Sparse PCA with Overlapping Supports: NP-Hardness and an FPTAS

Alberto Del Pia* Dekun Zhou†

November 11, 2026

Abstract

In sparse principal component analysis (SPCA) with several components, each must be sparse and all must be pairwise orthogonal. When the supports of the components overlap, these two requirements interact. For a covariance matrix of fixed rank, polynomial-time algorithms with optimality or approximation guarantees are known for one component, and for several components with disjoint supports or with a common support. We study the general model, with individual sparsity constants and arbitrary support overlaps. We are given a rational $n \times r$ matrix R and d sparsity constants $s_i \leq n$, and the goal is to find an $n \times d$ matrix X with orthonormal columns that maximizes $\left\|R^\top X\right\|_F^2$, such that column i has at most s_i non-zero entries. Here $RR^\top$ plays the role of the covariance matrix and has rank at most r. We show that SPCA is NP-hard even when $r = d = 2$, but admits a fully polynomial-time approximation scheme (FPTAS) when r and d are fixed. The returned solution has exactly orthonormal columns, and its entries are real algebraic numbers in an exact representation. For fixed ϵ, the number of steps of the algorithm is linear in n up to a polylogarithmic factor. Our approach reduces SPCA to a sparse version of the orthogonal Procrustes problem, that is, to maximizing a linear function over the feasible region of SPCA. We show that this problem is NP-hard even for $d = 2$, and we give an FPTAS for it when d is fixed. This FPTAS combines approximate dynamic programming over the rows, two-column updates that restore exact orthogonality, and an exact computation of the few variables for which these updates are not available.

Keywords: Sparse PCA · Procrustes problem · Approximation scheme · NP-hardness · Approximate dynamic programming.

1 Introduction

Principal component analysis (PCA) finds a few orthogonal directions, called *components*, that capture as much of the variance of a dataset as possible. Sparse PCA (SPCA), in addition, limits the number of features used by each component, making the components easier to interpret. With several components, sparsity and orthogonality interact through overlapping supports. Enforcing orthogonality between two components may disrupt their orthogonality with other components, while standard orthogonalization may violate the sparsity constraints.

Already for one component, SPCA is NP-hard and admits no fully polynomial-time approximation scheme (FPTAS) unless P = NP [26], and it is hard to approximate within any constant factor under stronger complexity assumptions [26, 11]. This motivates structural assumptions on the covariance matrix, such as a fixed rank. Fixed rank is also relevant for approximating SPCA on general covariance matrices. Indeed, for a single component, applying an algorithm for low-rank matrices to a low-rank approximation of the covariance matrix gives approximation guarantees that depend on the decay of the spectrum of the covariance matrix [28].

For a covariance matrix of fixed rank, polynomial-time algorithms with optimality or approximation guarantees are known for one component, and for several components with disjoint supports or with a

*Department of Industrial and Systems Engineering & Wisconsin Institute for Discovery, University of Wisconsin–Madison, Madison, WI, USA. E-mail: delpia@wisc.edu.

†E-mail: dekun.zhou6@gmail.com.

common support. Asteris, Papailiopoulos, and Karystinos [2] solve single-component SPCA exactly in polynomial time for fixed rank. For components with disjoint supports, orthogonality holds automatically. In this case, Asteris, Papailiopoulos, Kyrillidis, and Dimakis [3] give a $(1-\epsilon)$-approximation algorithm for every $\epsilon \in (0,1)$, whose running time is polynomial in the input size and in $1/\epsilon$ for fixed rank and number of components. Del Pia [15] solves this case exactly in polynomial time for fixed rank and number of components. For fixed rank, the same paper also solves exactly row-sparse PCA, where all components are supported on a common set of at most k rows, for a given integer k. Without a rank assumption, Cory-Wright and Pauphilet [14] study the general model with overlapping supports through semidefinite, Lagrangian, and combinatorial relaxations, and obtain computational methods with instance-specific optimality guarantees.

The general model with individual sparsity constants and arbitrary support overlaps remains open even for fixed rank. Indeed, restricting the rank simplifies the objective but does not remove the support-dependent orthogonality constraints. We ask whether fixing the rank and the number of components makes this model tractable. We show that the answer is negative for exact optimization and positive for approximation.

Notation. For a positive integer m, let $[m] := \{1, 2, \dots, m\}$. Denote the *support* of a vector x by $\operatorname{Supp}(x) := \{i : x_i \neq 0\}$, and let $\|x\|_0 := |\operatorname{Supp}(x)|$. Denote by $\mathbf{e}_i$ the i-th standard basis vector. We denote by $\mathbb{1}_J$ the indicator vector of a set J and by $\mathbb{1}_{\{P\}}$ the indicator of a condition P. For a matrix X, X_i denotes its i-th column and $X_{J,T}$ the submatrix with row indices J and column indices T. We denote by $\|\cdot\|_2$ the 2-norm for vectors and the spectral norm for matrices, and by $\|\cdot\|_F$ the Frobenius norm. We denote the trace by $\operatorname{tr}(\cdot)$, the identity matrix of order m by I_m, and a zero matrix by O.

The problem. Let $R \in \mathbb{Q}^{n\times r}$, let $d \leq n$ be the number of components, and let $s = (s_1, \dots, s_d) \in [n]^d$. The integers s_i are the *sparsity constants*. We study the problem

$$\mathsf{OPT}(R; s) := \max_{X \in \mathbb{R}^{n\times d}} \left\|R^\top X\right\|_F^2 \quad \text{s.t.} \quad X^\top X = I_d, \quad \|X_i\|_0 \leq s_i, \quad i \in [d]. \tag{SPCA}$$

Thus, s_i is the maximum number of non-zero entries of the i-th component X_i. We denote by $\mathcal{F}_s$ the feasible region of SPCA. Since $\left\|R^\top X\right\|_F^2 = \operatorname{tr}(X^\top C X)$ with $C := RR^\top$, the matrix C plays the role of the covariance matrix, and its rank is at most r.

Main results. Our first result shows that SPCA is NP-hard even for $r = d = 2$:

Theorem 1. *Given $R \in \mathbb{Z}^{n\times 2}$, $k \in [n]$, and $\tau_0 \in \mathbb{Z}$, it is NP-hard to decide whether $\mathsf{OPT}(R; (k,k)) \geq \tau_0$.*

Theorem 1 cannot be strengthened by lowering r or d. For $d = 1$ and fixed r, SPCA can be solved exactly in polynomial time [2, 15]. For $r = 1$, SPCA can be solved in polynomial time for every d (see Observation 11 in Appendix C).

We prove Theorem 1 by a reduction from Partition, which shows how overlapping supports create hardness. Let m be the number of integers in the given instance of Partition. In the constructed instance of SPCA, one component must have equal entries on $m+1$ rows and be orthogonal to a vector that encodes these integers. Hence the entries of this vector in these rows sum to zero, and these rows give a partition.

Our main result is an FPTAS.

Theorem 2. *For fixed r and d, there is an algorithm that, given $R \in \mathbb{Q}^{n\times r}$, $s \in [n]^d$, and $\epsilon \in \mathbb{Q} \cap (0,1)$, returns $X \in \mathcal{F}_s$ with $\left\|R^\top X\right\|_F^2 \geq (1-\epsilon)\mathsf{OPT}(R; s)$, in time polynomial in the encoding size of (R, s, ϵ) and in $1/\epsilon$.*

Since this reduction is from Partition, Theorems 1 and 2 give the same picture as for the knapsack problem: for fixed r and d, SPCA is NP-hard in the weak sense, and it admits an FPTAS.

The matrix returned by the algorithm of Theorem 2 has exactly orthonormal columns. All running times are in the bit model. The input is rational, and the entries of the output are real algebraic numbers. The algorithm outputs them exactly, in a standard finite encoding by an integer polynomial and an isolating interval (Appendix C.4). For fixed r and d, its bit complexity is at most $n\,\epsilon^{-\mathcal{O}(rd+d^3)} \log^{\mathcal{O}(d^2)}(2n/\epsilon)$ times a polynomial in the encoding size of (R, s, ϵ). For fixed r, d, and ϵ, the number of steps of the algorithm is nearly linear in n, that is, linear in n up to a polylogarithmic factor (see Section 3). Theorems 9 and 10 in Section 3 give the precise approximation guarantee and bit complexity.

The linear problem. Inspired by the auxiliary-variable and net approach of Asteris et al. [3, Lemma 9.2], we use Frobenius-norm duality to reduce SPCA to maximizing a linear function over $\mathcal{F}_s$. Specifically, by the Cauchy–Schwarz inequality,

$$\left\|R^\top X\right\|_F = \max\{\operatorname{tr}(W^\top R^\top X) : W \in \mathbb{R}^{r\times d},\ \|W\|_F \le 1\}.$$

Hence, up to a factor $1-\epsilon$, SPCA reduces to maximizing $(rd/\epsilon)^{\mathcal{O}(rd)}$ linear functions $\operatorname{tr}(A^\top X)$ over $\mathcal{F}_s$, with $A = RW$ for W in a grid. For $A \in \mathbb{Q}^{n\times d}$, let

$$\mathsf{OPT}_{\text{linear}}(A; s) := \max_{X\in\mathcal{F}_s} \operatorname{tr}(A^\top X). \tag{1}$$

For disjoint supports, the auxiliary problem of Asteris et al. [3] maximizes a sum of squared projections and reduces to bipartite matching. Our reduction instead uses the linear objective in (1). Problem (1) is a sparse version of the orthogonal Procrustes problem. Without the sparsity constraints, an optimal solution is obtained from a singular value decomposition of A (see [33, Theorem 4]), and Procrustes reformulations are used in algorithms for orthogonal sparse PCA [7]. With overlapping supports, the linear problem already contains the interaction between sparsity and orthogonality. Our next result shows that, unlike the classical problem, (1) is NP-hard even for $d = 2$.

Theorem 3. *Given $A \in \mathbb{Z}^{n\times 2}$, $k \in [n]$, and $\tau_0 \in \mathbb{Z}$, it is NP-hard to decide whether $\mathsf{OPT}_{\text{linear}}(A; (k,k)) \ge \tau_0$.*

The proof of Theorem 3 is similar to that of Theorem 1. In the exact algorithms of Del Pia [15], an auxiliary matrix with r rows plays a role similar to that of W, and the space of these matrices is partitioned exactly by a hyperplane arrangement. Here, by Theorem 3, the problem for a fixed W is already NP-hard, so the fundamental difficulty is not in enumeration of W, and we discretize this space directly. Our main technical result is an FPTAS for the linear problem when d is fixed.

Theorem 4. *For fixed d, there is an algorithm that, given $A \in \mathbb{Q}^{n\times d}$, $s \in [n]^d$, and $t \in \mathbb{Q} \cap (0,1)$, returns $X \in \mathcal{F}_s$ with $\operatorname{tr}(A^\top X) \ge (1-t)\mathsf{OPT}_{\text{linear}}(A; s)$, in time polynomial in the encoding size of (A, s, t) and in $1/t$.*

As for Theorem 2, for fixed d and t, the number of steps of this algorithm is linear in n up to a polylogarithmic factor. Theorem 4 holds for every $A \in \mathbb{Q}^{n\times d}$, so it gives an FPTAS for this sparse version of the Procrustes problem with a fixed number of columns, which may be of independent interest. In our algorithm for SPCA, fixing r is needed only to bound the number of matrices W in the grid. Theorem 2 follows from Theorem 4 and the reduction above.

Our techniques. The objective $\operatorname{tr}(A^\top X)$ and the sparsity constraints decompose over the rows of X, while the constraint $X^\top X = I_d$ couples all rows. The algorithm of Theorem 4 handles this coupling with three ingredients.

(a) *Approximate dynamic programming.* For every row, we choose the set of columns that may be non-zero in that row, called its row type, and the values of its entries. These choices are made by an approximate dynamic program (ADP) over a balanced binary tree whose leaves are the rows. Each state records rounded upper bounds on the number of rows in which each column may be non-zero, and the Gram matrix $X^\top X$ rounded relative to its trace. Since the tree has depth $\lceil \log_2 n \rceil$, the rounding errors accumulate over only logarithmically many levels. For fixed d and t, the number of states at each node is polylogarithmic in n. The candidate solutions at the root are approximately orthogonal.

(b) *Exact orthogonality by two-column updates.* Orthogonalizing the candidate solutions at the root, for example by the Gram–Schmidt process, would in general violate the sparsity constraints. Suppose that two columns i and l may both be non-zero in at least d common rows. In these rows, the other $d-2$ columns span a space of dimension at most $d-2$. Hence the vectors supported on these rows that are orthogonal to all the other columns form a space of dimension at least two. Changing columns i and l inside this space makes them orthogonal, does not change their inner products with the other columns, and does not decrease the objective (Proposition 8). A sequence of such updates, followed by a normalization, turns an approximately orthogonal candidate solution into a point in $\mathcal{F}_s$.

Table 1: Parameters of LinearOracle for accuracy t, where n is the number of rows and d the number of columns of A.

Parameter	Value	Role
δ	$t/(8(d+1))$	error allowed for $\left\|X^\top X - I_d\right\|_F$ at the root
D_d	$d(d-1)^2$	upper bound on the number D of hard variables
a	$\delta/(256(d+1)(D_d+1))$	upper bound on the side length of the boxes
J_p	$\lceil 2/a \rceil$	number of intervals covering $[-1,1]$
M	$\lceil 256d(d+1)/\delta \rceil$	inverse of relative mesh size in $\mathcal{V}$
J_0	$\lceil \log_2(256dn(d+1)/\delta) \rceil$	determines the smallest non-zero magnitude 2^{-J_0} in $\mathcal{V}$
H_{tree}	$\lceil \log_2 n \rceil$	height of the ADP tree
τ	$\delta/(128d^2)$	slack in the sparsity capacity usages
P_d	$d\binom{d}{2}$	upper bound on the number of kept rows in the analysis
K_0	$\lceil 16(P_d+1)/\tau \rceil$	largest capacity cap that is not rounded
α	$\tau/(8(H_{\text{tree}}+1))$	the geometric part of U_c has ratio $1+\alpha$
θ	$\delta/(128d^2(H_{\text{tree}}+1))$	relative width of the Gram matrix bins
κ_G	$\delta/(128n(H_{\text{tree}}+1))$	trace threshold of the small-trace bin

(c) *Hard pairs.* Pairs of columns with fewer than d common rows cannot be repaired in this way, and we call them hard. The number of variables involved in hard pairs is at most $d(d-1)^2$, independently of n. We fix the values of these variables before running the ADP. To this end, we cover the space of these variables by small boxes. In each box, we compute an exact solution of the orthogonality constraints of the hard pairs, if one exists, with algorithms from real algebraic geometry [6, 5].

Organization. Section 2 describes the algorithm for the linear problem, and Section 3 states the approximation guarantee, sketches its proof, and bounds the running time. Appendix A discusses further related work, Appendix B contains the pseudocode, and Appendix C contains all proofs.

2 The algorithm for the linear problem

This section describes the algorithm LinearOracle of Theorem 4 for the linear problem (1), with a fixed objective matrix $A \in \mathbb{Q}^{n\times d}$. Our algorithm for SPCA calls LinearOracle with accuracy $t = \epsilon/4$ and with $A = RW$ for every matrix W in an $(\epsilon/4)$-net of $\{W \in \mathbb{R}^{r\times d} : \|W\|_F \le 1\}$, and returns the best solution found (Algorithm 3 in Appendix B). Table 1 lists the parameters of LinearOracle for an accuracy t, and Appendix B gives the pseudocode of its subroutines.

2.1 Row types and column pairs

A *candidate solution* is a matrix in $\mathbb{R}^{n\times d}$ constructed by our algorithm, which may violate the constraint $X^\top X = I_d$. We decompose a candidate solution according to the set of columns that may be non-zero in each row.

Definition 5. *A* row type T *is a subset of* $[d]$*, indicating which entries of a row may be non-zero.*

Our algorithm fixes, row by row, which entries of a candidate solution may be non-zero, by assigning a row type to each row. The same row type can be assigned to several rows. There are only 2^d row types, a number that does not depend on n. For a candidate solution X whose j-th row is assigned a row type, we denote this row type by $\mathrm{RT}_j(X)$. For a candidate solution X with all its rows assigned row types, we classify the pairs of columns by the number of assigned row types that contain them.

Definition 6. *Let X be a candidate solution with all its rows assigned row types, and let* $\mathrm{CP}_d := \{\{i,l\} : 1 \le i < l \le d\}$ *be the set of* column pairs. *For every column pair $e \in \mathrm{CP}_d$, let* $m_e(X) := |\{j \in [n] : e \subseteq \mathrm{RT}_j(X)\}|$

be the number of rows whose assigned row type contains e. The hard pairs *are the pairs in* $H(X) := \{e \in \mathrm{CP}_d : m_e(X) < d\}$*, and the* simple pairs *are the pairs in* $S(X) := \mathrm{CP}_d \setminus H(X)$.

Our algorithm treats the rows whose assigned row types contain a hard pair separately, so we collect their row types in a list.

Definition 7. *For a candidate solution X with all its rows assigned row types, the* list of hard row types $\mathcal{L}(X)$ *is the multiset of the assigned row types of X that contain a hard pair. Let $b(X) := |\mathcal{L}(X)|$. The elements of $\mathcal{L}(X)$ have labels $1, \dots, b(X)$, one for each row whose assigned row type contains a hard pair, and we write $\mathcal{L}(X) = \{\!\{T^1, T^2, \dots, T^{b(X)}\}\!\}$.*

The rows whose assigned row types are in $\mathcal{L}(X)$ are the *hard rows*. The *hard variables* of X are the entries (j, i) such that $i \in e \subseteq \mathrm{RT}_j(X)$ for some $e \in H(X)$, and $D(X)$ denotes their number. Every hard row contains a hard pair, and every hard pair is contained in the assigned row types of at most $d-1$ rows. Hence $b(X) \le (d-1)\,|H(X)|$ and $D(X) \le 2(d-1)\,|H(X)| \le d(d-1)^2$ (Lemma 12 in Appendix C).

The two-column update of Section 2.2 makes a simple pair $\{i, l\}$ orthogonal using d rows whose assigned row types contain $\{i, l\}$, and a hard pair does not have that many rows. We therefore fix the values of the hard variables in advance. Our algorithm determines the hard pairs, the simple pairs, and the list of hard row types before the values of X. It first enumerates the sets H and S and the list $\mathcal{L}$, and then fixes the values of the hard variables. After that, it assigns the row types in $\mathcal{L}$ to b rows of the candidate solution, assigns row types without hard pairs to the remaining rows, and finally assigns values to the entries. Therefore, in the rest of the paper, we drop the dependence on X from the notation of this section.

2.2 Two-column updates for exact orthogonality

Let X be a candidate solution, and let $e = \{i, l\} \in S$ be a simple pair. By modifying at most $2d$ entries of the columns X_i and X_l, we obtain a matrix $\bar{X}$ such that (a) $\bar{X}_i^\top \bar{X}_l = 0$, (b) $\bar{X}_i^\top X_p = X_i^\top X_p$ and $\bar{X}_l^\top X_p = X_l^\top X_p$ for every $p \in [d] \setminus \{i, l\}$, and (c) $\mathrm{tr}(A^\top \bar{X}) \ge \mathrm{tr}(A^\top X)$. The modified entries lie in d rows whose assigned row types contain e. As explained in Section 1, the vectors supported on these rows that are orthogonal to all other columns form a space of dimension at least two. We change columns i and l only inside this space, which gives (b). To obtain (a) and (c), we choose the new entries by solving a problem of the form

$$\max\left\{\mathrm{tr}(\Gamma^\top Z) : Z \in \mathbb{R}^{h\times k},\ Z^\top Z = \Lambda\right\}, \tag{2}$$

where $\Gamma \in \mathbb{R}^{h\times k}$ with $k \le h$, and $\Lambda \in \mathbb{R}^{k\times k}$ is positive semidefinite. By von Neumann's trace inequality [27], the optimal value of (2) is $\mathrm{tr}(\sqrt{\Gamma\Lambda\Gamma^\top})$, and an optimal solution can be obtained from a singular value decomposition of $\Gamma\sqrt{\Lambda}$ (Proposition 13 in Appendix C). Algorithm 1 applies this update to all simple pairs, one at a time.

Algorithm 1: Repeated two-column updates

Input : Set of simple pairs S, candidate solution X, objective matrix A, and row types $\{\mathrm{RT}_j\}_{j\in[n]}$.
Output: $\bar{X}$ such that $\bar{X}^\top\bar{X}$ is diagonal and $\mathrm{tr}(A^\top\bar{X}) \ge \mathrm{tr}(A^\top X)$ under the assumptions of Proposition 8.

$Y \leftarrow X$, $G \leftarrow X^\top X$;
for each $e = \{i, l\} \in S$, *in lexicographic order* **do**
  Choose any d assigned rows containing e: $J_e \subseteq \{j : e \subseteq \mathrm{RT}_j\}$ with $|J_e| = d$;
  $g \leftarrow G_{il}$;
  if $g \neq 0$ **then**
    $D_e \leftarrow |g| \begin{pmatrix} 1 & -\mathrm{sgn}(g) \\ -\mathrm{sgn}(g) & 1 \end{pmatrix}$, $Z \leftarrow Y_{J_e,e}$, $B \leftarrow Y_{J_e,[d]\setminus e}$;
    Compute $N \in \mathbb{R}^{d\times 2}$ with $N^\top N = I_2$ and $B^\top N = O$;
    $K \leftarrow N^\top Z$, $Z_0 \leftarrow Z - NK$, $\Lambda_e \leftarrow K^\top K + D_e$, $\Gamma_e \leftarrow N^\top A_{J_e,e}$;
    $K_{\mathrm{new}} \leftarrow$ an optimal solution of (2) with $\Lambda = \Lambda_e$ and $\Gamma = \Gamma_e$, e.g., the one in Proposition 13 or in Lemma 23;
    $Y_{J_e,e} \leftarrow Z_0 + NK_{\mathrm{new}}$, $E_e \leftarrow (\mathbf{e}_i, \mathbf{e}_l) \in \mathbb{R}^{d\times 2}$, $G \leftarrow G + E_e D_e E_e^\top$;
return Y;

The next proposition shows that Algorithm 1 is correct.

Proposition 8. *Let $H := \mathrm{CP}_d \setminus S$. Let X satisfy $X_i^\top X_l = 0$ for every $\{i,l\} \in H$, and let every $e \in S$ satisfy $|\{j : e \subseteq \mathrm{RT}_j\}| \geq d$. Let $\bar{X}$ be the output of Algorithm 1 with input $(S, X, A, \{\mathrm{RT}_j\}_{j\in[n]})$. Then $\bar{X}^\top \bar{X}$ is diagonal with $(\bar{X}^\top \bar{X})_{ii} = (X^\top X)_{ii} + \sum_{l\neq i} \left|(X^\top X)_{il}\right|$ for every $i \in [d]$. Moreover, $\bar{X}_{jp} = X_{jp}$ for every $j \in [n]$ and $p \notin \mathrm{RT}_j$, and $\mathrm{tr}(A^\top \bar{X}) \geq \mathrm{tr}(A^\top X)$.*

Proof idea of Proposition 8. Consider the iteration of a simple pair $e = \{i,l\}$ with $g \neq 0$, and assume that $g = Y_i^\top Y_l$. Since $B^\top$ has $d-2$ rows, its kernel has dimension at least 2, so the matrix N exists. The iteration replaces the component NK of Z by NK_{new} and keeps Z_0. Since $B^\top N = O$, the matrix $B^\top Z$ does not change. Hence the inner products of columns i and l with the other columns do not change. Since $N^\top N = I_2$ and $N^\top Z_0 = O$, the matrix $Z^\top Z$ changes from $Z_0^\top Z_0 + K^\top K$ to $Z_0^\top Z_0 + \Lambda_e = Z_0^\top Z_0 + K^\top K + D_e$. By the choice of D_e, columns i and l become orthogonal, and their squared norms increase by $|g|$. For the objective, note that $\Lambda_e \succeq K^\top K$. Thus, by Proposition 13 and Weyl's monotonicity theorem, the optimal value of (2) with $\Lambda = \Lambda_e$ is at least the optimal value with $\Lambda = K^\top K$. Since K is feasible for the latter problem, we obtain $\mathrm{tr}(\Gamma_e^\top K_{\mathrm{new}}) \geq \mathrm{tr}(\Gamma_e^\top K)$, and the objective does not decrease. Finally, the entry (i,l) of the Gram matrix changes only in the iteration of $\{i,l\}$, where it becomes zero. The entries of the hard pairs are zero and never change, so $\bar{X}^\top \bar{X}$ is diagonal. The full proof is in Appendix C.2. □

2.3 Fixing the hard variables

Fix a set H of hard pairs and a list of hard row types $\mathcal{L} = \{\!\{T^1, \ldots, T^b\}\!\}$. For each $r \in [b]$, let $U_r := \cup_{e\in H, e\subseteq T^r} e$ be the set of indices in T^r that belong to a hard pair contained in T^r. The subroutine FixHardVariables (Algorithm 4 in Appendix B) computes values p_{ri}, for $r \in [b]$ and $i \in U_r$, that satisfy the orthogonality constraints of the hard pairs, and we write $p_r := (p_{ri})_{i\in U_r}$. The ADP decides later which rows receive them. The number of these variables is $D := \sum_{r=1}^b |U_r|$. Since every T^r contains a pair of H, and every pair of H is contained in at most $d-1$ of the sets $T^1, \ldots, T^b$ (Algorithm 2), the argument of Lemma 12 gives $D \leq 2(d-1)|H| \leq D_d$. A solution of the orthogonality equations in $[-1,1]^D$, such as the zero solution, can be far from the hard variables of a *target solution*, that is, an optimal solution of (1) that the algorithm aims to approximate. We therefore cover $[-1,1]$ by the J_p closed intervals $\Omega_\ell := [-1+2(\ell-1)/J_p,\ -1+2\ell/J_p]$, $\ell \in [J_p]$, each of length $2/J_p \leq a$, where a and J_p are given in Table 1. For each box $Q = \Omega_{\ell_1} \times \cdots \times \Omega_{\ell_D}$, we decide whether Q contains a tuple p satisfying

$$\sum_{r:\{i,k\}\subseteq T^r} p_{ri}p_{rk} = 0, \quad \forall\, \{i,k\} \in H, \quad \sum_{r=1}^{b} \sum_{i\in U_r} p_{ri}^2 \leq d. \tag{3}$$

Every $X \in \mathcal{F}_s$ satisfies $\mathrm{tr}(X^\top X) = d$, so its hard variables satisfy the last inequality. The bounded algebraic sampling algorithm of Basu, Leroy, and Roy [5] solves this semialgebraic feasibility problem exactly (see Lemma 27). If p^* denotes the hard variables of a target solution, then p^* lies in some box Q^*, and FixHardVariables returns a solution p^{Q^*} in this box, with $\left\|p^{Q^*} - p^*\right\|_2 \leq \sqrt{D}a$.

2.4 Approximate dynamic programming

Fix a set H of hard pairs, the set $S := \mathrm{CP}_d \setminus H$ of simple pairs, a list $\mathcal{L} = \{\!\{T^1, \ldots, T^b\}\!\}$ of hard row types, and a vector p returned by FixHardVariables. We call such a choice of $(H, \mathcal{L}, p)$ a *branch* of LinearOracle. For each branch, the subroutine SelectBranch (Algorithm 7 in Appendix B) assigns row types and variable values to the rows of a candidate solution through an approximate dynamic programming (ADP) tree. It keeps the values p of the hard variables. Its output consists of candidate solutions that satisfy the sparsity constraints and the orthogonality constraints of the hard pairs, and whose error $\left\|X^\top X - I_d\right\|_F$ is small.

The ADP tree is a complete binary tree of height $H_{\mathrm{tree}} := \lceil \log_2 n \rceil$ with $N_{\mathrm{leaf}} := 2^{H_{\mathrm{tree}}}$ leaves. Leaves $1, \ldots, n$ correspond to the rows. The remaining leaves are dummy leaves, which correspond to no row and contribute zero (Algorithm 6 in Appendix B). A *partial solution* at a node u is a matrix X_u whose rows in the subtree rooted at u are assigned row types and values, and whose remaining rows are zero. Let $G_u := X_u^\top X_u$ and $v_u := \mathrm{tr}(A^\top X_u)$. The *sparsity capacity usage* of column i in X_u is the number $c_{u,i}$ of rows of X_u whose

assigned row type contains i. Thus $c_{u,i}$ counts the rows in which column i is allowed to be non-zero, and it can exceed the number of non-zero entries assigned to column i. The *key* $\mathrm{key}(X_u)$ of X_u consists of the following terms:

- $\mathcal{L}_{\mathrm{occ}}$: the set of labels $r \in [b]$ such that the hard row type $T^r \in \mathcal{L}$ is assigned to a row in the subtree of u.
- $\{\mu_e : e \in S\}$: $\mu_e := \min\{m_{u,e}, d\}$, where $m_{u,e}$ is the number of rows in the subtree of u whose assigned row type contains the simple pair e.
- $\hat{c}$: a d-vector of upper bounds on the sparsity capacity usages $c_{u,i}$, defined below.
- $\hat{G}$: a bin label of the Gram matrix G_u, defined below.

At node u, we store a *dictionary* $\mathcal{D}_u$ that maps $\mathrm{key}(X_u)$ to the pair (v_u, G_u). For every key, the dictionary keeps the pair with the largest objective value found, and ties are broken by a fixed enumeration rule (Algorithm 5 in Appendix B). A *state* at u is an entry of $\mathcal{D}_u$, that is, a key together with its pair. There are $\prod_{i=1}^{d}(s_i + 1)$ possible vectors of sparsity capacity usages, and infinitely many possible Gram matrices G_u. Inspired by Bertsimas and Demir [9], we round both to reduce the number of states.

Capacity caps. We limit the number of possible sparsity capacity usages by rounding them up. We choose constants $K_0 > 0$ and $\alpha > 0$, whose values are given in Table 1. We define a finite set U_c as follows:

$$U_c := \begin{cases} \{0, 1, \ldots, n\}, & \text{if } K_0 \geq n, \\ \{0, 1, \ldots, K_0\} \cup \{K_0(1+\alpha)^l : l \in [l_{\max}]\}, & \text{otherwise,} \end{cases} \tag{4}$$

where $l_{\max}$ is the smallest positive integer such that $K_0(1+\alpha)^{l_{\max}} > n$. At a leaf u and for $i \in [d]$, the *capacity cap* of column i is

$$\hat{c}_{u,i} = R_c(c_{u,i}) := \min\{w \in U_c : w \geq c_{u,i}\}. \tag{5}$$

At the merge of two nodes, we add the two caps and round the sum up with R_c. We reject the merge if the sum, or its rounded value, is larger than s_i for some $i \in [d]$. As a result, for every node u, we have $c_{u,i} \leq \hat{c}_{u,i}$, but in general $\hat{c}_{u,i} \neq R_c(c_{u,i})$.

Gram matrix bins. We first use $\mathrm{tr}(G_u)$ to determine a scale, and then discretize every entry of G_u at the scale. More precisely, we choose constants $\kappa_G > 0$ and $\theta > 0$ as in Table 1, and define the following intervals:

$$I_{-1} := [0, \kappa_G) \cap [0, 4d), \quad I_0 := [\kappa_G, (1+\theta)\kappa_G) \cap [0, 4d), \quad \ldots$$
$$I_k := [(1+\theta)^k \kappa_G, (1+\theta)^{k+1} \kappa_G) \cap [0, 4d), \quad \ldots$$

Only finitely many of these intervals are nonempty. For each partial solution X_u with $\mathrm{tr}(G_u) < 4d$, we define the *label mapping* $\mathrm{lab}(G_u)$ as follows. If $\mathrm{tr}(G_u) \in I_{-1}$, then $\mathrm{lab}(G_u) := (-1, O)$. If $\mathrm{tr}(G_u) \in I_k$ for some $k \geq 0$, then

$$\mathrm{lab}(G_u) := \left(k, \left(\left\lfloor \frac{(G_u)_{il}}{\theta \cdot \kappa_G (1+\theta)^k} \right\rfloor\right)_{1 \leq i,l \leq d}\right).$$

The *bin label* of X_u is $\hat{G}(X_u) := \mathrm{lab}(G_u)$. Adding rows does not decrease the trace of the Gram matrix. Hence partial solutions with trace at least $4d$ cannot be extended to approximately orthogonal candidate solutions in our algorithm, and we discard them. The dictionaries store the exact Gram matrices, and the bin labels are used only to identify states. For a fixed k, the number of possible labels for each Gram matrix entry is $\mathcal{O}(1/\theta)$, as detailed in the workload analysis in Appendix C.4.

Leaves. For a row assigned to some row type T, we call the variables indexed by T that are not hard variables *free variables*. Their values are chosen from the finite set of rational numbers

$$\mathcal{V} := \{0, \pm 1\} \cup \left\{ \pm 2^{-k} \left(1 + \frac{\ell}{M}\right) : 1 \le k \le J_0, 0 \le \ell < M \right\}, \tag{6}$$

where M and J_0 are given in Table 1. The subintervals of $[2^{-k}, 2^{-k+1}]$ defined by $\mathcal{V}$ have length $2^{-k}/M$, so the spacing of $\mathcal{V}$ is relative to the magnitude of its elements. As a result, each entry z of a target solution with $|z| \ge 2^{-J_0}$ can be approximated by some $\hat{z} \in \mathcal{V}$ with error at most $|z|/M$. Every other entry can be rounded to 0, with error at most 2^{-J_0}. For a row type T, let $U_H(T) := \cup_{e \in H, e \subseteq T} e$. We call a row type T *ordinary* if it contains no hard pair, that is, if $U_H(T) = \emptyset$. At a leaf $j \le n$, the ADP considers three choices for row j (Algorithm 6 in Appendix B): (a) The row is unused. (b) The row is assigned a nonempty ordinary row type T, and the variables indexed by T take values in $\mathcal{V}$. (c) The row is assigned a hard row type $T^r \in \mathcal{L}$. Its hard variables, indexed by $U_r = U_H(T^r)$, take the values p_r, and its free variables, indexed by $T^r \setminus U_r$, take values in $\mathcal{V}$. Since their values p_r may differ, two copies of the same row type in $\mathcal{L}$ with different labels are distinct choices.

Merging and acceptance. At each internal node, SelectBranch combines every pair of states of the two children whose sets $\mathcal{L}_{\text{occ}}$ are disjoint and whose summed capacity caps do not exceed s. It adds the Gram matrices and the objective values, and it takes the union of the sets $\mathcal{L}_{\text{occ}}$. It also adds the counters μ_e up to d, rounds the summed caps with R_c, and updates the dictionary. At the root, a state is *accepted* if $\mathcal{L}_{\text{occ}} = [b]$, $\mu_e = d$ for every $e \in S$, $v \ge 0$, and $\|G - I_d\|_F \le \delta$. To reconstruct the candidate solution of an accepted state, every dictionary entry stores its leaf assignment or back-pointers to the two child states from which it was constructed.

2.5 The linear oracle

Algorithm 2 enumerates all branches. For every accepted root state of SelectBranch, it reconstructs the candidate solution Y, applies ExactFinalize (Algorithm 8 in Appendix B), and keeps the best matrix found.

Algorithm 2: LINEARORACLE

Input : Rational $A \in \mathbb{Q}^{n\times d}$, sparsity constants s, $t \in \mathbb{Q} \cap (0,1)$.
Output: $X \in \mathcal{F}_s$ such that $\operatorname{tr}(A^\top X) \geq (1-t)\mathsf{OPT}_{\text{linear}}(A;s)$.

$X_{\text{best}} \leftarrow \left(\begin{smallmatrix} I_d \\ O \end{smallmatrix}\right)$;
if $A = O$ **then**
  return X_{best};
if $d = 1$ **then**
  $x_{\text{best}} \leftarrow$ the restriction of A to s_1 entries of largest absolute value, normalized;
  return x_{best};
for each $H \subseteq \mathrm{CP}_d$ **do**
  $S \leftarrow \mathrm{CP}_d \setminus H$;
  for $b = 0, 1, \ldots, \min\{n, (d-1)|H|\}$ **do**
    for each *multiset* $\mathcal{L} = \{\!\{T^1, T^2, \ldots, T^b\}\!\}$ *of subsets of* $[d]$ **do**
      if *every* T^i *contains at least one* $e \in H$*, and every* $e \in H$ *is contained in at most* $d-1$ *of the sets* $T^1, \ldots, T^b$ **then**
        $D \leftarrow \sum_{i=1}^{b} \left|\cup_{e\in H, e\subseteq T^i} e\right|$;
        for each $p \in$ FIXHARDVARIABLES$(H, \mathcal{L}, D, t)$ **do**
          $\mathcal{R} \leftarrow$ SELECTBRANCH$(A, s, H, \mathcal{L}, p, t)$;
          for each *accepted root state* $q \in \mathcal{R}$ **do**
            Reconstruct its candidate solution Y and assigned row types RT_j;
            $X \leftarrow$ EXACTFINALIZE$(A, Y, (\mathrm{RT}_j), S)$;
            if $\operatorname{tr}(A^\top X) > \operatorname{tr}(A^\top X_{\text{best}})$ **then**
              $X_{\text{best}} \leftarrow X$;
return X_{best};

The subroutine EXACTFINALIZE first applies Algorithm 1 to Y. Since $\mathcal{L}_{\text{occ}} = [b]$ at the root and the sets $\mathcal{L}_{\text{occ}}$ of the two children are disjoint at every merge, each hard row type in $\mathcal{L}$ is assigned to exactly one row of Y. Moreover, ordinary row types contain no hard pair. Hence (3) gives $Y_i^\top Y_l = 0$ for every $\{i,l\} \in H$. Since $\mu_e = d$ for every $e \in S$ at the root, Proposition 8 applies, and the resulting matrix $\widetilde{Y}$ has orthogonal columns $\widetilde{Y}_i$. EXACTFINALIZE then divides each $\widetilde{Y}_i$ by its norm, and flips the sign of every $\widetilde{Y}_i$ with $A_i^\top \widetilde{Y}_i < 0$. The result is in $\mathcal{F}_s$, and the sign flips do not decrease the objective.

3 Approximation guarantee and running time

We now state the approximation guarantee of Algorithm 3 and sketch its proof. The full proof is in Appendix C.3. Then we bound the running time.

Theorem 9. *Let* $R \in \mathbb{Q}^{n\times r}$*,* $s \in [n]^d$*, and* $\epsilon \in \mathbb{Q} \cap (0,1)$*. Then Algorithm 3 returns* $X \in \mathcal{F}_s$ *satisfying* $\left\|R^\top X\right\|_F^2 \geq (1-\epsilon)\mathsf{OPT}(R;s)$.

Proof idea of Theorem 9. The grid of matrices W in Algorithm 3 contains a matrix within Frobenius distance $\epsilon/4$ of $R^\top X^* / \left\|R^\top X^*\right\|_F$, where X^* is an optimal solution of SPCA. Hence it suffices to show that LINEARORACLE returns $X \in \mathcal{F}_s$ with $\operatorname{tr}(A^\top X) \geq (1-t)\mathsf{OPT}_{\text{linear}}(A;s)$. By Proposition 8 and the acceptance tests at the root, every output of LINEARORACLE is in $\mathcal{F}_s$. For the objective value, let X^* now be an optimal solution of (1), and consider the branch whose hard pairs H and list $\mathcal{L}$ are those of X^*. We construct a chain of matrices

$$X^* \longrightarrow Z \longrightarrow Z_0 \longrightarrow Y \longrightarrow \widetilde{Y} \longrightarrow X,$$

in which each step loses only a small part of the objective value or of the orthogonality.

Slack. The matrix Z is obtained from X^* by deleting rows, that is, by setting them to zero. We keep the at most P_d rows that contain hard pairs or that we select to cover each simple pair d times. For every column

i with $s_i > K_0$, we then delete rows of small weight until at most $(1-4\tau)s_i$ rows of Z have a row type that contains i. The weight of a row is the sum of its share of the objective value and its share of $\|X^*\|_F^2$. Since $s_i > K_0$ is large compared with P_d/τ, the deleted rows have total weight $\mathcal{O}(d\tau)$. Thus Z loses at most a fraction $\mathcal{O}(d\tau)$ of the objective value, and $\left\|Z^\top Z - I_d\right\|_F = \mathcal{O}(d^2\tau)$. The slack $4\tau s_i$ absorbs the rounding of the capacity caps. Indeed, each merge multiplies a cap by at most $1+\alpha$, and over the H_{tree} levels this gives a factor at most $e^{\tau/8}$, which is smaller than $1/(1-4\tau)$. For $s_i \le K_0$, the caps are never rounded.

Discretization. The matrix Z_0 is obtained from Z by two changes. The hard variables of Z are replaced by the exact solution p that FixHardVariables returns for the box containing them, and p is within distance $\sqrt{D}a$ of these values. Each free variable is rounded to $\mathcal{V}$, with relative error at most $1/M$, or with absolute error at most 2^{-J_0} for very small entries. The difference $\operatorname{tr}(A^\top(Z_0 - Z))$ only involves entries of A in the support of X^*. By Lemma 14, these entries have Frobenius norm at most $d\mathsf{OPT}_{\text{linear}}(A;s)$. We obtain $\operatorname{tr}(A^\top Z_0) \ge (1-\delta/4)\mathsf{OPT}_{\text{linear}}(A;s)$ and $\left\|Z_0^\top Z_0 - I_d\right\|_F \le \delta/4$ (Proposition 15). Moreover, the columns of Z_0 corresponding to hard pairs are exactly orthogonal, and every simple pair is contained in the row types of at least d rows of Z_0.

Approximate dynamic programming. The rows of Z_0 give a sequence of states from the leaves to the root. At every node, KeepBest may replace a state by another state with the same key and a larger objective value. The two Gram matrices then lie in the same bin, so they differ by at most $d\theta\operatorname{tr}(G_u)+2\kappa_G$ in Frobenius norm. The trace-relative bins keep these errors proportional to the traces, which sum to less than $2d$ over the at most N_{leaf} states of a level. Summing over the $H_{\text{tree}}+1 = \mathcal{O}(\log n)$ levels, the total change of the Gram matrix is less than $\delta/8$. Hence the root has an accepted candidate Y with $\operatorname{tr}(A^\top Y) \ge \operatorname{tr}(A^\top Z_0)$ and $\left\|Y^\top Y - I_d\right\|_F < \delta$ (Proposition 17). Since the occupied labels of the two children are disjoint at every merge, each hard row type is used exactly once in Y, so the columns of Y corresponding to hard pairs are exactly orthogonal.

Exact orthogonality. Algorithm 1 turns Y into a matrix $\widetilde{Y}$ with orthogonal columns and $\operatorname{tr}(A^\top\widetilde{Y}) \ge \operatorname{tr}(A^\top Y)$. By Proposition 8, the squared norm of column i of $\widetilde{Y}$ is $(Y^\top Y)_{ii} + \sum_{l\neq i}\left|(Y^\top Y)_{il}\right|$, which is at most $1+(1+\sqrt{d-1})\delta$. Normalizing the columns and flipping their signs as in ExactFinalize gives $X \in \mathcal{F}_s$ and keeps at least a factor $1-d\delta$ of the objective. Since $\delta = t/(8(d+1))$, we obtain $\operatorname{tr}(A^\top X) \ge (1-t)\mathsf{OPT}_{\text{linear}}(A;s)$ (Proposition 18). □

The next theorem bounds the bit complexity of Algorithm 3. Here, real algebraic numbers are represented exactly, as described in Appendix C.4.

Theorem 10. *Let $R \in \mathbb{Q}^{n\times r}$, $s \in [n]^d$, and $\epsilon \in \mathbb{Q}\cap(0,1)$. Let B_{SPCA} be the encoding size of (R,s,ϵ), and let $U_\epsilon := 1+\log_2(n+1)+\log_2(r+d+1)+\log_2(1/\epsilon)$. Then there is an implementation of Algorithm 3 whose bit complexity T_{SPCA} satisfies*

$$T_{\text{SPCA}} \le F_{\text{SPCA}}(r,d)n\epsilon^{-(rd+d^3-d^2+4d+2+c)}(B_{\text{SPCA}}+1)^c U_\epsilon^{d^2+5d+4+4c},$$

where $F_{\text{SPCA}}(r,d)$ depends only on r and d, and c is an absolute constant.

Since $\log_2(n+1)$ and $\log_2(1/\epsilon)$ are at most B_{SPCA}, this bound is polynomial in B_{SPCA} and in $1/\epsilon$ for fixed r and d. Together with Theorem 9, Theorem 10 proves Theorem 2. The proof of Theorem 10 is in Appendix C.4.

For fixed r, d, and ϵ, the number of steps of Algorithm 3 is linear in n up to a polylogarithmic factor. This dependence on n comes from the ADP tree, the trace-relative Gram matrix bins, and the slack in the capacity caps, whereas a straightforward DP with exact capacity counts and a uniform Gram matrix grid can have $\mathcal{O}(n^{d+d(d+1)/2})$ states per row. The additional factors in the bit complexity bound account for the accuracy-dependent enumeration and discretization, as well as the cost of exact algebraic sampling, field operations, and comparisons. In particular, the degrees of the algebraic field extensions are bounded in terms of d, and the bit complexity of these operations is polynomial in the input encoding size. Further details are given in Appendix C.4.

Funding. A. Del Pia is partially supported by AFOSR Grant FA9550-23-1-0433 and ONR Grant N00014-25-1-2490. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Air Force Office of Scientific Research or the Office of Naval Research.

AI methodology. AI tools (Claude, developed by Anthropic, and ChatGPT, developed by OpenAI) were used for proofreading and suggesting improvements to the exposition. The authors take full responsibility for the content of the paper.

References

[1] Arash A. Amini and Martin J. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. In *2008 IEEE International Symposium on Information Theory*, pages 2454–2458. IEEE, 2008.

[2] Megasthenis Asteris, Dimitris S. Papailiopoulos, and George N. Karystinos. The sparse principal component of a constant-rank matrix. *IEEE Transactions on Information Theory*, 60(4):2281–2290, 2014.

[3] Megasthenis Asteris, Dimitris S. Papailiopoulos, Anastasios Kyrillidis, and Alexandros G. Dimakis. Sparse PCA via bipartite matchings. arXiv:1508.00625, 2015. Full version of the NIPS 2015 paper.

[4] Saugata Basu. Algorithms in real algebraic geometry: a survey. In Jean-Philippe Monnier, editor, *Real Algebraic Geometry*, volume 51 of *Panoramas et Synthèses*, pages 107–153. Société Mathématique de France, 2017.

[5] Saugata Basu, Richard Leroy, and Marie-Françoise Roy. A bound on the minimum of a real positive polynomial over the standard simplex. arXiv:0902.3304, 2009.

[6] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer, 2nd edition, 2006.

[7] Konstantinos Benidis, Ying Sun, Prabhu Babu, and Daniel P. Palomar. Orthogonal sparse PCA and covariance estimation via Procrustes reformulation. *IEEE Transactions on Signal Processing*, 64(23):6211–6226, 2016. doi:10.1109/TSP.2016.2605073.

[8] Dimitris Bertsimas, Ryan Cory-Wright, and Jean Pauphilet. Solving large-scale sparse PCA to certifiable (near) optimality. *Journal of Machine Learning Research*, 23(13):1–35, 2022.

[9] Dimitris Bertsimas and Ramazan Demir. An approximate dynamic programming approach to multidimensional knapsack problems. *Management Science*, 48(4):550–565, 2002.

[10] Dimitris Bertsimas and Driss Lahlou Kitane. Sparse PCA: a geometric approach. *Journal of Machine Learning Research*, 24(32):1–33, 2023.

[11] Siu On Chan, Dimitris S. Papailiopoulos, and Aviad Rubinstein. On the approximability of sparse PCA. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir, editors, *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 623–646. PMLR, 2016.

[12] Difei Cheng and Qiao Hu. A decomposition framework for certifiably optimal orthogonal sparse PCA. arXiv:2603.01144, 2026.

[13] Agniva Chowdhury, Petros Drineas, David P. Woodruff, and Samson Zhou. Approximation algorithms for sparse principal component analysis. arXiv:2006.12748, 2020.

[14] Ryan Cory-Wright and Jean Pauphilet. Sparse PCA with multiple components. *Operations Research*, 74(5):2650–2671, 2026. doi:10.1287/opre.2023.0598.

[15] Alberto Del Pia. Sparse PCA on fixed-rank matrices. *Mathematical Programming, Series A*, 198(1):139–157, 2023. doi:10.1007/s10107-022-01769-9.

[16] Alberto Del Pia and Dekun Zhou. A randomized algorithm for sparse PCA based on the basic SDP relaxation. arXiv:2507.09148, 2025.

[17] Alberto Del Pia, Dekun Zhou, and Yinglun Zhu. Efficient sparse PCA via block-diagonalization. In *The Thirteenth International Conference on Learning Representations (ICLR 2025)*, 2025.

[18] Yash Deshpande and Andrea Montanari. Sparse PCA via covariance thresholding. *Journal of Machine Learning Research*, 17(141):1–41, 2016.

[19] Santanu S. Dey, Rahul Mazumder, and Guanyi Wang. Using ℓ_1-relaxation and integer programming to obtain dual bounds for sparse PCA. *Operations Research*, 70(3):1914–1932, 2022. doi:10.1287/opre.2021.2153.

[20] Santanu S. Dey, Marco Molinaro, and Guanyi Wang. Solving sparse principal component analysis with global support. *Mathematical Programming, Series A*, 199(1–2):421–459, 2023. doi:10.1007/s10107-022-01857-w.

[21] William W. Hager, Dzung T. Phan, and Jiajie Zhu. Projection algorithms for nonconvex minimization with application to sparse principal component analysis. *Journal of Global Optimization*, 65:657–676, 2016.

[22] Serge Lang. *Algebra*, volume 211 of *Graduate Texts in Mathematics*. Springer, 3rd edition, 2002.

[23] Arjen K. Lenstra, Hendrik W. Lenstra, Jr., and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982. doi:10.1007/BF01457454.

[24] Yongchun Li and Weijun Xie. Exact and approximation algorithms for sparse principal component analysis. *INFORMS Journal on Computing*, 37(3):582–602, 2025. doi:10.1287/ijoc.2022.0372.

[25] Lester W. Mackey. Deflation methods for sparse PCA. In *Advances in Neural Information Processing Systems 21*, volume 21, pages 1017–1024, 2009.

[26] Malik Magdon-Ismail. NP-hardness and inapproximability of sparse PCA. *Information Processing Letters*, 126:35–38, 2017. doi:10.1016/j.ipl.2017.05.008.

[27] Leon Mirsky. A trace inequality of John von Neumann. *Monatshefte für Mathematik*, 79(4):303–306, 1975.

[28] Dimitris S. Papailiopoulos, Alexandros G. Dimakis, and Stavros Korokythakis. Sparse PCA through low-rank approximations. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 747–755. PMLR, 2013.

[29] Marcell János Uray. Algebraic number fields and the LLL algorithm. *Journal of Symbolic Computation*, 121:102261, 2024.

[30] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press, 3rd edition, 2013.

[31] Vincent Q. Vu, Juhee Cho, Jing Lei, and Karl Rohe. Fantope projection and selection: A near-optimal convex relaxation of sparse PCA. In *Advances in Neural Information Processing Systems*, volume 26, 2013.

[32] Xiao-Tong Yuan and Tong Zhang. Truncated power method for sparse eigenvalue problems. *Journal of Machine Learning Research*, 14:899–925, 2013.

[33] Hui Zou, Trevor Hastie, and Robert Tibshirani. Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15(2):265–286, 2006.

A Additional related work

In this section, we discuss additional related work.

Algorithms for single-component SPCA have been studied extensively. Existing approaches include greedy and iterative algorithms [32, 21, 24], regression-based formulations [33], and mixed-integer and convex optimization [8, 19, 24]. Chan, Papailiopoulos, and Rubinstein [11] give a polynomial-time algorithm with approximation factor $\max\{1/\sqrt{k}, n^{-1/3}\}$, where k is the sparsity constant and n is the dimension. Other approximation algorithms exploit low-rank structure [28], semidefinite relaxations [13, 16], and block-diagonal approximations [17].

Several methods compute multiple sparse components of a covariance matrix C. A common approach is *deflation* [25]. It computes a single sparse component x of C, replaces C by a deflated matrix, such as $(I - xx^\top)C(I - xx^\top)$, and iterates. Deflation can be combined with any algorithm for single-component SPCA. However, the resulting components need not be orthogonal, and so they may not be feasible for SPCA. In contrast, Bertsimas and Kitane [10] compute all components at once, with orthogonality enforced, when the components share the same support or have disjoint supports. Their method approximates the problem by binary linear optimization, and they report better solutions than deflation. Cheng and Hu [12] enforce sparsity and orthogonality together through Gram–Schmidt orthogonalization. They combine this approach with branch-and-bound to compute near-optimal solutions, and they speed it up by a decomposition based on block-diagonal approximations of the covariance matrix. For row-sparse PCA, where all components share a common support, Dey, Molinaro, and Wang [20] propose convex relaxations that provably approximate the feasible region, and convex integer programs that give dual bounds.

Another line of research studies statistical guarantees for SPCA, including estimation and support recovery under spiked covariance models [1, 18]. In particular, Vu, Cho, Lei, and Rohe [31] propose a Fantope relaxation with statistical guarantees for sparse principal subspace estimation.

B Pseudocode

This appendix gives the pseudocode of our algorithm for SPCA and of the subroutines of LinearOracle (Algorithm 2), which are described in Section 2. Algorithm 3 is our algorithm for SPCA. Algorithm 4 computes the values of the hard variables (Section 2.3). Algorithm 5 updates a dictionary with the data of a new partial solution, and Algorithm 6 initializes the dictionary of a leaf of the ADP tree (Section 2.4). Algorithm 7 runs the ADP for one branch of LinearOracle. It omits the back-pointers used to reconstruct the candidate solutions of accepted states. For a state q, we write $\mathcal{L}_{\text{occ}}(q)$, $\mu_e(q)$, $\hat{c}(q)$, $G(q)$, and $v(q)$ for its components. Algorithm 8 turns an accepted candidate solution into a point of $\mathcal{F}_s$ (Section 2.5).

Algorithm 3: $(1-\epsilon)$-algorithm for SPCA

Input : $R \in \mathbb{Q}^{n\times r}$, sparsity constants $s \in [n]^d$, $\epsilon \in \mathbb{Q} \cap (0,1)$.
Output: $X \in \mathcal{F}_s$ such that $\|R^\top X\|_F^2 \geq (1-\epsilon)\mathsf{OPT}(R; s)$.

$X_{\text{best}} \leftarrow \left(\begin{smallmatrix} I_d \\ O \end{smallmatrix}\right)$;
if $R = O$ **then**
  return X_{best};
$J_W \leftarrow \lceil 4rd/\epsilon \rceil$; $t \leftarrow \epsilon/4$;
for each $W \in J_W^{-1}\mathbb{Z}^{r\times d}$ *with* $\|W\|_F \leq 1$ **do**
  $A \leftarrow RW$;
  $X \leftarrow \text{LinearOracle}(A, s, t)$;
  if $\|R^\top X\|_F^2 > \|R^\top X_{\text{best}}\|_F^2$ **then**
    $X_{\text{best}} \leftarrow X$;
return X_{best};

Algorithm 4: FixHardVariables

Input : Set of hard pairs H, list of hard row types $\mathcal{L} = \{\{T^1, T^2, \ldots, T^b\}\}$, number of hard variables D, parameter $t \in \mathbb{Q} \cap (0,1)$.

Output: A finite list $\mathcal{P}$ of value assignments for the hard variables.

$\mathcal{P} \leftarrow \emptyset$;
if $D = 0$ **then**
  return $\{\emptyset\}$;
Set a and J_p as in Table 1;
Cover $[-1,1]$ by J_p closed intervals $\Omega_1, \ldots, \Omega_{J_p}$ of length at most a, as in Section 2.3;
for each *box* $Q = \Omega_{\ell_1} \times \cdots \times \Omega_{\ell_D} \subseteq [-1,1]^D$ **do**
  Decide whether the quadratic system given by $p \in Q$ and (3) is feasible;
  if *the system is feasible* **then**
    Find one solution p^Q, add p^Q to $\mathcal{P}$;
return $\mathcal{P}$;

Algorithm 5: KeepBest

Input : A dictionary $\mathcal{D}$, the data $\mathcal{L}_{\mathrm{occ}}, (\mu_e), \hat{c}, G, v$ of a candidate, sparsity constants s, and parameter $t \in \mathbb{Q} \cap (0,1)$

Output: The updated dictionary $\mathcal{D}$.

if $\hat{c}_i > s_i$ *for some* $i \in [d]$ ***or*** $\mathrm{tr}(G) \geq 4d$ **then**
  return $\mathcal{D}$;
$\hat{G} \leftarrow \mathrm{lab}(G)$ as described in Section 2.4, key $\leftarrow (\mathcal{L}_{\mathrm{occ}}, (\mu_e), \hat{c}, \hat{G})$;
if key $\notin \mathcal{D}$ **then**
  $\mathcal{D}[\mathrm{key}] \leftarrow (v, G)$;
else if $v >$ *the objective value recorded in* $\mathcal{D}[\mathrm{key}]$ **then**
  $\mathcal{D}[\mathrm{key}] \leftarrow (v, G)$;
return $\mathcal{D}$;

Algorithm 6: LeafInit

Input : Leaf index j, objective matrix A with rows a_j, sparsity constant vector s, set of hard pairs H, set of simple pairs S, list of hard row types $\mathcal{L} = \{\{T^1, \dots, T^b\}\}$, fixed hard-variable values p, and parameter $t \in \mathbb{Q} \cap (0,1)$.

Output: The initialized dictionary $\mathcal{D}_j$ for leaf j.

$\mathcal{D}_j \leftarrow \emptyset$;
Form the data $q_{\text{unused}} \leftarrow (\emptyset, 0, 0, O, 0)$ of an unused row;
$\mathcal{D}_j \leftarrow \text{KeepBest}(\mathcal{D}_j, q_{\text{unused}}, s, t)$;
if $j \le n$ **then**
  for each *nonempty* $T \subseteq [d]$ *with* $U_H(T) = \emptyset$ **do**
    for each $z \in \mathcal{V}^{|T|}$ **do**
      $y_T \leftarrow z$, $y_{[d]\setminus T} \leftarrow 0$, $\mathcal{L}_{\text{occ}} \leftarrow \emptyset$, $\hat{c} \leftarrow \mathbb{1}_T$, $\mu_e \leftarrow \mathbb{1}_{\{e \subseteq T\}}$ for every $e \in S$;
      $G \leftarrow yy^\top$, $v \leftarrow a_j y$;
      Form the candidate data $q \leftarrow (\mathcal{L}_{\text{occ}}, \mu, \hat{c}, G, v)$;
      $\mathcal{D}_j \leftarrow \text{KeepBest}(\mathcal{D}_j, q, s, t)$;
  for $r = 1, \dots, b$ **do**
    $U_r \leftarrow U_H(T^r)$, $F_r \leftarrow T^r \setminus U_r$;
    for each $z \in \mathcal{V}^{|F_r|}$ **do**
      $y_{U_r} \leftarrow p_r$, $y_{F_r} \leftarrow z$, $y_{[d]\setminus T^r} \leftarrow 0$, $\mathcal{L}_{\text{occ}} \leftarrow \{r\}$, $\hat{c} \leftarrow \mathbb{1}_{T^r}$, $\mu_e \leftarrow \mathbb{1}_{\{e \subseteq T^r\}}$ for every $e \in S$;
      $G \leftarrow yy^\top$, $v \leftarrow a_j y$;
      Form the candidate data $q \leftarrow (\mathcal{L}_{\text{occ}}, \mu, \hat{c}, G, v)$;
      $\mathcal{D}_j \leftarrow \text{KeepBest}(\mathcal{D}_j, q, s, t)$;
return $\mathcal{D}_j$;

Algorithm 7: SelectBranch

Input : Objective matrix A, sparsity constant vector s, set of hard pairs H, list of hard row types $\mathcal{L} = \{\{T^1, \dots, T^b\}\}$, fixed hard-variable values p, and parameter $t \in \mathbb{Q} \cap (0,1)$.

Output: A set $\mathcal{R}$ of accepted root states.

Set δ as in Table 1, $S \leftarrow \text{CP}_d \setminus H$;
$\mathcal{D}_j \leftarrow \text{LeafInit}(j, A, s, H, S, \mathcal{L}, p, t)$ for every $j \in [N_{\text{leaf}}]$;
for *each internal node* u*, from bottom to top* **do**
  Let u_L and u_R be the two children of u;
  $\mathcal{D}_u \leftarrow \emptyset$;
  for each $q_L \in \mathcal{D}_{u_L}$ *and* $q_R \in \mathcal{D}_{u_R}$ **do**
    if $\mathcal{L}_{\text{occ}}(q_L) \cap \mathcal{L}_{\text{occ}}(q_R) = \emptyset$ **then**
      $c_i^{\text{sum}} \leftarrow \hat{c}_i(q_L) + \hat{c}_i(q_R)$ for every $i \in [d]$;
      if $c_i^{\text{sum}} \le s_i$, $\forall i \in [d]$ **then**
        $\hat{c}_i \leftarrow R_c(c_i^{\text{sum}})$ according to (5) for every $i \in [d]$;
        $\mathcal{L}_{\text{occ}} \leftarrow \mathcal{L}_{\text{occ}}(q_L) \cup \mathcal{L}_{\text{occ}}(q_R)$, $\mu_e \leftarrow \min\{d, \mu_e(q_L) + \mu_e(q_R)\}$ for every $e \in S$;
        $v \leftarrow v(q_L) + v(q_R)$, $G \leftarrow G(q_L) + G(q_R)$;
        Form the data $q \leftarrow (\mathcal{L}_{\text{occ}}, \mu, \hat{c}, G, v)$ of the parent candidate;
        $\mathcal{D}_u \leftarrow \text{KeepBest}(\mathcal{D}_u, q, s, t)$;
$\mathcal{R} \leftarrow \emptyset$, $\mathcal{D}_{\text{root}} \leftarrow$ the dictionary at the root;
for each *state* q *in* $\mathcal{D}_{\text{root}}$ **do**
  if $\mathcal{L}_{\text{occ}}(q) = [b]$ ***and*** $\mu_e(q) = d$, $\forall e \in S$ ***and*** $v(q) \ge 0$ ***and*** $\|G(q) - I_d\|_F^2 \le \delta^2$ **then**
    $\mathcal{R} \leftarrow \mathcal{R} \cup \{q\}$;
return $\mathcal{R}$;

Algorithm 8: ExactFinalize

Input : Objective matrix A with columns A_i, candidate solution Y of an accepted root state of SelectBranch, assigned row types $(\mathrm{RT}_j)_{j=1}^n$, and set of simple pairs S.

Output: A matrix in $\mathcal{F}_s$.

$G^0 \leftarrow Y^\top Y$, $\widetilde{Y} \leftarrow$ Output of Algorithm 1 with input $(S, Y, A, (\mathrm{RT}_j)_{j=1}^n)$;

Divide each $\widetilde{Y}_i$ by $\left\|\widetilde{Y}_i\right\|_2 = \sqrt{\|G_i^0\|_1}$, where G_i^0 is the i-th column of G^0;

For each $i \in [d]$, flip the sign of $\widetilde{Y}_i$ if $A_i^\top \widetilde{Y}_i < 0$;

return $\widetilde{Y}$;

C Deferred proofs

C.1 Hardness results

In this section, we prove Theorems 1 and 3.

Theorem 1. *Given $R \in \mathbb{Z}^{n\times 2}$, $k \in [n]$, and $\tau_0 \in \mathbb{Z}$, it is NP-hard to decide whether* $\mathsf{OPT}(R;(k,k)) \geq \tau_0$.

Proof of Theorem 1. We reduce Partition to the decision version of SPCA. An instance of Partition is a multiset of positive integers $a_1, a_2, \ldots, a_m$, and the question is whether it can be split into two subsets with equal sums.

Construction. Let $M := \sum_{i=1}^m a_i$, and define the vector $b \in \mathbb{Z}^{2m+1}$ as follows. The first m entries of b are $2a_1, 2a_2, \ldots, 2a_m$, the $(m+1)$-th and $(m+2)$-th entries of b are both $-M$, and the remaining $m-1$ entries are zero. Let 1_{2m+1} be the all-one vector in $\mathbb{R}^{2m+1}$, and note that $1_{2m+1}^\top b = 2M - M - M = 0$. Let $B := 8m\|b\|_2^2$, $n := 2m+2$, and

$$R := \begin{pmatrix} 0 & B \\ Bb & 1_{2m+1} \end{pmatrix} \in \mathbb{Z}^{n\times 2}, \qquad RR^\top = \begin{pmatrix} B^2 & B1_{2m+1}^\top \\ B1_{2m+1} & B^2bb^\top + 1_{2m+1}1_{2m+1}^\top \end{pmatrix}.$$

Finally, let $k := m+2$ and $\tau_0 := B^2(\|b\|_2^2 + 1) + m + 1$. The instance (R, k, τ_0) can be computed in time polynomial in the encoding size of $a_1, \ldots, a_m$.

Let r_1 and r_2 be the two columns of R. Since $1_{2m+1}^\top b = 0$, the columns r_1 and r_2 are orthogonal. Moreover, $\|r_1\|_2^2 = B^2\|b\|_2^2$ and $\|r_2\|_2^2 = B^2 + 2m + 1$, and therefore

$$\operatorname{tr}(RR^\top) = \|r_1\|_2^2 + \|r_2\|_2^2 = B^2\|b\|_2^2 + B^2 + 2m + 1 = \tau_0 + m.$$

We show that $\mathsf{OPT}(R;(k,k)) \geq \tau_0$ if and only if the Partition instance is a yes-instance.

From a partition to a solution. Suppose that $J \subseteq [m]$ satisfies $\sum_{j\in J} a_j = M/2$. Since the a_j are positive, $1 \leq |J| \leq m-1$. Since b has $m-1$ zero entries, there is a set F of $m - |J|$ indices of zero entries of b. Let $I := J \cup \{m+1\} \cup F$. Then $|I| = m+1$ and $\sum_{i\in I} b_i = 2(M/2) - M = 0$. Let

$$x := \frac{1}{\sqrt{B^2+m+1}} \begin{pmatrix} B \\ 1_I \end{pmatrix}, \qquad y := \frac{1}{\|b\|_2} \begin{pmatrix} 0 \\ b \end{pmatrix},$$

and $X := (x, y)$. Both vectors have unit norm, and $\|x\|_0 = \|y\|_0 = m+2 = k$. Moreover, $x^\top y = 1_I^\top b / (\|b\|_2 \sqrt{B^2+m+1}) = 0$. Hence $X \in \mathcal{F}_{(k,k)}$. Moreover, $r_1^\top x = 0$, $r_2^\top x = \sqrt{B^2+m+1}$, $r_1^\top y = B\|b\|_2$, and $r_2^\top y = 0$. Therefore,

$$\left\|R^\top X\right\|_F^2 = (B^2+m+1) + B^2\|b\|_2^2 = \tau_0.$$

From a solution to a partition. Suppose that $\mathsf{OPT}(R;(k,k)) \geq \tau_0$. Let $X^* = (x^*, y^*)$ be an optimal solution, and let $P := x^*(x^*)^\top + y^*(y^*)^\top$ be the orthogonal projector onto $\operatorname{Span}(\{x^*, y^*\})$. Let $H := r_1r_1^\top/\|r_1\|_2^2 + r_2r_2^\top/\|r_2\|_2^2$ be the orthogonal projector onto $\operatorname{Span}(\{r_1, r_2\})$.

(a) We first show that the entries of P and H are close to each other. Since $\operatorname{tr}(PRR^\top) = \mathsf{OPT}(R;(k,k)) \geq \tau_0$ and $\operatorname{tr}(RR^\top) = \tau_0 + m$, we have $\operatorname{tr}((I-P)RR^\top) \leq m$. Since $I - P$ is positive semidefinite and $\|r_1\|_2^2, \|r_2\|_2^2 \geq B^2$, we obtain

$$\operatorname{tr}((I-P)H) = \frac{r_1^\top(I-P)r_1}{\|r_1\|_2^2} + \frac{r_2^\top(I-P)r_2}{\|r_2\|_2^2} \leq \frac{1}{B^2}\operatorname{tr}((I-P)RR^\top) \leq \frac{m}{B^2}.$$

Since $\|P\|_F^2 = \|H\|_F^2 = \operatorname{tr}(H) = 2$, we have $\|P-H\|_F^2 = 4 - 2\operatorname{tr}(PH) = 2\operatorname{tr}((I-P)H) \leq 2m/B^2$. Since $B = 8m\|b\|_2^2$ and $\sqrt{2m} < 2m$, every $i, j \in [n]$ satisfy

$$|P_{ij} - H_{ij}| \leq \|P-H\|_F \leq \frac{\sqrt{2m}}{B} < \frac{1}{4\|b\|_2^2}. \tag{7}$$

(b) Next, we determine the structure of the supports of x^* and y^*. Let $2 \le i, j \le m+3$. Then b_{i-1} and b_{j-1} are non-zero integers, and $B^2 > 8\|b\|_2^2$. Hence

$$|H_{ij}| = \left| \frac{b_{i-1}b_{j-1}}{\|b\|_2^2} + \frac{1}{B^2+2m+1} \right| \ge \frac{1}{\|b\|_2^2} - \frac{1}{B^2} > \frac{7}{8\|b\|_2^2},$$

and (7) gives $P_{ij} \neq 0$. Moreover, since $B^2 > 2m+1$, we have $H_{11} = B^2/(B^2+2m+1) > 1/2$. Since $\|b\|_2^2 \ge 2$, (7) gives $P_{11} = (x_1^*)^2 + (y_1^*)^2 > 1/4$. In particular, $x_1^* \neq 0$ or $y_1^* \neq 0$. Both sparsity constants equal k, so after swapping the two columns of X^* if necessary, we may assume that $x_1^* \neq 0$. After flipping the sign of x^* if necessary, we may also assume that $x_1^* > 0$.

We show that $y_1^* = 0$. Suppose not. Since $x_1^* \neq 0$, $n = 2m+2$, and $\|x^*\|_0 \le k = m+2$, there exists an index $i \in \{2, \dots, m+3\}$ such that $x_i^* = 0$. Since $y_1^* \neq 0$ and $\|y^*\|_0 \le k$, there exists an index $j \in \{2, \dots, m+3\}$ such that $y_j^* = 0$. Then $P_{ij} = x_i^* x_j^* + y_i^* y_j^* = 0$, a contradiction. Therefore, $y_1^* = 0$, and so $P_{11} = (x_1^*)^2$. We write

$$x^* = \begin{pmatrix} \sqrt{P_{11}} \\ z \end{pmatrix}, \qquad y^* = \begin{pmatrix} 0 \\ v \end{pmatrix},$$

where $z, v \in \mathbb{R}^{2m+1}$ satisfy $|\operatorname{Supp}(z)| \le m+1$, $z^\top v = 0$, $\|z\|_2^2 = 1 - P_{11}$, and $\|v\|_2 = 1$.

(c) Finally, we extract a partition from $\operatorname{Supp}(z)$. We first show that $z \neq 0$. Suppose not. Then $x^* = \mathbf{e}_1$, thus $(x^*)^\top RR^\top x^* = B^2$. Since r_1 and r_2 are orthogonal, the two non-zero eigenvalues of $RR^\top$ are $\|r_1\|_2^2$ and $\|r_2\|_2^2$. Since $\|b\|_2^2 \ge 2$ and $B^2 > 2m+1$, the largest eigenvalue of $RR^\top$ is $\|r_1\|_2^2 = B^2\|b\|_2^2$. Hence

$$\operatorname{tr}(PRR^\top) \le B^2 + B^2\|b\|_2^2 < \tau_0,$$

a contradiction.

Let α be the inner product of $z/\|z\|_2$ and $b/\|b\|_2$. Since $z/\|z\|_2$ and v are orthonormal, the squared norm of the projection of $b/\|b\|_2$ onto $\operatorname{Span}(\{z, v\})$ is $\alpha^2 + (b^\top v)^2/\|b\|_2^2 \le 1$. Hence $(b^\top v)^2 \le \|b\|_2^2(1-\alpha^2)$. Similarly, since $b/\|b\|_2$ and $1_{2m+1}/\sqrt{2m+1}$ are orthonormal, we have $(v^\top 1_{2m+1})^2 \le (2m+1)(1 - (v^\top b)^2/\|b\|_2^2)$. Using $B^2\|b\|_2^2 > 2m+1$, we obtain

$$\begin{aligned}
(y^*)^\top RR^\top y^* &= B^2(v^\top b)^2 + (v^\top 1_{2m+1})^2 \\
&\le 2m+1 + (B^2\|b\|_2^2 - 2m - 1)\frac{(v^\top b)^2}{\|b\|_2^2} \\
&\le 2m+1 + (B^2\|b\|_2^2 - 2m - 1)(1-\alpha^2) \\
&= B^2\|b\|_2^2 - (B^2\|b\|_2^2 - 2m - 1)\alpha^2.
\end{aligned}$$

Moreover, since $(b^\top z)^2 = \alpha^2\|b\|_2^2\|z\|_2^2 = \alpha^2\|b\|_2^2(1-P_{11})$, we have

$$\begin{aligned}
(x^*)^\top RR^\top x^* &= (r_1^\top x^*)^2 + (r_2^\top x^*)^2 \\
&= B^2\|b\|_2^2(1-P_{11})\alpha^2 + (B\sqrt{P_{11}} + 1_{2m+1}^\top z)^2.
\end{aligned}$$

Adding the two bounds, we obtain

$$\mathsf{OPT}(R; (k,k)) \le B^2\|b\|_2^2 + (B\sqrt{P_{11}} + 1_{2m+1}^\top z)^2 - (B^2\|b\|_2^2 P_{11} - 2m - 1)\alpha^2. \tag{8}$$

Since $B = 8m\|b\|_2^2$ and $P_{11} > 1/4$, we have $B^2\|b\|_2^2 P_{11} - 2m - 1 > 0$. Let $I' := \operatorname{Supp}(z)$. By the Cauchy–Schwarz inequality and since $|I'| \le m+1$, we have

$$(B\sqrt{P_{11}} + 1_{2m+1}^\top z)^2 \le (B^2 + |I'|)\|x^*\|_2^2 \le B^2 + m + 1. \tag{9}$$

Together with (8) and $\mathsf{OPT}(R; (k,k)) \ge \tau_0$, this gives

$$\tau_0 \le \mathsf{OPT}(R; (k,k)) \le \tau_0 - (B^2\|b\|_2^2 P_{11} - 2m - 1)\alpha^2.$$

Therefore $\alpha = 0$, that is, $z^\top b = 0$, and both inequalities in (9) hold with equality. Equality in the second inequality gives $|I'| = m+1$. Equality in the Cauchy–Schwarz inequality shows that $(\sqrt{P_{11}}, z_{I'})$ is a multiple of $(B, 1, \dots, 1)$, so $z_{I'}$ is a non-zero multiple of the all-one vector. Since $z^\top b = 0$, we obtain $\sum_{i \in I'} b_i = 0$.

The set I' cannot contain both $m+1$ and $m+2$. Otherwise, the entries $2a_1, \dots, 2a_m$ indexed by I' would have to sum to $2M$, so $[m] \subseteq I'$ and $|I'| \geq m+2$, a contradiction. The set I' also contains at least one of $m+1$ and $m+2$. Otherwise, since $|I'| = m+1$ and b has only $m-1$ zero entries, we have $I' \cap [m] \neq \emptyset$, and thus $\sum_{i \in I'} b_i \geq \min_{i \in [m]} 2a_i > 0$, a contradiction. Thus I' contains exactly one of $m+1$ and $m+2$, and $\bar{I} := I' \cap [m]$ satisfies

$$\sum_{i \in \bar{I}} 2a_i = M,$$

which gives the desired partition. □

The next observation shows that Theorem 1 cannot be extended to the case $r = 1$.

Observation 11. *Let $r = 1$ and $R = \rho \in \mathbb{Q}^n$, and let $K := \min\{n, \sum_{i=1}^d s_i\}$. Then $\mathsf{OPT}(\rho; s)$ is the sum of the K largest values among $\rho_1^2, \dots, \rho_n^2$, and an optimal solution can be computed exactly in polynomial time.*

Proof. For a set $P \subseteq [n]$, denote by ρ_P the vector that agrees with ρ on P and is zero elsewhere. Let $X \in \mathcal{F}_s$, and let U be the union of the supports of the columns of X. Since the columns of X are orthonormal, $\left\|\rho^\top X\right\|_F^2 = \left\|XX^\top \rho\right\|_2^2$ is the squared norm of the projection of ρ onto the column space of X. This space is contained in $\{x \in \mathbb{R}^n : \mathrm{Supp}(x) \subseteq U\}$, so $\left\|\rho^\top X\right\|_F^2 \leq \|\rho_U\|_2^2$. Since $|U| \leq K$, the value $\|\rho_U\|_2^2$ is at most the sum of the K largest values ρ_j^2.

Conversely, let U^* be a set of K indices with the largest values ρ_j^2. Since $d \leq n$ and $s_i \geq 1$ for all $i \in [d]$, we have $d \leq K \leq \sum_{i=1}^d s_i$. Hence we can partition U^* into nonempty sets $P_1, \dots, P_d$ with $|P_i| \leq s_i$ for all $i \in [d]$. For every $i \in [d]$, let $X_i := \rho_{P_i} / \|\rho_{P_i}\|_2$ if $\rho_{P_i} \neq 0$, and let $X_i := \mathbf{e}_j$ for some $j \in P_i$ otherwise. The columns of X have disjoint supports, so $X \in \mathcal{F}_s$. Moreover, $\left\|\rho^\top X\right\|_F^2 = \sum_{i=1}^d (\rho^\top X_i)^2 = \sum_{i=1}^d \|\rho_{P_i}\|_2^2 = \|\rho_{U^*}\|_2^2$. The set U^* and the matrix X can be computed in polynomial time, where each column of X is represented as a rational vector divided by the square root of a rational number. □

Next, we prove Theorem 3. As in the proof of Theorem 1, we reduce from Partition.

Theorem 3. *Given $A \in \mathbb{Z}^{n \times 2}$, $k \in [n]$, and $\tau_0 \in \mathbb{Z}$, it is NP-hard to decide whether $\mathsf{OPT}_{\mathrm{linear}}(A; (k,k)) \geq \tau_0$.*

Proof of Theorem 3. We reduce Partition to the decision version of the linear problem (1). Let $a_1, a_2, \dots, a_m$ be an instance of Partition, and let $M := \sum_{i=1}^m a_i$. After doubling all a_i, we may assume that M is even.

Construction. Let $q := \lceil \sqrt{m+2} \rceil$, $k := q^2$, $n := m+k+1$, $Q := \sum_{i=1}^m a_i^2 + (M/2)^2$, and $L_0 := Q+1$. We define $A = (A_1, A_2) \in \mathbb{Z}^{n \times 2}$ as follows. The first $m+k$ entries of A_1 are one, and the last entry of A_1 is zero. The first m entries of A_2 are $2a_1, 2a_2, \dots, 2a_m$, the $(m+1)$-th entry of A_2 is $-M$, the last entry of A_2 is $Q-1$, and the remaining $k-1$ entries of A_2 are zero. Finally, let $\tau_0 := q + L_0$. The instance (A, k, τ_0) can be computed in time polynomial in the encoding size of $a_1, \dots, a_m$.

We have $\|A_2\|_2^2 = 4Q + (Q-1)^2 = (Q+1)^2 = L_0^2$. Let $(x, y) \in \mathcal{F}_{(k,k)}$. Since the last entry of A_1 is zero and $\|x\|_0 \leq k$, the Cauchy–Schwarz inequality gives

$$\begin{aligned} A_1^\top x &\leq \left\|(A_1)_{\mathrm{Supp}(x)}\right\|_2 \|x\|_2 = \sqrt{|\mathrm{Supp}(x) \setminus \{n\}|} \leq \sqrt{k} = q, \\ A_2^\top y &\leq \|A_2\|_2 \|y\|_2 = L_0, \end{aligned} \tag{10}$$

where the vector $(A_1)_{\mathrm{Supp}(x)}$ agrees with A_1 on $\mathrm{Supp}(x)$ and is zero elsewhere. Hence $\mathsf{OPT}_{\mathrm{linear}}(A; (k,k)) \leq \tau_0$. We show that equality holds if and only if the Partition instance is a yes-instance.

From a partition to a solution. Suppose that $I \subseteq [m]$ satisfies $\sum_{i \in I} a_i = M/2$. Let F be a set of $k-1-|I|$ indices in $\{m+2, \dots, m+k\}$, which index zero entries of A_2. Let $S := I \cup \{m+1\} \cup F$, so that $|S| = k = q^2$ and $S \subseteq [m+k]$. Let $x := \mathbb{1}_S / q$, $y := A_2 / L_0$, and $X := (x, y)$. Both vectors have unit norm, $\|x\|_0 = k$, and $\|y\|_0 = m+2 \leq k$. Moreover, $x^\top y = (\sum_{i \in I} 2a_i - M)/(qL_0) = 0$. Hence $X \in \mathcal{F}_{(k,k)}$, and

$$\mathrm{tr}(A^\top X) = A_1^\top x + A_2^\top y = \frac{k}{q} + \frac{\|A_2\|_2^2}{L_0} = q + L_0 = \tau_0.$$

From a solution to a partition. Suppose that $\mathsf{OPT}_{\text{linear}}(A;(k,k)) \geq \tau_0$, and let $X^* = (x^*, y^*) \in \mathcal{F}_{(k,k)}$ satisfy $\operatorname{tr}(A^\top X^*) \geq \tau_0$. Then all inequalities in (10) hold with equality. Thus $S := \operatorname{Supp}(x^*)$ has cardinality k, $n \notin S$, $x^* = \mathbb{1}_S/q$, and $y^* = A_2/L_0$. Since $(x^*)^\top y^* = 0$, we have $\sum_{i\in S}(A_2)_i = 0$. Since $|S| = k$, $S \subseteq [m+k]$, and A_2 has exactly $k-1$ zero entries in $[m+k]$, we have $S \cap [m+1] \neq \emptyset$. If $m+1 \notin S$, then $\sum_{i\in S}(A_2)_i \geq \min_{i\in[m]} 2a_i > 0$, a contradiction. Hence $m+1 \in S$, and $I := S \cap [m]$ satisfies

$$0 = \sum_{i\in S}(A_2)_i = \sum_{i\in I} 2a_i - M.$$

Thus $\sum_{i\in I} a_i = M/2$, which gives the desired partition. □

C.2 Proofs for Section 2

In this section, we prove the bounds on the hard rows and hard variables, the solution of (2), and Proposition 8.

Lemma 12. *Let X be a candidate solution with all its rows assigned row types. Then*

$$b(X) \leq \sum_{e\in H(X)} m_e(X) \leq (d-1)\,|H(X)|\,,$$
$$D(X) \leq 2\sum_{e\in H(X)} m_e(X) \leq 2(d-1)\,|H(X)|\,.$$

Proof. Every hard row contains a hard pair, so it is counted in $m_e(X)$ for some $e \in H(X)$. This gives $b(X) \leq \sum_{e\in H(X)} m_e(X)$. Each occurrence of a hard pair in the assigned row type of a row gives at most two hard variables in that row. This gives $D(X) \leq 2\sum_{e\in H(X)} m_e(X)$. Since $m_e(X) \leq d-1$ for every $e \in H(X)$, the remaining two inequalities hold. □

Proposition 13. *Let $\Gamma \in \mathbb{R}^{h\times k}$ with $k \leq h$, and let $\Lambda \in \mathbb{R}^{k\times k}$ be positive semidefinite. Then the optimal value of (2) is $\operatorname{tr}(\sqrt{\Gamma\Lambda\Gamma^\top})$. Moreover, an optimal solution is $Z^* = UV^\top\sqrt{\Lambda}$, where $\Gamma\sqrt{\Lambda} = U\Sigma V^\top$ with $U \in \mathbb{R}^{h\times k}$, $U^\top U = I_k$, $\Sigma \in \mathbb{R}^{k\times k}$ diagonal and non-negative, and $V \in \mathbb{R}^{k\times k}$ orthogonal, i.e., a thin singular value decomposition including zero singular values.*

Proof of Proposition 13. We first show that every $Z \in \mathbb{R}^{h\times k}$ with $Z^\top Z = \Lambda$ can be written as $Z = \Psi\sqrt{\Lambda}$, where $\Psi \in \mathbb{R}^{h\times k}$ satisfies $\Psi^\top\Psi = I_k$. Indeed, let $\rho := \operatorname{rank}(Z) \leq \min\{k,h\} = k$. By the singular value decomposition, there are orthogonal matrices $P_1 \in \mathbb{R}^{h\times h}$ and $P_2 \in \mathbb{R}^{k\times k}$ such that

$$Z = P_1 \begin{pmatrix} \sigma_1 & & & \\ & \ddots & & \\ & & \sigma_\rho & \\ & & & O \end{pmatrix} P_2 \quad \xrightarrow{Z^\top Z = \Lambda} \quad \Lambda = P_2^\top \begin{pmatrix} \sigma_1^2 & & & \\ & \ddots & & \\ & & \sigma_\rho^2 & \\ & & & O \end{pmatrix} P_2,$$

which implies that $Z = P_1 \begin{pmatrix} I_k \\ O \end{pmatrix} P_2\sqrt{\Lambda}$. Conversely, every Ψ with $\Psi^\top\Psi = I_k$ gives a feasible $Z = \Psi\sqrt{\Lambda}$. Therefore,

$$\begin{aligned}\max_{Z^\top Z=\Lambda} \operatorname{tr}(\Gamma^\top Z) &= \max_{\Psi^\top\Psi=I_k} \operatorname{tr}(\Gamma^\top\Psi\sqrt{\Lambda}) = \max_{\Psi^\top\Psi=I_k} \operatorname{tr}(\sqrt{\Lambda}\Gamma^\top\Psi) \\ &= \left\|\sqrt{\Lambda}\Gamma^\top\right\|_* = \operatorname{tr}(\sqrt{\Gamma\Lambda\Gamma^\top}),\end{aligned}$$

where $\|\cdot\|_*$ is the nuclear norm. The third equality follows from von Neumann's trace inequality [27], applied after padding $\sqrt{\Lambda}\Gamma^\top$ and Ψ with zeros to square matrices, which gives $\operatorname{tr}(\sqrt{\Lambda}\Gamma^\top\Psi) \leq \left\|\sqrt{\Lambda}\Gamma^\top\right\|_* \|\Psi\|_2 \leq \left\|\sqrt{\Lambda}\Gamma^\top\right\|_*$, and from the fact that $\Psi = UV^\top$ attains equality, where $\Gamma\sqrt{\Lambda} = U\Sigma V^\top$ is the thin singular value decomposition in the statement. Indeed, $\operatorname{tr}(\sqrt{\Lambda}\Gamma^\top UV^\top) = \operatorname{tr}(V\Sigma U^\top UV^\top) = \operatorname{tr}(\Sigma) = \left\|\Gamma\sqrt{\Lambda}\right\|_*$. Hence $Z^* = UV^\top\sqrt{\Lambda}$ is an optimal solution. □

Next, we show Proposition 8, which states the correctness of Algorithm 1.

Proposition 8. *Let $H := \mathrm{CP}_d \setminus S$. Let X satisfy $X_i^\top X_l = 0$ for every $\{i,l\} \in H$, and let every $e \in S$ satisfy $|\{j : e \subseteq \mathrm{RT}_j\}| \geq d$. Let $\bar{X}$ be the output of Algorithm 1 with input $(S, X, A, \{\mathrm{RT}_j\}_{j\in[n]})$. Then $\bar{X}^\top \bar{X}$ is diagonal with $(\bar{X}^\top \bar{X})_{ii} = (X^\top X)_{ii} + \sum_{l\neq i} |(X^\top X)_{il}|$ for every $i \in [d]$. Moreover, $\bar{X}_{jp} = X_{jp}$ for every $j \in [n]$ and $p \notin \mathrm{RT}_j$, and $\mathrm{tr}(A^\top \bar{X}) \geq \mathrm{tr}(A^\top X)$.*

Proof of Proposition 8. Fix $e = \{i,l\} \in S$, and let $\hat{Y}$ and $\bar{Y}$ be the values of Y in Algorithm 1 immediately before and after the iteration of the for-loop corresponding to e.

(1) In the first part of the proof, we assume that $G = \hat{Y}^\top \hat{Y}$ at the beginning of the iteration, so that $g = \hat{Y}_i^\top \hat{Y}_l$. We show that $\bar{Y}^\top \bar{Y} = \hat{Y}^\top \hat{Y} + E_e D_e E_e^\top$. In particular, $\bar{Y}_i^\top \bar{Y}_l = 0$, the entries $\bar{Y}_i^\top \bar{Y}_i$ and $\bar{Y}_l^\top \bar{Y}_l$ increase by $|g|$, and no other entry of the Gram matrix changes. Since line 10 updates G by the same matrix $E_e D_e E_e^\top$, the equality $G = Y^\top Y$ is preserved, and by induction it holds at the beginning of every iteration. We also show that $\bar{Y}$ differs from $\hat{Y}$ only in entries (j,p) with $j \in J_e$ and $p \in e$, and that $\mathrm{tr}(A^\top \bar{Y}) \geq \mathrm{tr}(A^\top \hat{Y})$. If $g = 0$, then $\bar{Y} = \hat{Y}$ and there is nothing to show.

Assume $g \neq 0$, and assume without loss of generality that $e = \{1,2\}$ and that J_e consists of the first d rows. Then, with Z and B as in Algorithm 1, the matrix Y has the block form

$$Y = \begin{pmatrix} Z & B \\ Y_{[n]\setminus J_e, e} & Y_{[n]\setminus J_e, [d]\setminus e} \end{pmatrix},$$

and therefore

$$Y^\top Y = \begin{pmatrix} Z^\top Z + Y_{[n]\setminus J_e,e}^\top Y_{[n]\setminus J_e,e} & Z^\top B + Y_{[n]\setminus J_e,e}^\top Y_{[n]\setminus J_e,[d]\setminus e} \\ B^\top Z + Y_{[n]\setminus J_e,[d]\setminus e}^\top Y_{[n]\setminus J_e,e} & B^\top B + Y_{[n]\setminus J_e,[d]\setminus e}^\top Y_{[n]\setminus J_e,[d]\setminus e} \end{pmatrix}.$$

The iteration changes only the block Z. Hence, to compute $\bar{Y}^\top \bar{Y}$, it suffices to compute the new values of $Z^\top Z$ and $B^\top Z$. Let $\hat{Z} := \hat{Y}_{J_e,e}$ and $\bar{Z} := \hat{Z} - NN^\top \hat{Z} + NK_{\mathrm{new}}$ be the initial and updated values of Z in the iteration.

We first show that, on line 7 of Algorithm 1, the matrix N exists, that is, $\ker(B^\top)$ has dimension at least 2. Indeed, $B^\top \in \mathbb{R}^{(d-2)\times d}$, so $\mathrm{rank}(B^\top) \leq d-2$, and thus $\ker(B^\top) \subseteq \mathbb{R}^d$ has dimension at least $d - (d-2) = 2$. This also covers the case $d = 2$, where B has no columns and $\ker(B^\top) = \mathbb{R}^2$.

Next, we compute the new value of $Z^\top Z$. As in Algorithm 1, let $K := N^\top \hat{Z}$ and $Z_0 := \hat{Z} - NK$. Since $N^\top N = I_2$, we have $N^\top Z_0 = O$, and therefore

$$\begin{aligned} \bar{Z}^\top \bar{Z} &= (Z_0 + NK_{\mathrm{new}})^\top (Z_0 + NK_{\mathrm{new}}) = Z_0^\top Z_0 + K_{\mathrm{new}}^\top N^\top N K_{\mathrm{new}} \\ &= Z_0^\top Z_0 + K_{\mathrm{new}}^\top K_{\mathrm{new}} = Z_0^\top Z_0 + \Lambda_e \\ &= Z_0^\top Z_0 + K^\top K + D_e = \hat{Z}^\top \hat{Z} + D_e, \end{aligned}$$

where the last equality follows from $\hat{Z} = Z_0 + NK$ and $N^\top Z_0 = O$. Since the rows outside J_e do not change, we obtain $\bar{Y}_i^\top \bar{Y}_l = \hat{Y}_i^\top \hat{Y}_l + (D_e)_{12} = g - |g| \operatorname{sgn}(g) = 0$, and the diagonal entries $\bar{Y}_i^\top \bar{Y}_i$ and $\bar{Y}_l^\top \bar{Y}_l$ increase by $(D_e)_{11} = (D_e)_{22} = |g|$.

Then, we compute the new value of $B^\top Z$. By the choice of N, we have $B^\top N = O$, and therefore

$$B^\top \bar{Z} = B^\top (Z_0 + NK_{\mathrm{new}}) = B^\top Z_0 = B^\top \hat{Z}.$$

Therefore, except for the entries indexed by $e \times e$, the Gram matrix does not change, and $\bar{Y}^\top \bar{Y} = \hat{Y}^\top \hat{Y} + E_e D_e E_e^\top$. Moreover, only the entries of $Y_{J_e,e}$ change.

Finally, we show that $\mathrm{tr}(A^\top \bar{Y}) \geq \mathrm{tr}(A^\top \hat{Y})$. Since only Z changes, it suffices to compare the values of $\mathrm{tr}(A_{J_e,e}^\top Z)$. We claim that

$$\mathrm{tr}(A_{J_e,e}^\top NK_{\mathrm{new}}) \geq \mathrm{tr}(A_{J_e,e}^\top NK).$$

Indeed, by Proposition 13,

$$\mathrm{tr}(A_{J_e,e}^\top NK_{\mathrm{new}}) = \mathrm{tr}(\sqrt{N^\top A_{J_e,e}(K^\top K + D_e)A_{J_e,e}^\top N})$$

$$\geq \operatorname{tr}(\sqrt{N^\top A_{J_e,e} K^\top K A_{J_e,e}^\top N})$$
$$\geq \operatorname{tr}(A_{J_e,e}^\top N K),$$

where the equality holds by Proposition 13 with $\Gamma = \Gamma_e = N^\top A_{J_e,e}$ and $\Lambda = \Lambda_e$. The proposition applies because D_e has eigenvalues 0 and $2\,|g|$, and thus $\Lambda_e = K^\top K + D_e \succeq 0$. The first inequality follows from

$$N^\top A_{J_e,e}(K^\top K + D_e) A_{J_e,e}^\top N \succeq N^\top A_{J_e,e} K^\top K A_{J_e,e}^\top N$$

and from Weyl's monotonicity theorem: if $X \succeq Y \succeq 0$, then $\lambda_k(X) \geq \lambda_k(Y)$ for every k, and hence $\operatorname{tr}(\sqrt{X}) \geq \operatorname{tr}(\sqrt{Y})$. The second inequality follows from Proposition 13 with $\Lambda = K^\top K$, since K is feasible for (2) with this Λ. Therefore,

$$\operatorname{tr}(A_{J_e,e}^\top \bar{Z}) = \operatorname{tr}(A_{J_e,e}^\top (Z_0 + N K_{\text{new}})) \geq \operatorname{tr}(A_{J_e,e}^\top (Z_0 + NK)) = \operatorname{tr}(A_{J_e,e}^\top \hat{Z}).$$

(2) In the second part of the proof, we prove the statement. Initially $G = X^\top X$, so by part (1), $G = Y^\top Y$ at the beginning of every iteration. By part (1), the entry (i,l) of the Gram matrix changes only in the iteration of $\{i,l\}$, where it becomes zero. Hence the value g used in the iteration of $\{i,l\}$ equals $(X^\top X)_{il}$, and this iteration adds $\left|(X^\top X)_{il}\right|$ to the diagonal entries (i,i) and (l,l). At the end of Algorithm 1, the entry (i,l) is zero for every $\{i,l\} \in S$. It is also zero for every $\{i,l\} \in H$, since $(X^\top X)_{il} = 0$ and such entries never change. Thus $\bar{X}^\top \bar{X}$ is diagonal, and $(\bar{X}^\top \bar{X})_{ii} = (X^\top X)_{ii} + \sum_{l \neq i} \left|(X^\top X)_{il}\right|$, where the terms with $\{i,l\} \in H$ are zero. Each iteration changes only entries (j,p) with $j \in J_e$ and $p \in e \subseteq \mathrm{RT}_j$, so $\bar{X}_{jp} = X_{jp}$ whenever $p \notin \mathrm{RT}_j$. Finally, the objective value does not decrease in any iteration, so $\operatorname{tr}(A^\top \bar{X}) \geq \operatorname{tr}(A^\top X)$. □

C.3 Proof of Theorem 9

In this section, we prove Theorem 9. We first prove the approximation guarantee of LinearOracle (Algorithm 2). Afterwards, we use it to prove the guarantee for Algorithm 3, our main approximation algorithm for SPCA.

We first show that every output of LinearOracle is in $\mathcal{F}_s$. This is clear for the initial matrix X_{best} and for the output in the case $d = 1$. Consider an accepted root state of SelectBranch, with candidate solution Y and assigned row types $(\mathrm{RT}_j)_{j \in [n]}$. At the root, $\mathcal{L}_{\text{occ}} = [b]$, and the occupied labels of the two children are disjoint at every merge. Hence every hard row type $T^r \in \mathcal{L}$ is assigned to exactly one row, whose hard variables take the values p_r. Moreover, ordinary row types contain no pair of H. Therefore, for every $\{i,l\} \in H$, the only rows that contribute to $Y_i^\top Y_l$ are hard rows whose row type contains $\{i,l\}$, and (3) gives $Y_i^\top Y_l = \sum_{r:\{i,l\} \subseteq T^r} p_{ri} p_{rl} = 0$. At the root, we also have $\mu_e = d$ for every $e \in S$, so every simple pair is contained in at least d assigned row types. Thus, Proposition 8 applies to Y. By this proposition, the output $\widetilde{Y}$ of Algorithm 1 has orthogonal columns, and $\widetilde{Y}_{ji} = Y_{ji} = 0$ whenever $i \notin \mathrm{RT}_j$, since LeafInit sets the entries outside the row type to zero. The number of rows whose assigned row type contains i is at most $\hat{c}_i \leq s_i$, so $\left\|\widetilde{Y}_i\right\|_0 \leq s_i$. Moreover, $\left\|\widetilde{Y}_i\right\|_2^2 \geq G_{ii}^0 \geq 1 - \delta > 0$ by the root filter of SelectBranch, so the normalization in ExactFinalize (Algorithm 8) is well defined. The normalization and the sign flips change neither the supports nor the orthogonality of the columns. Therefore, the output of ExactFinalize, and hence of LinearOracle, is in $\mathcal{F}_s$. It remains to show that the output of LinearOracle has a sufficiently large objective value.

High-level idea. Let X^* be an optimal solution of (1). We show that there exists a matrix Z_0 that is close to X^* and whose branch of LinearOracle yields an accepted root candidate in SelectBranch (Algorithm 7). The output of LinearOracle is at least as good as the matrix that ExactFinalize obtains from this candidate, and this gives the approximation guarantee. We use the chain of matrices $X^* \to Z \to Z_0 \to Y \to \widetilde{Y} \to X$ described in the proof idea of Theorem 9 in Section 3.

Notation. If $A = O$, every matrix in $\mathcal{F}_s$ is optimal, and LinearOracle returns $X_{\text{best}} \in \mathcal{F}_s$. If $d = 1$, LinearOracle returns an optimal solution before it enumerates any branch. Thus, we assume that $A \neq O$ and $d \geq 2$. Then $\mathsf{OPT}_{\text{linear}}(A; s) > 0$. Indeed, if $A_{ji} \neq 0$, let $X_i := \operatorname{sign}(A_{ji}) \mathbf{e}_j$, and let the other columns of

X be distinct standard basis vectors $\mathbf{e}_k$ with $k \neq j$, multiplied by ± 1 so that $A_l^\top X_l \geq 0$. Then $X \in \mathcal{F}_s$ and $\operatorname{tr}(A^\top X) \geq |A_{ji}| > 0$. For the j-th row x_j^* of X^*, we define its *objective contribution* $q_j := a_j (x_j^*)^\top$, where a_j is the j-th row of A. Flipping the sign of a row x_j^* changes neither $(X^*)^\top X^*$ nor the supports, and it replaces q_j by $-q_j$. Hence the optimality of X^* gives $q_j \geq 0$ for every $j \in [n]$.

We next define the analogs for X^* of the objects in Section 2.1. For every row x_j^* of X^*, let $\mathrm{RT}_j^* := \operatorname{Supp}(x_j^*)$ be its row type. For a pair $e \in \mathrm{CP}_d$, let $m_e^* := \left|\{j : e \subseteq \mathrm{RT}_j^*\}\right|$. Let $H^* := \{e : m_e^* < d\}$ and $S^* := \mathrm{CP}_d \setminus H^*$ be the sets of hard and simple pairs, respectively.

The rows whose row types contain at least one pair in H^* give the list of hard row types $\mathcal{L}^* := \{\!\{(T^*)^1, (T^*)^2, \ldots, (T^*)^{b^*}\}\!\}$. For a hard row type $(T^*)^r$, the variables corresponding to $U_{H^*}((T^*)^r) = \cup_{e \in H^*, e \subseteq (T^*)^r} e$ are the hard variables, which are fixed by FixHardVariables. By Lemma 12, $b^* \leq \min\{n, (d-1)|H^*|\}$, and $D^* := \sum_{r=1}^{b^*} |U_{H^*}((T^*)^r)| \leq d(d-1)^2$. Moreover, every element of $\mathcal{L}^*$ contains a pair of H^*, and every $e \in H^*$ is contained in $m_e^* \leq d-1$ elements of $\mathcal{L}^*$. Hence LinearOracle enumerates the branch with $H = H^*$ and $\mathcal{L} = \mathcal{L}^*$. We label the hard rows of X^* so that the row with label r has the row type T^r of the multiset $\mathcal{L}$ enumerated by LinearOracle, and all labels below refer to this labeling.

Before constructing Z_0, we need a technical lemma.

Lemma 14. *Let $X^* \in \mathcal{F}_s$, let $\mathcal{I} := \{(j,i) : X_{ji}^* \neq 0\}$, and let $\bar{A} \in \mathbb{R}^{n \times d}$ be the matrix with $\bar{A}_{ji} = A_{ji}$ for $(j,i) \in \mathcal{I}$ and $\bar{A}_{ji} = 0$ otherwise. Then, $\left\|\bar{A}\right\|_F \leq d\mathsf{OPT}_{\mathrm{linear}}(A;s)$.*

Proof. Fix a column $i \in [d]$, and let $S_i^* := \operatorname{Supp}(X_i^*)$. Then $\left\|\bar{A}_i\right\|_2^2 = \sum_{j \in S_i^*} A_{ji}^2$. If $|S_i^*| \leq n - d + 1$, let $Q_i := S_i^*$. Otherwise, let $Q_i \subseteq S_i^*$ be a set of $n - d + 1$ indices j with the largest values A_{ji}^2. Since $|S_i^*| \leq n$ and $n \geq d$, we have

$$\sum_{j \in Q_i} A_{ji}^2 \geq \min\left\{\frac{n-d+1}{|S_i^*|}, 1\right\} \sum_{j \in S_i^*} A_{ji}^2 \geq \frac{n-d+1}{n} \sum_{j \in S_i^*} A_{ji}^2 \geq \frac{1}{d}\left\|\bar{A}_i\right\|_2^2. \tag{11}$$

We next show that $\mathsf{OPT}_{\mathrm{linear}}(A;s) \geq \left\|\bar{A}_i\right\|_2 / \sqrt{d}$. We may assume that $\bar{A}_i \neq 0$. Let $\tilde{a} \in \mathbb{R}^n$ agree with A_i on Q_i and be zero elsewhere. By (11), we have $\|\tilde{a}\|_2 \geq \left\|\bar{A}_i\right\|_2 / \sqrt{d} > 0$. Let $X_i := \tilde{a} / \|\tilde{a}\|_2$. Since $|Q_i| \leq n - d + 1$, there are at least $d-1$ indices outside Q_i. For $l \neq i$, let X_l be $\pm \mathbf{e}_j$ for distinct indices $j \notin Q_i$, with the sign chosen so that $A_l^\top X_l \geq 0$. Then $X^\top X = I_d$, $\|X_i\|_0 \leq |Q_i| \leq |S_i^*| \leq s_i$, and $\|X_l\|_0 = 1 \leq s_l$ for $l \neq i$. Hence $X \in \mathcal{F}_s$, and

$$\mathsf{OPT}_{\mathrm{linear}}(A;s) \geq \operatorname{tr}(A^\top X) = \|\tilde{a}\|_2 + \sum_{l \neq i} A_l^\top X_l \geq \|\tilde{a}\|_2 \geq \frac{1}{\sqrt{d}}\left\|\bar{A}_i\right\|_2.$$

Squaring and summing over $i \in [d]$, we obtain

$$d^2 \mathsf{OPT}_{\mathrm{linear}}(A;s)^2 \geq \sum_{i=1}^{d} \left\|\bar{A}_i\right\|_2^2 = \left\|\bar{A}\right\|_F^2.$$

Since $\mathsf{OPT}_{\mathrm{linear}}(A;s) \geq 0$, this gives $\left\|\bar{A}\right\|_F \leq d\mathsf{OPT}_{\mathrm{linear}}(A;s)$. □

In the next proposition, we show that there exists a matrix Z_0 that is close to X^* and passes all filters in SelectBranch, unless a better candidate replaces it.

Proposition 15. *For $H = H^*$ and $\mathcal{L} = \mathcal{L}^*$ in LinearOracle, there exist a vector p returned by FixHardVariables and a matrix Z_0 whose rows are among the enumeration of SelectBranch for the branch $(H^*, \mathcal{L}^*, p)$ such that*

$$\operatorname{tr}(A^\top Z_0) \geq \left(1 - \frac{\delta}{4}\right) \mathsf{OPT}_{\mathrm{linear}}(A;s), \tag{12}$$

$$\left\|Z_0^\top Z_0 - I_d\right\|_F \leq \frac{\delta}{4}. \tag{13}$$

Moreover: (i) $(Z_0)_i^\top (Z_0)_l = 0$ for every $\{i,l\} \in H^*$. (ii) Every labeled hard row type in $\mathcal{L}^*$ is assigned exactly once to a row of Z_0. (iii) Every simple pair $e \in S^*$ is contained in the assigned row types of at least d rows of Z_0. (iv) At every node of the ADP tree, the capacity caps computed along the rows of Z_0, and the sums of the caps of the two children, are at most s componentwise.

Proof. We first construct a matrix Z from X^*. We keep every hard row of X^* and, for each simple pair $e \in S^*$, any d rows whose row types contain e. The number of kept rows is at most $(d-1)\,|H^*| + d\,|S^*| \le P_d = d\binom{d}{2}$.

Next, we delete a small number of the remaining rows. We define the *weight* of each row $j \in [n]$ by $w_j := q_j / \operatorname{tr}(A^\top X^*) + \left\|x_j^*\right\|_2^2 / d$. Since $\sum_{j\in[n]} q_j = \operatorname{tr}(A^\top X^*)$ and $\|X^*\|_F^2 = d$, we have $\sum_{j\in[n]} w_j = 2$. Using the parameter τ in Table 1, we delete rows of small weight, column by column. For each column i, we proceed as follows.

- If $s_i \le K_0 = \lceil 16(P_d+1)/\tau \rceil$, we delete no row for this column.
- Otherwise, we define $t_i := \lfloor (1-4\tau)s_i \rfloor$, and compute the current sparsity capacity usage of this column c_i. If $c_i > t_i$, we delete the $c_i - t_i$ lowest-weight rows j that are currently undeleted and unkept and satisfy $i \in \mathrm{RT}_j^*$ (or equivalently, $j \in \mathrm{Supp}(X_i^*)$).

Deleting a row means setting the entire row to zero, not just its entry in column i, and replacing its assigned row type by $\emptyset$. We denote the resulting matrix by Z.

We claim that the cumulative weight of the removed rows never exceeds $10d\tau$. Indeed, consider a column i with $s_i > K_0$ and $c_i > t_i$. By the definitions of K_0 and t_i, we have $P_d + 1 < \tau s_i/16$, and $t_i \ge (1-4\tau)s_i - 1$. Since $\tau = \delta/(128d^2) < 1/32$, we have

$$c_i - P_d \ge t_i + 1 - P_d > (1-4\tau)s_i - P_d > \left(1 - 4\tau - \frac{\tau}{16}\right) s_i > 0.$$

Let C_i be the set of rows j that are currently undeleted and unkept and satisfy $i \in \mathrm{RT}_j^*$. Then $|C_i| \ge c_i - P_d$, and $c_i - t_i < c_i - P_d$ because $t_i \ge (1-4\tau)s_i - 1 > \tau s_i/16 - 1 > P_d$. Hence C_i contains enough rows to delete. The weights are non-negative, and we delete the $c_i - t_i$ rows of C_i of lowest weight. Therefore, the deleted weight is at most

$$\frac{c_i - t_i}{|C_i|} \sum_{j\in C_i} w_j \le 2 \cdot \frac{c_i - t_i}{c_i - P_d} \le 2 \cdot \frac{s_i - t_i}{s_i - P_d} \le 2 \cdot \frac{4\tau s_i + 1}{s_i - P_d} < 10\tau,$$

where the second inequality follows from $t_i > P_d$ and $c_i \le s_i$, the third inequality follows from $t_i \ge (1-4\tau)s_i - 1$, and the last inequality is equivalent to $1 + 5\tau P_d < \tau s_i$, which follows from $\tau s_i > 16(P_d+1)$. Since there are at most d columns, the removed rows have a total weight of at most $10d\tau$.

Since all hard rows and the selected d rows for each simple pair in S^* are kept, the deletion affects neither the hard constraints (3) nor the two-column updates of Algorithm 1. Since the deleted weight is at most $10d\tau$, the deleted rows satisfy $\sum_j q_j \le 10d\tau \operatorname{tr}(A^\top X^*)$ and $\sum_j \left\|x_j^*\right\|_2^2 \le 10d^2\tau$. Moreover, $Z^\top Z - I_d = -\sum_j (x_j^*)^\top x_j^*$, where both sums range over the deleted rows. Hence Z satisfies

$$\operatorname{tr}(A^\top Z) \ge (1 - 10d\tau)\mathsf{OPT}_{\mathrm{linear}}(A;s), \qquad \left\|Z^\top Z - I_d\right\|_F \le 10d^2\tau.$$

Since $\tau = \delta/(128d^2)$, we obtain

$$\operatorname{tr}(A^\top Z) \ge \left(1 - \frac{5\delta}{64d}\right) \mathsf{OPT}_{\mathrm{linear}}(A;s), \qquad \left\|Z^\top Z - I_d\right\|_F \le \frac{5\delta}{64}. \tag{14}$$

Next, let p^* be the values of the hard variables in Z. Then $p^* \in [-1,1]^{D^*}$, so p^* belongs to one of the boxes Q^* enumerated by FixHardVariables. Moreover, $\sum_{r,i}(p_{ri}^*)^2 \le \|X^*\|_F^2 = d$, and p^* satisfies the orthogonality constraints of the hard pairs in (3) exactly. Indeed, for $\{i,k\} \in H^*$, every row j with $X_{ji}^* X_{jk}^* \ne 0$ satisfies $\{i,k\} \subseteq \mathrm{RT}_j^*$, so it is a kept hard row and both entries are hard variables. Hence FixHardVariables returns some $p \in Q^*$ that satisfies (3). Let Z^H be the matrix obtained from Z by

replacing only the hard variables by the values in p, while leaving every free variable unchanged. Since every side of the box Q^* has length at most a, we have

$$\left\|Z^H - Z\right\|_F = \|p - p^*\|_2 \le \sqrt{D^*}a \le \sqrt{D_d}a, \tag{15}$$

where $D_d = d(d-1)^2$ is given in Table 1. Let $\kappa_v := 2^{-J_0}$. By the construction of $\mathcal{V}$, for every free variable $z \in [-1,1]$, we can choose $\hat{z} \in \mathcal{V}$ such that

$$|z - \hat{z}| \le \begin{cases} |z|/M, & \text{if } |z| \ge \kappa_v, \\ \kappa_v, & \text{otherwise.} \end{cases}$$

Indeed, if $|z| = 1$, we choose $\hat{z} = z$. If $\kappa_v \le |z| < 1$, we choose $k_z \in [J_0]$ such that $2^{-k_z} \le |z| < 2^{-k_z+1}$, let $\ell_z := \lfloor M(2^{k_z}|z| - 1)\rfloor$, and let $\hat{z} := \operatorname{sign}(z)2^{-k_z}(1 + \ell_z/M)$. Then $|z - \hat{z}| \le 2^{-k_z}/M \le |z|/M$. Otherwise, we choose $\hat{z} = 0$. We round every free variable of Z^H in this way to the finite set $\mathcal{V}$ defined in (6), and we denote the resulting matrix by Z_0.

We now show (13). Let F be the set of indices of the free variables, let $F_{\text{large}} := \{(j,i) \in F : \left|Z^H_{ji}\right| \ge \kappa_v\}$, and let $F_{\text{small}} := F \setminus F_{\text{large}}$. By the triangle inequality, and since $\sum_{(j,i)\in F}(Z^H_{ji})^2 \le \|X^*\|_F^2 = d$ and $|F| \le nd$, we have

$$\left\|Z_0 - Z^H\right\|_F \le \frac{1}{M} \cdot \sqrt{\sum_{(j,i)\in F_{\text{large}}} (Z^H_{ji})^2} + \kappa_v \cdot \sqrt{|F_{\text{small}}|} \le \frac{\sqrt{d}}{M} + \kappa_v\sqrt{nd}.$$

Together with (15) and the triangle inequality, this gives

$$\|Z_0 - Z\|_F \le \left\|Z_0 - Z^H\right\|_F + \left\|Z^H - Z\right\|_F \le \frac{\sqrt{d}}{M} + \kappa_v\sqrt{nd} + \sqrt{D_d}a.$$

Recall the choices of a, M, and J_0 in Table 1:

$$a = \frac{\delta}{256(d+1)(D_d+1)}, \quad M = \left\lceil \frac{256d(d+1)}{\delta} \right\rceil, \quad J_0 = \left\lceil \log_2\left(\frac{256dn(d+1)}{\delta}\right)\right\rceil.$$

These choices give

$$\max\left\{\frac{\sqrt{d}}{M}, \kappa_v\sqrt{nd}, \sqrt{D_d}a\right\} \le \frac{\delta}{256(d+1)},$$

and therefore $\|Z_0 - Z\|_F < \delta/(64(d+1))$. Since $\|Z\|_F \le \|X^*\|_F = \sqrt{d}$, we obtain

$$\begin{aligned} \left\|Z_0^\top Z_0 - Z^\top Z\right\|_F &= \left\|(Z_0 - Z)^\top(Z_0 - Z) + (Z_0 - Z)^\top Z + Z^\top(Z_0 - Z)\right\|_F \\ &\le \|Z_0 - Z\|_F^2 + 2\|Z\|_F\|Z_0 - Z\|_F \\ &< \left(\frac{\delta}{64(d+1)}\right)^2 + \frac{\delta}{32\sqrt{d+1}} < \frac{\delta}{16}. \end{aligned}$$

Together with (14), this gives

$$\left\|Z_0^\top Z_0 - I_d\right\|_F \le \left\|Z_0^\top Z_0 - Z^\top Z\right\|_F + \left\|Z^\top Z - I_d\right\|_F < \frac{\delta}{16} + \frac{5\delta}{64} < \frac{\delta}{4}.$$

Then, we show (12). The matrix $Z_0 - Z$ is supported in $\mathcal{I}$. Indeed, the hard variables of a hard row j lie in $\mathrm{RT}_j^* = \mathrm{Supp}(x_j^*)$, the free variables lie in RT_j^*, and deleted rows vanish in both matrices. Hence $\operatorname{tr}(A^\top(Z_0 - Z)) = \operatorname{tr}(A^\top(Z_0 - Z))$, and by the Cauchy–Schwarz inequality and Lemma 14,

$$\left|\operatorname{tr}(A^\top(Z_0 - Z))\right| \le d\mathsf{OPT}_{\text{linear}}(A; s)\|Z_0 - Z\|_F.$$

Together with (14), this gives

$$\begin{aligned}\operatorname{tr}(A^\top Z_0) &= \operatorname{tr}(A^\top Z) + \operatorname{tr}(A^\top (Z_0 - Z)) \\ &\geq \operatorname{tr}(A^\top Z) - d\mathsf{OPT}_{\text{linear}}(A; s) \cdot \frac{\delta}{64(d+1)} \\ &\geq \left(1 - \frac{5\delta}{64d} - \frac{\delta}{64}\right) \mathsf{OPT}_{\text{linear}}(A; s) \geq \left(1 - \frac{\delta}{4}\right) \mathsf{OPT}_{\text{linear}}(A; s).\end{aligned}$$

Finally, we show (i)–(iv). Properties (i)–(iii) hold by the construction of Z_0. It remains to show (iv). The transformations from Z to Z^H and from Z^H to Z_0 change only variable values, not the assigned row types. Hence the sparsity capacity usage $c_i(Z_0)$ of column i equals $c_i(Z)$. For a node u of the ADP tree, let $c_{u,i}$ be the number of rows of Z_0 in the subtree of u whose assigned row type contains i, and let $\hat{c}_{u,i}$ be the cap computed along these rows. At a leaf, $\hat{c}_{u,i} = c_{u,i} \in \{0, 1\}$, and at an internal node u with children u_L and u_R, $\hat{c}_{u,i} = R_c(\hat{c}_{u_L,i} + \hat{c}_{u_R,i}) \geq \hat{c}_{u_L,i} + \hat{c}_{u_R,i}$. Therefore, it suffices to show that $\hat{c}_{u,i} \leq s_i$ for every node u.

If $s_i \leq K_0 = \lceil 16(P_d + 1)/\tau \rceil$, then $R_c(c) = c$ for every integer $c \leq K_0$, so no rounding occurs as long as the caps are at most s_i. By induction on the height of u, we obtain $\hat{c}_{u,i} = c_{u,i} \leq c_i(Z_0) \leq s_i$.

If $s_i > K_0$, then $c_i(Z_0) \leq t_i \leq (1 - 4\tau)s_i$. Every sum of two caps is either an integer or larger than K_0. For every such sum $c \leq n$, we have $R_c(c) \leq (1 + \alpha)c$. Indeed, if $c \leq K_0$, then $R_c(c) = c$. Otherwise, $R_c(c) = K_0(1 + \alpha)^l$ for the smallest $l \geq 1$ with $K_0(1 + \alpha)^l \geq c$, and this l satisfies $K_0(1 + \alpha)^{l-1} < c$. By induction on the height $h(u)$ of u, we obtain

$$\hat{c}_{u,i} \leq (1 + \alpha)^{h(u)} c_{u,i} \leq (1 + \alpha)^{H_{\text{tree}}} c_i(Z_0) \leq e^{\alpha H_{\text{tree}}}(1 - 4\tau)s_i.$$

Since $\alpha = \tau/(8(H_{\text{tree}} + 1))$ and $1 - 4\tau \leq e^{-4\tau}$, the right-hand side is at most $e^{\tau/8 - 4\tau} s_i < s_i$, as desired. In the induction step, the induction hypothesis and this bound give $\hat{c}_{u_L,i} + \hat{c}_{u_R,i} < s_i \leq n$, so the merge is not rejected and the bound $R_c(c) \leq (1 + \alpha)c$ applies. □

KeepBest may replace the partial solutions given by the rows of Z_0 by other partial solutions with the same key and a larger objective value. In Proposition 17, we show that the resulting root candidate Y is still nearly orthogonal. We first need a technical lemma about the Gram matrix bins.

Lemma 16. *Let $G, G' \succeq 0$ satisfy $\operatorname{tr}(G), \operatorname{tr}(G') < 4d$ and $\operatorname{lab}(G) = \operatorname{lab}(G')$. If $\operatorname{lab}(G) = \operatorname{lab}(G') = (-1, O)$, then $\|G - G'\|_F < 2\kappa_G$. Otherwise, let $[u, (1 + \theta)u)$, with $u = \kappa_G(1 + \theta)^k$ for some integer $k \geq 0$, be the trace interval of the common label. Then $\|G - G'\|_F \leq d\theta u \leq d\theta \operatorname{tr}(G)$ and $\operatorname{tr}(G') \leq (1 + \theta) \operatorname{tr}(G)$.*

Proof. Recall the Gram matrix bins introduced in Section 2.4. Let M be a positive semidefinite matrix with vector of eigenvalues λ. Then $\|M\|_F = \|\lambda\|_2 \leq \|\lambda\|_1 = \operatorname{tr}(M)$. If both G and G' belong to the small-trace bin $(-1, O)$, then the triangle inequality gives

$$\|G - G'\|_F \leq \|G\|_F + \|G'\|_F \leq \operatorname{tr}(G) + \operatorname{tr}(G') < 2\kappa_G.$$

Now assume that the two matrices belong to the same bin with trace interval $[u, (1+\theta)u)$. Since they have the same bin label, we have $|G_{il} - G'_{il}| < \theta u$ for all $i, l \in [d]$. Thus $\|G - G'\|_F \leq d\|G - G'\|_\infty \leq d\theta u \leq d\theta \operatorname{tr}(G)$, where $\|\cdot\|_\infty$ denotes the largest absolute value of an entry. Finally, since both traces lie in $[u, (1 + \theta)u)$, we have $\operatorname{tr}(G') < (1 + \theta)u \leq (1 + \theta) \operatorname{tr}(G)$. □

Proposition 17. *Consider the branch of* LinearOracle *that contains the matrix Z_0 of Proposition 15. Then* SelectBranch *returns a root state whose candidate solution Y satisfies*

$$\operatorname{tr}(A^\top Y) \geq \operatorname{tr}(A^\top Z_0), \tag{16}$$

$$\left\|Y^\top Y - Z_0^\top Z_0\right\|_F < \frac{\delta}{8}. \tag{17}$$

Moreover: (i) For every $i \in [d]$, at most s_i rows of Y have an assigned row type that contains i. (ii) $Y_i^\top Y_l = 0$ for every $\{i, l\} \in H^*$. (iii) Every simple pair $e \in S^*$ is contained in the assigned row types of at least d rows of Y.

Proof. We follow the states of Z_0 through the ADP tree. At every leaf, we start from the state given by the corresponding row of Z_0, with its assigned row type and variable values. If KeepBest replaces this state by another state with the same key, we follow the state retained in the dictionary instead. At every internal node, we follow the merge of the two followed child states. The followed states have the same occupancy sets $\mathcal{L}_{\mathrm{occ}}$, capped counters μ_e, and capacity caps $\hat{c}$ as the states of Z_0, because these quantities are part of the key. Their partial solutions, Gram matrices, and Gram matrix bin labels may differ from those of Z_0. By Proposition 15 (ii) and (iv), no followed state is rejected by the occupancy test (line 7) or the capacity test (line 9) of SelectBranch, or by the capacity test in KeepBest.

We next show that the trace test in KeepBest rejects no followed state. Follow the states up to and including the first level, if any, at which one of them fails this test. At each reached node u, let q be the followed state immediately before this test. If q passes, let $\bar{q}$ be the representative finally retained for its key. Since $\mathrm{lab}(G_u(q)) = \mathrm{lab}(G_u(\bar{q}))$, Lemma 16 gives $\|G_u(q) - G_u(\bar{q})\|_F \le 2\kappa_G$ if $\mathrm{tr}(G_u(q)) < \kappa_G$, and $\|G_u(q) - G_u(\bar{q})\|_F \le d\theta\,\mathrm{tr}(G_u(q))$ otherwise. Moreover, Lemma 16 gives

$$\mathrm{tr}(G_u(\bar{q})) \le (1+\theta)\,\mathrm{tr}(G_u(q)) + \kappa_G. \tag{18}$$

For $\ell = 0, 1, \ldots, H_{\mathrm{tree}}$, let E_ℓ be the sum of the traces of the actual Gram matrices of the followed states immediately before KeepBest at level ℓ of the ADP tree, where the leaves are at level 0. Then $E_0 = \mathrm{tr}(Z_0^\top Z_0)$. By Proposition 15, $\left\|Z_0^\top Z_0 - I_d\right\|_F \le \delta/4$, and thus

$$E_0 = d + \mathrm{tr}(Z_0^\top Z_0 - I_d) \le d + \sqrt{d}\left\|Z_0^\top Z_0 - I_d\right\|_F \le d + \frac{\sqrt{d}\delta}{4}.$$

A merge adds the actual Gram matrices of the two child states. Since there are at most N_{leaf} followed states at every level, (18) gives

$$E_{\ell+1} \le (1+\theta)E_\ell + N_{\mathrm{leaf}}\kappa_G. \tag{19}$$

Applying (19) at most $H_{\mathrm{tree}} + 1$ times, we obtain, for every level ℓ,

$$\begin{aligned}
E_\ell &\le (1+\theta)^{H_{\mathrm{tree}}+1}\left(E_0 + (H_{\mathrm{tree}}+1)N_{\mathrm{leaf}}\kappa_G\right)\\
&\le (1+\theta)^{H_{\mathrm{tree}}+1}\left(d + \frac{\sqrt{d}\delta}{4} + (H_{\mathrm{tree}}+1)N_{\mathrm{leaf}}\kappa_G\right)\\
&\le \exp\{\theta(H_{\mathrm{tree}}+1)\}\cdot\left(d + \frac{\sqrt{d}\delta}{4} + (H_{\mathrm{tree}}+1)N_{\mathrm{leaf}}\kappa_G\right)\\
&= \exp\left\{\frac{\delta}{128d^2}\right\}\cdot\left(d + \frac{\sqrt{d}\delta}{4} + N_{\mathrm{leaf}}\cdot\frac{\delta}{128n}\right)\\
&\le \frac{128}{127}\cdot\frac{3d}{2} < 2d,
\end{aligned}$$

where the equality follows from the definitions $\theta = \delta/(128d^2(H_{\mathrm{tree}}+1))$ and $\kappa_G = \delta/(128n(H_{\mathrm{tree}}+1))$. The inequalities in the last line follow from $\delta < 1$, from $N_{\mathrm{leaf}} < 2n$, and from $\exp\{x\} \le 1/(1-x)$ for $x \in [0,1)$. Hence $\mathrm{tr}(G_u(q)) < 2d$ for every reached node u. Thus no followed state fails the trace test in KeepBest, and the followed states reach the root.

We now show (16) and (i)–(iii). Let q_u be the merged candidate at a node u before KeepBest, and let $\bar{q}_u$ be the representative retained for its key. Since KeepBest replaces a representative only by one with a larger objective value, $v(\bar{q}_u) \ge v(q_u)$. Moreover, $v(q_u) = v(\bar{q}_{u_L}) + v(\bar{q}_{u_R})$ for the children u_L and u_R of u. By induction on the height, $v(\bar{q}_u)$ is at least the objective value of the rows of Z_0 in the subtree of u, which gives (16) at the root. For (i), the number of rows of Y whose assigned row type contains i is at most $\hat{c}_i \le s_i$. For (ii), the root key has $\mathcal{L}_{\mathrm{occ}} = [b]$, and the occupied labels are disjoint at every merge, so each hard row type T^r is used exactly once, with hard variables equal to p_r. Since ordinary row types contain no pair of H^*, we obtain $Y_i^\top Y_l = \sum_{r:\{i,l\}\subseteq T^r} p_{ri}p_{rl} = 0$ for every $\{i,l\} \in H^*$ by (3). For (iii), the root key has $\mu_e = d$ for every $e \in S^*$, as the key of $\bar{Z}_0$ does.

We now show (17). Since the Gram matrices are added at every merge, we have $Y^\top Y = Z_0^\top Z_0 + \sum_u \big(G_u(\bar{q}_u) - G_u(q_u)\big)$, where the sum ranges over all nodes u. For every followed state at a node u, Lemma 16 gives

$$\|G_u(q_u) - G_u(\bar{q}_u)\|_F \le d\theta \operatorname{tr}(G_u(q_u)) + 2\kappa_G.$$

Since $E_\ell < 2d$ and $N_{\mathrm{leaf}} < 2n$, the sum over the followed states at level ℓ is at most $d\theta E_\ell + 2N_{\mathrm{leaf}}\kappa_G \le 3\delta/(64(H_{\mathrm{tree}}+1))$. Summing over the $H_{\mathrm{tree}}+1$ levels, we obtain

$$\left\|Y^\top Y - Z_0^\top Z_0\right\|_F \le (H_{\mathrm{tree}}+1)\cdot \frac{3\delta}{64(H_{\mathrm{tree}}+1)} < \frac{\delta}{8},$$

which gives us (17).

Finally, (13) and (17) give $\left\|Y^\top Y - I_d\right\|_F < \delta$. Hence the root state q of Y passes the filter $\|G(q) - I_d\|_F \le \delta$ in SelectBranch. Moreover, since $A \neq O$, we have $\mathsf{OPT}_{\mathrm{linear}}(A; s) > 0$. Hence

$$v(q) = \operatorname{tr}(A^\top Y) \ge \operatorname{tr}(A^\top Z_0) \ge (1-\delta/4)\mathsf{OPT}_{\mathrm{linear}}(A; s) > 0.$$

Therefore, q also passes the filter $v(q) \ge 0$, and SelectBranch returns q. □

We are now ready to prove the approximation guarantee of LinearOracle.

Proposition 18. *Let $A \in \mathbb{Q}^{n\times d}$, $s \in [n]^d$, and $t \in \mathbb{Q} \cap (0,1)$. Then* LinearOracle *returns $X \in \mathcal{F}_s$ satisfying*

$$\operatorname{tr}(A^\top X) \ge (1-t)\mathsf{OPT}_{\mathrm{linear}}(A; s).$$

Proof of Proposition 18. If $A = O$, then every feasible solution is optimal. If $d = 1$, then the normalized restriction of A_1 to s_1 entries of largest absolute value is optimal. Hence we assume that $d \ge 2$ and $A \neq O$.

Let Y be the root candidate of Proposition 17. By Propositions 15 and 17, we have $\operatorname{tr}(A^\top Y) \ge (1-\delta/4)\mathsf{OPT}_{\mathrm{linear}}(A; s) > 0$ and $\left\|Y^\top Y - I_d\right\|_F < \delta$.

Let $\widetilde{Y}$ be the output of Algorithm 1 in ExactFinalize. By Proposition 17 (ii) and (iii), Proposition 8 applies to Y. By Proposition 8, we have $\operatorname{tr}(A^\top \widetilde{Y}) \ge \operatorname{tr}(A^\top Y)$. Moreover, $\widetilde{Y}^\top \widetilde{Y}$ is diagonal, with diagonal entries $(Y^\top Y)_{ii} + \sum_{k\neq i} |(Y^\top Y)_{ik}|$ for $i \in [d]$. Since $\left\|Y^\top Y - I_d\right\|_F < \delta$, by the Cauchy–Schwarz inequality these entries are at most

$$(1+\delta) + \sum_{k\neq i} |(Y^\top Y)_{ik}| \le (1+\delta) + \sqrt{d-1}\,\delta.$$

Let $\rho := \sqrt{1 + (1+\sqrt{d-1})\delta}$, so that $\left\|\widetilde{Y}_i\right\|_2 \le \rho$ for every $i \in [d]$. After the normalization and the sign flips in ExactFinalize, the returned matrix X satisfies $A_i^\top X_i = \left|A_i^\top \widetilde{Y}_i\right| / \left\|\widetilde{Y}_i\right\|_2 \ge A_i^\top \widetilde{Y}_i/\rho$ for every $i \in [d]$. Summing over i, we obtain

$$\operatorname{tr}(A^\top X) \ge \frac{\operatorname{tr}(A^\top \widetilde{Y})}{\rho} \ge \frac{1}{\sqrt{1 + (1+\sqrt{d-1})\delta}} \cdot \left(1 - \frac{\delta}{4}\right) \mathsf{OPT}_{\mathrm{linear}}(A; s).$$

Since LinearOracle returns the best matrix it finds, it remains to show that $1-t \le (1-\delta/4)/\sqrt{1 + (1+\sqrt{d-1})\delta}$. Since $(1+x)^{-1/2} \ge 1 - x/2$ for every $x \ge 0$ and $\delta = t/(8(d+1))$, we obtain

$$\begin{aligned}\frac{1}{\sqrt{1 + (1+\sqrt{d-1})\delta}} \cdot \left(1 - \frac{\delta}{4}\right) &\ge \left(1 - \frac{(1+\sqrt{d-1})\delta}{2}\right) \cdot \left(1 - \frac{\delta}{4}\right) \\ &\ge (1 - d\delta) \cdot \left(1 - \frac{\delta}{4}\right)\end{aligned}$$

$$\geq \left(1 - \frac{t}{8}\right)\left(1 - \frac{t}{32d}\right)$$
$$\geq 1 - \left(\frac{1}{8} + \frac{1}{32d}\right) t \geq 1 - t.$$

□

Next, we prove the approximation guarantee of Algorithm 3.

Theorem 9. *Let $R \in \mathbb{Q}^{n\times r}$, $s \in [n]^d$, and $\epsilon \in \mathbb{Q} \cap (0,1)$. Then Algorithm 3 returns $X \in \mathcal{F}_s$ satisfying $\left\|R^\top X\right\|_F^2 \geq (1-\epsilon)\mathsf{OPT}(R;s)$.*

Proof of Theorem 9. If $\mathsf{OPT}(R;s) = 0$, every matrix in $\mathcal{F}_s$ is optimal. Thus, we assume that $\mathsf{OPT}(R;s) > 0$. Let X^* be an optimal solution of SPCA, and let $W^* := R^\top X^* / \left\|R^\top X^*\right\|_F$, so that $\|W^*\|_F = 1$. Let J_W and $t = \epsilon/4$ be as in Algorithm 3, and let $m := rd$. Let $W \in J_W^{-1}\mathbb{Z}^{r\times d}$ be defined by

$$W_{ij} := \frac{\operatorname{sign}(W^*_{ij})\lfloor J_W \left|W^*_{ij}\right| \rfloor}{J_W}, \quad i \in [r],\ j \in [d].$$

Since $|W_{ij}| \leq \left|W^*_{ij}\right|$ and $\left|W_{ij} - W^*_{ij}\right| \leq 1/J_W$, we have $\|W\|_F \leq 1$ and $\|W - W^*\|_F \leq \sqrt{m}/J_W \leq t$. Hence Algorithm 3 enumerates W. By the Cauchy–Schwarz inequality,

$$\begin{aligned}\mathsf{OPT}_{\mathrm{linear}}(RW;s) &\geq \operatorname{tr}((RW)^\top X^*) = \operatorname{tr}(W^\top R^\top X^*)\\ &\geq \operatorname{tr}((W^*)^\top R^\top X^*) - \|W - W^*\|_F \left\|R^\top X^*\right\|_F\\ &\geq (1-t)\sqrt{\mathsf{OPT}(R;s)}.\end{aligned}$$

Let X_W be the output of LinearOracle for the input (RW, s, t). Since $\|W\|_F \leq 1$, the Cauchy–Schwarz inequality and Proposition 18 give

$$\begin{aligned}\left\|R^\top X_W\right\|_F &\geq \operatorname{tr}((RW)^\top X_W) \geq \left(1 - \frac{\epsilon}{4}\right)\mathsf{OPT}_{\mathrm{linear}}(RW;s)\\ &\geq \left(1 - \frac{\epsilon}{4}\right)^2 \sqrt{\mathsf{OPT}(R;s)},\end{aligned}$$

and thus

$$\left\|R^\top X_W\right\|_F^2 \geq \left(1 - \frac{\epsilon}{4}\right)^4 \mathsf{OPT}(R;s) \geq (1-\epsilon)\,\mathsf{OPT}(R;s),$$

where the last inequality follows from $(1-x)^4 \geq 1 - 4x$ for $x \in [0,1]$. Since Algorithm 3 returns the best matrix it finds, its output satisfies the same bound. □

C.4 Bit complexity

In this section, we prove Theorem 10, which we restate below. Real algebraic numbers are represented exactly, as described after Definition 19.

Theorem 10. *Let $R \in \mathbb{Q}^{n\times r}$, $s \in [n]^d$, and $\epsilon \in \mathbb{Q} \cap (0,1)$. Let B_{SPCA} be the encoding size of (R, s, ϵ), and let $U_\epsilon := 1 + \log_2(n+1) + \log_2(r+d+1) + \log_2(1/\epsilon)$. Then there is an implementation of Algorithm 3 whose bit complexity T_{SPCA} satisfies*

$$T_{\mathrm{SPCA}} \leq F_{\mathrm{SPCA}}(r,d)\, n \epsilon^{-(rd+d^3-d^2+4d+2+c)} (B_{\mathrm{SPCA}} + 1)^c U_\epsilon^{d^2+5d+4+4c},$$

where $F_{\mathrm{SPCA}}(r,d)$ depends only on r and d, and c is an absolute constant.

We first bound the running time of one call to LinearOracle, and then multiply this bound by the number of matrices W enumerated by Algorithm 3. We first count the arithmetic operations performed by the ADP and by LinearOracle, and then we bound the bit complexity of these operations.

C.4.1 Workload analysis.

Let $q := \binom{d}{2}$, which is an upper bound on $|H|$. Recall that $D_d = d(d-1)^2$ is an upper bound on the number of hard variables in one branch of LinearOracle. Let $q_G := \binom{d+1}{2}$ be the number of entries in the lower triangular part of a $d \times d$ Gram matrix, and let B_{in} be the encoding size of the input (A, s, t), including the dimensions.

Number of choices of H and $\mathcal{L}$. There are at most 2^q sets H of hard pairs. By Lemma 12, the number of hard rows satisfies $b \le b_d := (d-1)q$. For each b, there are at most 2^{db_d} lists $\mathcal{L} = \{\{T^1, T^2, \dots, T^b\}\}$. Therefore, the total number of choices in lines 7–11 of LinearOracle is at most

$$E_d := (b_d + 1)2^{q+db_d}.$$

Number of boxes Q. Since $a = \delta/(256(d+1)(D_d+1))$ and $\delta = t/(8(d+1))$ in Table 1, the number of boxes is at most

$$J_p^{D_d} = \left\lceil \frac{4096(d+1)^2(D_d+1)}{t} \right\rceil^{D_d} \le f_{\text{branches}}(d) t^{-D_d} \tag{20}$$

for some function f_{branches} that depends only on d. For every box, the subroutine FixHardVariables decides whether a quadratic system is feasible and, if so, computes one solution. This system has at most D_d variables and a number of constraints that depends only on d. Hence the cost of one such call is a factor that depends only on d times a fixed polynomial in the bit length of the coefficients (see Lemma 27).

Number of choices of free variables at a leaf. SelectBranch first initializes the leaves. The running time of this step is governed by the number of choices of free variables at a leaf. Recall that $M = \lceil 256d(d+1)/\delta \rceil$, and $J_0 = \lceil \log_2(256dn(d+1)/\delta) \rceil = \mathcal{O}(\log(n+1) + \log(d+1) + \log(1/t))$. Let

$$U := 1 + \log_2(n+1) + \log_2(d+1) + \log_2(1/t).$$

By the definition of $\mathcal{V}$, we have

$$|\mathcal{V}| \le 2MJ_0 + 3 \le \frac{4097d(d+1)^2}{t} J_0 + 3 \le f_{\text{free}}(d) \cdot \frac{U}{t}$$

for some function f_{free}. At a leaf of the ADP tree, there are at most 2^d ordinary row types and at most b_d hard row types in $\mathcal{L}$. The row is either unused, or it has one of these row types and at most d free variables with values in $\mathcal{V}$. Hence the number of leaf candidates is at most

$$L_{\text{leaf}} \le 1 + (2^d + b_d)\,|\mathcal{V}|^d \le f_{\text{leaf}}(d) t^{-d} U^d \tag{21}$$

for some function f_{leaf}.

Next, we bound the number of states at the internal nodes.

Number of capacity caps. Recall that $K_0 = \lceil 16(P_d+1)/\tau \rceil$ and $\alpha = \tau/(8(H_{\text{tree}}+1))$ in Table 1, with $P_d = dq$, $\tau = \delta/(128d^2)$, and $H_{\text{tree}} = \lceil \log_2(n) \rceil$. Since $\log(1+x) \ge x/2$ for $x \in [0,1]$, we have

$$\log_{1+\alpha}(n) \le \frac{\log(n+1)}{\log(1+\alpha)} \le \frac{2\log(n+1)}{\alpha} \le \frac{16384d^2(d+1)(H_{\text{tree}}+1)U}{t}.$$

Since U_c consists of the integers $0, \dots, K_0$ and a geometric sequence above K_0, we obtain

$$|U_c| \le K_0 + 1 + \log_{1+\alpha}(n) + 1 \le \frac{16385(P_d + 1 + (H_{\text{tree}}+1)U)d^2(d+1)}{t} + 2.$$

Therefore, the number of vectors $\hat{c}$ of capacity caps is at most

$$|U_c|^d \le f_{\text{sparsity}}(d) t^{-d} U^{2d} \tag{22}$$

for some function $f_{\text{sparsity}}(d)$.

Number of Gram matrix bins. Since Gram matrices are symmetric, a Gram matrix bin label only needs to store the lower triangular part of the rounded Gram matrix. Recall that $\theta = \delta/(128d^2(H_{\text{tree}}+1)) < 1$ and $\kappa_G = \delta/(128n(H_{\text{tree}}+1))$. We first count the number L_G of intervals I_k with $k \geq 0$ that contain a trace smaller than $4d$. We have

$$L_G \leq 1 + \frac{\log(4d) - \log(\kappa_G)}{\log(1+\theta)} \leq 1 + \frac{2\log\left(\frac{4d}{\kappa_G}\right)}{\theta}.$$

We claim that, within each trace interval $[(1+\theta)^k\kappa_G, (1+\theta)^{k+1}\kappa_G)$, each lower triangular entry has at most $7/\theta$ possible bins. Indeed, since $G \succeq 0$, we have $|G_{il}| \leq \operatorname{tr}(G) < (1+\theta)^{k+1}\kappa_G$, and thus the number of possible bins is at most

$$\begin{aligned} 2\left\lceil \frac{(1+\theta)^{k+1}\kappa_G}{\theta(1+\theta)^k\kappa_G} \right\rceil + 1 = 2\left\lceil \frac{1+\theta}{\theta} \right\rceil + 1 &\leq 2\left(\frac{1}{\theta} + 1 + 1\right) + 1 \\ &\leq \frac{7}{\theta} = \frac{7168(d+1)d^2(H_{\text{tree}}+1)}{t}. \end{aligned}$$

Therefore, the total number M_G of Gram matrix bin labels satisfies

$$M_G \leq L_G \cdot \left(\frac{7}{\theta}\right)^{q_G} + 1 \leq f_{\text{bin}}(d) t^{-q_G-1} U^{q_G+2} \tag{23}$$

for some function f_{bin}.

Number of states at one internal node. Recall that a dictionary key consists of a set $\mathcal{L}_{\text{occ}}$, a vector of capped counters $(\mu_e)_{e\in S}$, a vector $\hat{c}$ of capacity caps, and a Gram matrix bin label $\text{lab}(G)$. There are at most 2^b possible sets $\mathcal{L}_{\text{occ}}$ and at most $(d+1)^q$ possible vectors $(\mu_e)_{e\in S}$. Moreover, there are at most $|U_c|^d$ possible vectors $\hat{c}$ and at most M_G possible labels $\text{lab}(G)$. Thus, the number of states at one node is at most

$$K_{\text{state}} \leq 2^b (d+1)^q \, |U_c|^d \, M_G. \tag{24}$$

Since $b \leq (d-1)q$, the first two factors depend only on d. Substituting (22) and (23) gives

$$K_{\text{state}} \leq f_{\text{state}}(d) t^{-(q_G+d+1)} U^{q_G+2d+2}$$

for some function f_{state}.

Number of state pairs at a merge. When two child nodes are merged, the number K_{merge} of state pairs satisfies

$$K_{\text{merge}} \leq K_{\text{state}}^2 \leq f_{\text{state}}^2(d) t^{-2(q_G+d+1)} U^{2q_G+4d+4}. \tag{25}$$

ADP workload for a fixed p. Fix H, b, $\mathcal{L}$, and a vector p of values of the hard variables in LinearOracle. The ADP tree has $N_{\text{leaf}} < 2n$ leaves and $N_{\text{leaf}} - 1$ internal nodes. At each leaf, LeafInit enumerates at most L_{leaf} candidates and processes each of them with KeepBest. Hence the leaves require at most $N_{\text{leaf}} L_{\text{leaf}}$ processing steps. At each internal node, SelectBranch processes every pair of states of its two children, that is, at most K_{merge} pairs. Hence the internal nodes require at most $(N_{\text{leaf}} - 1)K_{\text{merge}}$ processing steps. Finally, each of the at most K_{state} root states is checked against the acceptance conditions. Therefore, for a fixed p, the total ADP workload is

$$\begin{aligned} W_{\text{ADP}} &\leq 1 + N_{\text{leaf}} L_{\text{leaf}} + (N_{\text{leaf}} - 1)K_{\text{merge}} + K_{\text{state}} \\ &\leq f_{\text{bw}}(d) n t^{-(d^2+3d+2)} U^{d^2+5d+4} \end{aligned} \tag{26}$$

for some function f_{bw}, where we used $2q_G = d^2 + d$.

LinearOracle workload. For every box Q, FixHardVariables solves one quadratic system exactly and, if the system is feasible, returns one vector p. Therefore,

$$W_{\text{LinearOracle}} \le E_d J_p^{D_d}(1 + N_{\text{leaf}} L_{\text{leaf}} + (N_{\text{leaf}} - 1)K_{\text{merge}} + (4nd + 1)K_{\text{state}}),$$

where the leading 1 accounts for the call to FixHardVariables for the box. The term $4ndK_{\text{state}}$ is an upper bound that accounts, for each accepted root state, for reconstructing its candidate solution by backtracking to the n leaves and for ExactFinalize. Substituting the estimates of E_d and $J_p^{D_d}$, and absorbing all factors that depend only on d into a function $f_{\text{lin}}(d)$, we obtain

$$W_{\text{LinearOracle}} \le f_{\text{lin}}(d) n t^{-D_d - d^2 - 3d - 2} U^{d^2 + 5d + 4}. \tag{27}$$

Next, we bound the bit complexity of each subroutine, and then of Algorithm 3.

C.4.2 Bit complexity analysis.

To bound the bit complexity of Algorithm 3, we need to address two issues. The first is how to represent irrational algebraic numbers. The second is the bit complexity of arithmetic operations on these representations. Irrational numbers arise only in ExactFinalize and FixHardVariables, which solve (2) and (3), respectively.

A standard way to represent a real root of a polynomial with integer coefficients is to store the polynomial together with an interval that contains exactly one of its roots (see, e.g., [4, Section 2.2.2]). For example, the polynomial $x^2 - 2$ and the *root isolating interval* $(1, 2)$ specify $\sqrt{2}$ exactly, because $\sqrt{2}$ is the only root of the polynomial in this interval. To store the number $2 + 3\sqrt{2}$, we store $\sqrt{2}$ as above, together with its *rational coordinate coefficients* 2 and 3. We use the following standard notions from field theory, see for example [22, Chapter V].

Definition 19. *Let $\gamma \in \mathbb{R}$ be a root of a non-zero polynomial with integer coefficients. We denote by $\mathbb{K} := \mathbb{Q}(\gamma)$ the smallest subfield of $\mathbb{R}$ that contains $\mathbb{Q}$ and γ. In other words, $\mathbb{K}$ is the set of numbers obtained from rational numbers and γ by addition, subtraction, multiplication, and division by non-zero numbers. Let $f \in \mathbb{Q}[x]$ be an irreducible polynomial with $f(\gamma) = 0$. We call f a* defining polynomial *of $\mathbb{K}$. The* degree $[\mathbb{K} : \mathbb{Q}]$ *of the field extension is the degree of f.*

We represent $\mathbb{K}$ by a defining polynomial f and a rational interval that contains γ and no other root of f, and we represent each element of $\mathbb{K}$ by its rational coordinate coefficients. For $m := [\mathbb{K} : \mathbb{Q}]$, each element of $\mathbb{K}$ is stored as $\sum_{j=0}^{m-1} q_j \gamma^j$ with $q_j \in \mathbb{Q}$. The *representation size* of a finite set of elements of $\mathbb{K}$ is the largest encoding size of the coefficients of f, of the endpoints of the interval, and of the rational coordinate coefficients of these elements. We first bound the bit complexity of computations that satisfy the following assumption.

Assumption 20. *All input entries belong to one field $\mathbb{K} = \mathbb{Q}(\gamma) \subseteq \mathbb{R}$ with $[\mathbb{K} : \mathbb{Q}] \le \Delta$. Let $B_{\text{enc}} \ge 1$ be an upper bound on the representation size of the input entries. Every divisor is non-zero, and every square root is the non-negative root of a non-negative number. The computation has at most N_0 input entries and performs at most L additions, subtractions, multiplications, divisions, and equality or sign tests, and at most κ square-root computations.*

For example, the computation $x_1 x_2 + x_3$ with $x_1, x_2, x_3 \in \mathbb{K}$ has $N_0 = 3$, $L = 2$, and $\kappa = 0$.

Lemma 21. *Consider a computation that satisfies Assumption 20 and keeps every scalar result. Let T be its bit complexity, let $\mathbb{K}_{\text{out}}$ be a field that contains all scalar results, and let B_{out} be the representation size of the scalar results in $\mathbb{K}_{\text{out}}$. Then $\mathbb{K}_{\text{out}}$ can be chosen so that*

$$[\mathbb{K}_{\text{out}} : \mathbb{Q}] \le \Delta 2^{\kappa}, \quad T \le F(\Delta, N_0, L, \kappa)(B_{\text{enc}} + 1)^c,$$
$$B_{\text{out}} \le C(\Delta, N_0, L, \kappa)(B_{\text{enc}} + 1),$$

for some parameter-dependent constants F and C, and for an absolute constant $c > 0$.

Proof. The computation consists of steps, each of which is an arithmetic operation, a comparison, or a square root. Before each step, let $\mathbb{K}$ be the current field, which contains all kept values. Let $m := [\mathbb{K} : \mathbb{Q}]$, let u be the number of kept values, and let b_{cur} be the representation size of the kept values. We first show that the bit complexity T_{step} of each step and the largest encoding size b_{new} after the step satisfy

$$T_{\mathrm{step}} \le F_{\mathrm{step}}(m,u)(b_{\mathrm{cur}}+1)^c, \quad b_{\mathrm{new}} \le C_{\mathrm{step}}(m,u)(b_{\mathrm{cur}}+1), \tag{28}$$

for some parameter-dependent constants F_{step} and C_{step}, and for an absolute constant $c > 0$.

(i) *Arithmetic and comparison in* $\mathbb{K}$. By [29, Section 2.2], arithmetic operations and comparisons in $\mathbb{K}$ have bit complexity $C(m)(b_{\mathrm{cur}}+1)^c$, and the encoding sizes of the results are at most $C(m)(b_{\mathrm{cur}}+1)$. An equality test has the same bounds, since it compares the coordinate coefficients reduced modulo the defining polynomial.

(ii) *Field extension after adjoining a root.* If the square root of an element of $\mathbb{K}$ does not belong to $\mathbb{K}$, we extend $\mathbb{K}$. The multiplication table and univariate representation procedures in [6, Sections 12.2 and 12.4] compute a primitive element of the new field, of degree at most $2m$. If necessary, the factorization algorithm of Lenstra, Lenstra, and Lovász [23] computes its irreducible defining polynomial. The real-root isolation and sign determination procedures in [6, Sections 10.2 and 10.3] identify the intended real root. The univariate representation then expresses the u kept values in the new field. The bit complexity of these procedures is $F(m,u)(b_{\mathrm{cur}}+1)^c$, and the encoding sizes of the results are at most $C(m,u)(b_{\mathrm{cur}}+1)$.

We now combine the steps. Arithmetic operations and comparisons stay in the current field, and adjoining a square root multiplies the degree of the field extension by at most two. Hence $[\mathbb{K}_{\mathrm{out}} : \mathbb{Q}] \le \Delta 2^\kappa$. Let $J := L + \kappa$, $m_{\max} := \Delta 2^\kappa$, and $u_{\max} := \max\{1, N_0 + J\}$. Then every step has $m \le m_{\max}$ and $u \le u_{\max}$. Let $F_{\max}$ be the maximum of $F_{\mathrm{step}}(m,u)$ over $m \in [m_{\max}]$ and $u \in [u_{\max}]$, and let $C_{\max}$ be one plus the maximum of $C_{\mathrm{step}}(m,u)$ over the same range. Let b_j be the largest encoding size after j steps. Then (28) gives

$$b_{j+1} + 1 \le C_{\max}(b_j + 1) \le \cdots \le C_{\max}^{j+1}(b_0 + 1) \le C_{\max}^{j+1}(B_{\mathrm{enc}} + 1).$$

Since there are at most J steps, (28) gives the statement with $C(\Delta, N_0, L, \kappa) := C_{\max}^J$ and $F(\Delta, N_0, L, \kappa) := JF_{\max}C_{\max}^{cJ}$. □

Next, we analyze the running time of ExactFinalize. We start with the computation of N in Algorithm 1.

Lemma 22. *Let $d \ge 2$, let $\mathbb{K} \subseteq \mathbb{R}$ be an algebraic number field, and let $B \in \mathbb{K}^{d\times(d-2)}$ be as in Algorithm 1. One can construct $N \in \mathbb{R}^{d\times 2}$ such that $N^\top N = I_2$ and $B^\top N = O$, with entries in an extension of $\mathbb{K}$ obtained by adjoining at most two square roots, using $\mathcal{O}(d^3)$ field operations and comparisons.*

Proof. If $d = 2$, we set $N := I_2$. Assume that $d > 2$. In the proof of Proposition 8, we showed that $\ker(B^\top)$ has dimension at least 2. Gaussian elimination computes two linearly independent vectors $w_1, w_2 \in \ker(B^\top) \cap \mathbb{K}^d$ with $\mathcal{O}(d^3)$ field operations. Let $\hat{w}_2 := w_2 - (w_1^\top w_2 / w_1^\top w_1) w_1$, so that $\hat{w}_2 \ne 0$ and $\hat{w}_2^\top w_1 = 0$. Let $N_1 := w_1/\sqrt{w_1^\top w_1}$, $N_2 := \hat{w}_2/\sqrt{\hat{w}_2^\top \hat{w}_2}$, and $N := (N_1, N_2)$. Then $N^\top N = I_2$ and $B^\top N = O$, and the construction uses two square roots. □

Next, we consider the computation of K_{new}.

Lemma 23. *Let $\mathbb{K} \subseteq \mathbb{R}$ be an algebraic number field, and let $\Gamma, \Lambda \in \mathbb{K}^{2\times 2}$ with $\Lambda \succeq 0$. An exact optimal solution to (2) can be computed using $\mathcal{O}(1)$ field operations and comparisons and at most three square roots.*

Proof. In the proof of Proposition 13, we showed that every feasible solution of (2) has the form $\Psi\sqrt{\Lambda}$, where Ψ is a 2×2 orthogonal matrix. Hence it suffices to compute $\sqrt{\Lambda}$ and an optimal Ψ. If $\Lambda = O$, we set $K_{\mathrm{new}} := O$. Assume that $\Lambda \ne O$. Let $a := \sqrt{\det(\Lambda)}$, $b := \sqrt{\mathrm{tr}(\Lambda) + 2a} > 0$, and $P := (\Lambda + aI_2)/b$. By the Cayley–Hamilton theorem, $\Lambda^2 - \mathrm{tr}(\Lambda)\Lambda + \det(\Lambda) I_2 = O$, and therefore

$$P^2 = \frac{\Lambda^2 + 2a\Lambda + a^2 I_2}{b^2} = \frac{\mathrm{tr}(\Lambda)\Lambda + 2a\Lambda}{b^2} = \Lambda.$$

Since $a \geq 0$, we have $P \succeq 0$, and thus $\sqrt{\Lambda} = P$.

Next, we determine Ψ. Every 2×2 orthogonal matrix has one of the two forms $\Psi_+ = \begin{pmatrix} x & -y \\ y & x \end{pmatrix}$ or $\Psi_- = \begin{pmatrix} x & y \\ y & -x \end{pmatrix}$, where $x^2 + y^2 = 1$. We choose Ψ to maximize $\operatorname{tr}(\Gamma^\top \Psi \sqrt{\Lambda}) = \operatorname{tr}(\sqrt{\Lambda}\Gamma^\top \Psi)$. Write $M := \sqrt{\Lambda}\Gamma^\top = (m_{ij})$, and define

$$v_+ := \begin{pmatrix} m_{11} + m_{22} \\ m_{12} - m_{21} \end{pmatrix}, \quad v_- := \begin{pmatrix} m_{11} - m_{22} \\ m_{12} + m_{21} \end{pmatrix}.$$

Since $\operatorname{tr}(M\Psi_+) = (x, y)v_+$ and $\operatorname{tr}(M\Psi_-) = (x, y)v_-$, by the Cauchy–Schwarz inequality the maximum of (2) is $\max\{\|v_+\|_2, \|v_-\|_2\}$. If $\|v_+\|_2^2 = \|v_-\|_2^2 = 0$, we take $\Psi = I_2$. Otherwise, assume without loss of generality that $\|v_+\|_2^2 \geq \|v_-\|_2^2$. Then we set $\Psi = \Psi_+$ with $(x, y)^\top = v_+ / \|v_+\|_2$.

Finally, we set $K_{\text{new}} := \Psi\sqrt{\Lambda}$. The computation uses $\mathcal{O}(1)$ field operations and comparisons, and at most three square roots. □

We now bound the degree of the field extension and the encoding sizes of the scalar results of ExactFinalize.

Lemma 24. *Consider the computation* ExactFinalize*, with the notation of Assumption 20. Let $Y^0 := Y$ be the input candidate solution and $G^0 := Y^\top Y$. Let $B_{\text{tot-enc}} \geq 1$ be an upper bound on the total encoding size of the coefficients of the defining polynomial of $\mathbb{K}$, of the endpoints of its root isolating interval, and of the rational coordinate coefficients of the entries of A, Y^0, and G^0. Let $\mathbb{K}_{\text{out}}$ be a field that contains all scalar results of* ExactFinalize*, and let B_{out} be the representation size of these scalar results in $\mathbb{K}_{\text{out}}$. Then $\mathbb{K}_{\text{out}}$ can be chosen so that*

$$[\mathbb{K}_{\text{out}} : \mathbb{Q}] \leq \Delta 2^{5\binom{d}{2}+d}, \quad B_{\text{out}} \leq C_{\text{ExactFinalize}}(d, \Delta)(B_{\text{tot-enc}} + 1),$$

for some parameter-dependent constant $C_{\text{ExactFinalize}}$.

Proof. We first bound $[\mathbb{K}_{\text{out}} : \mathbb{Q}]$. By Lemmas 22 and 23, each pair $e \in S$ in Algorithm 1 uses at most five square roots. The normalization of the columns uses d further square roots. Hence Lemma 21 gives $[\mathbb{K}_{\text{out}} : \mathbb{Q}] \leq \Delta 2^{5\binom{d}{2}+d}$.

Every scalar value produced by the local computations of ExactFinalize is obtained through arithmetic operations or taking square roots. In particular, each iteration of Algorithm 1 involves only a number of scalar inputs, arithmetic operations, and square roots bounded in terms of d. We can defer the field representation update of any stored value until it is needed for a computation, using the embeddings between successive fields. Since the number of iterations is $\mathcal{O}(d^2)$, Lemma 21 gives encoding sizes at most $C(d, \Delta)(B_{\text{tot-enc}} + 1)$ for all locally computed values in Algorithm 1. After the local updates, the representations of the remaining entries are updated into the final field $\mathbb{K}_{\text{out}}$. As each such update involves only a number of operations depending on d and Δ, the same encoding-size bound holds for all output entries, with a constant that does not depend on n.

Finally, since $B_{\text{tot-enc}}$ is the total encoding size of the inputs, the encoding sizes of $\left\|\widetilde{Y}_i\right\|_2 = \sqrt{\|G_i^0\|_1}$ and $A_i^\top \widetilde{Y}_i$ are at most the desired upper bound. Here $\widetilde{Y}$ denotes the matrix after Algorithm 1 and before the normalization. Indeed, since G_i^0 has d entries, Lemma 21 bounds the encoding size of this term by $C_{2-\text{norm}}(d, \Delta)(B_{\text{tot-enc}} + 1)$ for some parameter-dependent constant $C_{2-\text{norm}}$. It remains to upper bound the encoding size of $A_i^\top \widetilde{Y}_i$. Let I_i denote the set of rows in column i modified by Algorithm 1. Then $|I_i| \leq d(d-1)$, and

$$A_i^\top \widetilde{Y}_i = A_i^\top Y_i^0 + \sum_{j \in I_i} A_{ji}(\widetilde{Y}_{ji} - Y_{ji}^0).$$

Since $B_{\text{tot-enc}}$ is the total input encoding size, computing the first term on the right-hand side by clearing denominators and adding its at most n terms gives an encoding size of at most $C_{\text{sum}}(\Delta)(B_{\text{tot-enc}} + 1)$, for some constant C_{sum} that depends only on Δ. The sum on the right-hand side has only $\mathcal{O}(d^2)$ terms, whose encoding sizes are bounded by Lemma 21 and the analysis above. The normalization and the sign tests satisfy the same bound. □

Proposition 25. *Let $B_{\mathrm{tot-enc}}$ be as in Lemma 24. Then there is an implementation of* ExactFinalize *whose bit complexity $T_{\textsc{ExactFinalize}}$ satisfies*

$$T_{\textsc{ExactFinalize}} \le nF_{\textsc{ExactFinalize}}(d,\Delta)(B_{\mathrm{tot-enc}}+1)^c$$

for some parameter-dependent constant $F_{\textsc{ExactFinalize}}$ and an absolute constant c.

Proof. By Lemmas 22 and 23, the iterations of Algorithm 1 use $\mathcal{O}(d^5)$ field operations and comparisons and at most $5\binom{d}{2}$ square roots. Therefore, by Lemma 21, their bit complexity $T_{\mathrm{iterations}}$ satisfies

$$T_{\mathrm{iterations}} \le F(d,\Delta)(B_{\mathrm{tot-enc}}+1)^c.$$

The remaining computations of ExactFinalize require $\mathcal{O}(nd^2)$ scalar operations and updates of field representations. By Lemma 24, their inputs have encoding sizes at most $C_{\textsc{ExactFinalize}}(d,\Delta)(B_{\mathrm{tot-enc}}+1)$, and the degree of the field extension is at most $\Delta 2^{5\binom{d}{2}+d}$. As in the proof of Lemma 24, the partial sums of the n-term sums in these computations have encoding sizes at most $C(d,\Delta)(B_{\mathrm{tot-enc}}+1)$. Hence, by Lemma 21 applied to each addition, multiplication, and field update separately, each of these scalar computations has bit complexity $F(d,\Delta)(B_{\mathrm{tot-enc}}+1)^c$. Absorbing d^2 into $F(d,\Delta)$, the remaining computations have bit complexity $nF(d,\Delta)(B_{\mathrm{tot-enc}}+1)^c$. Adding the two bounds concludes the proof. □

Next, we bound the bit complexity of FixHardVariables. The main step is to solve the system (3) in each box Q. We reduce this problem to finding a zero of a non-negative multivariate quartic polynomial, and then apply a bounded algebraic sampling algorithm.

Proposition 26. *Let $k \ge 1$, and let $p \in \mathbb{Z}[x_1,x_2,\dots,x_k]$ be a polynomial of degree at most ν with coefficients of encoding size at most $B_{\mathrm{coef}} \ge 1$. Assume that $p \ge 0$ on $\mathbb{R}^k$ and that $Z(p) := \{x \in \mathbb{R}^k : p(x) = 0\}$ is bounded. Then there is an algorithm that returns exact points of $Z(p)$ that meet every connected component of $Z(p)$. In particular, it returns no point if $Z(p) = \emptyset$. Let T_{samp} be its bit complexity, let Δ_{samp} be the largest degree of the field extension of a returned point, and let B_{samp} be the largest representation size of a returned point. Then*

$$T_{\mathrm{samp}} \le F_{\mathrm{samp}}(k,\nu)(B_{\mathrm{coef}}+1)^c, \quad \Delta_{\mathrm{samp}} \le (\nu+2)(\nu+1)^{k-1},$$
$$B_{\mathrm{samp}} \le C_{\mathrm{samp}}(k,\nu)(B_{\mathrm{coef}}+1),$$

for some parameter-dependent constants F_{samp} and C_{samp}, and an absolute constant c.

Proof. We apply the bounded algebraic sampling algorithm [6, Algorithm 12.15], as recalled in [5, Algorithm 2.8]. This algorithm deforms p using an infinitesimal, computes finitely many critical points of the deformed polynomial, and takes their limits to obtain points of $Z(p)$. These points meet every bounded connected component of $Z(p)$. Since $Z(p)$ is bounded, they meet every connected component. In particular, the returned list is empty when $Z(p) = \emptyset$.

Each returned point is specified by a real univariate representation

$$(f, g_0, g_1, \dots, g_k, \pi),$$

where π is a *Thom encoding* that identifies a real root γ of f, and the point is

$$x = \left(\frac{g_1(\gamma)}{g_0(\gamma)}, \dots, \frac{g_k(\gamma)}{g_0(\gamma)}\right), \qquad g_0(\gamma) \ne 0.$$

We next bound the sizes of these representations. Let $N := (\nu+2)(\nu+1)^{k-1}$. By the analysis of [5, Algorithm 2.8], the degrees of $f, g_0, \dots, g_k$ are at most N, and the bitsizes of their coefficients are bounded by a function of k and ν that is linear in B_{coef}. In particular, these encoding sizes are at most $C_0(k,\nu)(B_{\mathrm{coef}}+1)$ for some parameter-dependent constant C_0. By the complexity analysis of this algorithm in [6, Chapter 12], the number of arithmetic operations depends only on k and ν, and these operations are performed on integers whose bitsizes are bounded in the same way. Hence the algorithm has bit complexity at most $F_1(k,\nu)(B_{\mathrm{coef}}+1)^{c_1}$, for some parameter-dependent constant F_1 and an absolute constant c_1.

Finally, we convert the real univariate representations to the representation of Definition 19. For the tuple $(f, g_0, \dots, g_k, \pi)$, let γ be the real root of f specified by π. By polynomial factorization over $\mathbb{Q}$ [23], real-root isolation, and sign determination [6, Section 10.2], we can compute the primitive irreducible factor $m \in \mathbb{Z}[x]$ of f such that $m(\gamma) = 0$, and a rational interval isolating γ.

Since $g_0(\gamma) \neq 0$ and m is irreducible, we have $\gcd(g_0, m) = 1$. The extended Euclidean algorithm computes $u, v \in \mathbb{Q}[x]$ with $ug_0 + vm = 1$. Hence $g_i = ug_0g_i + vmg_i$, and so $g_i(\gamma)/g_0(\gamma) = u(\gamma)g_i(\gamma)$. Consequently, each coordinate $x_i = g_i(\gamma)/g_0(\gamma)$ equals $r_i(\gamma)$, where $r_i = (ug_i) \bmod m$ and $\deg r_i < \deg m$. Hence all coordinates belong to the same selected real field $\mathbb{K} = \mathbb{Q}(\gamma)$, with $[\mathbb{K} : \mathbb{Q}] = \deg m \leq N$.

Polynomial factorization over $\mathbb{Q}$ [23], real-root isolation [6, Section 10.2], and the extended Euclidean algorithm over $\mathbb{Q}[x]$ [30, Theorem 6.52] all have bit complexity polynomial in the degrees and in the coefficient encoding sizes. Since all degrees are $\mathcal{O}(N)$, the resulting encoding sizes of the coefficients and of the root isolating interval are at most $C(k, \nu)(B_{\text{coef}} + 1)$, and the total bit complexity of these steps is at most $F(k, \nu)(B_{\text{coef}} + 1)^c$, for an absolute constant c. Enlarging the constants concludes the proof. □

In the next lemma, we bound the bit complexity of solving (3) in one box. Recall from Table 1 that $J_p = \lceil 4096(d+1)^2(D_d+1)/t \rceil$, and let $B_{J_p} := 1 + \lceil \log_2(J_p + 1) \rceil$ be its encoding size.

Lemma 27. *For each box Q, the system (3) with $p \in Q$ can be solved exactly in at most $F_{\text{hard}}(d)B_{J_p}^c$ bit operations, for some parameter-dependent constant F_{hard} and an absolute constant c.*

Proof. We denote the variables in (3) by $p_1, p_2, \dots, p_D$. We write each orthogonality constraint in (3) as $g_e(p) = 0$ for $e \in H$, and the norm constraint as $\sum_{j=1}^D p_j^2 \leq d$. For a box $Q = \Omega_{\ell_1} \times \cdots \times \Omega_{\ell_D}$, the constraint $p \in Q$ can be written as

$$p_j \geq \frac{2(\ell_j - 1)}{J_p} - 1 =: l_j, \quad p_j \leq \frac{2\ell_j}{J_p} - 1 =: u_j, \quad \forall j \in [D].$$

Solving (3) in Q is equivalent to finding a zero of the following non-negative quartic polynomial in $3D+1$ variables:

$$h(p, \alpha, \beta, \eta) := \sum_{e \in H} g_e^2(p) + \left(d - \eta^2 - \sum_{j=1}^D p_j^2 \right)^2 + J_p^2 \sum_{j=1}^D \left[(\alpha_j^2 + l_j - p_j)^2 + (\beta_j^2 + p_j - u_j)^2 \right].$$

Indeed, a vector p can be extended to a zero (p, α, β, η) of h if and only if $g_e(p) = 0$ for all $e \in H$, $\sum_{j=1}^D p_j^2 \leq d$, and $p_j \in [l_j, u_j]$ for all $j \in [D]$. Moreover, the zero set of h is bounded. Indeed, if $h = 0$, then $p \in Q$ and $\eta^2 \leq d$. Moreover, $\alpha_j^2 = p_j - l_j$ and $\beta_j^2 = u_j - p_j$, so $\alpha_j^2 + \beta_j^2 = u_j - l_j = 2/J_p$.

Since $J_p l_j$ and $J_p u_j$ are integers and $|l_j|, |u_j| \leq 1$, the polynomial h has integer coefficients of absolute value at most $d^2 + 4(D+1)J_p^2$, and its number of monomials depends only on d. Thus, the total encoding size of the coefficients is at most $C_h(d)B_{J_p}$, for some parameter-dependent constant C_h. Proposition 26 with $\nu = 4$ and $k = 3D + 1$ concludes the proof. □

We now bound the bit complexity of FixHardVariables.

Proposition 28. *Let B_t be the encoding size of $t \in \mathbb{Q} \cap (0, 1)$. There is an implementation of* FixHardVariables *whose bit complexity $T_{\text{FixHardVariables}}$ satisfies*

$$T_{\text{FixHardVariables}} \leq F_{\text{FixHardVariables}}(d) \left[(B_t + 1)^c + J_p^{D_d} B_{J_p}^c \right]$$

for some parameter-dependent constant $F_{\text{FixHardVariables}}$ and an absolute constant c. Moreover, each returned solution p belongs to a field $\mathbb{K}_p$ such that

$$[\mathbb{K}_p : \mathbb{Q}] \leq 6 \cdot 5^{3D}, \quad B_p \leq C_p(d)B_{J_p},$$

where B_p is the representation size of p in $\mathbb{K}_p$, and C_p is a parameter-dependent constant.

Proof. Computing J_p has bit complexity $F_1(d)(B_t+1)^c$, for some parameter-dependent constant F_1 and an absolute constant c. By (20), the number of boxes is at most $J_p^{D_d}$, since $D \le D_d = d(d-1)^2$. Together with Lemma 27, this gives the bound on $T_{\textsc{FixHardVariables}}$. The bounds on $[\mathbb{K}_p : \mathbb{Q}]$ and B_p follow from Proposition 26 with $\nu = 4$ and $k = 3D+1$, applied to the polynomial h in the proof of Lemma 27. □

Next, we bound the bit complexity of the ADP, starting with one branch.

Lemma 29. *For a fixed branch of* LinearOracle*, let B_{Gv} be the representation size of the Gram entries G_{il} and objective values v produced in* SelectBranch*. Then*

$$B_{\mathrm{Gv}} \le C_{\mathrm{Gv}}(d)(U + B_{\mathrm{in}})$$

for some parameter-dependent constant C_{Gv}.

Proof. Fix one branch and one output p of FixHardVariables for this branch. By Proposition 28, the vector p belongs to a field $\mathbb{K}_p = \mathbb{Q}(\gamma)$ whose degree $[\mathbb{K}_p : \mathbb{Q}]$ is bounded by a function of d, and the encoding sizes of the coordinates of p are at most $C_p(d)B_{J_p} \le C(d)U$. The ADP performs only arithmetic operations and comparisons, so no further field extension is needed.

To bound the encoding sizes of G_{il} and v, let Q_A be the product of the denominators of the entries of A, and let $Q_V = M2^{J_0}$ be a common denominator of the numbers in free-variable finite set (6). Then

$$\log_2(Q_A) \le \mathcal{O}(B_{\mathrm{in}}), \qquad \log_2(Q_V) \le C(d)U,$$

for some parameter-dependent constant C. By Proposition 28 and arithmetic in $\mathbb{K}_p$, the rational coordinate coefficients of all fixed hard-variable values and their pairwise products admit a common denominator Q_H with $\log_2(Q_H) \le C(d)U$. We claim that the contributions to the Gram matrix and objective values have a common denominator

$$Q_* = Q_A Q_V^2 Q_H, \qquad \log_2(Q_*) \le C(d)(B_{\mathrm{in}} + U),$$

and that the encoding sizes of the numerators satisfy the same bound. Indeed, every Gram entry and every objective value has the form

$$G_{il} = \sum_{j \in I} y_{ji} y_{jl}, \qquad v = \sum_{j \in I} \sum_{i=1}^{d} A_{ji} y_{ji},$$

for some subset $I \subseteq [n]$, where y_{ji} denotes the (j,i) entry of the partial solution, so Q_* is a common denominator. These sums have at most nd terms, so the encoding sizes of their numerators increase by at most $\mathcal{O}(\log(nd))$, which is absorbed into $C(d)(B_{\mathrm{in}} + U)$. Since a merge combines disjoint sets of rows and KeepBest only selects among candidates, these bounds hold throughout the branch. □

Proposition 30. *There is an implementation of* SelectBranch *whose bit complexity $T_{\textsc{SelectBranch}}$ satisfies*

$$T_{\textsc{SelectBranch}} \le F_{\textsc{SelectBranch}}(d) n t^{-(d^2+3d+2+c)} (B_{\mathrm{in}}+1)^c U^{d^2+5d+4+3c}$$

for some parameter-dependent constant $F_{\textsc{SelectBranch}}$ and an absolute constant c.

Proof. Fix one branch and one output p of FixHardVariables for this branch. By Proposition 28, the vector p belongs to a field $\mathbb{K}_p$ whose degree $[\mathbb{K}_p : \mathbb{Q}]$ is bounded by a function of d. Every value computed by the ADP belongs to $\mathbb{K}_p$. By Lemma 29, every Gram entry G_{il} and objective value v has encoding size at most $C_{\mathrm{Gv}}(d)(B_{\mathrm{in}} + U)$, also after merges and dictionary updates.

It remains to account for the rational discretization thresholds. Let $l_{\max}$ be as in (4), and let $k_{\max}$ be the smallest integer k with $\kappa_G(1+\theta)^k \ge 4d$. These are the largest exponents in the capacity and Gram-bin thresholds, and they satisfy

$$\max\{l_{\max}, k_{\max}\} \le C_{\exp}(d) t^{-1} U^2$$

for some parameter-dependent constant C_{exp}. Indeed, by Table 1, we have $H_{\mathrm{tree}} = \lceil \log_2(n) \rceil$, $\delta = t/(8(d+1))$, and $\tau = \delta/(128d^2)$. The bound follows from $\log(1+x) \geq x/2$ for $x \in (0,1)$. Since the encoding sizes of K_0, $1+\alpha$, $1+\theta$, and κ_G are at most $C_{\mathrm{base}}(d)(B_{\mathrm{in}}+U)$ for some parameter-dependent constant C_{base}, all thresholds have encoding sizes at most

$$B_{\mathrm{thre}} := C_{\mathrm{thre}}(d) t^{-1} (B_{\mathrm{in}} + 1) U^3$$

for some parameter-dependent constant C_{thre}. Consequently, arithmetic in $\mathbb{K}_p$, comparisons, bin assignments, and dictionary operations can each be implemented in at most

$$F_{\mathrm{work}}(d)(B_{\mathrm{thre}} + 1)^c$$

bit operations, for some parameter-dependent constant F_{work} and an absolute constant c, by Lemma 21. Here, by means of binary search, a Gram matrix bin assignment or a sparsity capacity usage rounding consists of a number of sign tests that is logarithmic in $k_{\max}$, $|U_c|$, and $1/\theta$. Each sign test compares an element of $\mathbb{K}_p$ with a precomputed rational threshold and creates no new values, so Lemma 21 applies to each test with $L = 1$. The number of these tests is at most $C_{\mathrm{test}}(d)U$ for some parameter-dependent constant C_{test}, and it is absorbed by enlarging F_{work} and c. The construction of $\mathcal{V}$ and of the thresholds is absorbed into the same bound. Finally, by (26), the ADP performs at most $W_{\mathrm{ADP}} \leq f_{\mathrm{bw}}(d) n t^{-(d^2+3d+2)} U^{d^2+5d+4}$ operations. Multiplying this bound by the bit complexity of one operation concludes the proof. □

Next, we bound the bit complexity of LinearOracle. We first bound the cost of comparing the objective values of two candidates that may lie in different fields.

Lemma 31. *Let $x \in \mathbb{K}_1$ and $y \in \mathbb{K}_2$ with $[\mathbb{K}_1 : \mathbb{Q}] \leq \Delta_1$, $[\mathbb{K}_2 : \mathbb{Q}] \leq \Delta_2$. Let B_{val} be an upper bound on the representation sizes of x in $\mathbb{K}_1$ and of y in $\mathbb{K}_2$. Then the sign of $x - y$ can be determined with bit complexity*

$$F_{\mathrm{cmp}}(\Delta_1, \Delta_2)(B_{\mathrm{val}} + 1)^c$$

for some parameter-dependent constant F_{cmp} and an absolute constant c.

Proof. The compositum $\mathbb{K}_1\mathbb{K}_2$, that is, the smallest field that contains $\mathbb{K}_1$ and $\mathbb{K}_2$, has degree at most $\Delta_1\Delta_2$ [22, Chapter V, Exercise 7]. One can compute a primitive element of $\mathbb{K}_1\mathbb{K}_2$ via resultants, its minimal polynomial via factorization [23], and the representations of x and y in $\mathbb{K}_1\mathbb{K}_2$, as in the proof of Lemma 21 (see also [6, Chapter 12]). Then, by sign determination [6, Section 10.3], the sign of $x - y$ can be determined exactly in $F_{\mathrm{cmp}}(\Delta_1, \Delta_2)(B_{\mathrm{val}} + 1)^c$ bit operations, for some parameter-dependent constant F_{cmp} and an absolute constant c. □

Proposition 32. *Let B_{in} be the encoding size of the input (A, s, t), including the dimensions. There is an implementation of* LinearOracle *whose bit complexity $T_{\textsc{LinearOracle}}$ satisfies*

$$T_{\textsc{LinearOracle}} \leq F_{\textsc{LinearOracle}}(d) n t^{-(d^3-d^2+4d+2+c)} (B_{\mathrm{in}} + 1)^c U^{d^2+5d+4+3c}$$

for some parameter-dependent constant $F_{\textsc{LinearOracle}}$ and an absolute constant c.

Proof. The cases $A = O$ and $d = 1$ can be handled within the stated bound. Thus, we assume that $A \neq O$ and $d \geq 2$. By (27), the number of computational tasks of LinearOracle is at most

$$W_{\textsc{LinearOracle}} \leq f_{\mathrm{lin}}(d) n t^{-\beta_d} U^{\gamma_d},$$

where $\beta_d := D_d + d^2 + 3d + 2 = d^3 - d^2 + 4d + 2$ and $\gamma_d := d^2 + 5d + 4$. This bound includes the calls to FixHardVariables for all boxes Q and the calls to ExactFinalize for all accepted root candidates.

We first bound the input encoding size of ExactFinalize. Consider an accepted candidate solution Y reconstructed at the root of the ADP. By construction, Y contains at most D_d fixed hard-variable entries, while all remaining entries belong to the finite rational set $\mathcal{V}$. By Proposition 28, the fixed hard-variable values belong to a common algebraic field whose degree is bounded in terms of d, and their encoding sizes are at most $C_p(d) B_{J_p} \leq C_{\mathrm{hv}}(d) U$ for some parameter-dependent constant C_{hv}. Moreover, every element of $\mathcal{V}$ has encoding size at most $C_V(d) U$ for some parameter-dependent constant C_V. By Lemma 29, the entries

of $G^0 = Y^\top Y$ have encoding sizes at most $C_{\mathrm{Gv}}(d)(B_{\mathrm{in}} + U)$. Since B_{in} is the total encoding size of the numerical input and A has n rows, we have $n \le B_{\mathrm{in}}$. Therefore, the total input encoding size required by ExactFinalize satisfies

$$B_{\mathrm{tot-enc}} \le C_0(d)(B_{\mathrm{in}} + nU) \le C_1(d)(B_{\mathrm{in}} + 1)U,$$

for some parameter-dependent constants C_0 and C_1.

Next, by Propositions 25, 28, and 30, together with Lemmas 29 and 31, the bit complexity of each computational task is at most

$$T_{\mathrm{task}} \le F_{\mathrm{task}}(d)(B_{\mathrm{task}} + 1)^c,$$

for some parameter-dependent constant F_{task} and an absolute constant c, where B_{task} is the input encoding size of the task. By the proofs of these propositions, we have

$$B_{\mathrm{task}} \le C_{\mathrm{task}}(d)t^{-1}(B_{\mathrm{in}} + 1)U^3$$

for some parameter-dependent constant C_{task}. Indeed, by Proposition 28, the sampled hard variables have encoding sizes bounded by $C_2(d)U$. By Lemma 29, the Gram entries and objective values in the ADP have encoding sizes bounded by $C_3(d)(B_{\mathrm{in}} + U)$. Moreover, the proof of Proposition 30 upper bounds the encoding sizes of the discretization thresholds by $C_4(d)t^{-1}(B_{\mathrm{in}} + 1)U^3$. By the bound on $B_{\mathrm{tot-enc}}$ above, the input of ExactFinalize has total encoding size at most $C_1(d)(B_{\mathrm{in}} + 1)U$. Here, C_2, C_3, and C_4 are parameter-dependent constants. These bounds give the bound on B_{task}.

Finally, by Propositions 28 and 25, each candidate lies in a field whose degree is bounded in terms of d. By Lemma 31, comparisons between different candidates use temporary field extensions, while the better candidate is retained in its original field. Therefore, these comparisons do not accumulate field extensions. We conclude the proof by $T_{\textsc{LinearOracle}} \le W_{\textsc{LinearOracle}} T_{\mathrm{task}}$. □

Since $n \le B_{\mathrm{in}}$ and U is logarithmic in n, d, and $1/t$, the bound in Proposition 32 is polynomial in B_{in} and $1/t$ for fixed d. Together with Proposition 18, this proves Theorem 4.

Finally, we prove Theorem 10.

Proof of Theorem 10. Algorithm 3 enumerates at most

$$M_W \le \left(2\left\lceil \frac{4rd}{\epsilon} \right\rceil + 1\right)^{rd} \le \left(\frac{8rd+3}{\epsilon}\right)^{rd}$$

matrices W. For each W, the entries have the common denominator $J_W = \lceil 4rd/\epsilon \rceil$. To bound the total encoding size B_A of $A = RW$, write each entry of R as $R_{jk} = a_{jk}/b_{jk}$, where $a_{jk} \in \mathbb{Z}$ and b_{jk} is a positive integer. Since every entry of W has denominator J_W, we can write

$$A_{ji} = \frac{1}{J_W} \sum_{k=1}^{r} \frac{a_{jk} w_{ki}}{b_{jk}},$$

where $w_{ki} = J_W W_{ki} \in \mathbb{Z}$ and $|w_{ki}| \le J_W$. Summing the encoding sizes over all rows and columns gives

$$B_A \le C_0(r,d)\left(B_{\mathrm{SPCA}} + n(1 + \log_2(J_W + 1))\right) \le C_{\mathrm{SPCA}}(r,d)(B_{\mathrm{SPCA}} + 1)U_\epsilon,$$

for some parameter-dependent constants C_0 and C_{SPCA}, where the last inequality follows from $n \le B_{\mathrm{SPCA}}$ and $\log_2(J_W + 1) \le C_1(r,d)U_\epsilon$ for some parameter-dependent constant C_1. Hence the input (A, s, t) of LinearOracle has encoding size $B_{\mathrm{in}} \le B_A + B_{\mathrm{SPCA}} + \mathcal{O}(U_\epsilon)$.

Applying Proposition 32 with $t = \epsilon/4$, each call to LinearOracle has bit complexity at most

$$F'_{\textsc{LinearOracle}}(r,d)n\epsilon^{-(d^3-d^2+4d+2+c)}(B_{\mathrm{SPCA}} + 1)^c U_\epsilon^{d^2+5d+4+4c},$$

for some parameter-dependent constant $F'_{\textsc{LinearOracle}}$ and an absolute constant c. Computing RW, obtaining the returned objectives, and comparing them exactly using Lemma 31 are dominated by this bound. The

value $\left\|R^\top X\right\|_F^2$ is computed from the rd sums $\sum_{j=1}^{n} R_{jk}X_{ji}$, each with n terms. Clearing denominators as in the proof of Lemma 29, their encoding sizes are polynomial in B_{SPCA} and U_ϵ, because $n \le B_{\mathrm{SPCA}}$. Indeed, by Proposition 28 and Lemma 24, the degrees of the field extensions are bounded in terms of d, and by Lemma 24 and the bound on $B_{\mathrm{tot-enc}}$ in the proof of Proposition 32, the encoding sizes are bounded by a polynomial in B_{SPCA} and U_ϵ. Thus, after enlarging the parameter-dependent constant and the absolute constant c if necessary, these additional computations are absorbed into the stated bound. Multiplying by M_W concludes the proof. □

Dependence on n. The proof of Theorem 10 gives a finer bound on the dependence on n. For fixed r and d, Algorithm 3 processes at most $F(r,d)\, n\, \epsilon^{-(rd+d^3-d^2+4d+2)} U_\epsilon^{d^2+5d+4}$ states, merges, and finalization steps, by (27) and the bound on the number of matrices W. Each of these steps uses a number of arithmetic operations and sign tests, in a number field whose degree is bounded in terms of d, that is at most a function of d times U_ϵ. The number of square-root extensions in ExactFinalize and of exact comparisons of candidates in different fields is polylogarithmic in n. Thus the number of steps is linear in n up to a polylogarithmic factor. For comparison, a straightforward dynamic program that processes the rows one at a time would need exact sparsity capacity usages of order n^d and a Gram grid of mesh of order $\epsilon/(nd^2)$, since its rounding errors accumulate over all n rows. This gives a number of states of order $n^{d+d(d+1)/2}$ at each row.