handbooks in operations research and management science*, 12:321–391, 2005.
- [15] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Springer Science & Business Media, 2012.
- [16] Haotian Jiang. Minimizing convex functions with integral minimizers. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 976–985. SIAM, 2021.

- [17] Christos H Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM)*, 55(3):1–29, 2008.
- [18] Albert Xin Jiang and Kevin Leyton-Brown. Polynomial-time computation of exact correlated equilibrium in compact games. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 119–126, 2011.
- [19] Blake Woodworth and Nathan Srebro. Open problem: The oracle complexity of convex optimization with limited memory. In *Conference on Learning Theory*, pages 3202–3210. PMLR, 2019.
- [20] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.
- [21] Jacob Steinhardt and John Duchi. Minimax rates for memory-bounded sparse linear regression. In *Proceedings of The 28th Conference on Learning Theory*, pages 1564–1587. PMLR, 2015.
- [22] Vatsal Sharan, Aaron Sidford, and Gregory Valiant. Memory-sample tradeoffs for linear regression with small error. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, page 890–901. Association for Computing Machinery, 2019.
- [23] Ioannis Mitliagkas, Constantine Caramanis, and Prateek Jain. Memory limited, streaming pca. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’13, page 2886–2894, Red Hook, NY, USA, 2013. Curran Associates Inc.
- [24] Jacob Steinhardt, Gregory Valiant, and Stefan Wager. Memory, communication, and statistical queries. In *29th Annual Conference on Learning Theory*, pages 1490–1516. PMLR, 2016.
- [25] Gavin Brown, Mark Bun, Vitaly Feldman, Adam Smith, and Kunal Talwar. When is memorization of irrelevant training data necessary for high-accuracy learning? In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, page 123–132. Association for Computing Machinery, 2021.
- [26] Gavin Brown, Mark Bun, and Adam Smith. Strong memory lower bounds for learning natural models. In *Proceedings of Thirty Fifth Conference on Learning Theory*, pages 4989–5029. PMLR, 2022.
- [27] Dana Moshkovitz and Michal Moshkovitz. Mixing implies lower bounds for space bounded learning. In *Proceedings of the 2017 Conference on Learning Theory*, pages 1516–1566. PMLR, 2017.
- [28] Dana Moshkovitz and Michal Moshkovitz. Entropy Samplers and Strong Generic Lower Bounds For Space Bounded Learning. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*, volume 94 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 28:1–28:20. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.
- [29] Paul Beame, Shayan Oveis Gharan, and Xin Yang. Time-space tradeoffs for learning finite functions from random evaluations, with applications to polynomials. In *Proceedings of the 31st Conference On Learning Theory*, pages 843–856. PMLR, 2018.
- [30] Sumegha Garg, Ran Raz, and Avishay Tal. Extractor-based time-space lower bounds for learning. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, page 990–1002. Association for Computing Machinery, 2018.

- [31] Gillat Kol, Ran Raz, and Avishay Tal. Time-space hardness of learning sparse parities. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, page 1067–1080. Association for Computing Machinery, 2017.
- [32] Ran Raz. A time-space lower bound for a large class of learning problems. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 732–742, 2017. doi: 10.1109/FOCS.2017.73.
- [33] A.S. Nemirovsky, D.B. Yudin, and E.R. Dawson. *Problem Complexity and Method Efficiency in Optimization*. A Wiley-Interscience publication. Wiley, 1983. ISBN 978-0471103455.
- [34] Blake E Woodworth and Nati Srebro. Tight complexity bounds for optimizing composite objectives. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [35] Blake E. Woodworth and Nathan Srebro. Lower bound for randomized first order convex optimization. *arXiv: Optimization and Control*, 2017.
- [36] Jorge Nocedal. Updating quasi-newton matrices with limited storage. *Mathematics of Computation*, 35(151):773–782, 1980.
- [37] Dong C. Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1):503–528, August 1989.
- [38] Adrian S. Lewis and Michael L. Overton. Nonsmooth optimization via quasi-Newton methods. *Mathematical Programming*, 141(1):135–163, October 2013.
- [39] Arkadi Nemirovski. On parallel complexity of nonsmooth convex optimization. *Journal of Complexity*, 10(4):451–463, 1994.
- [40] Eric Balkanski and Yaron Singer. Parallelization does not accelerate convex optimization: Adaptive lower bounds for non-smooth convex minimization. *arXiv preprint arXiv:1808.03880*, 2018.
- [41] Sébastien Bubeck, Qijia Jiang, Yin-Tat Lee, Yuanzhi Li, and Aaron Sidford. Complexity of highly parallel non-smooth convex optimization. *Advances in neural information processing systems*, 32, 2019.
- [42] Uriel Feige and Gideon Schechtman. On the optimality of the random hyperplane rounding technique for max cut. *Random Structures & Algorithms*, 20(3):403–440, 2002.

## A Concentration bounds

The following result gives concentration bounds for the norm of the projection of a random unit vector onto linear subspaces.

**Proposition 20.** *Let  $P$  be a projection in  $\mathbb{R}^d$  of rank  $r$  and let  $\mathbf{x} \in \mathbb{R}^d$  be a random vector sampled uniformly on the unit sphere  $\mathbf{x} \sim \mathcal{U}(S^{d-1})$ . Then, for every  $t > 0$ ,*

$$\max \left\{ \mathbb{P} \left( \|P(\mathbf{x})\|^2 - \frac{r}{d} \geq t \right), \mathbb{P} \left( \|P(\mathbf{x})\|^2 - \frac{r}{d} \leq -t \right) \right\} \leq e^{-dt^2}.$$

Further, if  $r = 1$  and  $d \geq 2$ ,

$$\mathbb{P} \left( \|P(\mathbf{x})\| \geq \sqrt{\frac{t}{d-1}} \right) \leq 2\sqrt{t}e^{-t/2}.$$

**Proof** First, by isometry, we can assume that  $P$  is the projection onto the coordinate vectors  $\mathbf{e}_1, \dots, \mathbf{e}_r$ . Then, let  $\mathbf{y} \sim \mathcal{N}(0, 1)$  be a normal vector. Note that  $\mathbf{x} = \frac{\mathbf{y}}{\|\mathbf{y}\|} \sim \mathcal{U}(S^{d-1})$ . Further,

$$\|\mathbf{x}\|^2 \geq \frac{r}{d} + t \iff \left(1 - \frac{r}{d} - t\right) \sum_{i=1}^r y_i^2 \geq \left(\frac{r}{d} + t\right) \sum_{i=r+1}^d y_i^2.$$

Now note that  $Z_1 = \sum_{i=1}^r y_i^2$  and  $Z_2 = \sum_{i=r+1}^d y_i^2$  are two independent random chi squared variables of parameters  $r$  and  $d - r$  respectively. Recalling that the moment generating function of  $Z \sim \chi^2(k)$  is  $\mathbb{E}[e^{sZ}] = (1 - 2s)^{-k/2}$  for  $s < 1/2$ . Therefore, for any

$$-\frac{1}{2(r/d + t)} < s < \frac{1}{2(1 - r/d - t)}, \quad (13)$$

one has

$$\begin{aligned} \mathbb{P} \left( \|P(\mathbf{x})\|^2 - \frac{r}{d} \geq t \right) &\leq \mathbb{E} \left[ \exp \left( s \left(1 - \frac{r}{d} - t\right) Z_1 - s \left(\frac{r}{d} + t\right) Z_2 \right) \right] \\ &= \frac{[1 - 2s(1 - \frac{r}{d} - t)]^{-r/2}}{[1 - 2s(\frac{r}{d} + t)]^{-(d-r)/2}}. \end{aligned}$$

Now let  $s = \frac{1}{2} \left( \frac{1-r/d}{1-r/d-t} - \frac{r/d}{r/d+t} \right)$ , which satisfies Eq (13). The previous equation readily yields

$$\mathbb{P} \left( \left| \|P(\mathbf{x})\|^2 - \frac{r}{d} \right| \geq t \right) \leq \exp \left( -\frac{d}{2} d_{KL} \left( \frac{r}{d}; \frac{r}{d} + t \right) \right) \leq e^{-dt^2}.$$

In the last inequality we used Pinsker's inequality  $d_{KL}(r/d; r/d + t) \geq 2\delta(\mathcal{B}(r/d), \mathcal{B}(d/r + t))^2 = 2t^2$ , where  $\mathcal{B}(q)$  is the Bernoulli distribution of parameter  $q$ . Replacing  $P$  with  $Id - P$  and  $r$  with  $d - r$  gives the other inequality

$$\mathbb{P} \left( \|P(\mathbf{x})\|^2 - \frac{r}{d} \leq -t \right) \leq e^{-dt^2}.$$

This gives first claim. For the second claim, supposing that  $r = 1 < d$ , from the above equation, we have

$$\mathbb{P} \left( \|P(\mathbf{x})\|^2 \geq \frac{t}{d} \right) \leq \exp \left( -\frac{d}{2} d_{KL} \left( \frac{1}{d}; \frac{t}{d} \right) \right) = \sqrt{t} \left( \frac{1 - \frac{t}{d}}{1 - \frac{1}{d}} \right)^{(d-1)/2} \leq \sqrt{2t} e^{-t(d-1)/(2d)}.$$

Thus,

$$\mathbb{P} \left( \|P(\mathbf{x})\|^2 \geq \frac{t}{d-1} \right) \leq \sqrt{\frac{2(d-1)}{d}} \sqrt{t} e^{-t/2},$$

which ends the proof of the proposition. ■

Next, we need the following lemma which gives a concentration inequality for discretized samples in  $\mathcal{D}_d$  and approximately perpendicular to  $k \leq d/3 - 1$  vectors.

**Lemma 21.** Let  $0 \leq k \leq d/3 - 1$  and  $\mathbf{x}_1, \dots, \mathbf{x}_k \in B_d(0, 1)$  be  $k$  orthonormal vectors in the unit ball, and  $\mathbf{x} \in B_d(0, 1)$ . Denote by  $\mu$  the distribution on the unit sphere corresponding to the uniform distribution  $\mathbf{y} \sim \mathcal{U}(S^{d-1} \cap \{\mathbf{w} \in \mathbb{R}^d : |\mathbf{x}_i^\top \mathbf{w}| \leq d^{-3}, \forall i \leq k\})$ . Let  $\mathbf{y} \sim \mu$ . Then, for  $t \geq 2$ ,

$$\mathbb{P} \left( |\mathbf{x}^\top \mathbf{y}| \geq \sqrt{\frac{t}{d}} + \frac{1}{d^2} \right) \leq 2\sqrt{t}e^{-t/3}.$$

Further, let  $\delta \leq 1$  and  $\mathbf{z} = \phi_\delta(\mathbf{y})$ . Then for  $t \geq 4$ ,

$$\mathbb{P} \left( |\mathbf{x}^\top \mathbf{z}| \geq \sqrt{\frac{t}{d}} + \frac{1}{d^2} + \delta \right) \leq 2\sqrt{t}e^{-t/3}.$$

**Proof** We use the same notations as above and denote by  $\mathcal{E} = \{|\mathbf{x}_i^\top \mathbf{y}| \leq d^{-3}, \forall i \leq k\}$  the event considered and  $\mathbf{y} \sim \mu$ . We decompose  $\mathbf{y} = \alpha_1 \mathbf{x}_1 + \dots + \alpha_k \mathbf{x}_k + \mathbf{y}'$ , where  $\mathbf{y}' \in \text{Span}(\mathbf{x}_i, i \leq k)^\perp := E$ . Now note that  $\frac{\mathbf{y}'}{\|\mathbf{y}'\|}$  is a uniformly random unit vector in  $E$ . As a result, using Proposition 20, we obtain for any  $t \geq 2$ ,

$$\begin{aligned} \mathbb{P} \left( |\mathbf{x}^\top \mathbf{y}'| \geq \sqrt{\frac{t}{d-k-1}} \right) &= \mathbb{P} \left( |P_E(\mathbf{x})^\top \mathbf{y}'| \geq \sqrt{\frac{t}{d-k-1}} \right) \\ &\leq 2\sqrt{t}e^{-t/2}. \end{aligned}$$

Also, because by definition of  $\mu$ , we have  $|\alpha_i| \leq d^{-3}$  for all  $i \leq k$ , we obtain  $|\mathbf{x}^\top \mathbf{y}| \leq \frac{k}{d^3} + |\mathbf{x}^\top \mathbf{y}'| \leq \frac{1}{d^2} + |\mathbf{x}^\top \mathbf{y}'|$ . As a result, using the fact that  $d - k - 1 \geq 2d/3$ , the previous equation shows that

$$\mathbb{P} \left( |\mathbf{x}^\top \mathbf{y}| \geq \sqrt{\frac{3t}{2d}} + \frac{1}{d^2} \right) \leq \mathbb{P} \left( |\mathbf{x}^\top \mathbf{y}'| \geq \sqrt{\frac{t}{d-k-1}} \right) \leq 2\sqrt{t}e^{-t/2}.$$

Next, we use the fact that  $\|\mathbf{z} - \mathbf{y}\| = \|\phi_\delta(\mathbf{y}) - \mathbf{y}\| \leq \delta$  to obtain

$$\mathbb{P} \left( |\mathbf{x}^\top \mathbf{z}| \geq \sqrt{\frac{t}{d}} + \frac{1}{d^2} + \delta \right) \leq \mathbb{P} \left( |\mathbf{x}^\top \mathbf{y}| \geq \sqrt{\frac{t}{d}} + \frac{1}{d^2} \right) \leq 2\sqrt{t}e^{-t/3}.$$

This ends the proof of the lemma. ■

## B An improved result on robustly-independent vectors

The following lemma serves the same purpose as [1, Lemma 34]. Namely, from successful vectors of the Game 2, it allows to recover an orthonormal basis that is still approximately in the nullspace of  $\mathbf{A}$ . The following version gives a stronger version that improves the dependence in  $d$  of our chosen parameters.

**Lemma 22.** Let  $\delta \in (0, 1]$  and suppose that we have  $r \leq d$  unit norm vectors  $\mathbf{y}_1, \dots, \mathbf{y}_r \in \mathbb{R}^d$ . Suppose that for any  $i \leq k$ ,

$$\|P_{\text{Span}(\mathbf{y}_j, j < i)^\perp}(\mathbf{y}_i)\| \geq \delta.$$

Let  $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_r]$  and  $s \geq 2$ . There exists  $\lceil r/s \rceil$  orthonormal vectors  $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_{\lceil r/s \rceil}]$  such that for any  $\mathbf{a} \in \mathbb{R}^d$ ,

$$\|\mathbf{Z}^\top \mathbf{a}\|_\infty \leq \left( \frac{\sqrt{d}}{\delta} \right)^{s/(s-1)} \|\mathbf{Y}^\top \mathbf{a}\|_\infty.$$

**Proof** Let  $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_r)$  be the orthonormal basis given by the Gram-Schmidt decomposition of  $\mathbf{y}_1, \dots, \mathbf{y}_r$ . By definition of the Gram-Schmidt decomposition, we can write  $\mathbf{Y} = \mathbf{BC}$  where  $\mathbf{C}$  is an upper-triangular matrix. Further, its diagonal is exactly  $\text{diag}(\|P_{\text{Span}(\mathbf{y}_{l'}, l' < l)^\perp}(\mathbf{y}_l)\|, l \leq r)$ . Hence,

$$\det(\mathbf{Y}) = \det(\mathbf{C}) = \prod_{l \leq r} \|P_{\text{Span}(\mathbf{y}_{l'}, l' < l)^\perp}(\mathbf{y}_l)\| \geq \delta^r.$$

We now introduce the singular value decomposition  $\mathbf{Y} = \mathbf{U} \text{diag}(\sigma_1, \dots, \sigma_r) \mathbf{V}^\top$ , where  $\mathbf{U} \in \mathbb{R}^{d \times r}$  and  $\mathbf{V} \in \mathbb{R}^{r \times r}$  have orthonormal columns, and  $\sigma_1 \geq \dots \geq \sigma_r$ . Next, for any vector  $\mathbf{z} \in \mathbb{R}^d$ , since the columns of  $\mathbf{Y}$  have unit norm,

$$\|\mathbf{Y}\mathbf{z}\|_2 \leq \sum_{l \leq r} |z_l| \|\mathbf{y}_l\|_2 \leq \|\mathbf{z}\|_1 \leq \sqrt{d} \|\mathbf{z}\|_2.$$

In the last inequality we used Cauchy-Schwartz. Therefore, all singular values of  $\mathbf{Y}$  are upper bounded by  $\sigma_1 \leq \sqrt{d}$ . Thus, with  $r' = \lceil r/s \rceil$

$$\delta^r \leq \det(\mathbf{Y}) = \prod_{l=1}^r \sigma_l \leq d^{(r'-1)/2} \sigma_{r'}^{r-r'+1} \leq d^{r/2s} \sigma_{r'}^{(s-1)r/s},$$

so that  $\sigma_{r'} \geq \delta^{s/(s-1)} / d^{1/(2s)}$ . We are ready to define the new vectors. We pose for all  $i \leq r'$ ,  $\mathbf{z}_i = \mathbf{u}_i$  the  $i$ -th column of  $\mathbf{U}$ . These correspond to the  $r'$  largest singular values of  $\mathbf{Y}$  and are orthonormal by construction. Then, for any  $i \leq r'$ , we also have  $\mathbf{z}_i = \mathbf{u}_i = \frac{1}{\sigma_i} \mathbf{Y} \mathbf{v}_i$  where  $\mathbf{v}_i$  is the  $i$ -th column of  $\mathbf{V}$ . Hence, for any  $\mathbf{a} \in \mathbb{R}^d$ ,

$$|\mathbf{z}_i^\top \mathbf{a}| = \frac{1}{\sigma_i} |\mathbf{v}_i^\top \mathbf{Y}^\top \mathbf{a}| \leq \frac{\|\mathbf{v}_i\|_1}{\sigma_i} \|\mathbf{Y}^\top \mathbf{a}\|_\infty \leq \frac{d^{1/2+1/(2s)}}{\delta^{s/(s-1)}} \|\mathbf{Y}^\top \mathbf{a}\|_\infty.$$

This ends the proof of the lemma. ■