

Learning Arithmetic Formulas in the Presence of Noise: A General Framework and Applications to Unsupervised Learning

Pritam Chandra

Microsoft Research

t-pchandra@microsoft.com

Ankit Garg

Microsoft Research

garga@microsoft.com

Neeraj Kayal

Microsoft Research

neeraka@microsoft.com

Kunal Mittal

Princeton University

kmittal@cs.princeton.edu

Tanmay Sinha

Microsoft Research

t-tsinha@microsoft.com

November 13, 2023

Abstract

We present a general framework for designing efficient algorithms for unsupervised learning problems, such as mixtures of Gaussians and subspace clustering. Our framework is based on a meta algorithm that learns arithmetic formulas in the presence of noise, using lower bounds. This builds upon the recent work of Garg, Kayal and Saha (FOCS '20), who designed such a framework for learning arithmetic formulas without any noise. A key ingredient of our meta algorithm is an efficient algorithm for a novel problem called *Robust Vector Space Decomposition*. We show that our meta algorithm works well when certain matrices have sufficiently large smallest non-zero singular values. We conjecture that this condition holds for smoothed instances of our problems, and thus our framework would yield efficient algorithms for these problems in the smoothed setting.

Contents

1	Introduction	1
1.1	Overview - Arithmetic Formula Learning algorithm.	5
1.2	Overview - Vector Space Decomposition algorithm	8
1.3	Application 1: Subspace Clustering.	11
1.4	Application 2: Learning Mixtures of Gaussians	14
1.5	Conclusion and Future Directions	17
2	Preliminaries	19
3	Robust Recovery from Scaling Maps (RRSM)	21
3.1	RRSM: Algorithm for the Noiseless Case	22
3.2	RRSM: Algorithm for the Robust Case	22
4	Robust Vector Space Decomposition (RVSD)	25
4.1	RVSD: Algorithm for the Noiseless Case	26
4.2	RVSD: Algorithm for the Robust Case	27
4.3	RVSD: Using a Common Tuple of Operators on a Larger Space	29
A	Linear Algebra and Probability	33
A.1	Matrices and Subspaces	33
A.1.1	Matrix Norms	33
A.1.2	Pseudo-Inverse	33
A.1.3	Distances Between Subspaces	33
A.2	Matrix Perturbation Bounds	34
A.2.1	Perturbation bounds for Singular values and Singular vectors	34
A.2.2	Perturbation bounds for Eigenvalues and Eigenvectors	36
A.2.3	Perturbation bounds for Pseudo-Inverse	38
A.3	Anti-Concentration of Gaussian Linear Forms	38
B	Subspace Clustering	39
B.1	Robust Recovery from Symmetric Tensor Power	39
B.1.1	RRSTP: Noiseless Case	39
B.1.2	RRSTP: Robust Case	40
B.2	Subspace Clustering: Noiseless Case	42
B.2.1	Structure of The Adjoint Algebra	43
B.2.2	Completing the Proof	46
B.3	Subspace Clustering: Robust Case	46
B.3.1	Closeness of Subspaces	48
B.3.2	Using RVSD and RRSTP	49
B.4	Singular Value Analysis for The Adjoint Algebra	50
B.5	Smoothed Analysis of Subspace Clustering	51
C	Learning Mixtures of Gaussians	54
D	Learning Arithmetic Circuits in the Presence of Noise	58
D.1	Applying Robust Vector Space Decomposition	60
D.2	Recovering the Polynomials	61

E	Analysis of the RRSM Algorithm	62
E.1	Notation	62
E.2	Condition Number Relations	63
E.3	Canonical Decomposition and Random Sampling	63
E.4	Bounding the Spectral Gap	64
E.5	Analysis of the Algorithm	65
E.5.1	Action on The Space of Scaling Maps	65
E.5.2	Perturbation Bound	66
E.5.3	Recovering the Components	67
E.5.4	Runtime Analysis	67
E.6	Direct Sum Property	68
F	Analysis of the RVSD Algorithm	68
F.1	Notation	69
F.2	Perturbation Bound on the Adjoint Algebra	69
F.3	Applying RRSM and Recovering The Component Subspaces	71
F.4	Runtime Analysis	71
G	Singular Values of the Adjoint Algebra Operator for Subspace Clustering	72
G.1	Adjoint Algebra Operators corresponding to Partial Derivatives on Tensored Spaces	72
G.1.1	Matrix Representations of Derivatives and Shifts	73
G.1.2	Adjoint Algebra Operator in a Special Inner Product	74
G.1.3	Singular Values of the Φ Operator	77
G.2	Adjoint Algebra Operator for Subspace Clustering	80

1 Introduction

Unsupervised learning involves discovering hidden patterns and structure in data without using any labels or direct human supervision. Here we consider data that has a nice mathematical structure or is generated from a mathematically well-defined distribution. An example of the former is when the data points can be grouped into meaningful clusters based on some similarity patterns and the goal is to find the underlying clusters. An example of the latter is mixture modeling, which assumes that the data is generated from a mixture of succinctly described probability distributions, such as Gaussian distributions, and the goal is to learn the parameters of these distributions from samples. A general framework for solving many unsupervised learning problems is the method of moments, which leverages the statistical moments¹ of the data to infer the underlying structure or the underlying parameters of the model. For many unsupervised learning problem scenarios wherein the underlying data has some nice mathematical structure, the moments of the data are well-defined functions of the parameters. Heuristic arguments then suggest that the converse should typically hold, i.e. the parameters of the structure/distribution are *typically* uniquely determined by a few low order moments of the data. In this broad direction, the main challenge then is to design algorithms to (approximately) recover the underlying parameters from the (empirical) moments². We further want the algorithm to be *efficient*, noise-tolerant (i.e. work well even when the moments are known only approximately rather than exactly) and are even *outlier-tolerant* (i.e. work well even when a few data points do not conform to the underlying structure/distribution). But even the simplest problems in this area tend to be NP-hard and remain so even when there is no noise and no outliers. So one cannot realistically hope for an algorithm with provable worst-case guarantees. But what one can hope are algorithms that are guaranteed to typically work well, i.e. either for *random* problem instances or even more desirably for instances chosen in a smoothed fashion. Accordingly, many different algorithms have been designed for each such problem in unsupervised learning with varying levels of efficiency, noise-tolerance, outlier-tolerance and provable guarantees. In this work we give a *single meta-algorithm* that applies to many such unsupervised learning problems. The starting point of our work is the observation that many such problems reduce to the task of learning an appropriate subclass of arithmetic formulas.

Connecting unsupervised learning to arithmetic complexity. We now give a few more details of how such a reduction works for the setting in which the data points are drawn from a distribution having a nice mathematical structure. Let \mathcal{D} be a distribution over points in \mathbb{R}^n . We introduce n formal variables (x_1, x_2, \dots, x_n) and denote it as \mathbf{x} . For a suitably chosen integer $d \geq 2$, form a degree- d polynomial $f(\mathbf{x})$

¹Recall that moments are measures of the shape and variability of a data set. They are used to describe the the location and dispersion of the data. When the dataset consists of a collection of points

$$A = \{\mathbf{a}_i = (a_{i1}, a_{i2}, \dots, a_{in}) \in \mathbb{R}^n \mid i \in [N]\},$$

some examples of (low-order) moments are $\mathbb{E}_{\mathbf{a}_i \in A} [a_{i1}]$, $\mathbb{E}_{\mathbf{a}_i \in A} [a_{i1} \cdot a_{i2}]$, etc.

²In scenarios where the data is a finite sample drawn from a distribution \mathcal{D} over \mathbb{R}^n , the empirical moments (which can be very easily and efficiently computed) are estimates of, but *not* equal to, the true underlying moments.

which encodes the d -th order moments³ of the distribution in some suitable way. For example, in some applications the coefficient of a monomial of $f(\mathbf{x})$ is simply (a canonically scaled version of) the corresponding moment. At this point, such a formal polynomial is a mere bookkeeping device for the d -th order moments of \mathcal{D} . For many nice, well-structured distributions such as mixtures of Gaussians, however this polynomial (or variants thereof) turns out to have a remarkable property - it can be computed/represented by a small arithmetic formula! Special cases of this remarkable phenomenon were noted earlier when it was observed that many problems in (unsupervised) learning reduce to the problem of learning set-multilinear depth-three formulas, better known as tensor decomposition. Such connection(s) inspired a whole body of work on tensor decomposition with applications including independent component analysis, learning Hidden Markov Models, learning special cases of mixtures of Gaussians, latent Dirichlet allocation, dictionary learning, etc. (cf. the surveys [Vij20,KB09,AGH⁺14]).

Noise-tolerance. Notice however that we are given a finite set of points sampled from the distribution \mathcal{D} , so we do not have the (d -th order) moments of \mathcal{D} exactly but only approximately. Thus for such applications we need the algorithm for learning arithmetic formulas to also be *noise-tolerant*, i.e. given a polynomial $\tilde{f}(\mathbf{x})$ that is *close to*⁴ a polynomial $f(\mathbf{x})$ that has a small arithmetic formula ϕ , we want to learn/reconstruct a arithmetic formula $\tilde{\phi}$ from the same subclass as that of ϕ whose output polynomial is close to $\tilde{f}(\mathbf{x})$ (and therefore to $f(\mathbf{x})$ as well). Recently, [GKS20] gave a meta-algorithm for learning many different subclasses of formulas including the ones relevant for unsupervised learning (assuming that certain nondegeneracy conditions hold). But it has one important shortcoming that was also pointed out in [BDJ⁺22]: the techniques of [GKS20] were algebraic and it was unclear if they could handle noise arising out of the fact that the moments are known only approximately and not exactly. Qualitatively, our main result builds upon and suitably adapts the algorithm [GKS20] to make it noise-tolerant. Quantitatively, in the noisy setting, we provide bounds on the quality of the output of our algorithm that depend on singular values of certain matrices that underlie the algorithm. We expect that for most applications, the relevant singular values would be well-behaved for random instances and maybe even for smoothed/perturbed worst-case instances. If so, our algorithm would work and yield good quality outputs on such instances. Accordingly, we then go on to analyze the singular values of the relevant matrices pertaining to subspace clustering⁵. We also expect (suitable adaptations of) our algorithm to be tolerant to the presence of a few outliers but we do not pursue this direction here and leave it for future work.

Illustrative example - mixtures of Gaussians. Let us make the above discussion concrete via the example of learning mixtures of Gaussians which in itself is a very well-studied problem with history going back to more than a hundred years. Suppose we

³We are making the mild assumption here that the d -th order moments of \mathcal{D} are bounded.

⁴Under a natural notion of distance between a pair of polynomials akin to Euclidean distance between the coefficient vectors of the pair of polynomials - see section 2.

⁵A recent work [BHKX22] analyzed the singular values of matrices arising in a related (but also different) algorithm that was tailor-made for the mixtures of (zero-mean) Gaussians and verified that for random instances the singular values are indeed well-behaved.

are given a dataset consisting of a finite set of points $A \subset \mathbb{R}^n$

$$A = \{\mathbf{a}_i = (a_{i1}, a_{i2}, \dots, a_{in}) \in \mathbb{R}^n \mid i \in [N]\}. \quad (1)$$

The points are drawn independently at random from an unknown mixture of s Gaussians $\mathcal{D} := \sum_{i=1}^s w_i \mathcal{N}(\boldsymbol{\mu}_i, \Sigma_i)$, which means that the i -th component of the mixture has weight⁶ $w_i \in [0, 1]$, mean $\boldsymbol{\mu}_i \in \mathbb{R}^n$ and covariance matrix $\Sigma_i \in \mathbb{R}^{n \times n}$. Our goal is to estimate the parameters w_i and $\boldsymbol{\mu}_i$ and Σ_i ($i \in [s]$) from the given samples/data A . Let $\mathbf{x} = (x_1, x_2, \dots, x_n)$ be a tuple of formal variables and consider the polynomial $f(\mathbf{x}) := \mathbb{E}_{\mathbf{a} \sim \mathcal{D}} [\langle \mathbf{x}, \mathbf{a} \rangle^d]$. It is (a scalar multiple of) a *slice* of the formal moment generating function defined as $\mathbb{E}_{\mathbf{a} \sim \mathcal{D}} [\exp(\langle \mathbf{x}, \mathbf{a} \rangle)]$. Notice that the coefficients of a given monomial (over \mathbf{x}) in $f(\mathbf{x})$ equals the corresponding moment of the distribution (upto some canonical scaling). Then in this case, $f(\mathbf{x})$ has the following small formula⁷:

$$f(\mathbf{x}) = \sum_{i \in [s]} w_i G_d(\ell_i(\mathbf{x}), Q_i(\mathbf{x})),$$

where $\ell_i(\mathbf{x}) := \langle \boldsymbol{\mu}_i, \mathbf{x} \rangle$, $Q_i(\mathbf{x}) := \frac{1}{2} \mathbf{x}^T \Sigma_i \mathbf{x}$ and G_d is a fixed bivariate polynomial depending on d . In the zero-mean case (i.e. when $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2 = \dots = \boldsymbol{\mu}_s = \mathbf{0}$), the formula for $f(\mathbf{x})$ is

$$f(\mathbf{x}) = \sum_{i \in [s]} \frac{d!}{(d/2)!} w_i Q_i(\mathbf{x})^{d/2}$$

when d is even (and 0 if d is odd). In this way, if the sample size was infinite (or equivalently that if we knew the true moments of the distribution), learning mixtures of Gaussians would reduce to the problem of learning/reconstructing the subclass of arithmetic formulas indicated by the rhs of the above expression for $f(\mathbf{x})$. But we don't have access to the exact moments. Using the empirical moments, we can get hold of an approximate version of f ,

$$\tilde{f}(\mathbf{x}) := \mathbb{E}_{\mathbf{a} \sim A} [\langle \mathbf{x}, \mathbf{a} \rangle^d] = \frac{1}{N} \sum_{i \in [N]} [\langle \mathbf{x}, \mathbf{a}_i \rangle^d].$$

We will have $\tilde{f}(\mathbf{x}) = f(\mathbf{x}) + \eta(\mathbf{x})$ for a noise polynomial $\eta(\mathbf{x})$ whose magnitude will be inversely proportional to square root of the number of samples N . In this way, learning mixtures of Gaussians reduces to the problem of reconstructing the indicated subclass of arithmetic formulas in the presence of noise.

Learning arithmetic formulas in the presence of noise - problem formulation. The above discussion motivates us to consider the problem of learning (arbitrary subclasses of) arithmetic formulas in the presence of noise. In many practical settings the output gate of the underlying formula is a (generalized⁸) addition gate so that the problem can be formulated as follows. We are given a polynomial $\tilde{f}(\mathbf{x})$ of the form $\tilde{f}(\mathbf{x}) := T_1(\mathbf{x}) + \dots + T_s(\mathbf{x}) + \eta(\mathbf{x})$, for *structured* polynomials $T_i(\mathbf{x})$'s and a noise polynomial $\eta(\mathbf{x})$. Our goal is to approximately recover each summand $T_i(\mathbf{x})$. For example,

⁶The weights satisfy $\sum_{i \in [s]} w_i = 1$.

⁷This formula for $f(\mathbf{x})$ can be inferred from the fact that for a single Gaussian distribution $N(\boldsymbol{\mu}, \Sigma)$, its moment generating function is in fact equal to $\exp(\mathbf{x}^T \cdot \boldsymbol{\mu} + \frac{1}{2} \mathbf{x}^T \cdot \Sigma \cdot \mathbf{x})$.

⁸A generalized addition gate can compute any fixed linear combination of its inputs.

for the case of mixture of spherical Gaussians we would have $T_i(\mathbf{x}) = w_i \cdot \langle \boldsymbol{\mu}_i, \mathbf{x} \rangle^3$ (see Remark 1). For the case of mixture of zero-mean Gaussians we would have $T_i(\mathbf{x}) = w_i \cdot Q_i(\mathbf{x})^{d/2}$ and so on. In the noiseless setting, i.e. when $\eta(\mathbf{x}) = 0$, the paper [GKS20] designed a meta-algorithm applicable to learning many interesting subclasses using a general framework exploiting lower bound techniques in arithmetic complexity theory. The algorithm worked under certain relatively mild non-degeneracy assumptions. However, their algorithm had some algebraic components and it was not clear how to design an algorithm in the noisy case when the noise polynomial $\eta(\mathbf{x})$ is non-zero.⁹ Our main contribution is that we show how to modify the general framework in [GKS20] to the noisy setting. We also show how to use this framework to design efficient algorithms for two well studied problems in unsupervised learning: mixtures of (zero mean) Gaussians and subspace clustering.

Remark 1. (a). **Simpler reductions.** *The ability to handle arbitrary subclasses of arithmetic formulas not only yields a common (meta) algorithm that applies to a wide variety of problems in unsupervised learning but it also often makes the reductions simpler. For example, in the discussion above the reduction of learning mixtures of arbitrary Gaussians to learning the appropriate subclass of arithmetic formulas is perhaps simpler than the reduction of learning mixtures of spherical Gaussians¹⁰ to tensor decomposition. We sketch this reduction now. Consider $f(\mathbf{x}) := \mathbb{E}_{\mathbf{a} \sim \mathcal{D}} \left[\langle \mathbf{x}, \mathbf{a} \rangle^3 - 3 \left(\sum_{i \in [n]} x_i^2 \right) \cdot \langle \mathbf{x}, \mathbf{a} \rangle \right]$. When \mathcal{D} is a mixture of spherical Gaussians, expanding and simplifying this expression, we can get that $f(\mathbf{x}) = \sum_{i=1}^s w_i \langle \boldsymbol{\mu}_i, \mathbf{x} \rangle^3$.*

(b). **Mixtures of general Gaussians.** *We expect that our algorithm can be extended to general mixtures of Gaussians (different means and/or covariance matrices) but its analysis will likely get much more cumbersome, so we avoid this more general case for the sake of simplicity.*

(c). **Handling outliers.** *The ability to handle arbitrary subclasses of arithmetic formulas can also allow the algorithm to be tolerant to the presence of outliers. To see this, consider the case of zero-mean Gaussians and suppose that the given set of data points A contains a subset $\hat{A} \subset A$ of outliers of size $\hat{N} \ll N$. In that case the empirical moment polynomial $\tilde{f}(\mathbf{x})$ would have the following structure:*

$$\tilde{f}(\mathbf{x}) = \frac{N - \hat{N}}{N} \cdot \frac{d!}{(d/2)!} \cdot \left(\sum_{i \in [s]} w_i Q_i(\mathbf{x})^{d/2} \right) + \frac{1}{N} \cdot \left(\sum_{\mathbf{a}_j \in \hat{A}} (\mathbf{x} \cdot \mathbf{a}_j)^d \right) + \eta(\mathbf{x}).$$

We expect that our algorithm can be adapted to learn the class of formulas corresponding to the right side of the above expression however the analysis of such an algorithm can get cumbersome. For the sake of keeping the length of this paper to within reasonable bounds, we do not do the analysis of the outlier tolerance of our algorithm.

(d). **Other mixtures models.** *The connection between learning mixtures of Gaussians and learning an appropriate subclass of arithmetic formulas arose out of the fact that (any*

⁹In most settings, one would like the running time of the algorithm to be inverse polynomial in the magnitude of the noise, to have a polynomial dependence on the number of samples in the final learning problem.

¹⁰A spherical Gaussian is one where the covariance matrix Σ_i is the identity matrix.

slice of) the moment generating function of a multivariate Gaussian has a simple algebraic expression. For some other distributions also the (slices of) moment generating function or some other related function like the cumulant generating function or the characteristic function have a nice algebraic expression and we can expect our approach to be applicable for such mixtures also.

- (e). **Mixtures of structured point sets and those sampled from probability distributions.** Consider a set of points $A \subset \mathbb{R}^n$ that can be partitioned into two subsets $A = A_1 \uplus A_2$ such that A_1 is some structured set of points (such as being contained in the union of a small number of low-dimensional subspaces for example) and A_2 is chosen from some mixture model (such as being chosen from a mixture of Gaussians for example). When say the moment polynomials of both the structured set A_1 and the sampled set A_2 admit small formulas from a tractable subclass of formulas, we can expect our methods to apply. In particular, we expect (a suitable adaptation of) our algorithm to be to handle the case where points in A_1 are chosen from a union of low-dimensional subspaces in an non-degenerate way without conforming to any nice distribution and points in A_2 conform to a (mixture of) Gaussians. We leave the task of handling such mixed datasets and analyzing the relevant algorithms as a possible direction for future work.
- (f). **Potential application - Topic Modeling.** It turns out that there are some other problems in unsupervised learning which reduce to robustly learning an appropriate subclass of arithmetic formulas. We expect that a suitable instantiation/adaptation of our algorithm should apply for these applications but we do not pursue these applications here and leave it as a direction for future work. One such problem is called topic modelling. It is known that learning some simple topic models reduce to tensor decomposition. It turns out that learning some general topic models as proposed in [Wal06] reduce to the problem of learning set-multilinear formulas of larger depth.
- (g). **Potential application - Learning (Mixtures of) Polynomial Transformations.** Another such application is the problem of learning polynomial transformations as studied in [CLLZ23] which also reduces to learning a certain subclass of arithmetic formulas¹¹. The generality of our approach makes us expect that it should apply to this task also as well as to its generalizations like learning **mixtures of polynomial transformations**. We do not pursue this potential application here but leave it as a direction for future work.

1.1 Overview - Arithmetic Formula Learning algorithm.

Background. Arithmetic formulas are a natural model of computing polynomials using the basic operations of addition (+) and multiplication (\times). A natural problem about arithmetic formulas is that of learning: given a polynomial $f(\mathbf{x})$ ¹², find the smallest (or somewhat small) arithmetic formula computing $f(\mathbf{x})$. We consider formulas in their alternating normal form: i.e. the formula consists of alternating layers of addition and multiplication gates. The learning problem boils down to recovering

¹¹The work of [CLLZ23] does not state it this way but this can be inferred from the observations underlying their work.

¹²There are various input models all of which lead to interesting questions. Some of the common ones are as a black box or described explicitly as a list of coefficients.

the polynomials computed at each child of a node v given the polynomial computed at v . When v is a multiplication node then generically, the polynomials computed at its children are irreducible¹³ in which case the efficient multivariate polynomial factorization algorithm of Kaltofen and Trager [KT90] recovers the children’s outputs. Even when there is noise, the robust factorization algorithm of [KMYZ08] can recover the factors approximately¹⁴. Thus the main challenge is to recover the children of addition gates. This connects us to the problem discussed in the previous section with the *structured* polynomials being the polynomials computed at the children gates. In the noiseless setting, a meta algorithm for this problem was given in [GKS20]. We provide a meta algorithm in the noisy case and show worst-case bounds on the quality of the output in terms of singular values of certain matrices¹⁵. The abstract problem is as follows. Given a polynomial $\tilde{f}(\mathbf{x})$ that can be expressed as

$$\tilde{f}(\mathbf{x}) = T_1(\mathbf{x}) + T_2(\mathbf{x}) + \dots + T_s(\mathbf{x}) + \eta(\mathbf{x}), \quad (2)$$

where T_i ’s are *structured* polynomials and the noise/perturbation polynomial $\eta(\mathbf{x})$ has *small norm*, can we approximately recover the T_i ’s via an efficient algorithm?

Learning from lower bounds. [GKS20] showed how the linear maps used in the known arithmetic formula lower bound proofs could be used to recover the T_i ’s in the noiseless ($\eta = 0$) setting, *assuming that appropriate non-degeneracy conditions hold*. [GKS20] observed that the assumption that the T_i ’s are structured can effectively be operationalized via the existence of a known set of linear maps \mathcal{L} from the vector space of polynomials to some appropriate vector space W_1 such that $\dim(\langle \mathcal{L} \cdot T_i \rangle)$ is¹⁶ *small* for every simple polynomial T_i . When we apply such a set of linear operator \mathcal{L} to (2) with $\eta = 0$, we get:

$$\langle \mathcal{L} \cdot \tilde{f}(\mathbf{x}) \rangle \subseteq \langle \mathcal{L} \cdot T_1(\mathbf{x}) \rangle + \langle \mathcal{L} \cdot T_2(\mathbf{x}) \rangle + \dots + \langle \mathcal{L} \cdot T_s(\mathbf{x}) \rangle. \quad (3)$$

[GKS20] observe that generically two things tend to happen.

Assumption 1.1. First blessing of dimensionality¹⁷. *If $\sum_{i \in [s]} \dim(\langle \mathcal{L} \cdot T_i(\mathbf{x}) \rangle) \ll \dim(W_1)$ then almost surely (over the independent random choice of the T_i ’s), it holds that the subspaces $\langle \mathcal{L} \cdot T_i(\mathbf{x}) \rangle$ form a direct sum, i.e.*

$$\dim(\langle \mathcal{L} \cdot T_1(\mathbf{x}) \rangle + \langle \mathcal{L} \cdot T_2(\mathbf{x}) \rangle + \dots + \langle \mathcal{L} \cdot T_s(\mathbf{x}) \rangle) = \sum_{i \in [s]} \dim(\langle \mathcal{L} \cdot T_i(\mathbf{x}) \rangle).$$

¹³Random multivariate polynomials are almost surely irreducible and with that as intuition, one expects the output of a formula with output being an addition gate to almost surely be an irreducible polynomial when the underlying field constants are chosen randomly. However proving this can be technically involved for any given subclass of formulas.

¹⁴The work of [KT90] aims to devise a factorization algorithm that is empirically as robust as possible and does not contain theoretical bounds on how much the output factors get perturbed as a function of the noise added to a true factorization. Nevertheless such a bound can be inferred from their work. The bound would depend on the appropriate singular values of an instance-dependent matrix called the Ruppert matrix that comes up in their algorithm.

¹⁵The matrices whose singular values are used to bound the quality of the output depend on the input instance as well as on the choice of linear operators used to instantiate our framework

¹⁶Here, $\langle S \rangle$ denotes the \mathbb{R} -linear span of a set S that consists of vectors or linear maps. Also $\mathcal{L} \cdot T_i$ denotes the set of vectors obtained by applying each linear map in \mathcal{L} to T_i .

¹⁷The intuition is that (pseudo)-randomly chosen small-dimensional subspaces of a large-dimensional ambient space should form a direct sum.

Assumption 1.2. Second blessing of dimensionality¹⁸. If $\sum_{i \in [s]} \dim(\langle \mathcal{L} \cdot T_i(\mathbf{x}) \rangle) \ll \dim(\langle \mathcal{L} \rangle)$ then almost surely (over the independent random choice of the T_i 's), it holds that for all $i \in [s]$:

$$\langle \mathcal{L} \cdot (T_1(\mathbf{x}) + \dots + T_s(\mathbf{x})) \rangle \supseteq \langle \mathcal{L} \cdot T_i(\mathbf{x}) \rangle.$$

Under these nondegeneracy assumptions we then have (for $\eta = 0$):

$$U \stackrel{\text{def}}{=} \langle \mathcal{L} \cdot \tilde{f}(\mathbf{x}) \rangle = \langle \mathcal{L} \cdot T_1(\mathbf{x}) \rangle \oplus \langle \mathcal{L} \cdot T_2(\mathbf{x}) \rangle \oplus \dots \oplus \langle \mathcal{L} \cdot T_s(\mathbf{x}) \rangle. \quad (4)$$

We observe that in the noisy case, on input \tilde{f} , finding the best $(\dim(U))$ -rank subspace through the set of points $(\mathcal{L} \circ \tilde{f})$ yields a subspace \tilde{U} that is pretty close to U (lemma D.1 gives quantitative bounds). Coming back to the noiseless case, [GKS20] then observe that linear maps constructed for the purpose of proving lower bounds also yield a set of linear maps \mathcal{B} such that

$$V \stackrel{\text{def}}{=} \langle \mathcal{B} \cdot U \rangle = \langle \mathcal{B} \cdot U_1 \rangle \oplus \dots \oplus \langle \mathcal{B} \cdot U_s \rangle, \quad (5)$$

where $U_i \stackrel{\text{def}}{=} \langle \mathcal{L} \cdot T_i(\mathbf{x}) \rangle$. This motivated the following problem which they call *Vector Space Decomposition*. Given a set of linear maps \mathcal{B} between two vector spaces U and V , find a (maximal) decomposition $U = U_1 \oplus \dots \oplus U_s$, $V = V_1 \oplus \dots \oplus V_s$ s.t. $\mathcal{B} \cdot U_i \subseteq V_i$ for all $i \in [s]$. In most applications, such a decomposition turns out to be unique (up to some obvious symmetries like permuting the subspaces) and hence an algorithm for vector space decomposition finds the intended decomposition.

Reduction to Vector Space Decomposition in the noisy setting. We then formulate and give an algorithm for a robust/noise-tolerant version of vector space decomposition. But there is an important difficulty that crops up in trying to use the problem of robust vector space decomposition as formulated below to the setting of learning arithmetic formulas in the presence of noise. \mathcal{B} is a collection of maps from W_1 to W_2 where $U \subseteq W_1, V \subseteq W_2$ and $\mathcal{B} \cdot U$ equals V . However, $\mathcal{B} \cdot \tilde{U}$ will typically *not* be contained in \tilde{V} . In fact the dimension of the image of \tilde{U} under the action of \mathcal{B} (denoted $\dim(\langle \mathcal{B} \cdot \tilde{U} \rangle)$) will typically be much larger than the dimension of \tilde{V} (denoted $\dim(\tilde{V})$). To overcome this difficulty, our idea is to compose maps in \mathcal{B} with the projection¹⁹ map to \tilde{V} to obtain a tuple of maps $\tilde{\mathcal{B}}$ from \tilde{U} to \tilde{V} . In general such a composition can completely spoil the structure of the set of maps \mathcal{B} but our conceptual insight here is that in this situation, one can set up a natural correspondence between $\text{Lin}(U, V)$ and $\text{Lin}(\tilde{U}, \tilde{V})$ that can be used to infer that the projection-composed maps $\tilde{\mathcal{B}}$ are *slight perturbations* of the corresponding maps in \mathcal{B} (lemma 4.1 gives quantitative bounds).

¹⁸The intuition is that in most applications when the underlying dimension $n = |\mathbf{x}|$ is large enough then the dimension of the set of operators \mathcal{L} is large relative to $\dim(\langle \mathcal{L} \cdot T_i(\mathbf{x}) \rangle)$ for any i . In such a situation if the T_i 's are chosen generically then \mathcal{L} tends to contain many operators that kill all the other T_j 's (for $j \neq i$) so that $\langle \mathcal{L} \cdot f(\mathbf{x}) \rangle$ tends to contain each of the subspaces $\langle \mathcal{L} \cdot T_i(\mathbf{x}) \rangle$.

¹⁹Projection to \tilde{V} here implicitly uses a decomposition of the ambient space W_2 into \tilde{V} and its orthogonal complement (defined via some canonical inner product on W_2 that is clear from context). It is the unique map in $\text{Lin}(W_2, W_2)$ which is identity on \tilde{V} and whose kernel is the orthogonal complement of \tilde{V} .

This insight gives us the reduction. Then, the robust vector space decomposition algorithm yields a decomposition

$$\tilde{U} = \tilde{U}_1 \oplus \tilde{U}_2 \oplus \dots \oplus \tilde{U}_s \quad (6)$$

where $\tilde{U}_1, \tilde{U}_2, \dots, \tilde{U}_s$ are slightly perturbed versions of $\langle \mathcal{L} \cdot T_1(\mathbf{x}) \rangle, \langle \mathcal{L} \cdot T_2(\mathbf{x}) \rangle, \dots, \langle \mathcal{L} \cdot T_s(\mathbf{x}) \rangle$ respectively (corollary 4.1 gives quantitative bounds). In particular this implies that for each $L \in \mathcal{L}$ we can obtain a vector close to $L \cdot T_1(\mathbf{x})$ by projecting²⁰ $L \cdot \tilde{f}(\mathbf{x})$ to \tilde{U}_1 . This implies that we can approximately recover $T_1(\mathbf{x})$ itself via an appropriate pseudo-inverse computation. Similarly, we can recover all the $T_i(\mathbf{x})$'s up to some error (Theorem 14 gives quantitative bounds). Before stating the quantitative bound on this error (Theorem 1) let us discuss the subroutine of robust vector space decomposition which is perhaps of interest in itself and might have wider applicability.

1.2 Overview - Vector Space Decomposition algorithm

We refer to the noise-tolerant version of vector space decomposition as *Robust Vector Space Decomposition (RVSD)*. The setting is the following: let W_1 and W_2 be vector spaces, and let $U = U_1 \oplus \dots \oplus U_s \subseteq W_1$ and $V = V_1 \oplus \dots \oplus V_s \subseteq W_2$ be subspaces. Let $\mathcal{B} = (B_1, B_2, \dots, B_m)$ be an m -tuple of linear operators, with each $B_j : U \rightarrow V$ being a linear map from U to V . Suppose that, under the action of \mathcal{B} , each U_i is mapped inside V_i ; that is, for each $i \in [s]$, it holds that $\langle \mathcal{B} \cdot U_i \rangle \subseteq V_i$. We consider the problem of recovering the U_i 's approximately given noisy access to U, V and \mathcal{B} . Specifically²¹

Robust Vector Space Decomposition (RVSD). We are given as input the integer s , two vector spaces $\tilde{U} \subseteq W_1$ and $\tilde{V} \subseteq W_2$, and a m -tuple of operators $\tilde{\mathcal{B}} = (\tilde{B}_1, \tilde{B}_2, \dots, \tilde{B}_m)$ from \tilde{U} to \tilde{V} , such that $\text{dist}(\tilde{U}, U)$, $\text{dist}(\tilde{V}, V)$ and $\text{dist}(\tilde{\mathcal{B}}, \mathcal{B})$ ²² are "small." Our goal is to *efficiently* find an s -tuple $\tilde{U} = (\tilde{U}_1, \tilde{U}_2, \dots, \tilde{U}_s)$ of subspaces in $\tilde{U} \subseteq W_1$, such that (upto a reordering of the components) for each $i \in [s]$, $\text{dist}(\tilde{U}_i, U_i)$ is "small"²³.

Now we give some rough ideas that go behind our Robust Vector Space Decomposition algorithm. For more details, the reader is referred to Section 4. Let us first consider the noiseless setting, in which we are given an integer s , the vector spaces $U \subseteq W_1, V \subseteq W_2$, and a m -tuple of operators $\mathcal{B} = (B_1, \dots, B_m)$ from $U \rightarrow V$; the goal is to find a decomposition $U = U_1 \oplus \dots \oplus U_s$ and $V = V_1 \oplus \dots \oplus V_s$, such that each U_i is mapped into V_i under the action of \mathcal{B} , i.e.

$$U = U_1 \oplus U_2 \oplus \dots \oplus U_s \quad \text{and} \quad V = V_1 \oplus V_2 \oplus \dots \oplus V_s, \quad \langle \mathcal{B} \cdot U_i \rangle \subseteq V_i \quad \forall i \in [s]. \quad (7)$$

²⁰Projection to \tilde{U}_1 here refers to using the decomposition given by (6). It is applying the unique map in $\text{Lin}(\tilde{U}, \tilde{U})$ which is identity on \tilde{U}_1 and whose kernel is $(\tilde{U}_2 \oplus \tilde{U}_3 \oplus \dots \oplus \tilde{U}_s)$.

²¹As discussed above, it is often the case that a set of operators (B_1, \dots, B_m) , with each $B_i : W_1 \rightarrow W_2$, satisfying the above property are exactly known. In this case, we can instantiate the Robust Vector Space Decomposition problem with suitable projections of these operators on the set of linear maps from $U \rightarrow V$, and $\tilde{U} \rightarrow \tilde{V}$ respectively. For more details, the reader is referred to Section 4.3.

²²In the formulation here, the distance $\text{dist}(\tilde{\mathcal{B}}, \mathcal{B})$ is defined by extending all operators to map W_1 into W_2 .

²³As we note in Remark 6, our algorithms can be used to find (V_1, \dots, V_s) approximately as well, but we omit that here since our applications do not need it.

The adjoint algebra and its properties. Based on [Qia18, CIK97], [GKS20] defined a notion called the *adjoint algebra*²⁴ whose structure can be used to understand (the potentially many) decompositions. Let us recall this notion.

Definition 1.1. Adjoint algebra The adjoint algebra, corresponding to the vector spaces U, V , and the tuple of operators \mathcal{B} , denoted $\text{Adj}_{U,V}(\mathcal{B})$ is defined to be the set of all tuples of linear maps (D, E) , with $D : U \rightarrow U, E : V \rightarrow V$, such that $B_j \cdot D = E \cdot B_j$ for all $j \in [m]$.

Observe that the adjoint algebra always contains the space of scaling maps²⁵: that is, the set of maps $D : U \rightarrow U, E : V \rightarrow V$ such that D (resp. E) simply scales each U_i (resp. V_i) by some scalar λ_i , for each $i \in [s]$. We observe that in most applications these maps are all that the adjoint algebra contains, and in this case, there is a simple algorithm to solve the vector space decomposition, and the obtained decomposition is unique:

Proposition 1.1 (Proposition A.3²⁶ in [GKS20]). *Suppose that U, V admit a decomposition into direct sum of s spaces under the action of \mathcal{B} as in (7). If $\dim(\text{Adj}_{U,V}(\mathcal{B})) = s$, then it holds that:*

1. $\text{Adj}_{U,V}(\mathcal{B})$ equals the set of scaling maps (as defined above) and,
2. The decomposition given by (7) is the unique irreducible decomposition, i.e. if

$$U = \hat{U}_1 \oplus \hat{U}_2 \oplus \dots \oplus \hat{U}_{\hat{s}} \text{ and } V = \hat{V}_1 \oplus \hat{V}_2 \oplus \dots \oplus \hat{V}_{\hat{s}}, \quad \hat{s} \geq s,$$

and

$$\langle \mathcal{B} \cdot \hat{U}_i \rangle \subseteq \hat{V}_i, \quad \forall i \in [\hat{s}],$$

then $\hat{s} = s$ and upto reordering if necessary, $\hat{U}_i = U_i$ and $\hat{V}_i = V_i$ for all $i \in [s]$.

Noiseless algorithm. Note that given \mathcal{B} (and U, V) computing $\text{Adj}_{U,V}(\mathcal{B})$ is easy and simply involves solving for D and E that satisfy the linear constraints specified in definition 1.1. Further under the assumption that $\text{Adj}_{U,V}(\mathcal{B})$ equals the set of scaling maps (this we refer to as strong uniqueness), the required subspaces U_1, U_2, \dots, U_s can be obtained as the eigenspaces corresponding to distinct eigenvalues of the linear map $D : U \rightarrow U$ which is the component of a random element (D, E) of $\text{Adj}_{U,V}(\mathcal{B})$.

Making the algorithm robust. There is a relatively straightforward way to make this algorithm robust: we use the maps in $\tilde{\mathcal{B}}$ to compute a vector space²⁷ that is in some sense an approximation to the original adjoint algebra. Finally, we recover the U_i 's approximately as (the sum of a few) eigenspaces of suitably chosen elements of this approximate adjoint algebra. In the noiseless setting it suffices to chose random elements of the adjoint algebra but in the noisy setting this does not work very well.

²⁴The adjoint algebra is a generalization of the notion of the centralizer algebra in matrix/group theory to the case when the image space of the set of linear maps is different from the domain space.

²⁵This observation is due to [CIK97] and forms the starting point of the [GKS20] algorithm for vector space decomposition.

²⁶This proposition is a special case of the more general proposition A.3 in [GKS20] wherein the blocks of $\text{Adj}_{U,V}(\mathcal{B})$ consist of scalar matrices only.

²⁷This space is typically not closed under multiplication and so does not form an algebra.

This is because the error incurred in the recovery of an eigenvector/eigenspace of an operator is inversely related to the corresponding eigengap(s) (see lemma A.8). Simply picking a random element of the adjoint algebra $\text{Adj}_{U,V}(\mathcal{B})$ leads to a rather small eigengap and we therefore incur a rather large error both theoretically and practically (i.e. in both the worst case noise scenario and the random noise scenarios). Our insight here is that the multiplicative structure of the adjoint algebra can be exploited to find operators in it with (some) large eigengaps and this yields an algorithm that is more robust. Indeed, our initial experiments suggest that the resulting algorithm when applied to tensor decomposition empirically performs better (in terms of error in the output) than any of the known algorithms for tensor decomposition. The details and quantitative bounds are provided in section 4.

Our Results. The noise-tolerance and performance of our meta-algorithm is captured by the following theorem which bounds the error incurred in terms of various parameters involved.

Theorem 1 (Learning Noisy Arithmetic Circuits, Informal version of Theorem 14). *Let $f(\mathbf{x}) = T_1(\mathbf{x}) + \dots + T_s(\mathbf{x})$ be a polynomial such that each $T_i \in \mathbb{R}[\mathbf{x}]^d$ belongs to a circuit class \mathcal{C} that admits operators \mathcal{L} and \mathcal{B} satisfying the following properties:*

- \mathcal{L} consists of linear maps $L : \mathbb{R}[\mathbf{x}]^d \rightarrow W_1$ such that $U \stackrel{\text{def}}{=} \langle \mathcal{L} \cdot f \rangle = U_1 \oplus \dots \oplus U_s$, $\dim(U) = d_U$, where $U_i \stackrel{\text{def}}{=} \langle \mathcal{L} \cdot T_i \rangle$.
- \mathcal{B} consists of linear maps $B : W_1 \rightarrow W_2$ satisfying $V \stackrel{\text{def}}{=} \langle \mathcal{B} \cdot \mathcal{L} \cdot f \rangle = V_1 \oplus \dots \oplus V_s$, $\dim(V) = d_V$, where $V_i \stackrel{\text{def}}{=} \langle \mathcal{B} \cdot \mathcal{L} \cdot T_i \rangle$.
- The decomposition of (U, V) under \mathcal{B} is strongly unique, i.e. $\dim(\text{Adj}_{U,V}(\mathcal{B})) = s$.

We also need the robust versions of the above assumptions and that \mathcal{L} and \mathcal{B} are appropriately normalized. Let M, N be matrices with columns $L \cdot f, L \in \mathcal{L}$ and $B \cdot L \cdot f, B \in \mathcal{B}, L \in \mathcal{L}$ respectively. Suppose that the d_U^{th} and the d_V^{th} largest singular values of M and N , respectively, are bounded from below by some $\sigma > 0$. Similarly, for an appropriate operator corresponding to the adjoint algebra, we need an appropriate singular value lower bounded by σ .

Let $\tilde{f}(\mathbf{x}) = f(\mathbf{x}) + \eta(\mathbf{x})$ be a polynomial such that $\|\eta\| \leq \epsilon^{28}$. Then, there is an efficient algorithm, which on input \tilde{f} , recovers $\tilde{T}_1, \tilde{T}_2, \dots, \tilde{T}_s$, such that for any $\delta > 0$, with probability at least $1 - \delta$, (upto reordering) for each $i \in [s]$ it holds that

$$\left\| T_i - \tilde{T}_i \right\| \leq \text{poly}(s, d, d_U, d_V, 1/\delta, 1/\sigma) \cdot \epsilon.$$

Remark 2. 1) **Error for random noise.** *The above bound on the output error is for the case when the noise $\eta(\mathbf{x})$ is chosen in an adversarial (i.e. worst-case) fashion, subject of course to the indicated upper bound on its norm. In practice $\eta(\mathbf{x})$ often behaves like a random vector so that the output error is in practice significantly less²⁹ than the worst-case bound in the above theorem. Our intuition is that when $\eta(\mathbf{x})$ is random the output*

²⁸Under an appropriate norm called the Bombieri norm as defined in Section 2. The Bombieri norm is a suitably scaled version of the ℓ_2 norm that has many desirable properties including being invariant under a unitary transformation of the underlying variables.

²⁹This situation is reminiscent of the well-studied spiked tensor problem in machine learning which can be thought of as a very special case of our problem.

error should be less by a factor of $\text{poly}(\dim(\langle \mathcal{L} \rangle))$ compared to when $\eta(\mathbf{x})$ is adversarially chosen. We leave it as a potential direction for future investigation.

- 2) **Noise-tolerance.** As noted earlier, our initial experiments indicate that for the well-studied special case of tensor decomposition our algorithm seems to be more noise-tolerant than existing algorithms. We remark here that for subspace clustering, one can have a somewhat different reduction to vector space decomposition which also incorporates the affinity-based information to obtain a more noise-tolerant clustering algorithm. It might be interesting to do an empirical comparison of noise-tolerance of (such adaptations of) our algorithm to existing algorithms for various applications of interest.
- 3) **Running Time.** The algorithm boils down to computing singular value decompositions and/or pseudoinverses of certain matrices and thus its running time³⁰ is upper bounded by the cube of the dimension of the largest vector space involved.
- 4) We suggest a potential way to speed up the above algorithm in section 1.5.

1.3 Application 1: Subspace Clustering.

Subspace clustering is the following problem - we are given a set of N points $A = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_N\} \subseteq \mathbb{R}^n$ that admit a partition

$$A = A_1 \uplus A_2 \uplus \dots \uplus A_s,$$

such that the points in each A_j ($j \in [s]$) span a low-dimensional (relative to the number of points in A_j) space $\langle A_j \rangle$. The goal is to find such a partition.

Even for $n = 3$, this problem is NP-hard in the worst case [MT82]. Despite this, it has been intensely studied and we refer the reader to the surveys [PHL04], [QXCK23] and the references therein. Most state of the art techniques rely on constructing an affinity matrix, which measures how likely two points are to be in the same subspace, followed by spectral clustering using the affinity matrix. Most such algorithms have little theoretical analysis about the robustness and recovery guarantees.

A non-degeneracy condition and a reduction. Suppose now that the span of the A_j 's satisfy the following non-degeneracy condition: they form a direct sum, i.e.

$$\langle A \rangle = \langle A_1 \rangle \oplus \langle A_2 \rangle \oplus \dots \oplus \langle A_s \rangle. \quad (8)$$

We will see that in this case subspace clustering reduces to vector space decomposition in the following way. For a point $\mathbf{a} = (a_1, a_2, \dots, a_n) \in \mathbb{R}^n$, let $\mathbf{a} \cdot \mathbf{x} \in \mathbb{R}[\mathbf{x}]$ denote the linear form $a_1x_1 + a_2x_2 + \dots + a_nx_n$, in the formal variables $\mathbf{x} = (x_1, x_2, \dots, x_n)$. For $d \geq 1$ we denote by $A^{\otimes d}$ the set $\{(\mathbf{a} \cdot \mathbf{x})^d : \mathbf{a} \in A\} \subseteq \mathbb{R}[\mathbf{x}]^{=d}$. Consider the space of first-order partial differential operators $\mathcal{B} = \partial^{=1}$ acting on the subspace of polynomials

$$\langle A^{\otimes 2} \rangle = \langle (\mathbf{a}_1 \cdot \mathbf{x})^2, (\mathbf{a}_2 \cdot \mathbf{x})^2, \dots, (\mathbf{a}_N \cdot \mathbf{x})^2 \rangle \subseteq \mathbb{R}[\mathbf{x}]^{=2}.$$

The image space is then

$$\langle A^{\otimes 1} \rangle = \langle (\mathbf{a}_1 \cdot \mathbf{x}), (\mathbf{a}_2 \cdot \mathbf{x}), \dots, (\mathbf{a}_N \cdot \mathbf{x}) \rangle \subseteq \mathbb{R}[\mathbf{x}]^{=1}.$$

³⁰This is in the model where operations over real numbers are of unit cost. A more precise bound on the running time in terms of the dimensions of the various relevant vector spaces can be

Note that our non-degeneracy condition can be restated as saying that

$$\langle A^{\otimes 1} \rangle = \langle A_1^{\otimes 1} \rangle \oplus \langle A_2^{\otimes 1} \rangle \oplus \dots \oplus \langle A_s^{\otimes 1} \rangle.$$

This implies that the subspaces $\langle A_j^{\otimes 2} \rangle$ also form a direct sum. It is also easily seen that the image of each $\langle A_j^{\otimes 2} \rangle$ under $\mathcal{B} = \partial^1$ is precisely $\langle A_j^{\otimes 1} \rangle$. Thus the vector space $\langle A^{\otimes 2} \rangle$ admits a decomposition under the action of \mathcal{B} . Furthermore, under the additional mild assumption that each $\langle A_j^{\otimes 2} \rangle$ is *indecomposable* under the action of \mathcal{B} it turns out (using Corollary B.1) that the decomposition is unique and thus the subspace clustering problem reduces to the problem of vector space decomposition.

A weaker non-degeneracy condition. Note that the non-degeneracy condition given by (8) is rather restrictive - it implies in particular that the number of subspaces s cannot exceed n , the dimension of the ambient space. We can get a weaker non-degeneracy condition by considering the action of first-order partial differential operators $\mathcal{B} = \partial^1$ on the space $\langle A^{\otimes d} \rangle$ instead (for some suitable choice of $d \geq 2$). The image space is then $\langle A^{\otimes(d-1)} \rangle$. As before, under the (now weaker) non-degeneracy condition that

$$\langle A^{\otimes(d-1)} \rangle = \langle A_1^{\otimes(d-1)} \rangle \oplus \langle A_2^{\otimes(d-1)} \rangle \oplus \dots \oplus \langle A_s^{\otimes(d-1)} \rangle,$$

the vector space $\langle A^{\otimes d} \rangle$ admits a decomposition under the action of \mathcal{B} . Furthermore, as before, under the additional mild assumption that each $\langle A_j^{\otimes d} \rangle$ is indecomposable under the action of \mathcal{B} it turns out (Corollary B.1) that the decomposition is unique and thus the subspace clustering problem reduces to the problem of vector space decomposition (see Theorem 9).

Robust subspace clustering. The robust or noisy version of the subspace clustering problem is the following. Given a set of points $\tilde{A} = \{\tilde{\mathbf{a}}_1, \tilde{\mathbf{a}}_2, \dots, \tilde{\mathbf{a}}_N\} \subseteq \mathbb{R}^n$ suppose that each point $\tilde{\mathbf{a}}_i$ is *close to* an (unknown point) $\mathbf{a}_i \in \mathbb{R}^n$ such that the resulting set of points $A = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_N\} \subseteq \mathbb{R}^n$ can be clustered using s subspaces, i.e.

$$A = A_1 \uplus A_2 \uplus \dots \uplus A_s,$$

where each A_j spans a *low-dimensional* subspace $\langle A_j \rangle$. The computational task is to *approximately* recover each subspace $\langle A_j \rangle$, that is, output $\tilde{\mathbf{W}} = (\tilde{W}_1, \tilde{W}_2, \dots, \tilde{W}_s)$ such that (upto reordering) each \tilde{W}_j is close to $\langle A_j \rangle$ for each $j \in [s]$. We can reduce this problem to the robust vector space decomposition as follows. Let $m_d \stackrel{\text{def}}{=} \dim(\langle A^{\otimes d} \rangle)$ and $m_{d-1} \stackrel{\text{def}}{=} \dim(\langle A^{\otimes(d-1)} \rangle)$. Given \tilde{A} we algorithmically compute the best fitting subspace \tilde{U} (resp. \tilde{V}) of dimension m_d (resp. m_{d-1}) to $\tilde{A}^{\otimes d}$ (resp. to $\tilde{A}^{\otimes(d-1)}$). It turns out then that \tilde{U} (resp. \tilde{V}) is *close to* $\langle A^{\otimes d} \rangle$ (resp. to $\langle A^{\otimes(d-1)} \rangle$) (Lemmas B.7, B.8 give the quantitative bounds). Applying the robust version of vector space decomposition on $(\tilde{U}, \tilde{V}, \mathcal{B})$, the subspaces that we obtain are close to $\langle A_j^{\otimes d} \rangle$ ($j \in [s]$) and from these we can, in turn, also approximately recover $\langle A_j \rangle$ (Proposition B.1), as required. This yields the following theorem.

Theorem 2 (Robust Subspace Clustering, Informal version of Theorem 10). Let $A = \{\mathbf{a}_1, \dots, \mathbf{a}_N\} \subseteq \mathbb{R}^n$ be a finite set of N points of unit norm, which can be partitioned as $A = A_1 \uplus \dots \uplus A_s$, where each $\langle A_i \rangle$ is a subspace of dimension at most t .

Let $d \geq 2$ be an integer, let $\mathbf{U} = (U_1, \dots, U_s)$ (resp. $\mathbf{V} = (V_1, \dots, V_s)$) be an s -tuple of subspaces with $U_j = \langle A_j^{\otimes d} \rangle$ (resp. $V_j = \langle A_j^{\otimes d-1} \rangle$) for each $j \in [s]$. Let $U = \langle \mathbf{U} \rangle$ (resp. $V = \langle \mathbf{V} \rangle$) have dimension m_d (resp. m_{d-1}).

Suppose that:

- $U = U_1 \oplus \dots \oplus U_s$, $V = V_1 \oplus \dots \oplus V_s$, and for each $j \in [s]$, it holds that $\dim(U_j) = \binom{\dim(\langle A_j \rangle) + d - 1}{d}$, $\dim(V_j) = \binom{\dim(\langle A_j \rangle) + d - 2}{d-1}$.
- σ_A is the minimum of $\sigma_{m_d}(M_{A,d})$ and $\sigma_{m_{d-1}}(M_{A,d-1})$, where $M_{A,d}$ (resp. $M_{A,d-1}$) is the matrix whose columns are the polynomials $(\mathbf{a}_i \cdot \mathbf{x})^d$ (resp. $(\mathbf{a}_i \cdot \mathbf{x})^{d-1}$) (see Definition B.4).
- $\kappa(\mathbf{U})$ denotes the condition number of the tuple of subspaces \mathbf{U} (see Section 2).
- $\sigma_{-(s+1)}(\mathfrak{A})$ is the $(s+1)^{\text{th}}$ smallest singular value of the adjoint algebra map (see Definition 4.2), corresponding to the action of $\mathcal{B} = (B_1, \dots, B_n)$ on \mathbf{U}, \mathbf{V} , where B_i corresponds to the operator ∂_{x_i} .

Let $\tilde{A} = \{\tilde{\mathbf{a}}_1, \tilde{\mathbf{a}}_2, \dots, \tilde{\mathbf{a}}_N\} \subseteq \mathbb{R}^n$ be a set of unit norm vectors such that $\|\mathbf{a}_i - \tilde{\mathbf{a}}_i\|_2 \leq \epsilon$ for each $i \in [N]$. Then, there is an algorithm, which on input \tilde{A} , runs in time $\text{poly}(N, n^d)$, and recovers subspaces $(\tilde{W}_1, \tilde{W}_2, \dots, \tilde{W}_s)$, such that with probability at least $1 - \delta$, (upto reordering) for each $j \in [s]$ it holds that

$$\text{dist}(\tilde{W}_j, \langle A_j \rangle) \leq \text{poly}\left(t, N, d, s, 1/\delta, \kappa(\mathbf{U}), 1/\sigma_A, 1/\sigma_{-(s+1)}(\mathfrak{A})\right) \cdot \epsilon.$$

We show how to lower bound $\sigma_{-(s+1)}(\mathfrak{A})$ (Theorem 11). The main technical component is an inductive argument to analyze singular values of basic adjoint operators, which is inspired by the inductive argument in recent works on analyzing eigenvalues for random walks on simplicial complexes (e.g. [ALGV19]). Next we see what our algorithms would yield in the smoothed case and state some explicit conjectures about singular values of relevant smoothed matrices.

Smoothed analysis of subspace clustering. We first describe the input model. For simplicity, we assume that each of the subspaces have the same dimension (equal to t).

1. **Perturbation model for subspaces.** We have a tuple of s hidden subspaces of \mathbb{R}^n , $\mathbf{W} = (W_1, W_2, \dots, W_s)$, each of dimension t . Let $P_1, P_2, \dots, P_s \in \mathbb{R}^{n \times t}$ be matrices with orthonormal columns, such that the column span of P_i is W_i . Each subspace W_i is perturbed by perturbing P_i by a random (and independent) Gaussian matrix $G_i \sim \mathcal{N}(0, \rho^2/n)^{n \times t}$. Let $\hat{P}_i = P_i + G_i$, and $\hat{W}_1, \hat{W}_2, \dots, \hat{W}_s$ be the column spans of $\hat{P}_1, \hat{P}_2, \dots, \hat{P}_s$ respectively.
2. **Perturbation models for points from each subspace.** Sample (possibly adversarially) sets of points A_1, A_2, \dots, A_s from $\hat{W}_1, \hat{W}_2, \dots, \hat{W}_s$ respectively, of unit norm. For each $i \in [s]$, perturb each point in A_i with respect to \hat{W}_i to get the

set of points \hat{A}_i . Formally, this means perturbing points in A_i by $\hat{B}_i \cdot v$, where \hat{B}_i is an $n \times t$ matrix describing an orthonormal basis for \hat{W}_i and $v \sim \mathcal{N}(0, \rho^2/t)^t$ (independently generated for each point), and normalizing. Let $\hat{A} = \hat{A}_1 \cup \hat{A}_2 \cup \dots \cup \hat{A}_s$.

3. **Adding noise.** For each $\mathbf{a} \in \hat{A}$, add noise (possibly adversarially) and normalize to get a unit norm point \mathbf{a}' such that $\|\mathbf{a} - \mathbf{a}'\|_2 \leq \epsilon$. We are given as input \hat{A}' , the set of noise-added points.

Given the set of points \hat{A}' , the goal is to recover subspaces $\tilde{\mathbf{W}} = (\tilde{W}_1, \tilde{W}_2, \dots, \tilde{W}_s)$ such that $\text{dist}(\tilde{\mathbf{W}}, \hat{\mathbf{W}})$ is small. Next we state a couple of conjectures about minimum singular values of smoothed random matrices that we encounter:

Conjecture 1.1. *Let $\mathbf{v}_{i1}, \dots, \mathbf{v}_{it}$ be an orthonormal basis for \hat{W}_i generated as above. Define the linear forms $\ell_{ij}(\mathbf{x}) = \langle \mathbf{v}_{ij}, \mathbf{x} \rangle$. Consider the $\binom{n+d-1}{d} \times s \binom{t+d-1}{d}$ matrix M where the columns are divided into s chunks and in the i^{th} chunk, the columns are all the monomials of degree d in the polynomials $\ell_{i1}, \dots, \ell_{it}$. Also suppose $s \binom{t+d-1}{d} \leq (1 - \delta) \binom{n+d-1}{d}$ for a constant $\delta > 0$. Then for constant d , with high probability, $\sigma_{s \binom{t+d-1}{d}}(M) \geq \text{poly}(\rho, 1/n)$.*

Conjecture 1.2. *Consider arbitrary vectors $v_1, \dots, v_s \in \mathbb{R}^t$ of unit norm and their smoothed versions $\hat{v}_1, \dots, \hat{v}_s$, where $\hat{v}_i = v_i + g_i$, $g_i \sim \mathcal{N}(0, \rho^2/t)^t$ (and then further normalized to unit norm). Consider the $s \times \binom{t+d-1}{d}$ matrix M where the i^{th} row contains the polynomial $\langle \hat{v}_i, \mathbf{x} \rangle^d$. Suppose $s \geq (1 + \delta) \binom{t+d-1}{d}$ for a constant $\delta > 0$. Then for constant d , with high probability, $\sigma_{\binom{t+d-1}{d}}(M) \geq \text{poly}(\rho, 1/t)$.*

Theorem 3 (Smoothed analysis of subspace clustering, Theorem 12 restated). *Suppose Conjectures 1.1 and 1.2 are true. Then for constant d , Algorithm 5 on input $(\hat{A}', d, s, m_d, m_{d-1})$ outputs $\tilde{\mathbf{W}} = (\tilde{W}_1, \dots, \tilde{W}_s)$ such that with high probability,*

$$\text{dist}(\tilde{W}_j, \hat{W}_j) \leq \text{poly}(n, t, 1/\rho) \cdot \epsilon$$

Regarding the two conjectures, Conjecture 1.2 is closely linked to the paper [BCPV19]. There they considered the setting where $s \leq (1 - \delta) \binom{t+d-1}{d}$ and proved a similar lower bound for $\sigma_s(M)$. In both the settings there is slack, so it is plausible that the techniques of [BCPV19] can be adapted to prove Conjecture 1.2. But we don't know how to do that. In Conjecture 1.1, the matrix M is such that both the rows and columns share random variables. Most of the smoothed analysis till now focuses on matrices where either rows or columns have different sets of variables involved, and this makes it amenable to the leave-one-out distance method. Still, in Conjecture 1.1, the sharing of variables is not completely arbitrary. One can divide rows into chunks so that different chunks have different sets of variables. However, even this setting seems to require new techniques to analyze.

1.4 Application 2: Learning Mixtures of Gaussians

In this section we will see how the problem of computing the parameters of a mixture of Gaussians reduces to (several instances of) vector space decomposition.

Reduction to a special case of formula learning. It is implicit in [GHK15] that learning a mixture of s zero-mean Gaussians reduces to robustly expressing a given homogeneous polynomial $p(\mathbf{x})$ as a sum of s powers of quadratics, i.e.

$$p(\mathbf{x}) = p_1(\mathbf{x})^d + p_2(\mathbf{x})^d + \dots + p_s(\mathbf{x})^d, \quad (9)$$

where the p_i 's are homogeneous quadratic polynomials. Following the ideas in [GKS20], we give a direct reduction³¹ to vector space decomposition as follows.

Obtaining a vector space that is the direct sum of unknown spaces. Following [GKS20], we apply partial derivatives followed by a *random projection* to obtain a vector space that is a direct sum of s unknown subspaces, one corresponding to each $p_i(\mathbf{x})$. Specifically, let \mathcal{L} be the set of operators corresponding to taking k -th order partial derivatives followed by a *random restriction*³². Applying \mathcal{L} to both sides of equation (9), we get

$$\langle \mathcal{L} \cdot p(\mathbf{x}) \rangle \subseteq \langle \mathcal{L} \cdot p_1(\mathbf{x})^d \rangle + \langle \mathcal{L} \cdot p_2(\mathbf{x})^d \rangle + \dots + \langle \mathcal{L} \cdot p_s(\mathbf{x})^d \rangle,$$

It turns out that (Lemma C.1) under relatively mild nondegeneracy conditions on the choice of the p_i 's, the vector space sum on the right hand side of the above equation is actually a direct sum and the containment is actually an equality, i.e.

$$\langle \mathcal{L} \cdot p(\mathbf{x}) \rangle = \langle \mathcal{L} \cdot p_1(\mathbf{x})^d \rangle \oplus \langle \mathcal{L} \cdot p_2(\mathbf{x})^d \rangle \oplus \dots \oplus \langle \mathcal{L} \cdot p_s(\mathbf{x})^d \rangle.$$

We now carefully choose another set of operators \mathcal{B} such that the subspace $U \stackrel{\text{def}}{=} \langle \mathcal{L} \cdot p(\mathbf{x}) \rangle$ admits a *unique* decomposition under the action of \mathcal{B} .

Choice of \mathcal{B} . The set of operators \mathcal{L} maps polynomials in \mathbf{x} to polynomials in a subset of variables $\mathbf{y} \subseteq \mathbf{x}$. Under the above mentioned nondegeneracy conditions, it also turns out that for each $i \in [s]$, $\langle \mathcal{L} \cdot p_i(\mathbf{x})^d \rangle$ is of the form $U_i \stackrel{\text{def}}{=} \langle \mathbf{y}^k \cdot q_i(\mathbf{y})^{d-k} \rangle \subseteq \mathbb{R}[\mathbf{y}]^{=(2d-k)}$. With this in mind, we choose \mathcal{B} as the following set of operators: first order partial derivatives followed by multiplication³³ by polynomials of degree 1. In detail: \mathcal{B} consists of $|\mathbf{y}|^2$ operators with the (i, j) -th operator ($i, j \in [|\mathbf{y}|]$) being

$$B_{ij} : \mathbb{R}[\mathbf{y}]^{=(2d-k)} \mapsto \mathbb{R}[\mathbf{y}]^{=(2d-k)}, \quad B_{ij} \cdot q(\mathbf{y}) = y_j \cdot (\partial_{y_i} q(\mathbf{y})) \text{ for any } q(\mathbf{y}) \in \mathbb{R}[\mathbf{y}]^{=(2d-k)}.$$

It turns out that for any $i \in [s]$, under the action of \mathcal{B} , the image of

$$U_i \stackrel{\text{def}}{=} \langle \mathbf{y}^k \cdot q_i(\mathbf{y})^{d-k} \rangle \text{ is the subspace } V_i \stackrel{\text{def}}{=} \langle \mathbf{y}^{=(k+2)} \cdot q_i(\mathbf{y})^{d-k-1} \rangle$$

and that the U_i 's and V_j 's form direct sums (Lemma C.2). Furthermore, under mild non-degeneracy conditions such a decomposition is unique (Corollary C.1) implying that our vector space U has a unique decomposition into s subspaces under the action

³¹In [GKS20], there is an additional "multi-gcd" step which we avoid here.

³²We can think of a random projection as keeping a subset $\mathbf{y} \subseteq \mathbf{x}$ of the variables alive and setting the rest to zero.

³³The relevant literature on arithmetic formula lower bounds would refer to the set of operators \mathcal{B} as *shifted partials* and denote it by $\mathbf{y}^{-1} \cdot \partial_{\mathbf{y}}^{-1}$.

of \mathcal{B} . Lastly, from each U_i we can recover the corresponding $q_i(\mathbf{y})$ which is a restriction of $p_i(\mathbf{x})$ to a chosen subspace. Any polynomial can be recovered from its restriction to a small number of chosen subspaces and we use this to recover each $p_i(\mathbf{x})$ ($i \in [s]$), as required. In this way, the problem of learning mixtures of Gaussians reduces to robust vector space decomposition.

Robust version. Our general algorithm for learning arithmetic circuits with noise (Theorem 1) can be used to make the above algorithm robust. We will also need to use the algorithm of [BHKX22] in this case to combine the various projections of p_i 's. Our algorithm will depend on condition numbers of certain matrices which can be deduced from the operators used in the above algorithm. Lemmas C.1, C.2 and C.3 show that at least the ranks of these matrices are as expected. Lemmas C.1 and C.2 are from [GKS20]. Lemma C.3 is new and is the main technical contribution for this section, and shows that the relevant adjoint algebra is of the correct dimension. Also [BHKX22] analyze similar matrices corresponding to Lemmas C.1 and C.2 and prove the required condition number bounds in the fully random case. For the singular values of the adjoint operator (robustification of Lemma C.3), we believe similar techniques as Theorem 11 should work to give us a bound but the setting is more challenging and we don't know how to prove a bound here yet.

Comparison to [GKS20] and [BHKX22]. The algorithms of [GKS20, BHKX22] for learning mixtures of Gaussians roughly proceed as follows (for simplicity, we only consider the noiseless case here).

Given a polynomial $p(x) = \sum_{i=1}^s p_i(\mathbf{x})^d$, where each p_i is a quadratic polynomial:

1. Apply a set of operators \mathcal{L} to $p(\mathbf{x})$, where \mathcal{L} corresponds to taking some k -th order partial derivatives followed by a random restriction: as described before, each $\langle \mathcal{L} \cdot p_i(\mathbf{x}) \rangle$ is of the form $\langle \mathbf{y}^k \cdot q_i(\mathbf{y})^{d-k} \rangle \subseteq \mathbb{R}[\mathbf{y}]^{=(2d-k)}$. This step is essentially the same in both [GKS20, BHKX22].

We note however that [BHKX22] actually do not work under the non-degeneracy condition of the spaces $\langle \mathcal{L} \cdot p_i(\mathbf{x}) \rangle$'s forming a direct sum, and instead explicitly characterize the structure of the intersections $\langle \mathcal{L} \cdot p_i(\mathbf{x}) \rangle \cap \langle \mathcal{L} \cdot p_j(\mathbf{x}) \rangle$. This allows them to deal with a broader range of parameters compared to [GKS20].

2. The next step is a "multi-gcd" step, which is used to find the vector space $\langle q_1(\mathbf{y})^{d-k} \rangle + \dots + \langle q_s(\mathbf{y})^{d-k} \rangle$. This step is already present in the algorithm of [GKS20], however [BHKX22] give a significantly simpler algorithm for this step, along with an analysis for the robust version of this step.
3. The next step, which is in some sense the "main part" of the algorithm, is where the two algorithms [GKS20] and [BHKX22] differ:
 - (a) The algorithm of [GKS20] considers another application of k -th order partial derivatives + random restriction on this vector space, and uses vector space decomposition with respect to this set of operators. This allows them to recover the component polynomials.
 - (b) The algorithm of [BHKX22] follows the approach in [GHK15], and does a "desymmetrization + tensor-decomposition" step. This roughly enables

them to convert the sum of polynomials to a sum of tensors, and then apply standard tensor decomposition methods to obtain the required components.

4. The final step is to repeat the above procedure multiple times, using a different random restriction each time, and then aggregating the obtained $q_i(\mathbf{y})$'s into $p_i(\mathbf{x})$, as described before.

Our algorithm essentially follows the same first and final step as both these algorithms. It significantly deviates from the two algorithms in Steps 2 and 3:

1. While we follow the same vector space decomposition paradigm as [GKS20], our algorithm completely eliminates the use of the multi-gcd step. Instead, we use a very simple set of operators, namely order one partial derivatives + order one shifts, directly on the vector space $\langle \mathcal{L} \cdot p(\mathbf{x}) \rangle$. Hence, our approach provides a much more direct reduction to vector space decomposition.
2. In comparison to [BHKX22], we first eliminate the use of the multi-gcd step, and further we do not go through the desymmetrization step at all. Instead, our framework of vector space decomposition allows us to deal with symmetric polynomials throughout the algorithm; this inherently seems much more natural since the inputs and outputs all deal only with polynomials (symmetric tensors).

Finally, we note the the above described simplification allows us to obtain a much better range of parameters compared to [GKS20], whereas we still expect them to be slightly worse than [BHKX22].

1.5 Conclusion and Future Directions

In this work we showed how to adapt the algorithm of [GKS20] for learning subclasses of arithmetic formulas to make it noise-tolerant. This turns out to have a number of applications arising out of the remarkable fact that in these applications, a suitably defined polynomial formed out of the statistics of the data has a small arithmetic formula. We feel that our approach has the potential to give algorithms which are fast, noise-tolerant, outlier-tolerant and come with provable guarantees³⁴ for many such applications and is therefore worthy of further investigation. We now pose some problems that might encourage or guide such further study.

Making the vector space decomposition algorithm faster. Consider a set of operators \mathcal{B} mapping a real vector space U to another real vector space V . Our algorithm for decomposition of U (and V) under the action of \mathcal{B} involved computations with the adjoint algebra which entailed working in the vector spaces of linear maps $\text{Lin}(U, U)$ and $\text{Lin}(V, V)$. These spaces of linear maps have larger dimension than that of U and V themselves and consequently, our approach for decomposing U has running time pertaining to the cost of doing linear algebra over spaces of dimension $(\dim(U)^2 + \dim(V)^2)$. Let us first make an observation. Suppose that the decomposition induced by \mathcal{B} , namely:

$$U = U_1 \oplus U_2 \oplus \dots \oplus U_s, \quad V = V_1 \oplus V_2 \oplus \dots \oplus V_s$$

³⁴For most such applications the worst-case instances are intractable so the best we can hope for are algorithms whose performance can be bounded using singular values of certain instance-dependent matrices.

had the property that the U_i 's (respectively also the V_i 's) were orthogonal complements of each other (under some canonical inner product on the spaces U and V). Consider the collection of linear maps $\mathcal{L} \subseteq \text{Lin}(U, U)$ defined as $\mathcal{L} := \{B_j^T \cdot B_i : B_i, B_j \in \mathcal{B}\}$. Then each U_i is an invariant subspace (i.e. an eigenspace) of every operator in \mathcal{L} . In such a situation we typically expect the following simple algorithm to work: simply pick three random maps $B_1, B_2, B_3 \in \langle \mathcal{B} \rangle$ and compute³⁵ $L := B_2^T \cdot B_1$ and $M := B_3^T \cdot B_1$. Then for each eigenvector u of L , compute the span of the orbit of u under the action of M . The distinct subspaces so obtained should typically give us the required subspaces U_1, U_2, \dots, U_s . Clearly such an algorithm, *when it works*, would be much faster. We expect that for most applications, the above algorithm should work but we don't know.

Problem 1. For problems such as subspace clustering and learning mixtures of Gaussians, if the relevant U_i 's (respectively also the V_i 's) are orthogonal to each other, does the above algorithm correctly recover the U_i 's?

Problem 2. Whats the best way to make this algorithm noise-tolerant?

Finally, in situations where the U_i 's (resp the V_i 's) are not orthogonal to each other we can clearly make them so by using appropriate inner products on U and V . But how do we find such an inner product? We expect the operator scaling algorithm of [GGdOW20] to yield such an inner product(!)

Problem 3. For (noiseless) subspace clustering, does the operator scaling algorithm of [GGdOW20] applied on the relevant \mathcal{B} yield inner products under which the relevant subspaces are orthogonal?

Mixture of Gaussians. As mentioned in remark 1(b) earlier, we expect that our algorithm can be extended to handle general mixtures of Gaussians with differing means and covariance matrices. Let us formally state this as an open problem.

Problem 4. Random instances of general mixtures of Gaussians. Let $n, s \geq n$ be integers. For $i \in [s]$ suppose that we pick $\mu_i \in \mathbb{R}^n$ and covariance matrices $\Sigma_i \in \mathbb{R}^{n \times n}$ independently at random³⁶. Let $\mathcal{D} := \sum_{i=1}^s \frac{1}{s} \cdot \mathcal{N}(\mu_i, \Sigma_i)$ be the equi-weighted mixture of Gaussians with the above randomly chosen parameters. Design an efficient algorithm that given samples from D recovers the μ_i 's and Σ_i 's approximately.

Our work as well as that of [BHKX22] leave open the problem of doing a smoothed analysis of the corresponding algorithm for mixtures of zero-mean Gaussians. To encourage this direction of research, let us state this explicitly in the form of a conjecture.

Conjecture 1.3. Smoothed analysis of our algorithm for mixture of zero-mean Gaussians. *Our algorithm efficiently recovers the unknown parameters for smoothed instances of mixtures of zero-mean Gaussians.*

Handling outliers and other applications. In Remark 1, we conjectured that our approach/framework should enable the design of efficient algorithms that can handle outliers and also be useful for many more applications in unsupervised learning. It would be nice to have concrete results in such directions.

³⁵The linear maps B_2^T, B_3^T from V to U are defined using the canonical inner products on these two spaces.

³⁶Any reasonable distribution would do but for concreteness say we pick $\mu_i \sim \mathcal{N}(0, I_n)$ and we pick $\Sigma_i = B^T \cdot B$, where $B \sim \mathcal{N}(0, 1)^{n \times n}$

2 Preliminaries

We shall use $[n]$ to denote the set $\{1, 2, \dots, n\}$.

Matrices, Norms, Pseudo-inverse. Let $\mathbb{R}^{m \times n}$ denote the space of $m \times n$ matrices over \mathbb{R} . Given an $m \times n$ matrix M , we denote by $\|M\|_2$ its operator norm, and by $\|M\|_F$ its Frobenius norm. We shall use M^\dagger to denote its Moore-Penrose pseudo-inverse, and we define its condition number as $\kappa(M) \stackrel{\text{def}}{=} \|M\|_2 \|M^\dagger\|_2$. We mention some relevant properties of matrix norms and the pseudo-inverse in Section A.1.

Vector Spaces, Linear Operators, Projection Maps. Every vector space V that we will deal with in this work will be a finite dimensional real vector space that comes equipped with an inner product³⁷ denoted $\langle \cdot, \cdot \rangle_V$, or simply $\langle \cdot, \cdot \rangle$ when the underlying vector space is clear from context. For two vector spaces U and V , $\text{Lin}(U, V)$ shall denote the set of linear maps from U to V .

Inner products on U and V can be used to generate an inner product on $\text{Lin}(U, V)$, called the Hilbert–Schmidt inner product: For $A, B \in \text{Lin}(U, V)$, we define $\langle A, B \rangle_{\text{Lin}(U, V)} = \sum_{i \in [\dim(U)]} \langle A \cdot \mathbf{e}_i, B \cdot \mathbf{e}_i \rangle_V$ where $\mathbf{e}_1, \dots, \mathbf{e}_{\dim(U)}$ is an orthonormal basis of U with respect to $\langle \cdot, \cdot \rangle_U$. It is shown easily that this inner product is independent of the choice of the orthonormal basis, and it matches the usual Frobenius inner product on the space of $\dim(U) \times \dim(V)$ matrices, when A and B are represented as matrices under a choice of orthonormal basis for U and V .

For any linear map $A \in \text{Lin}(U, V)$, we shall use $\|A\|_2$ to denote its operator norm, and $\|A\|_F = \langle A, A \rangle_{\text{Lin}(U, V)}^{1/2}$ to denote its norm under the Hilbert-Schmidt inner product.

For any linear map $A \in \text{Lin}(U, V)$, we use $\sigma_n(A)$ to denote the n^{th} largest singular value of A . We also use $\sigma_{-n}(A)$ to denote the n^{th} smallest singular value.

If $U \subseteq V$ is a subspace, then $U^\perp \subseteq V$ shall denote the subspace that is an the orthogonal complement of U , and $\text{Proj}_U \in \text{Lin}(V, V)$ shall denote the projection onto U , i.e.

$$U^\perp \stackrel{\text{def}}{=} \{\mathbf{w} \in V : \langle \mathbf{u}, \mathbf{w} \rangle = 0\} \subseteq V \quad \text{and}$$

$$\text{Proj}_U \cdot \mathbf{v} = \mathbf{u} \quad \text{where } \mathbf{u} \in U, \mathbf{w} \in U^\perp \text{ are the unique vectors such that } \mathbf{v} = \mathbf{u} + \mathbf{w}.$$

Tuples of subspaces. Let $\mathbf{U} = (U_1, U_2, \dots, U_s)$, $U_i \subseteq V$, be an s -tuple of subspaces. $\langle \mathbf{U} \rangle$ shall denote the span of the constituent subspaces, i.e.

$$\langle \mathbf{U} \rangle \stackrel{\text{def}}{=} U_1 + U_2 + \dots + U_s.$$

We will be interested in recovering the constituents of a subspace tuple \mathbf{U} using operators acting on $\langle \mathbf{U} \rangle$. Towards this end, we fix some relevant terminology.

Associated matrices, independent tuples of subspaces, condition numbers. Let $\mathbf{U} = (U_1, U_2, \dots, U_s)$ be an s -tuple of subspaces of an n -dimensional vector space V , and $d_i = \dim(U_i)$ and $d = \sum_{i \in [s]} d_i$. We will say that an $n \times d$ matrix M is a \mathbf{U} -associated matrix if and only if

³⁷Very often, V is a space of homogeneous multivariate polynomials in which case the inner product is the Bombieri inner product.

the first set of d_1 columns of M forms an orthonormal basis for U_1 , the next set of d_2 columns forms an orthonormal basis for U_2 and so on and the last set of d_s columns forms an orthonormal basis of U_s . Any two \mathbf{U} -associated matrices are equal up to right multiplication by an orthogonal matrix and so the rank and condition number of all matrices associated to \mathbf{U} are the same.

We will say that $\mathbf{U} = (U_1, U_2, \dots, U_s)$ is an *independent* tuple of subspaces if

$$\langle \mathbf{U} \rangle = U_1 \oplus U_2 \oplus \dots \oplus U_s.$$

In particular, \mathbf{U} is an independent tuple of subspaces if and only if any \mathbf{U} -associated matrix has full rank. Motivated by this, we define a measure of the "robustness" of independence of the U_i 's, called the condition number $\kappa(\mathbf{U})$ of the tuple \mathbf{U} , as the condition number $\kappa(M)$ of a \mathbf{U} -associated matrix M .

Distances between Subspaces. The distance $\text{dist}(U, V)$ between subspaces $U, V \subseteq W$ will be defined by

$$\text{dist}(U, V) = \|\text{Proj}_U - \text{Proj}_V\|_2,$$

and correspondingly we will say that these two subspaces are ϵ -close if $\text{dist}(U, V) \leq \epsilon$. In particular, observe that for any two subspaces $U, V \subseteq W$, it holds that $\text{dist}(U, V) \in [0, 1]$, and that $\text{dist}(U, V) = 1$ if $\dim(U) \neq \dim(V)$.

The distance between two s -tuples $\mathbf{U} = (U_1, U_2, \dots, U_s)$ and $\mathbf{V} = (V_1, V_2, \dots, V_s)$ of subspaces of a vector space W , is defined as $\text{dist}(\mathbf{U}, \mathbf{V}) \stackrel{\text{def}}{=} \max_{i \in [s]} \text{dist}(U_i, V_i)$.

Tuples of operators. Let $\mathcal{B} = (B_1, B_2, \dots, B_m) \in (\text{Lin}(U, V))^m$ be an m -tuple of linear operators. For $\mathbf{u} \in U$, $\langle \mathcal{B} \cdot \mathbf{u} \rangle$ shall denote the space spanned by $\{B_i \cdot \mathbf{u} : i \in [m]\}$. Similarly, for a subspace $U' \subseteq U$, $\langle \mathcal{B} \cdot U' \rangle$ shall denote the space spanned by $\{B_i \cdot \mathbf{u} : \mathbf{u} \in U', i \in [m]\}$.

Definition 2.1. Corresponding to any such m -tuple $\mathcal{B} = (B_1, \dots, B_m)$ of operators, we shall associate a linear map $\hat{\mathcal{B}} \in \text{Lin}(U, V^m)$, given by

$$\hat{\mathcal{B}} \cdot \mathbf{u} = (B_1 \cdot \mathbf{u}, \dots, B_m \cdot \mathbf{u}).$$

Further, we shall define the norm $\|\mathcal{B}\|_2$ of the tuple of operators, to be the norm $\|\hat{\mathcal{B}}\|_2$. More generally, we define the i -th singular value of \mathcal{B} to be the i -th singular value of $\hat{\mathcal{B}}$, and $\kappa(\mathcal{B}) = \|\hat{\mathcal{B}}\|_2 \cdot \|\hat{\mathcal{B}}^\dagger\|_2$.

Note that it holds trivially that $\|\mathcal{B}\|_2 \leq \sqrt{m} \cdot \max_{i \in [m]} \|B_i\|_2$.

Vector Spaces of Homogenous Polynomials. We denote by $\mathbb{R}[\mathbf{x}]^d$ the space of degree d homogenous polynomials in n variables $\mathbf{x} = (x_1, \dots, x_n)$. Let \mathbb{N}_d^n denote the set of multi-indices i.e. the set of n -tuples of non-negative integers $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$ such that $|\alpha| = \alpha_1 + \dots + \alpha_n = d$. Then, for $\alpha \in \mathbb{N}_d^n$, we use \mathbf{x}^α to denote the monomial $x_1^{\alpha_1} x_2^{\alpha_2} \dots x_n^{\alpha_n} \in \mathbb{R}[\mathbf{x}]^d$. These monomials form a basis of $\mathbb{R}[\mathbf{x}]^d$, and we have $\dim \mathbb{R}[\mathbf{x}]^d = |\mathbb{N}_d^n| = \binom{n+d-1}{d}$.

We endow $\mathbb{R}[\mathbf{x}]^d$ with the Bombieri inner product, defined on the monomials as follows:

$$\langle \mathbf{x}^\alpha, \mathbf{x}^\beta \rangle_B = \begin{cases} \frac{\alpha!}{d!} & \text{if } \alpha = \beta, \\ 0 & \text{otherwise.} \end{cases}$$

where $\boldsymbol{\alpha}! = \alpha_1! \alpha_2! \cdots \alpha_n!$. The Bombieri basis i.e. the orthonormal basis with respect to this inner product is the basis of scaled monomials: $p_{\boldsymbol{\alpha}}(\mathbf{x}) = \sqrt{\frac{d!}{\boldsymbol{\alpha}!}} \mathbf{x}^{\boldsymbol{\alpha}}$, $\boldsymbol{\alpha} \in \mathbb{N}_d^n$.

We use the following properties of the Bombieri inner product. For any homogeneous polynomials p, q , and any $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, it holds:

- $\|p \cdot q\|_B \leq \|p\|_B \|q\|_B$.
- $\langle (\mathbf{a} \cdot \mathbf{x})^d, (\mathbf{b} \cdot \mathbf{x})^d \rangle_B = \langle \mathbf{a}, \mathbf{b} \rangle^d$, where $\mathbf{a} \cdot \mathbf{x} = \sum_{i \in [n]} a_i x_i \in \mathbb{R}[\mathbf{x}]^d$.

Lemma 2.1. Let $p \in \mathbb{R}[\mathbf{x}]^d$ be any polynomial, where $\mathbf{x} = (x_1, \dots, x_n)$. Then, we have that

$$\sum_{i=1}^n \|\partial_{x_i} p\|_B^2 = d^2 \|p\|_B^2.$$

Proof. Let $p(\mathbf{x}) = \sum_{\boldsymbol{\alpha} \in \mathbb{N}_d^n} c_{\boldsymbol{\alpha}} \mathbf{x}^{\boldsymbol{\alpha}}$. Then, $\|p\|_B^2 = \sum_{\boldsymbol{\alpha}} c_{\boldsymbol{\alpha}}^2 \cdot \frac{\boldsymbol{\alpha}!}{d!}$, and

$$\sum_{i \in [n]} \|\partial_{x_i} p\|_B^2 = \sum_{i \in [n]} \sum_{\boldsymbol{\alpha}: \alpha_i > 0} (\alpha_i c_{\boldsymbol{\alpha}})^2 \cdot \frac{\boldsymbol{\alpha}! / \alpha_i}{(d-1)!} = \sum_{i \in [n]} \sum_{\boldsymbol{\alpha}} c_{\boldsymbol{\alpha}}^2 \cdot \frac{\boldsymbol{\alpha}!}{d!} \cdot (d \alpha_i) = d^2 \cdot \sum_{\boldsymbol{\alpha}} c_{\boldsymbol{\alpha}}^2 \cdot \frac{\boldsymbol{\alpha}!}{d!}. \quad \square$$

3 Robust Recovery from Scaling Maps (RRSM)

In this section, we look at a special case of the robust vector space decomposition problem, involving linear maps that correspond to scaling the component subspaces. This shall later be used in our algorithm for the general robust vector decomposition problem as a sub-routine.

Let $\mathbf{U} = (U_1, U_2, \dots, U_s)$ be an *independent* s -tuple of subspaces in W and let $U = \langle \mathbf{U} \rangle \subseteq W$.

Definition 3.1. Space of Scaling Maps. The space of scaling maps $S(\mathbf{U}) \subseteq \text{Lin}(U, U)$ is defined as:

$$S(\mathbf{U}) \stackrel{\text{def}}{=} \{A \in \text{Lin}(U, U) : \exists \lambda_1, \lambda_2, \dots, \lambda_s \in \mathbb{R} \text{ such that } \forall i \in [s], \mathbf{u}_i \in U_i \text{ we have } A \cdot \mathbf{u}_i = \lambda_i \mathbf{u}_i\}.$$

In other words, in a basis of U obtained by concatenating the bases of U_i 's, the space $S(\mathbf{U})$ consists of block diagonal matrices wherein the i -th diagonal block is a scalar multiple of the identity matrix of size $\dim(U_i)$. For ease of notation, we will use S to denote $S(\mathbf{U})$.

We are interested in the following problem:

Problem 5. Robust Recovery from Scaling Maps (RRSM). We are given as input the integer s , a vector space $\tilde{U} \subseteq W$, and a space $\tilde{S} \subseteq \text{Lin}(\tilde{U}, \tilde{U})$ of linear operators on \tilde{U} . It is known that $\text{dist}(\tilde{U}, U)$ and $\text{dist}(\tilde{S}, S(\mathbf{U}))$ are "small," and our goal is to *efficiently* find an s -tuple $\tilde{\mathbf{U}} = (\tilde{U}_1, \tilde{U}_2, \dots, \tilde{U}_s)$ of subspaces in $\tilde{U} \subseteq W$, such that (upto reordering) $\text{dist}(\tilde{\mathbf{U}}, \mathbf{U})$ is "small."

In the above formulation, $\text{dist}(\tilde{S}, S)$ is defined as follows: We extend the space S (resp. \tilde{S}) to be a subspace of $\text{Lin}(W, W)$, by extending each $A \in S$ (resp. $A \in \tilde{S}$) to be zero on U^\perp (resp. \tilde{U}^\perp). Then, the distance between \tilde{S} and S is defined using the Hilbert-Schmidt inner product on $\text{Lin}(W, W)$.

3.1 RRSM: Algorithm for the Noiseless Case

First, we consider the noiseless case of Problem 5, in which $\tilde{U} = U$ and $\tilde{S} = S$ are exactly known, and we wish to recover U_1, \dots, U_s exactly.

In this case it is trivially easy to recover the constituent U_i 's (up to permutation):

Exact Algorithm 1: pick a random $A \in S$, diagonalize it and output the eigenspaces corresponding to distinct eigenvalues of A . For a random A , eigenvalues corresponding to distinct U_i 's will be distinct with high probability, and the algorithm answers correctly.

Next, we will give another (slightly more complicated) algorithm for this exact case. We expect the robust version of this algorithm to have better tolerance to noise (when the noise is random) compared to the algorithm described above.

Definition 3.2. Projection Maps. For each $i \in [s]$, the projection map $P_i \in \text{Lin}(U, U)$ is defined as

$$P_i \cdot (\mathbf{u}_1 + \dots + \mathbf{u}_s) = \mathbf{u}_i,$$

where $\mathbf{u}_j \in U_j$ for each $j \in [s]$. Further, we define the map $\hat{M} : \mathbb{R}^s \rightarrow \text{Lin}(U, U)$ by

$$\hat{M}(\lambda_1, \dots, \lambda_s) = \sum_{i=1}^s \lambda_i P_i.$$

Note that the maps P_i 's do not correspond to orthogonal projections if the spaces U_i 's are not orthogonal to each other.

Observe that:

1. The space of scaling maps $S = \text{span}(P_1, \dots, P_s)$ in $\text{Lin}(U, U)$.
2. Recovering $\mathbf{U} = (U_1, \dots, U_s)$ is equivalent to recovering (P_1, \dots, P_s) .

Next, given any $A \in S$, we consider the action (by left multiplication) of this map A on the space S itself:

Definition 3.3. Given any map $A \in S \subseteq \text{Lin}(U, U)$, we define the map $\hat{A} \in \text{Lin}(S, S)$ by $\hat{A} \cdot B = A \cdot B$ for all $B \in S$.

This map \hat{A} is well-defined since the space S is closed under composition of maps, and, if $A = \sum_{i=1}^s \lambda_i P_i$, then \hat{A} has eigenvalues $\lambda_1, \dots, \lambda_s$ with eigenvectors P_1, \dots, P_s respectively.

Exact Algorithm 2: pick a random $A \in S$, compute the map $\hat{A} \in \text{Lin}(S, S)$, and diagonalize it. With high probability, its eigenvectors are P_1, \dots, P_s (appropriately scaled), and the spaces U_1, \dots, U_s are the images of these maps.

3.2 RRSM: Algorithm for the Robust Case

In the robust case, we are given a subspace \tilde{U} which is "close" to U , and a space $\tilde{S} \subseteq \text{Lin}(\tilde{U}, \tilde{U})$ which is "close" to S , and we wish to *efficiently* recover a tuple of subspaces $\tilde{\mathbf{U}}$ close to \mathbf{U} . Formally, for the time complexity analysis, we shall assume that the input to the algorithm is given as follows: Let $\dim(W) = n$, $\dim(U) = d$, $\dim(S) = s \leq d$. The

vector space $\tilde{U} \subseteq W$ is given as nd field elements, consisting of an orthonormal basis of \tilde{U} with respect to some fixed orthonormal basis of W . The vector space $\tilde{S} \subseteq \text{Lin}(\tilde{U}, \tilde{U})$ is given as sd^2 field elements, consisting of an orthonormal basis of \tilde{S} with respect to the above orthonormal basis of \tilde{U} . The total input size is $N = nd + sd^2$.

Throughout this section, we shall let $d^* = \max_{i \in [s]} \dim(U_i)$ and $d_* = \min_{i \in [s]} \dim(U_i)$. Before giving an algorithm, we note that the performance of our algorithm will depend on how "well-separated" the component U_i 's are. For this purpose, we shall be interested in two condition numbers, namely $\kappa(\mathbf{U})$ and $\kappa(\hat{M})$ (see Definition 3.2). These satisfy the following relations:

Lemma 3.1.

$$\|\hat{M}\|_2 \leq \kappa(\mathbf{U}) \cdot \sqrt{d^*}, \quad \kappa(\hat{M}) \leq \kappa(\mathbf{U}) \cdot \sqrt{\frac{d^*}{d_*}}.$$

Proof. We defer the proof of this Lemma to Section E.2. □

We note that while the inequalities in the above lemma may be tight in the worst-case, we expect $\kappa(\hat{M}) \ll \kappa(\mathbf{U})$ in practice, since it sort of measures the "average separation" between the component subspaces U_i 's.

Next, we shall give robust versions of both the exact case-algorithms. We note that the performance of the robust version of the first exact-case algorithm depends on $\kappa(\mathbf{U})$, whereas that of the second case exact-case algorithm depends on $\kappa(\hat{M})$. As stated above, we expect $\kappa(\hat{M}) \ll \kappa(\mathbf{U})$, and so the first algorithm is expected to be worse than the second algorithm. For this reason, we will only analyze the second algorithm formally in this work. It also turns out that the second algorithm is technically a bit easier to analyze: as we will see, it only requires perturbation bounds on eigenvectors corresponding to simple eigenvalues (see Lemma A.8).

Robust Algorithm 1: pick a "random" (suitably defined) map $\tilde{A} \in \tilde{S}$, "cluster" the eigenvalues which are close together into s clusters, and output the eigenspaces corresponding to each cluster. We expect that if S and \tilde{S} are "sufficiently close" then \mathbf{U} and $\tilde{\mathbf{U}}$ are "fairly close".

Next, we will describe the robust version of the second algorithm for the exact case. First, we give an (approximate) analogue of Definition 3.3.

Definition 3.4. Given any map $\tilde{A} \in \tilde{S} \subseteq \text{Lin}(\tilde{U}, \tilde{U})$, we define the map $\hat{A} \in \text{Lin}(\tilde{S}, \tilde{S})$ by $\hat{A} \cdot \tilde{B} = \text{Proj}_{\tilde{S}}(\tilde{A} \cdot \tilde{B})$, where $\text{Proj}_{\tilde{S}} : \text{Lin}(\tilde{U}, \tilde{U}) \rightarrow \tilde{S}$ is the orthogonal projection onto \tilde{S} .

Robust Algorithm 2: pick a random $\tilde{A} \in \tilde{S}$, compute the map $\hat{A} \in \text{Lin}(\tilde{S}, \tilde{S})$, and diagonalize it. Let its eigenvectors be $\tilde{P}_1, \dots, \tilde{P}_s \in \tilde{S}$ ³⁸. For each $i \in [s]$, let $\tilde{U}_i \subseteq \tilde{U}$ be the span of the left singular vectors of the map \tilde{P}_i , with singular values "not too small." Output $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_s)$.

The above algorithm is formally described as Algorithm 1 and it gets the following guarantees:

³⁸ We will show in the analysis that with high probability the eigenvalues of \hat{A} are real and distinct, and hence the eigenvectors lie in the real vector space \tilde{S} .

Theorem 4. Robust Recovery from Scaling Maps. Let $\mathbf{U} = (U_1, \dots, U_s)$ be an independent s -tuple of subspaces in a vector space W , and let $U = \langle \mathbf{U} \rangle \subseteq W$. Let $S = S(\mathbf{U}) \subseteq \text{Lin}(U, U)$ be the space of scaling maps as defined in Definition 3.1, and let the map \hat{M} be as defined in Definition 3.2.

Let $\tilde{U} \subseteq W$ and $\tilde{S} \subseteq \text{Lin}(\tilde{U}, \tilde{U})$ be vector spaces, and let $\tau \in (0, 1)$ be such that:

1. $\text{dist}(\tilde{U}, U) < 1$.
2. $\text{dist}(\tilde{S}, S) \leq \epsilon < 1$, where the distance is measured after extending both S, \tilde{S} to subspaces of $\text{Lin}(W, W)$.
3. The parameter τ satisfies $\frac{1}{3} \cdot \frac{1}{\|\hat{M}\|_2} < \tau \leq \frac{2}{3} \cdot \frac{1}{\|\hat{M}\|_2}$.

Then, for any $\delta > 0$, Algorithm 1, on input $(W, \tilde{U}, \tilde{S}, \tau)$, runs in time $O(s^3 + s^2 d^\omega + s d^3 + s d^2 n) = O(N^{5/3})$, and outputs an s -tuple $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_s)$ of subspaces in \tilde{U} , such that with probability at least $1 - \delta$, it holds (upto reordering) that for each $i \in [s]$,

$$\begin{aligned} \text{dist}(U_i, \tilde{U}_i) &\leq 300 \cdot \kappa(\hat{M}) \cdot \|\hat{M}\|_2^2 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \frac{\epsilon}{\delta} \\ &\leq 300 \cdot \sqrt{\frac{d^{*3}}{d_*}} \cdot \kappa(\mathbf{U})^3 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \frac{\epsilon}{\delta}. \end{aligned}$$

Proof. We defer the proof of the Theorem to Section E. □

Remark 3. We notice that our algorithm uses an auxiliary parameter $\tau \in (0, 1)$ and the correctness of the algorithm depends on τ lying in a correct range. This is fine for our purposes, since in applications one can usually check the correctness of the final solution obtained; so it is possible to simply iterate over τ , halving it in each iteration, and checking the solution obtained for correctness. The number of iterations is at most logarithmic in the condition number: Theorem 4 shows that a valid τ is encountered in at most $O(\log_2 \|\hat{M}\|_2) = O(\log_2(\kappa(\mathbf{U}) \cdot d^*))$ iterations. As the error bounds in Theorem 4 depend on $\kappa(\mathbf{U})$, we will be mostly interested in the case where this condition number is not too large, and hence the runtime blow up is small. More formally, since we require $\kappa(\mathbf{U})^3 \cdot \epsilon \ll 1$, we will have a blow up of at most $O(\log_2 1/\epsilon)$.

Note that in Algorithm 1, there is another natural way to get \tilde{U}_i once the map \tilde{P}_i is known: we can simply let \tilde{U}_i be the span of the left singular vectors of the map \tilde{P}_i , corresponding to the $\dim(U_i)$ largest singular vectors. However, it turns out that in applications, each of the s dimensions $\dim(U_i)$'s may not be known. Hence, we use a single threshold parameter $\tau \in (0, 1)$, and just consider all the singular vectors corresponding to singular value at least τ .

Algorithm 1 RRSM: Robust Recovery From Scaling Maps.

Input: $(W, \tilde{U}, \tilde{S}, \tau)$, $\tilde{U} \subseteq W$ is a subspace of vector space W , and $\tilde{S} \subseteq \text{Lin}(\tilde{U}, \tilde{U})$ is a subspace of $\dim(\tilde{S}) = s$, and $\tau \in (0, 1)$.

Assumptions: $\mathbf{U} = (U_1, U_2, \dots, U_s)$ is an independent s -tuple of subspaces in W , and $U = \langle \mathbf{U} \rangle$, $S = S(\mathbf{U})$ are such that:

1. $\text{dist}(\tilde{U}, U) < 1$.
2. $\text{dist}(\tilde{S}, S) \leq \epsilon < 1$, where the distance is measured after extending both to subspaces of $\text{Lin}(W, W)$.

Output: s -tuple $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_s)$ of subspaces in $\tilde{U} \subseteq W$ such that $\text{dist}(\tilde{\mathbf{U}}, \mathbf{U})$ is small.

- 1: Pick a random element $\tilde{A} \in \tilde{S}$ with $\|\tilde{A}\|_F = 1$ as follows:
 - (a) Let $\tilde{\mathbf{s}}_1, \dots, \tilde{\mathbf{s}}_s$ be any orthonormal basis of \tilde{S} .
 - (b) Pick $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_s) \in \mathbb{R}^s$, with each $\alpha_i \sim \mathcal{N}(0, 1)$ chosen independently.
 - (c) Let $\tilde{A} = \sum_{i=1}^s \frac{\alpha_i}{\|\boldsymbol{\alpha}\|_2} \cdot \tilde{\mathbf{s}}_i$.
 - 2: Compute the map $\hat{A} \in \text{Lin}(\tilde{S}, \tilde{S})$, as in Definition 3.4.
 - 3: Compute the eigen-decomposition of \hat{A} : Suppose that it has eigenvectors $\tilde{P}_1, \dots, \tilde{P}_s \in \tilde{S}$, with $\|\tilde{P}_i\|_F = 1$ for each $i \in [s]$ ³⁸.
 - 4: For each $i \in [s]$, let $\tilde{U}_i \subseteq \tilde{U}$ be the span of all left singular vectors of \tilde{P}_i , with singular value at least τ .
 - 5: Output $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_s)$.
-

Finally, we also show that if the error in Theorem 4 is small, then the recovered subspaces form a direct sum.

Proposition 3.1. *Let $U = U_1 \oplus \dots \oplus U_s \subseteq W$ be the direct sum of s -subspaces of a vector space W . Let $\tilde{U} \subseteq W$ be a subspace such that $\text{dist}(U, \tilde{U}) < 1$, and let $\tilde{U}_1, \dots, \tilde{U}_s \subseteq \tilde{U}$ be such that for each $i \in [s]$, $\text{dist}(\tilde{U}_i, U_i) \leq \gamma < 1$. If $2\gamma\sqrt{s} \cdot \kappa(\mathbf{U}) < 1$, then $\tilde{U} = \tilde{U}_1 \oplus \dots \oplus \tilde{U}_s$.*

Proof. We defer the proof to Section E.6. □

4 Robust Vector Space Decomposition (RVSD)

Let W_1 and W_2 be real vector spaces, and let $\mathbf{U} = (U_1, U_2, \dots, U_s)$ be an independent s -tuple of subspaces in W_1 , and let $\mathbf{V} = (V_1, V_2, \dots, V_s)$ be an independent s -tuple of subspaces in W_2 . Let $U \stackrel{\text{def}}{=} \langle \mathbf{U} \rangle \subseteq W_1$ and $V \stackrel{\text{def}}{=} \langle \mathbf{V} \rangle \subseteq W_2$, and let $\mathcal{B} = (B_1, B_2, \dots, B_m) \in (\text{Lin}(U, V))^m$ be an m -tuple of linear operators from U to V . Suppose that each U_i is mapped inside V_i under the action of \mathcal{B} , that is, for each $i \in [s]$, it holds that $\langle \mathcal{B} \cdot U_i \rangle \subseteq V_i$.

We are interested in the following problem:

Problem 6. Robust Vector Space Decomposition (RVSD). We are given as input the integer s , two vector spaces $\tilde{U} \subseteq W_1$ and $\tilde{V} \subseteq W_2$ such that $\text{dist}(\tilde{U}, U)$ and $\text{dist}(\tilde{V}, V)$

are "small," and a tuple of operators $\tilde{\mathcal{B}} = (\tilde{B}_1, \tilde{B}_2, \dots, \tilde{B}_m) \in \text{Lin}(\tilde{U}, \tilde{V})^m$, such that $\tilde{\mathcal{B}}$ is close to \mathcal{B} . Our goal is to *efficiently* find an s -tuple $\tilde{\mathbf{U}} = (\tilde{U}_1, \tilde{U}_2, \dots, \tilde{U}_s)$ of subspaces in $\tilde{U} \subseteq W_1$, such that (upto a common reordering of the components) $\text{dist}(\tilde{\mathbf{U}}, \mathbf{U})$ is "small."

Remark 4. Note that in applications, it is usually sufficient to only find the tuple \mathbf{U} approximately and so we frame our problem in this form. If one wishes to find \mathbf{V} approximately, our algorithm can easily be extended to do that (see Remark 6).

In the above formulation, we formally define closeness between $\tilde{\mathcal{B}}$ and \mathcal{B} as follows:

Definition 4.1. Let $\hat{B} \in \text{Lin}(U, V^m)$ (resp. $\hat{\tilde{B}} \in \text{Lin}(\tilde{U}, \tilde{V}^m)$) be the map corresponding to \mathcal{B} (resp. $\tilde{\mathcal{B}}$) defined as in Definition 2.1. We say that \mathcal{B} and $\tilde{\mathcal{B}}$ are ϵ -close if $\|\hat{B} - \hat{\tilde{B}}\|_2 \leq \epsilon \cdot \|\hat{B}\|_2 = \epsilon \|\mathcal{B}\|_2$, where the difference is taken by viewing $\hat{B}, \hat{\tilde{B}}$ as elements of $\text{Lin}(W_1, W_2^m)$ (by defining them to be zero on U^\perp, \tilde{U}^\perp respectively).

The exact version of the problem (where $\tilde{U} = U, \tilde{V} = V, \tilde{\mathcal{B}} = \mathcal{B}$, and we wish to find \mathbf{U}, \mathbf{V} exactly) first appeared in in [GKS20], where they give an algorithm to solve it efficiently (in some special cases), and use it to recover individual components in a "sum of powers of low-degree polynomials." Building on the ideas in [GKS20], we give a very general algorithm to solve this robust version of the vector space decomposition problem.

4.1 RVSD: Algorithm for the Noiseless Case

We look at the exact version of RVSD where the spaces $\tilde{U} = U, \tilde{V} = V, \tilde{\mathcal{B}} = \mathcal{B}$ are exactly known. We follow the sketch defined in Section 1.1.

Definition 4.2. Map Corresponding to the Adjoint Algebra. The *adjoint algebra map* corresponding to the subspaces U and V and the operator tuple \mathcal{B} , denoted by $\mathfrak{A}_{U,V}(\mathcal{B}) : \text{Lin}(U, U) \times \text{Lin}(V, V) \rightarrow \text{Lin}(U, V)^m$, is defined as

$$\mathfrak{A}_{U,V}(\mathcal{B}) \cdot (D, E) = (B_1 D - E B_1, \dots, B_m D - E B_m).$$

For ease of notation, we shall simply use \mathfrak{A} to denote $\mathfrak{A}_{U,V}(\mathcal{B})$.

Definition 4.3 (Adjoint algebra, [CIK97, GKS20, Qia18]). The adjoint algebra, corresponding to the vector spaces U, V , and the tuple of operators \mathcal{B} , denoted $\text{Adj}_{U,V}(\mathcal{B}) \subseteq \text{Lin}(U, U) \times \text{Lin}(V, V)$, is defined to be the null space of the map \mathfrak{A} , given by

$$\text{Adj}_{U,V}(\mathcal{B}) \stackrel{\text{def}}{=} \{(D, E) : B_j \cdot D = E \cdot B_j \text{ for all } j \in [m]\}.$$

For ease of notation, we shall simply use Adj to denote $\text{Adj}_{U,V}(\mathcal{B})$. Also, we shall use $\text{Adj}_1 \subseteq \text{Lin}(U, U)$ to denote the projection of Adj onto the "U-part," formally defined as:

$$\text{Adj}_1 \stackrel{\text{def}}{=} \{D : \exists E \text{ such that } (D, E) \in \text{Adj}\}.$$

Definition 4.4. We define $\mathbf{U} \times \mathbf{V}$ to be the s -tuple of subspaces given by $\mathbf{U} \times \mathbf{V} = (U_1 \times V_1, \dots, U_s \times V_s)$, which satisfies $\langle \mathbf{U} \times \mathbf{V} \rangle = U \times V$.

Observe that the space of scaling maps $S(\mathbf{U} \times \mathbf{V})$ (see Definition 3.1) is always contained in the adjoint algebra Adj , that is, $S(\mathbf{U} \times \mathbf{V}) \subseteq \text{Adj}$. The next proposition shows that if the two are equal, then the decomposition of U and V into the s components is unique, and further indecomposable.

Proposition 4.1 (Proposition A.3 in [GKS20]). *If $\dim(\text{Adj}) = s$, or equivalently, if $\text{Adj} = S(\mathbf{U} \times \mathbf{V})$, then:*

The decomposition $U = U_1 \oplus \dots \oplus U_s, V = V_1 \oplus \dots \oplus V_s$ is the unique irreducible decomposition satisfying $\langle \mathcal{B} \cdot U_i \rangle \subseteq V_i$ for each $i \in [s]$. That is, if

$$U = \hat{U}_1 \oplus \hat{U}_2 \oplus \dots \oplus \hat{U}_{\hat{s}} \quad \text{and} \quad V = \hat{V}_1 \oplus \hat{V}_2 \oplus \dots \oplus \hat{V}_{\hat{s}}, \quad \hat{s} \geq s,$$

and

$$\forall i \in [\hat{s}] \quad \langle \mathcal{B} \cdot \hat{U}_i \rangle \subseteq \hat{V}_i,$$

then $\hat{s} = s$ and upto reordering if necessary, $\hat{U}_i = U_i$ and $\hat{V}_i = V_i$ for all $i \in [s]$.

Also, the above implies that $\text{Adj}_1 = S(\mathbf{U})$, and that $V_i = \langle \mathcal{B} \cdot U_i \rangle$ for each $i \in [s]$.

Based on the above proposition, the following simple algorithm recovers the U_i 's and the V_i 's under the assumption that $\text{Adj} = S(\mathbf{U} \times \mathbf{V})$:

Exact Algorithm 1: Compute the adjoint algebra Adj (by solving a system of linear equations). Run the exact algorithm for RRSB (see Section 3.1) on Adj with respect to the space $U \times V$; suppose that the output is $T_1, \dots, T_s \subseteq U \times V$. For each $i \in [s]$, let U_i, V_i be the projection of T_i on the " U, V -parts" respectively.

We also give a slightly different variation of the above algorithm, which turns out to be easier to analyze in the robust case.

Exact Algorithm 2: Compute the adjoint algebra Adj , and then compute Adj_1 . Run the exact algorithm for RRSB on Adj_1 with respect to the space U ; the output is (U_1, \dots, U_s) . For each $i \in [s]$, compute $V_i = \langle \mathcal{B} \cdot U_i \rangle$.

4.2 RVSD: Algorithm for the Robust Case

Next, we give a robust version of the (second) exact-case algorithm. We are given a subspaces $\tilde{U} \approx U, \tilde{V} \approx V$, and an m -tuple of operators $\tilde{\mathcal{B}} \approx \mathcal{B}$, and we wish to recover a tuple of subspaces $\tilde{\mathbf{U}} \approx \mathbf{U}$. Formally, for the time complexity analysis, we shall assume that the input to the algorithm is given as follows: Let $\dim(W_1) = n_1, \dim(W_2) = n_2, \dim(U) = d_1, \dim(V) = d_2$. The vector space $\tilde{U} \subseteq W_1$ (resp. $\tilde{V} \subseteq W_2$) is given as $n_1 d_1$ (resp. $n_2 d_2$) field elements, consisting of an orthonormal basis of \tilde{U} (resp. \tilde{V}), with respect to some fixed orthonormal basis of W_1 (resp. W_2). The m -tuple of operators $\tilde{\mathcal{B}} = (\tilde{B}_1, \dots, \tilde{B}_m) \in \text{Lin}(\tilde{U}, \tilde{V})^m$ is given as $md_1 d_2$ field elements, with each \tilde{B}_i given as a matrix with respect to the above orthonormal basis of \tilde{U} and \tilde{V} . The total input size is $N = n_1 d_1 + n_2 d_2 + md_1 d_2$.

Definition 4.5. (Approximate) Map Corresponding to the Adjoint Algebra. The (approximate) *adjoint algebra map* corresponding to the subspaces \tilde{U} and \tilde{V} and the operator tuple $\tilde{\mathcal{B}}$, denoted by $\tilde{\mathfrak{A}}_{\tilde{U}, \tilde{V}}(\tilde{\mathcal{B}}) : \text{Lin}(\tilde{U}, \tilde{U}) \times \text{Lin}(\tilde{V}, \tilde{V}) \rightarrow \text{Lin}(\tilde{U}, \tilde{V})^m$, is defined as

$$\tilde{\mathfrak{A}}_{\tilde{U}, \tilde{V}}(\tilde{\mathcal{B}}) \cdot (D, E) = (\tilde{B}_1 D - E \tilde{B}_1, \dots, \tilde{B}_m D - E \tilde{B}_m).$$

For ease of notation, we shall use $\tilde{\mathfrak{A}}$ to denote $\tilde{\mathfrak{A}}_{\tilde{U}, \tilde{V}}(\tilde{\mathcal{B}})$.

Definition 4.6. (Approximate) Adjoint Algebra. The (approximate) *adjoint algebra* corresponding to the subspaces \tilde{U} and \tilde{V} and the operator tuple $\tilde{\mathcal{B}}$, denoted by $\widetilde{\text{Adj}}_{U,V}(\mathcal{B}) \subseteq \text{Lin}(\tilde{U}, \tilde{U}) \times \text{Lin}(\tilde{V}, \tilde{V})$, is defined to be the vector space spanned by the right singular vectors of the map $\tilde{\mathcal{A}}$, corresponding to the s smallest singular values. Note that the relevant inner product on $\text{Lin}(\tilde{U}, \tilde{U}) \times \text{Lin}(\tilde{V}, \tilde{V})$ is the direct sum of the two inner products in the natural way.

For ease of notation, we shall simply use $\widetilde{\text{Adj}}$ to denote $\widetilde{\text{Adj}}_{U,V}(\mathcal{B})$. Note that by definition, the dimension of $\widetilde{\text{Adj}}$ is equal to s .

Further, we define $\widetilde{\text{Adj}}_1 \stackrel{\text{def}}{=} \left\{ D : \exists E \text{ such that } (D, E) \in \widetilde{\text{Adj}} \right\}$.

Based on the above definitions, we next give an algorithm for the RVSD problem. It relies on the problem of Robust Recovery from Scaling Maps (RRSM), which we discussed in Section 3.

Robust Algorithm: Compute the map $\tilde{\mathcal{A}}$ and the adjoint algebras $\widetilde{\text{Adj}}$ and $\widetilde{\text{Adj}}_1$. Run RRSM algorithm on $\widetilde{\text{Adj}}_1$ with respect to the space $\tilde{U} \subseteq W_1$; let the output be $(\tilde{U}_1, \dots, \tilde{U}_s)$. Output $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_s)$.

The above algorithm is formally described as Algorithm 2 and it gets the following guarantees:

Theorem 5. Robust Vector Space Decomposition. Let $\mathbf{U} = (U_1, \dots, U_s)$ and $\mathbf{V} = (V_1, \dots, V_s)$ be independent s -tuple of subspaces in vector spaces W_1 and W_2 respectively, and let $U = \langle \mathbf{U} \rangle, V = \langle \mathbf{V} \rangle$. Let $\mathcal{B} = (B_1, \dots, B_m) \in \text{Lin}(U, V)^m$ be an m -tuple of operators such that for each $i \in [s], \langle \mathcal{B} \cdot U_i \rangle \subseteq V_i$.

Let the map $\hat{M} : \mathbb{R}^s \rightarrow \text{Lin}(U, U)$ corresponding to the space U be defined as in Definition 3.2. Let $\sigma_{-(s+1)}(\mathcal{A})$ denote the $(s+1)^{\text{th}}$ -smallest singular value of the adjoint-algebra map \mathcal{A} (see Definition 4.2), and let $\|\mathcal{B}\|_2$ be defined as in Definition 2.1.

Suppose that $\tilde{U} \subseteq W_1, \tilde{V} \subseteq W_2, \tilde{\mathcal{B}} \in \text{Lin}(\tilde{U}, \tilde{V})^m$, and $\tau \in (0, 1)$ are such that:

1. $\text{dist}(U, \tilde{U}) \leq \epsilon_1 < 1, \text{dist}(V, \tilde{V}) \leq \epsilon_2 < 1$, and that the operator tuples $\mathcal{B}, \tilde{\mathcal{B}}$ are ϵ -close (see Definition 4.1).
2. $\dim(\text{Adj}) = s$ (see Definition 4.3).
3. The parameter τ satisfies $\frac{1}{3} \cdot \frac{1}{\|\hat{M}\|_2} < \tau \leq \frac{2}{3} \cdot \frac{1}{\|\hat{M}\|_2}$.

Then, for any $\delta > 0$ we have that Algorithm 2, on input $(W_1, W_2, s, \tilde{U}, \tilde{V}, \tilde{\mathcal{B}}, \tau)$, runs in time $O((md_1d_2 \cdot (d_1^2 + d_2^2))^3 + sd_1^2n_1) = O(N^6)$, and outputs an s -tuple $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_s)$ of subspaces in \tilde{U} , such that with probability at least $1 - \delta$, it holds (upto reordering) that for each $i \in [s]$,

$$\begin{aligned} \text{dist}(U_i, \tilde{U}_i) &\leq 1800 \cdot \kappa(\hat{M}) \cdot \|\hat{M}\|_2^2 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \frac{\epsilon + \epsilon_1 + \epsilon_2}{\delta} \cdot \frac{\|\mathcal{B}\|_2}{\sigma_{-(s+1)}(\mathcal{A})} \\ &\leq 1800 \cdot \sqrt{\frac{d^{*3}}{d_*}} \cdot \kappa(\mathbf{U})^3 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \frac{\epsilon + \epsilon_1 + \epsilon_2}{\delta} \cdot \frac{\|\mathcal{B}\|_2}{\sigma_{-(s+1)}(\mathcal{A})}. \end{aligned}$$

Furthermore, if $2\gamma\sqrt{s} \cdot \kappa(\mathbf{U}) < 1$, then $\tilde{\mathbf{U}} = \tilde{U}_1 \oplus \dots \oplus \tilde{U}_s$, where γ denotes the above error bound.

Proof. We defer the proof of the above theorem to Section F. \square

Remark 5. Note that the assumption that the integer s is given to the algorithm as input is merely for simplicity. If s is not known, its value can be determined by looking at the singular values of the map $\tilde{\mathfrak{A}}$ corresponding to the adjoint algebra: the smallest s singular values are usually very close to zero, and are "much smaller" than the $(s + 1)^{\text{th}}$ smallest singular value. Alternatively, we can just iterate over s , since in applications we can usually check the correctness of the final solution obtained.

Also, the algorithm having access to the correct value of τ can be handled using the iteration strategy mentioned in Remark 3.

Remark 6. Note that if we wish to find $\mathbf{V} = (V_1, \dots, V_s)$ approximately as well, one can run the RRSMD algorithm on $(W_1 \times W_2, \tilde{U} \times \tilde{V}, \tilde{\text{Adj}}, \tau)$, to recover approximate versions of $U_1 \times V_1, \dots, U_s \times V_s$. In this case, the error bounds will be the same as in Theorem 5, with the definition of the map \hat{M} changed appropriately to $\hat{M} : \mathbb{R}^s \rightarrow \text{Lin}(U \times V, U \times V)$.

Algorithm 2 RVSD: Robust Vector Space Decomposition.

Input: $(W_1, W_2, s, \tilde{U}, \tilde{V}, \tilde{\mathcal{B}}, \tau)$, where s is a positive integer, $\tilde{U} \subseteq W_1, \tilde{V} \subseteq W_2$ are subspaces of vector spaces W_1, W_2 respectively, $\tilde{\mathcal{B}} = (\tilde{B}_1, \dots, \tilde{B}_m) \in \text{Lin}(\tilde{U}, \tilde{V})^m$ is an m -tuple of linear operators, and $\tau \in (0, 1)$.

Assumptions: $\mathbf{U} = (U_1, \dots, U_s)$ and $\mathbf{V} = (V_1, \dots, V_s)$ are independent s -tuples of subspaces in W_1, W_2 respectively, and $U = \langle \mathbf{U} \rangle, V = \langle \mathbf{V} \rangle$, and $\mathcal{B} = (B_1, \dots, B_m) \in \text{Lin}(U, V)^m$ are such that:

1. For each $i \in [s]$, it holds that $\langle \mathcal{B} \cdot U_i \rangle \subseteq V_i$.
2. $\text{dist}(U, \tilde{U}) \leq \epsilon_1 < 1$ and $\text{dist}(V, \tilde{V}) \leq \epsilon_2 < 1$.
3. \mathcal{B} and $\tilde{\mathcal{B}}$ are ϵ -close, according to Definition 4.1.

Output: s -tuple of subspaces $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_s)$ in $\tilde{U} \subseteq W_1$ such that $\text{dist}(\tilde{\mathbf{U}}, \mathbf{U})$ is small.

- 1: Compute a singular value decomposition of the adjoint algebra map $\tilde{\mathfrak{A}}$ as defined in Definition 4.5.
 - 2: Compute the approximate adjoint algebra $\tilde{\text{Adj}} \subseteq \text{Lin}(\tilde{U}, \tilde{U}) \times \text{Lin}(\tilde{V}, \tilde{V})$ and $\tilde{\text{Adj}}_1 \subseteq \text{Lin}(\tilde{U}, \tilde{U})$, as defined in Definition 4.6.
 - 3: Run Robust Recovery from Scaling Maps (RRSM, Algorithm 1) on $(W_1, \tilde{U}, \tilde{\text{Adj}}_1, \tau)$, and let $(\tilde{U}_1, \dots, \tilde{U}_s)$ be the s -tuple of subspaces in $\tilde{U} \subseteq W_1$ it outputs.
 - 4: Output $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_s)$.
-

4.3 RVSD: Using a Common Tuple of Operators on a Larger Space

It is often the case in applications that a tuple of operators $\mathcal{B} = (B_1, \dots, B_m) \in \text{Lin}(W_1, W_2)^m$ are known exactly, and these satisfy $\langle \mathcal{B} \cdot U_i \rangle \subseteq V_i$ for each $i \in [s]$. In this case, we work with the relevant projections of these operators to $\text{Lin}(U, V)$ and $\text{Lin}(\tilde{U}, \tilde{V})$ respectively.

Definition 4.7. Projected Tuple of Operators. Let $\mathcal{B} = (B_1, \dots, B_m) \in \text{Lin}(W_1, W_2)^m$ be an m -tuple of operators from W_1 to W_2 . For subspaces $U \subseteq W_1, V \subseteq W_2$, we define

the operator tuple $\mathcal{C} = (C_1, \dots, C_m) \in \text{Lin}(W_1, W_2)^m$ as follows: For each $j \in [m]$,

$$C_j = \text{Proj}_V \cdot B_j \cdot \text{Proj}_U,$$

where $\text{Proj}_U : W_1 \rightarrow W_1$, $\text{Proj}_V : W_2 \rightarrow W_2$ are the orthogonal projection maps onto U, V respectively. Observe that each C_j maps the space U into V , and is the zero map on U^\perp .

Lemma 4.1. *Let $\mathcal{B} = (B_1, \dots, B_m) \in \text{Lin}(W_1, W_2)^m$ be an m -tuple of operators from W_1 to W_2 , and let $U, \tilde{U} \subseteq W_1$ and $V, \tilde{V} \subseteq W_2$ be subspaces satisfying $\text{dist}(U, \tilde{U}) \leq \epsilon_1$ and $\text{dist}(V, \tilde{V}) \leq \epsilon_2$. Let \mathcal{C} (resp. $\tilde{\mathcal{C}}$) be the projected tuple of operators with respect to U, V (resp. \tilde{U}, \tilde{V}) according to Definition 4.7.*

Then, $\|\mathcal{C}\|_2 \leq \|\mathcal{B}\|_2$, and \mathcal{C} and $\tilde{\mathcal{C}}$ are ϵ -close (see Definition 4.1), for $\epsilon = (\epsilon_1 + \epsilon_2) \cdot \frac{\|\mathcal{B}\|_2}{\|\mathcal{C}\|_2}$.

Proof. By Definition 2.1 and Definition 4.7, we have for any $\mathbf{w} \in W_1$,

$$\left\| (\hat{\mathcal{C}} - \hat{\tilde{\mathcal{C}}}) \cdot \mathbf{w} \right\|_2 = \left\| \left((\text{Proj}_V \cdot B_i \cdot \text{Proj}_U - \text{Proj}_{\tilde{V}} \cdot B_i \cdot \text{Proj}_{\tilde{U}}) \cdot \mathbf{w} \right)_{i=1}^m \right\|_2$$

Then, using $\|\text{Proj}_U - \text{Proj}_{\tilde{U}}\|_2 \leq \epsilon_1$, $\|\text{Proj}_V - \text{Proj}_{\tilde{V}}\|_2 \leq \epsilon_2$, and the triangle inequality, we get that

$$\left\| \hat{\mathcal{C}} - \hat{\tilde{\mathcal{C}}} \right\|_2 \leq (\epsilon_1 + \epsilon_2) \cdot \|\mathcal{B}\|_2.$$

Similarly, the fact $\|\mathcal{C}\|_2 \leq \|\mathcal{B}\|_2$ follows easily. \square

Now, Lemma 4.1, combined with Theorem 5, immediately gives us the following corollary:

Corollary 4.1. *Suppose that a tuple of operators $\mathcal{B} = (B_1, \dots, B_m) \in \text{Lin}(W_1, W_2)^m$ is known exactly, where for each $i \in [s]$, $\langle \mathcal{B} \cdot \mathbf{U}_i \rangle \subseteq V_i$. Then, under the assumptions of Theorem 5, Algorithm 2 can recover the tuple \mathbf{U} of subspaces, upto error*

$$\begin{aligned} \gamma &= 3600 \cdot \kappa(\hat{M}) \cdot \|\hat{M}\|_2^2 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \frac{\epsilon_1 + \epsilon_2}{\delta} \cdot \frac{\|\mathcal{B}\|_2}{\sigma_{-(s+1)}(\mathfrak{A})} \\ &\leq 3600 \cdot \sqrt{\frac{d^{*3}}{d_*}} \cdot \kappa(\mathbf{U})^3 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \frac{\epsilon_1 + \epsilon_2}{\delta} \cdot \frac{\|\mathcal{B}\|_2}{\sigma_{-(s+1)}(\mathfrak{A})}, \end{aligned}$$

where $\|\mathcal{B}\|_2$ is as in Definition 2.1 (with respect to W_1 and W_2), and the adjoint algebra map \mathfrak{A} is defined by the projection of operators in \mathcal{B} onto $\text{Lin}(U, V)$ (see Definition 4.7).

Furthermore, if $2\gamma\sqrt{s} \cdot \kappa(\mathbf{U}) < 1$, then $\tilde{U} = \tilde{U}_1 \oplus \dots \oplus \tilde{U}_s$.

References

- [AGH⁺14] Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M. Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15(1):2773–2832, jan 2014.
- [ALGV19] Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials II: High-dimensional walks and an FPRAS for counting bases of a matroid. In *STOC'19—Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 1–12, 2019.

- [BCPV19] Aditya Bhaskara, Aidao Chen, Aidan Perreault, and Aravindan Vijayaraghavan. Smoothed analysis in unsupervised learning via decoupling. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 582–610. IEEE, 2019.
- [BDJ⁺22] Ainesh Bakshi, Ilias Diakonikolas, He Jia, Daniel M. Kane, Pravesh K. Kothari, and Santosh S. Vempala. Robustly learning mixtures of k arbitrary Gaussians. In *STOC '22—Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1234–1247. ACM, New York, 2022.
- [BF60] F. L. Bauer and C. T. Fike. Norms and exclusion theorems. *Numer. Math.*, 2:137–141, 1960.
- [BHKX22] Mitali Bafna, Jun-Ting Hsieh, Pravesh K Kothari, and Jeff Xu. Polynomial-time power-sum decomposition of polynomials. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 956–967. IEEE, 2022.
- [CIK97] Alexander L. Chistov, Gábor Ivanyos, and Marek Karpinski. Polynomial time algorithms for modules over finite dimensional algebras. In *Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation, ISSAC '97, Maui, Hawaii, USA, July 21-23, 1997*, pages 68–74, 1997.
- [CLLZ23] Sitan Chen, Jerry Li, Yuanzhi Li, and Anru R. Zhang. Learning polynomial transformations via generalized tensor decompositions. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 1671–1684. ACM, 2023.
- [Ell69] W. J. Ellison. A ‘waring’s problem’ for homogeneous forms. *Mathematical Proceedings of the Cambridge Philosophical Society*, 65(3):663–672, 1969.
- [GGdOW20] Ankit Garg, Leonid Gurvits, Rafael Mendes de Oliveira, and Avi Wigderson. Operator scaling: Theory and applications. *Found. Comput. Math.*, 20(2):223–290, 2020.
- [GHK15] Rong Ge, Qingqing Huang, and Sham M. Kakade. Learning mixtures of gaussians in high dimensions. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC*, pages 761–770, 2015. Open source version at <https://arxiv.org/abs/1503.00424>.
- [GKS20] Ankit Garg, Neeraj Kayal, and Chandan Saha. Learning sums of powers of low-degree polynomials in the non-degenerate case. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 889–899. IEEE, 2020. Open source version at <https://arxiv.org/abs/2004.06898>.
- [Jor75] Camille Jordan. Essai sur la géométrie à n dimensions. *Bull. Soc. Math. France*, 3:103–174, 1875.

- [KB09] Tamara G. Kolda and Brett W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.
- [KMYZ08] Erich Kaltofen, John P. May, Zhengfeng Yang, and Lihong Zhi. Approximate factorization of multivariate polynomials using singular value decomposition. *Journal of Symbolic Computation*, 43(5):359–376, 2008.
- [KT90] Erich Kaltofen and Barry M. Trager. Computing with polynomials given by black boxes for their evaluations: Greatest common divisors, factorization, separation of numerators and denominators. *Journal of Symbolic Computation*, 9(3):301–320, 1990.
- [Lat05] R. Latala. Some estimates of norms of random matrices. *Proceedings of the American Mathematical Society*, 133:1273–1282, 2005.
- [LM00] B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Ann. Statist.*, 28(5):1302–1338, 2000.
- [MT82] Nimrod Megiddo and Arie Tamir. On the complexity of locating linear facilities in the plane. *Operations Research Letters*, 1(5):194–197, 1982.
- [PHL04] Lance Parsons, Ehtesham Haque, and Huan Liu. Subspace clustering for high dimensional data: A review. *SIGKDD Explor. Newsl.*, 6(1):90–105, jun 2004.
- [Qia18] Youming Qiao. Block diagonalization for adjoint action. Private communication, 2018.
- [QXCK23] Wentao Qu, Xianchao Xiu, Huangyue Chen, and Lingchen Kong. A survey on high-dimensional subspace clustering. *Mathematics*, 11(2), 2023.
- [SGS90] G. W. Stewart and Ji Guang Sun. *Matrix Perturbation Theory*. Academic Press, 1990.
- [Sza91] Stanislaw J. Szarek. Condition numbers of random matrices. *J. Complex.*, 7(2):131–149, 1991.
- [Vij20] Aravindan Vijayaraghavan. Efficient tensor decompositions. In Tim Roughgarden, editor, *Beyond the Worst-Case Analysis of Algorithms*, pages 424–444. Cambridge University Press, 2020.
- [Wal06] Hanna Wallach. Topic modeling: Beyond bag-of-words. In *ICML 2006 - Proceedings of the 23rd International Conference on Machine Learning*, volume 2006, pages 977–984, 01 2006.
- [Wed72] Per-Åke Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, pages 99–111, 1972.
- [Wed73] Per-Åke Wedin. Perturbation theory for pseudo-inverses. *Nordisk Tidskr. Informationsbehandling (BIT)*, 13:217–232, 1973.
- [Wey12] Hermann Weyl. Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung). *Math. Ann.*, 71(4):441–479, 1912.

A Linear Algebra and Probability

A.1 Matrices and Subspaces

A.1.1 Matrix Norms

The following are some easy to verify properties of the Frobenius norm:

Proposition A.1.

1. (Sub-multiplicativity) For any $m \times n$ matrix M , and $n \times p$ matrix N , it holds $\|MN\|_F \leq \|M\|_F \cdot \|N\|_2 \leq \|M\|_F \cdot \|N\|_F$.
2. Let the $n \times n$ matrix M have eigenvalues $\lambda_1, \dots, \lambda_n$. Then, $\|M\|_F^2 \geq \sum_{i \in [n]} |\lambda_i|^2$.

A.1.2 Pseudo-Inverse

Recall that for a matrix M , we use M^\dagger to denote its Moore-Penrose pseudo-inverse. This satisfies the following easy to check properties:

Proposition A.2. For any $m \times n$ matrix M with $\text{rank}(M) = n \leq m$, it holds that:

1. $M^\dagger = (M^\top M)^{-1} M^\top$, where M^\top is the transpose of M .
2. $M^\dagger M$ equals the identity matrix of size n .
3. $M^\dagger (M^\dagger)^\top = (M^\top M)^{-1}$.
4. $\|M^\dagger\|_2 = 1/\sigma_n(M)$, where $\sigma_n(M)$ denotes the n^{th} largest (or the smallest non-zero) singular value of M .

We will need also the following lemma about pseudo-inverses.

Lemma A.1. Let A, B be $n \times r$ matrices be matrices with $\text{rank}(A) = \text{rank}(B) = r \leq n$, and with the same column space. Then, $A^\dagger B \in \mathbb{R}^{r \times r}$ is invertible, and $(A^\dagger B)^{-1} = B^\dagger A$.

Proof. Fix any $\mathbf{x} \in \mathbb{R}^r$. Then, there is a unique $\mathbf{y} \in \mathbb{R}^r$ such that $A\mathbf{x} = B\mathbf{y}$, and it holds that

$$(A^\dagger B) \cdot (B^\dagger A) \cdot \mathbf{x} = (A^\dagger B)(B^\dagger B)\mathbf{y} = (A^\dagger B) \cdot \mathbf{y} = A^\dagger A\mathbf{x} = \mathbf{x}.$$

Hence, $(A^\dagger B) \cdot (B^\dagger A)$ must be the identity matrix. \square

A.1.3 Distances Between Subspaces

For any pair of subspaces in \mathbb{R}^n , we can find a nice basis that relates how they are situated with respect to each other in \mathbb{R}^n , as follows:

Theorem 6. Canonical Decomposition and Angles Between Subspaces. ([\[Jor75\]](#); see [Theorem I.5.2 in \[SGS90\]](#))

Let $U, \tilde{U} \subseteq \mathbb{R}^n$ be two subspaces, each of dimension r , and let $k = r - \dim(U \cap \tilde{U})$. Then, there exists an orthonormal basis $\mathbf{e}_1, \dots, \mathbf{e}_k, \mathbf{e}_{k+1}, \dots, \mathbf{e}_r, \mathbf{f}_1, \dots, \mathbf{f}_k, \mathbf{h}_1, \dots, \mathbf{h}_{n-(r+k)}$ of \mathbb{R}^n , and angles $\frac{\pi}{2} \geq \theta_1 \geq \dots \geq \theta_k > 0$, such that:

1. $\mathbf{e}_1, \dots, \mathbf{e}_r$ form an orthonormal basis of U .

2. $\tilde{\mathbf{e}}_1, \dots, \tilde{\mathbf{e}}_r$ form an orthonormal basis of \tilde{U} , where for each $i \in [r]$,

$$\tilde{\mathbf{e}}_i = \begin{cases} \cos(\theta_i) \cdot \mathbf{e}_i + \sin(\theta_i) \cdot \mathbf{f}_i, & i \leq k \\ \mathbf{e}_i, & o/w \end{cases}.$$

Furthermore, it holds that:

1. $\text{dist}(U, \tilde{U}) = \|\text{Proj}_U - \text{Proj}_{\tilde{U}}\|_2 = \sin(\theta_1)$.

2. Let $\alpha_1, \dots, \alpha_r \in \mathbb{R}$ be such that $\sum_{i=1}^k |\alpha_i|^2 = 1$, and let $u = \sum_{i=1}^r \alpha_i \cdot \mathbf{e}_i \in U$, $\tilde{u} = \sum_{i=1}^r \alpha_i \cdot \tilde{\mathbf{e}}_i \in \tilde{U}$. Then, it holds that $\|u - \tilde{u}\|_2 \leq 2 \sin\left(\frac{\theta_1}{2}\right) \leq 2 \sin(\theta_1) = 2 \cdot \text{dist}(\tilde{U}, U)$.

Note that for consistency, we define $\theta_1 = 0$ if $k = 0$.

Remark 7. Theorem 6 implies that, given subspaces U, V of \mathbb{R}^n of dimensions r , there are orthogonal matrices P, Q such that $U = \langle P \rangle$, $V = \langle Q \rangle$ and $P^T Q = C$, where C is a diagonal matrix with entries $C_{ii} = \cos(\theta_i)$. In the case when the dimensions of the subspaces are different, C is a rectangular matrix of shape $\dim(U) \times \dim(V)$, the principal diagonal having elements $\cos(\theta_1), \dots, \cos(\theta_r)$ where $r = \min(\dim(U), \dim(V))$.

We will need also the following lemma:

Lemma A.2. Let $U, \tilde{U} \subseteq \mathbb{R}^n$ be subspaces of dimension d such that for each $\tilde{\mathbf{u}} \in \tilde{U}$, $\|\tilde{\mathbf{u}}\| = 1$, $\text{dist}(\tilde{\mathbf{u}}, U) \leq \eta$. Then, $\text{dist}(U, \tilde{U}) \leq \eta$.

Proof. Consider the Canonical Decomposition as in Theorem 6. We have that

$$\text{dist}(U, \tilde{U}) = \sin(\theta_1) = \|\tilde{\mathbf{e}}_1 - \cos(\theta_1) \cdot \mathbf{e}_1\|_2 = \text{dist}(\tilde{\mathbf{e}}_1, U) \leq \eta. \quad \square$$

A.2 Matrix Perturbation Bounds

A.2.1 Perturbation bounds for Singular values and Singular vectors

If we *slightly* perturb an $m \times n$ matrix A to obtain $\tilde{A} = A + E$, how "far" are the singular values and singular spaces of \tilde{A} from those of A ? A satisfactory answer (which is asymptotically the best possible in the worst case) is provided by a theorem due to Wedin.

Lemma A.3. (Weyl's Inequality [Wey12]) Let A, E be $m \times n$ matrices with $m \geq n$, and let $\tilde{A} = A + E$. Then, for each $i \in [n]$, it holds that

$$|\sigma_i(\tilde{A}) - \sigma_i(A)| \leq \|E\|_2,$$

where σ_i denotes the i^{th} largest singular value.

Lemma A.4. Perturbation of singular spaces (Wedin [Wed72]; see Theorem V.4.4 in [SGS90]). Let A, E be $m \times n$ matrices with $m \geq n$, and let $\tilde{A} = A + E$, and suppose their singular value decomposition is as follows:

$$A = [U_1 \ U_2 \ U_3] \begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} V_1^\top \\ V_2^\top \end{bmatrix}, \quad \tilde{A} = [\tilde{U}_1 \ \tilde{U}_2 \ \tilde{U}_3] \begin{bmatrix} \tilde{\Sigma}_1 & 0 \\ 0 & \tilde{\Sigma}_2 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \tilde{V}_1^\top \\ \tilde{V}_2^\top \end{bmatrix}$$

where the first block of the above decomposition corresponds to the top r singular values and the second block to the bottom $(n - r)$ singular values.

Let $\delta > 0$ be such that $\delta \leq \min_{i \in [r]} (\Sigma_1)_{i,i} - \max_{j \in [n-r]} (\Sigma_2)_{j,j}$. Then, it holds that

$$\text{dist}(\langle U_1 \rangle, \langle \tilde{U}_1 \rangle) \leq \frac{2 \|E\|_2}{\delta}, \quad \text{dist}(\langle V_1 \rangle, \langle \tilde{V}_1 \rangle) \leq \frac{2 \|E\|_2}{\delta}.$$

Proof. Theorem V.4.4 in [SGS90] shows that the above distances are bounded by $\frac{\|E\|_2}{\eta}$, where $\eta \stackrel{\text{def}}{=} \min_{i \in [r]} (\Sigma_1)_{i,i} - \max_{j \in [n-r]} (\Sigma_2)_{j,j}$. The result then follows from Weyl's inequality (Lemma A.3), which says that for each $i \in [r]$, $|(\Sigma_1)_{i,i} - (\tilde{\Sigma}_1)_{i,i}| \leq \|E\|_2$, and so $\eta \geq \delta - \|E\|_2$.

We remark that result over \mathbb{C} (as is proven in [SGS90]) also implies the same over \mathbb{R} . \square

As an immediate corollary we have:

Corollary A.1. (lemma G.5 in [GHK15]) Let A, E be $m \times n$ matrices with $m \geq n$. Suppose that A has rank r and the smallest (non-zero) singular value of A is given by $\sigma_r(A)$. Let S, \tilde{S} (resp. T, \tilde{T}) be the subspaces spanned by the top r right (resp. left) singular vectors of A and $\tilde{A} = A + E$ respectively. Then we have:

$$\text{dist}(S, \tilde{S}) \leq \frac{2 \|E\|_2}{\sigma_r(A)}, \quad \text{dist}(T, \tilde{T}) \leq \frac{2 \|E\|_2}{\sigma_r(A)}.$$

The above bounds are hold true for even when the perturbation matrix E is chosen in a worst-case/adversarial fashion and can be rather pessimistic. In many applications, E is more like a random matrix in which the case the perturbation of singular values and spaces will be significantly less.

Lemma A.5. (cf. [Lat05, Sza91]) For a random real matrix $E \sim (\mathcal{N}(0, \rho^2))^{m \times n}$, we have that almost surely

$$\|E\|_2 = \Theta(\rho \cdot \sqrt{m+n}).$$

Lemma A.6. For $i \in \{1, 2\}$, let $S_i \subseteq \mathbb{R}^{n_i}$ be a subspace of dimensions $s_i \leq n_i$ and let $\text{Proj}_{S_i} : \mathbb{R}^{n_i} \mapsto \mathbb{R}^{n_i}$ be the corresponding projection map. Then for a random matrix $E \sim (\mathcal{N}(0, \rho^2))^{n_2 \times n_1}$, we have that almost surely

$$\left\| \text{Proj}_{S_2} \cdot E \cdot \text{Proj}_{S_1} \right\|_2 = \Theta(\rho \cdot \sqrt{s_1 + s_2}).$$

Proof. The probability density function for random Gaussian matrices is invariant under orthogonal change of basis of either the domain space or the target space. Consequently we can assume without loss of generality that S_1 (resp. S_2) is spanned by the first s_1 (resp. s_2) canonical unit vectors of \mathbb{R}^{n_1} (resp. of \mathbb{R}^{n_2}). Then $\left\| (\text{Proj}_{S_2} \cdot E \cdot \text{Proj}_{S_1}) \right\|_2 = \|B\|_2$, where $B \in \mathbb{R}^{s_2 \times s_1}$ is the top-left $(s_2 \times s_1)$ -dimensional submatrix of E . Thus $B \sim (\mathcal{N}(0, \rho^2))^{s_2 \times s_1}$ and hence the conclusion follows from an application of lemma A.5. \square

A.2.2 Perturbation bounds for Eigenvalues and Eigenvectors

If we *slightly* perturb a matrix $A \in \mathbb{R}^{n \times n}$ to obtain $\tilde{A} = A + E$, how "far" are the eigenvalues and eigenvectors of \tilde{A} from those of A ? We show quantitative bounds when all eigenvalues of A are simple (have multiplicity 1), and $\|E\|_2$ is *small*.

First, we state a result that shows that the eigenvalues of any complex matrix vary continuously with the error $\|E\|_2$. For $\lambda \in \mathbb{C}, \epsilon > 0$, we define $\mathcal{D}(\lambda, \epsilon) \stackrel{\text{def}}{=} \{\zeta \in \mathbb{C} : |\zeta - \lambda| \leq \epsilon\}$.

Theorem 7. (See Theorem IV.1.1 in [SGS90]) Let $A, E \in \mathbb{C}^{n \times n}$ and $\tilde{A} = A + E$. Let λ be an eigenvalue of A with algebraic multiplicity m . Then, for any (small enough) $\epsilon > 0$, there exists a $\delta > 0$, such that if $\|E\|_2 < \delta$, the disk $\mathcal{D}(\lambda, \epsilon)$ contains exactly m eigenvalues of \tilde{A} .

Proof Idea. The theorem follows by noticing that the characteristic polynomial of any matrix is a continuous function of the matrix entries, and then applying Rouché's Theorem. \square

Theorem 8. (Bauer-Fike [BF60]) Let $A = X \cdot \Lambda \cdot X^{-1} \in \mathbb{C}^{n \times n}$ be a diagonalizable matrix with $X \in \mathbb{C}^{n \times n}$, and $\Lambda \in \mathbb{C}^{n \times n}$ diagonal. Let $E \in \mathbb{C}^{n \times n}$ and $\tilde{A} = A + E$. Then, for each eigenvalue $\tilde{\lambda}$ of \tilde{A} , there is an eigenvalue λ of A such that $|\tilde{\lambda} - \lambda| \leq \kappa(X) \cdot \|E\|_2$, where $\kappa(X) = \|X\|_2 \cdot \|X^{-1}\|_2$ is the condition number of X .

Lemma A.7. Perturbation of Eigenvalues. Let $A = X \cdot \Lambda \cdot X^{-1} \in \mathbb{C}^{n \times n}$ be a diagonalizable matrix, with $X \in \mathbb{C}^{n \times n}$, and $\Lambda = \text{diag}(\underbrace{\lambda_1, \dots, \lambda_1}_{m_1 \text{ times}}, \dots, \underbrace{\lambda_k, \dots, \lambda_k}_{m_k \text{ times}})$ and $\sum_{i=1}^k m_i = n$.

Let $\delta = \min_{i,j \in [k], i \neq j} |\lambda_i - \lambda_j| > 0$, and let $\kappa(X) = \|X\|_2 \cdot \|X^{-1}\|_2$ be the condition number of X .

Let $E \in \mathbb{C}^{n \times n}$ be such that $\kappa(X) \cdot \|E\|_2 < \delta/2$, and let $\tilde{A} = A + E$. Then,

1. The eigenvalues of \tilde{A} can be grouped into k groups $\tilde{\lambda}_{1,1}, \dots, \tilde{\lambda}_{1,m_1}, \dots, \tilde{\lambda}_{k,1}, \dots, \tilde{\lambda}_{k,m_k}$ such that for each $i \in [k], j \in [m_i]$, it holds that $|\tilde{\lambda}_{i,j} - \lambda_i| \leq \kappa(X) \cdot \|E\|_2 < \delta/2$.
2. Suppose that all of $A, E, X, \Lambda \in \mathbb{R}^{n \times n}$ are real matrices. Let $i \in [n]$ be such that λ_i has multiplicity 1 (that is $m_i = 1$). Then, it holds that $\tilde{\lambda}_{i,1}$ is real.

Proof. 1. For each $i \in [k]$, let $\mathcal{D}_i = \mathcal{D}(\lambda_i, \kappa(X) \cdot \|E\|_2) \subsetneq \mathcal{D}(\lambda_i, \delta/2)$. We know that the \mathcal{D}_i 's are disjoint from each other, and for each $i \in [k]$, \mathcal{D}_i contains exactly m_i eigenvalues of $A_0 = A$. By the Bauer-Fike theorem (Theorem 8), we know that for each $\tau \in [0, 1]$, each eigenvalue of $A_\tau \stackrel{\text{def}}{=} A + \tau E$ lies in $\cup_{i \in [k]} \mathcal{D}_i$. Then, by Theorem 7, it must hold that for each $\tau \in [0, 1]$, each \mathcal{D}_i contains exactly m_i eigenvalues of A_τ , since the eigenvalues can "never jump" from one \mathcal{D}_i to another. We leave the formal details of the last sentence to the reader.

2. Suppose that all of $A, E, X, \Lambda \in \mathbb{R}^{n \times n}$. Without loss of generality, we assume that λ_1 is the eigenvalue multiplicity 1. It holds that $|\tilde{\lambda}_{1,1} - \lambda_1| < \delta/2$, and for each $i \in [k], i \neq 1, j \in [m_i]$, that

$$|\tilde{\lambda}_{i,j} - \lambda_1| \geq |\lambda_i - \lambda_1| - |\lambda_i - \tilde{\lambda}_{i,j}| > \delta - \delta/2 = \delta/2.$$

That is, $\tilde{\lambda}_{1,1}$ is the unique eigenvalue of \tilde{A} that is within $\delta/2$ distance of λ_1 . But, the complex conjugate $\bar{\tilde{\lambda}}_{1,1}$ is also an eigenvalue of \tilde{A} satisfying $|\bar{\tilde{\lambda}}_{1,1} - \lambda_1| = |\tilde{\lambda}_{1,1} - \lambda_1| < \delta/2$. Hence, $\tilde{\lambda}_{1,1} = \bar{\tilde{\lambda}}_{1,1} \in \mathbb{R}$. \square

Lemma A.8. Perturbation of Eigenvectors of Simple Eigenvalues for Real Matrices. Let $A = X \cdot \Lambda \cdot X^{-1} \in \mathbb{R}^{n \times n}$ be a real diagonalizable matrix, with $X \in \mathbb{R}^{n \times n}$, and where $\Lambda = \text{diag}(\underbrace{\lambda_1, \dots, \lambda_1}_{m_1 \text{ times}}, \dots, \underbrace{\lambda_k, \dots, \lambda_k}_{m_k \text{ times}}) \in \mathbb{R}^{n \times n}$ with $\sum_{i=1}^k m_i = n$. Let $\delta =$

$\min_{i,j \in [k], i \neq j} |\lambda_i - \lambda_j|$, and let $\kappa(X) = \|X\|_2 \cdot \|X^{-1}\|_2$ be the condition number of X .

Let $E \in \mathbb{R}^{n \times n}$ be such that $\kappa(X) \cdot \|E\|_2 < \delta/2$, and let $\tilde{A} = A + E$. Then,

1. The eigenvalues of \tilde{A} can be grouped into k groups $\tilde{\lambda}_{1,1}, \dots, \tilde{\lambda}_{1,m_1}, \dots, \tilde{\lambda}_{k,1}, \dots, \tilde{\lambda}_{k,m_k} \in \mathbb{C}$ such that for each $i \in [k], j \in [m_i]$, it holds that $|\tilde{\lambda}_{i,j} - \lambda_i| \leq \kappa(X) \cdot \|E\|_2$.
2. Let $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^n$ be the columns of X , and suppose that λ_1 is a simple eigenvalue of A (that is, it has multiplicity $m_1 = 1$), with \mathbf{x}_1 as the corresponding eigenvector.

Let $\tilde{\lambda}_1 \stackrel{\text{def}}{=} \tilde{\lambda}_{1,1}$ be the unique eigenvalue of \tilde{A} such that $|\tilde{\lambda}_1 - \lambda_1| < \delta/2$. Then, $\tilde{\lambda}_1$ is real. Further, suppose that $\tilde{\mathbf{x}} \in \mathbb{R}^n, \|\tilde{\mathbf{x}}\|_2 = 1$ is an eigenvector of \tilde{A} with eigenvalue $\tilde{\lambda}_1$. Then, there exists $\zeta \in \{-1, 1\}$ such that

$$\left\| \tilde{\mathbf{x}} - \zeta \cdot \frac{\mathbf{x}_1}{\|\mathbf{x}_1\|_2} \right\|_2 \leq \frac{4\kappa(X) \|E\|_2}{\delta}.$$

Proof. The first part follows from first part of Lemma A.7. Similarly, in the second part, the fact that $\tilde{\lambda}_1$ is real follows from the second part in Lemma A.7.

Now, suppose that $\tilde{\mathbf{x}} \in \mathbb{R}^n, \|\tilde{\mathbf{x}}\|_2 = 1$ is an eigenvector of \tilde{A} with eigenvalue $\tilde{\lambda}_1$. Let $\boldsymbol{\alpha} \in \mathbb{R}^n$ be such that $\tilde{\mathbf{x}} = X\boldsymbol{\alpha} = \sum_{i=1}^n \alpha_i \mathbf{x}_i$. Then, we know

$$\tilde{\lambda}_1 X\boldsymbol{\alpha} = \tilde{\lambda}_1 \tilde{\mathbf{x}} = \tilde{A}\tilde{\mathbf{x}} = (A + E)X\boldsymbol{\alpha} = X\Lambda\boldsymbol{\alpha} + E\tilde{\mathbf{x}}.$$

Rearranging, we get

$$(\Lambda - \tilde{\lambda}_1 I)\boldsymbol{\alpha} = -X^{-1}E\tilde{\mathbf{x}}.$$

We know that for each $i \in [k], i \neq 1$,

$$|\lambda_i - \tilde{\lambda}_1| \geq |\lambda_i - \lambda_1| - |\lambda_1 - \tilde{\lambda}_1| > \delta/2 > 0.$$

Setting $P = \text{diag}\left(0, \frac{1}{\Lambda_{2,2} - \tilde{\lambda}_1}, \dots, \frac{1}{\Lambda_{n,n} - \tilde{\lambda}_1}\right) \in \mathbb{R}^{n \times n}$, we know $\|P\|_2 \leq \frac{2}{\delta}$, and

$$\tilde{\mathbf{x}} - \alpha_1 \mathbf{x}_1 = X\left(\boldsymbol{\alpha} - [\alpha_1, 0, \dots, 0]^\top\right) = -XPX^{-1}E\tilde{\mathbf{x}},$$

and

$$\|\tilde{\mathbf{x}} - \alpha_1 \mathbf{x}_1\|_2 \leq \left\| XPX^{-1}E\tilde{\mathbf{x}} \right\|_2 \leq \frac{2\kappa(X) \|E\|_2}{\delta} \stackrel{\text{def}}{=} \gamma.$$

We can assume that $\gamma < 1$, or else the statement we wish to prove is true trivially by triangle inequality. This, in particular implies that $\alpha_1 \neq 0$, and

$$\begin{aligned} \left\| \tilde{\mathbf{x}} - \frac{\alpha_1 \mathbf{x}_1}{\|\alpha_1 \mathbf{x}_1\|_2} \right\|_2 &\leq \|\tilde{\mathbf{x}} - \alpha_1 \mathbf{x}_1\|_2 + \left\| \alpha_1 \mathbf{x}_1 - \frac{\alpha_1 \mathbf{x}_1}{\|\alpha_1 \mathbf{x}_1\|_2} \right\|_2 \\ &= \|\tilde{\mathbf{x}} - \alpha_1 \mathbf{x}_1\|_2 + |1 - \|\alpha_1 \mathbf{x}_1\|_2| \\ &\leq 2 \|\tilde{\mathbf{x}} - \alpha_1 \mathbf{x}_1\|_2 \\ &\leq 2\gamma. \end{aligned}$$

Choosing $\zeta = \alpha_1 / |\alpha_1|$, this proves the desired result. \square

A.2.3 Perturbation bounds for Pseudo-Inverse

Lemma A.9. (Wedin [Wed73]; see Theorem III.3.9 in [SGS90]) Let A, E be $m \times n$ matrices with $m \geq n$, and let $\tilde{A} = A + E$. If $\text{rank}(A) = \text{rank}(\tilde{A}) = n$, then

$$\|A^\dagger - \tilde{A}^\dagger\|_2 \leq \sqrt{2} \|A^\dagger\|_2 \|\tilde{A}^\dagger\|_2 \|E\|_2.$$

Corollary A.2. Let A, E be $m \times n$ matrices with $m \geq n$, and let $\tilde{A} = A + E$. If $\text{rank}(A) = n$ and $\|E\|_2 \leq \sigma_n(A)/2$, then we have

$$\|A^\dagger - \tilde{A}^\dagger\|_2 \leq 3 \|A^\dagger\|_2^2 \|E\|_2.$$

Proof. By Weyl's inequality (Lemma A.3), we know $\sigma_n(\tilde{A}) \geq \sigma_n(A) - \|E\|_2 \geq \sigma_n(A)/2 > 0$. Hence, $\text{rank}(\tilde{A}) = n$, and $\|\tilde{A}^\dagger\|_2 = 1/\sigma_n(\tilde{A}) \leq 2/\sigma_n(A)$. Plugging this into Lemma A.9, we get the desired result. \square

A.3 Anti-Concentration of Gaussian Linear Forms

We show an anti-concentration bound for linear forms in independent Gaussian random variables.

Lemma A.10. Let $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$ be such that for each $i \in [n]$, $x_i \sim \mathcal{N}(0, 1)$ is chosen independently, and let $\mathbf{y} = \mathbf{x} / \|\mathbf{x}\|_2$. Then, for any $\mathbf{a} \in \mathbb{R}^n$, $\mathbf{a} \neq 0$ and any $\delta \geq 0$, it holds that

$$\Pr \left[\left| \mathbf{a}^\top \mathbf{y} \right| \leq \frac{\delta \|\mathbf{a}\|_2}{6\sqrt{n + \ln \frac{1}{\delta}}} \right] \leq \delta.$$

Proof. Let $\mathbf{a} \in \mathbb{R}^n$, $\mathbf{a} \neq 0$. We can assume $\|\mathbf{a}\|_2 = 1$.

1. Note that $\mathbf{a}^\top \mathbf{x} \sim \mathcal{N}(0, 1)$, and so for any $\eta \geq 0$,

$$\Pr \left[\left| \mathbf{a}^\top \mathbf{x} \right| \leq \eta \right] = \Pr_{g \sim \mathcal{N}(0,1)} [|g| \leq \eta] = \frac{1}{\sqrt{2\pi}} \int_{-\eta}^{\eta} e^{-t^2} dt \leq \frac{1}{\sqrt{2\pi}} \cdot 2\eta \leq \eta.$$

2. By concentration bounds on chi-square random variables [LM00], we have for any $c \geq 1$ that

$$\Pr \left[\|\mathbf{x}\|_2 > 3c\sqrt{n} \right] \leq e^{-c^2 n}.$$

Combining the above inequalities, we get

$$\begin{aligned} \Pr \left[\left| \mathbf{a}^\top \mathbf{y} \right| \leq \frac{\delta}{6\sqrt{n + \ln \frac{1}{\delta}}} \right] &\leq \Pr \left[\left| \mathbf{a}^\top \mathbf{x} \right| \leq \frac{\delta}{2} \right] + \Pr \left[\|\mathbf{x}\|_2 > 3\sqrt{n} \cdot \sqrt{1 + \frac{1}{n} \ln \frac{1}{\delta}} \right] \\ &\leq \frac{\delta}{2} + \delta e^{-n} \leq \delta. \end{aligned} \quad \square$$

B Subspace Clustering

In this section, we will consider the problem of Subspace Clustering. Following the outline in Section 1.3, we will show a reduction to vector space decomposition. In Section B.1, we solve a problem which we call Robust Recover from Symmetric Tensor Power, which we later use as a subroutine in our subspace clustering algorithm. Then, in Section B.2 and Section B.3 we analyze the noiseless and the robust case of subspace clustering.

B.1 Robust Recovery from Symmetric Tensor Power

For any set $A \subseteq \mathbb{R}^n$, and any $d \in \mathbb{N}$, we define the set $A^{\otimes d} \stackrel{\text{def}}{=} \{(\mathbf{a} \cdot \mathbf{x})^d : \mathbf{a} \in A\} \subseteq \mathbb{R}[\mathbf{x}]^d$, where $\mathbf{x} = (x_1, \dots, x_n)$ are formal variables. We consider the following problem:

Robust Recovery from Symmetric Tensor Power (RRSTP). Let $A = \{\mathbf{a}_1, \dots, \mathbf{a}_N\} \subseteq \mathbb{R}^n$ be a set of N points, and let $d \in \mathbb{N}$. We are given as input a subspace $\tilde{U} \subseteq \mathbb{R}[\mathbf{x}]^d$ such that $\text{dist}(\tilde{U}, \langle A^{\otimes d} \rangle)$ is "small," and our goal is to *efficiently* find a subspace $\tilde{V} \subseteq \mathbb{R}[\mathbf{x}]^1$ such that $\text{dist}(\tilde{V}, \langle A \rangle)$ is "small."

Note that we are working with the Bombieri inner product over $\mathbb{R}[\mathbf{x}]^d$ and $\mathbb{R}[\mathbf{x}]^{(d-1)}$. Informally speaking, the goal is to (approximately) recover the subspace $\langle A \rangle$, when given as input the subspace $\langle A^{\otimes d} \rangle$ (approximately).

For the rest of this section, we fix a set $A = \{\mathbf{a}_1, \dots, \mathbf{a}_N\}$. Let $V \stackrel{\text{def}}{=} \langle A \rangle$ and $r = \dim(V)$, and let $U \stackrel{\text{def}}{=} \langle A^{\otimes d} \rangle$ and $R = \dim(U)$.

B.1.1 RRSTP: Noiseless Case

We first consider the noiseless case, in which the input is the space U . Let $(\mathbf{u}_1, \dots, \mathbf{u}_R)$ be an arbitrary orthonormal basis of the space U .

Definition B.1. We define the partial derivative map $T : \mathbb{R}^n \rightarrow (\mathbb{R}[\mathbf{x}]^{(d-1)})^R$ by

$$T(\mathbf{c}) = \left(\sum_{i=1}^n c_i \cdot \partial_{x_i} \mathbf{u}_1, \dots, \sum_{i=1}^n c_i \cdot \partial_{x_i} \mathbf{u}_R \right) = (\partial_{\mathbf{c}} \mathbf{u}_1, \dots, \partial_{\mathbf{c}} \mathbf{u}_R).$$

Lemma B.1.

$$\ker(T)^\perp = \langle A \rangle.$$

Proof. We show the equivalent fact that $\ker(T) = \langle A \rangle^\perp$, which is implied by the following:

$$\begin{aligned}
\mathbf{c} \in \ker(T) &\iff \partial_{\mathbf{c}} \mathbf{u}_i = 0 \text{ for all } i \in [R] \\
&\iff \partial_{\mathbf{c}} \mathbf{u} = 0 \text{ for all } \mathbf{u} \in U \\
&\iff \partial_{\mathbf{c}} (\mathbf{a}_i \cdot \mathbf{x})^d = 0 \text{ for all } i \in [N] \\
&\iff \mathbf{c} \cdot \mathbf{a}_i = 0 \text{ for all } i \in [N] \\
&\iff \mathbf{c} \in \langle A \rangle^\perp. \quad \square
\end{aligned}$$

The above lemma gives a natural *algorithm for the noiseless case*: Compute the map T (using any orthonormal basis of $\langle A^{\otimes d} \rangle$), and output $\ker(T)^\perp$. Further this algorithm is efficient: if the input is given as an orthonormal basis of $\langle A^{\otimes d} \rangle$, of size $R \cdot \binom{n+d-1}{d} \leq R \cdot n^d$, the algorithm runs in time $\text{poly}(n^d)$.

B.1.2 RRSTP: Robust Case

Next, we provide a robust version of the above algorithm. We are given as input a vector space $\tilde{U} \subseteq \mathbb{R}[\mathbf{x}]^{=d}$ of dimension R , such that $\text{dist}(\tilde{U}, U)$ is small. Let $(\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_R)$ be an arbitrary orthonormal basis of \tilde{U} . We define a noisy version of the partial derivative operator in Definition B.1.

Definition B.2. We define the partial derivative map $\tilde{T} : \mathbb{R}^n \rightarrow (\mathbb{R}[\mathbf{x}]^{=(d-1)})^R$ by

$$\tilde{T}(\mathbf{c}) = \left(\sum_{i=1}^n c_i \cdot \partial_{x_i} \tilde{\mathbf{u}}_1, \dots, \sum_{i=1}^n c_i \cdot \partial_{x_i} \tilde{\mathbf{u}}_R \right) = (\partial_{\mathbf{c}} \tilde{\mathbf{u}}_1, \dots, \partial_{\mathbf{c}} \tilde{\mathbf{u}}_R).$$

Our algorithm is then formally described as Algorithm 3.

Algorithm 3 Robust Recovery from Symmetric Tensor Power.

Input: \tilde{U} is a subspace of dimension R , where $\tilde{U} \subseteq \mathbb{R}[\mathbf{x}]^{=d}$ with $d \in \mathbb{N}$ and $\mathbf{x} = (x_1, \dots, x_n)$.

Assumptions: There is a set $A \subseteq \mathbb{R}^n$ of size N , such that $\text{dist}(\tilde{U}, \langle A^{\otimes d} \rangle) \leq \epsilon$.

Output: Subspace $\tilde{V} \subseteq \mathbb{R}^n$ such that $\text{dist}(\tilde{V}, \langle A \rangle)$ is small.

- 1: Let $(\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_R)$ be an orthonormal basis for \tilde{U} (with respect to the Bombieri inner product).
 - 2: Let $\tilde{T} : \mathbb{R}^n \rightarrow (\mathbb{R}[\mathbf{x}]^{=(d-1)})^R$ be the (directional-derivative) map defined as in Definition B.2.
 - 3: Let r be such that $R = \binom{r+d-1}{d}$.
 - 4: Let $\tilde{V} \subseteq \mathbb{R}^n$ be the space spanned by the right singular vectors of \tilde{T} , corresponding to the top r singular values.
 - 5: Output \tilde{V} .
-

We shall prove that the algorithm gets the following guarantees. We shall work with the extra assumption that $R = \binom{r+d-1}{d}$, which for example is satisfied when N is large and the set A is chosen in some random manner.

Proposition B.1. Let $d \in \mathbb{N}$, and let $A = \{\mathbf{a}_1, \dots, \mathbf{a}_N\} \subseteq \mathbb{R}^n$ be a set of N points. Let $V = \langle A \rangle$ be of dimension r , and let $U = \langle A^{\otimes d} \rangle \subseteq \mathbb{R}[\mathbf{x}]^{=d}$ be of dimension R . Suppose that:

1. $R = \binom{r+d-1}{d}$.
2. $\tilde{U} \subseteq \mathbb{R}[\mathbf{x}]^{=d}$ is a subspace such that $\text{dist}(\tilde{U}, U) \leq \epsilon < 1$.

Then, Algorithm 3, on input \tilde{U} , runs in time $\text{poly}(n^d)$, and outputs a subspace $\tilde{V} \in \mathbb{R}^n$ such that

$$\text{dist}(\tilde{V}, V) \leq 4\epsilon\sqrt{r}.$$

Before we prove our main proposition, we show that the choice of basis does not affect the above algorithm in any way.

Lemma B.2. The singular value decomposition of the operator T (resp. \tilde{T}) does not depend on the choice of the orthonormal basis $(\mathbf{u}_1, \dots, \mathbf{u}_R)$ (resp. $(\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_R)$) for the vector space U (resp. \tilde{U}).

Proof. For any $\mathbf{c}, \mathbf{d} \in \mathbb{R}^n$, we have that

$$\langle T\mathbf{c}, T\mathbf{d} \rangle = \sum_{k=1}^R \langle \partial_{\mathbf{c}}\mathbf{u}_k, \partial_{\mathbf{d}}\mathbf{u}_k \rangle_B.$$

This is the Hilbert-Schmidt inner product between $\partial_{\mathbf{c}}$ and $\partial_{\mathbf{d}}$ over the vector space $\text{Lin}(U, \mathbb{R}[\mathbf{x}]^{=(d-1)})$, which does not depend on the choice of orthonormal basis of U . Hence, the singular value decomposition of T is independent of this choice of basis as well.

The same proof shows the result for \tilde{U} as well. \square

We will also need the following singular value lower bound.

Lemma B.3. Let T be as defined in Definition B.1 with respect to an arbitrary orthonormal basis $(\mathbf{u}_1, \dots, \mathbf{u}_R)$ for U . If $R = \binom{r+d-1}{d}$, then it holds that

$$\sigma_r(T) = d \cdot \sqrt{\frac{R}{r}}.$$

Proof. Let $\mathbf{c} \in \mathbb{R}^n, \|\mathbf{c}\|_2 = 1$ be such that $\mathbf{c} \in \ker(T)^\perp = \langle A \rangle = V$ (see Lemma B.1). We wish to lower bound

$$\langle T(\mathbf{c}), T(\mathbf{c}) \rangle = \sum_{i=1}^R \|\partial_{\mathbf{c}}\mathbf{u}_i\|_B^2.$$

Note that by Lemma B.2, the singular value does not depend on the choice of $(\mathbf{u}_1, \dots, \mathbf{u}_R)$. Hence, we will choose a convenient basis to work with.

First, consider an orthonormal basis $(\mathbf{v}_1 = \mathbf{c}, \dots, \mathbf{v}_r)$ of V . Then, based on the variables $\mathbf{x} = (x_1, \dots, x_n)$, we define the variables $\mathbf{y} = (y_1, \dots, y_r)$ by $y_i = \mathbf{v}_i \cdot \mathbf{x}$ for each $i \in [r]$. This allows us to view $\langle A^{\otimes d} \rangle = U \subseteq \mathbb{R}[\mathbf{x}]^{=d}$ in a natural way: for each polynomial $p(\mathbf{x}) \in U$, there is a corresponding polynomial $q(\mathbf{y}) \in \mathbb{R}[\mathbf{y}]^{=d}$ such that $p(\mathbf{x}) = q(\mathbf{v}_1 \cdot \mathbf{x}, \dots, \mathbf{v}_r \cdot \mathbf{x})$. Furthermore, since the Bombieri norm is preserved under isometries, it holds that $\|p(\mathbf{x})\|_B = \|q(\mathbf{y})\|_B$ (where the norms are with respect to the spaces $\mathbb{R}[\mathbf{x}]^{=d}$ and $\mathbb{R}[\mathbf{y}]^{=d}$ respectively).

Now, since $\dim(U) = R = \binom{r+d-1}{d}$, we know $U = \mathbb{R}[\mathbf{y}]^{=d}$, and so we can choose the orthonormal basis $(q_\alpha(\mathbf{y}) = \sqrt{\frac{d!}{\alpha!}} \mathbf{y}^\alpha)_{\alpha \in \mathbb{N}_d^r}$. Then,

$$\begin{aligned}
\langle T(\mathbf{v}_1), T(\mathbf{v}_1) \rangle &= \sum_{\alpha \in \mathbb{N}_d^r} \|\partial_{\mathbf{v}_1} q_\alpha(\mathbf{v}_1 \cdot \mathbf{x}, \dots, \mathbf{v}_r \cdot \mathbf{x})\|_B^2 \\
&= \sum_{\alpha \in \mathbb{N}_d^r} \|\partial_{y_1} q_\alpha(\mathbf{y})\|_B^2 \\
&= \sum_{\alpha \in \mathbb{N}_d^r: \alpha_1 > 0} \frac{d!}{\alpha!} \cdot \alpha_1^2 \cdot \frac{\alpha!}{\alpha_1 \cdot (d-1)!} \\
&= d \cdot \sum_{\alpha \in \mathbb{N}_d^r} \alpha_1 = d \cdot \frac{Rd}{r}. \quad \square
\end{aligned}$$

Proof of Proposition B.1. Observe that by Lemma B.2, we can work with any orthonormal basis for the vector space \tilde{U} , and the corresponding operator \tilde{T} . By the Canonical Decomposition (Theorem 6), we can choose a basis $(\mathbf{u}_1, \dots, \mathbf{u}_R)$ for U , and $(\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_R)$ for \tilde{U} , such that $\|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_B \leq 2\epsilon$ for each $i \in [R]$. Let T and \tilde{T} be the operators as defined in Definition B.1 and Definition B.2 with respect to these basis, and let M and \tilde{M} be the matrices corresponding to these operators. Then, by Lemma 2.1, we have

$$\left\| \tilde{M}\mathbf{c} - M\mathbf{c} \right\|_F^2 = \sum_{i=1}^R \sum_{j=1}^n \left\| \partial_{x_j} (\tilde{\mathbf{u}}_i - \mathbf{u}_i) \right\|_B^2 = d^2 \cdot \sum_{i=1}^R \|\tilde{\mathbf{u}}_i - \mathbf{u}_i\|_B^2 \leq 4\epsilon^2 d^2 R.$$

Now by Corollary A.1 and Lemma B.3, we get

$$\text{dist}(\tilde{V}, V) \leq \frac{2 \cdot \left\| \tilde{M}\mathbf{c} - M\mathbf{c} \right\|_F}{\sigma_r(M)} \leq \frac{2 \cdot 2\epsilon d \sqrt{R}}{d \sqrt{\frac{R}{r}}} = 4\epsilon \sqrt{r}.$$

Runtime: We observe that $R \leq \dim(\mathbb{R}[\mathbf{x}]^{=d}) = \binom{n+d-1}{d} = \text{poly}(n^d)$. Hence, the map \tilde{T} and its singular value decomposition can also be computed in time $\text{poly}(n^d)$. \square

B.2 Subspace Clustering: Noiseless Case

We begin by considering the noiseless version of the subspace clustering problem. Recall that we are given a set of N points $A = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_N\} \subseteq \mathbb{R}^n$, which admit a partition

$$A = A_1 \uplus A_2 \uplus \dots \uplus A_s,$$

such that the points in each A_j span a *low-dimensional* space $\langle A_j \rangle$. Our goal is to find this partition.

For each $i \in [N]$, we define $\ell_i \in \mathbb{R}[\mathbf{x}]^{=1}$ as the linear form $\ell_i(\mathbf{x}) = (\mathbf{a}_i \cdot \mathbf{x})$ in the formal variables $\mathbf{x} = (x_1, \dots, x_n)$. For any $d \in \mathbb{N}$, we define the set $A^{\otimes d} \stackrel{\text{def}}{=} \{(\mathbf{a} \cdot \mathbf{x})^d : \mathbf{a} \in A\} \subseteq \mathbb{R}[\mathbf{x}]^{=d}$. Proceeding as in Section 1.3, we devise an algorithm (Algorithm 4) for this problem, via a reduction to vector space decomposition.

Algorithm 4 Subspace Clustering: Noiseless Case.

Input: (A, d) where $A \subseteq \mathbb{R}^n$ is a set of size N , and $d \geq 2$ is a positive integer.

Assumptions: The set A admits a partition $A = A_1 \uplus A_2 \uplus \dots \uplus A_s$ such that each A_i spans a *low-dimensional* subspace.

Output: The partition (A_1, \dots, A_s) .

- 1: Compute the spaces $U = \langle A^{\otimes d} \rangle, V = \langle A^{\otimes(d-1)} \rangle$ and the tuple of operators $\mathcal{B} = (B_1, \dots, B_n) \in \text{Lin}(U, V)^n$, where B_i corresponds to the operator $\partial_{x_i} : U \rightarrow V$.
 - 2: Run RVSD Algorithm (Noiseless Case) on (U, V, \mathcal{B}) , and obtain the spaces $\langle A_j^{\otimes d} \rangle$, for each $j \in [s]$.
 - 3: For each $j \in [s]$, run RRSTP Algorithm (Noiseless Case; see Section B.1.1) on $\langle A_j^{\otimes d} \rangle$ to obtain $\langle A_j \rangle$.
 - 4: For each $j \in [s]$, compute $A_j = A \cap \langle A_j \rangle$.
 - 5: Output (A_1, \dots, A_s) .
-

Next, we state our assumptions, and then analyze our algorithm.

Definition B.3. (Subspace Clustering: Non-degeneracy conditions) We say that the partition $A = A_1 \uplus A_2 \uplus \dots \uplus A_s$, is non-degenerate with respect to the integer $d \geq 2$ if the following conditions are satisfied:

1.
$$\langle A^{\otimes(d-1)} \rangle = \langle A_1^{\otimes(d-1)} \rangle \oplus \langle A_2^{\otimes(d-1)} \rangle \oplus \dots \oplus \langle A_s^{\otimes(d-1)} \rangle.$$
2. For each $j \in [s]$, the space $\langle A_j^{\otimes d} \rangle$ is irreducible with respect to the action of first order partials $\partial^{\neq 1}$. That is, we cannot write $\langle A_j^{\otimes d} \rangle = U_{j,1} \oplus U_{j,2}$ and $\langle A_j^{\otimes(d-1)} \rangle = V_{j,1} \oplus V_{j,2}$, with all $U_{j,1}, U_{j,2}, V_{j,1}, V_{j,2}$ non-zero, such that $\partial^{\neq 1}$ maps $U_{j,1}$ into $V_{j,1}$ and $U_{j,2}$ into $V_{j,2}$.

Theorem 9. Let $A \subseteq \mathbb{R}^n$ be set of size N , and $d \geq 2$ be an integer. Let

$$A = A_1 \uplus A_2 \uplus \dots \uplus A_s$$

be a partition of A that it is non-degenerate with respect to d (see Definition B.3). Then, on input (A, d) , Algorithm 4 runs in time $\text{poly}(N, n^d)$, and outputs the partition (A_1, \dots, A_s) .

In the remainder of this section, we shall analyze our algorithm and prove the above theorem.

B.2.1 Structure of The Adjoint Algebra

In this section, we analyze Step 2 of Algorithm 4.

Fix some $d \in \mathbb{N}$, such that the non-degeneracy conditions in Definition B.3 are satisfied. Let

$$U \stackrel{\text{def}}{=} \langle A^{\otimes d} \rangle = \langle \ell_1^d, \ell_2^d, \dots, \ell_N^d \rangle \subseteq \mathbb{R}[\mathbf{x}]^{\neq d}$$

and

$$V \stackrel{\text{def}}{=} \langle A^{\otimes(d-1)} \rangle = \langle \ell_1^{d-1}, \ell_2^{d-1}, \dots, \ell_N^{d-1} \rangle \subseteq \mathbb{R}[\mathbf{x}]^{=(d-1)}.$$

Let $\mathcal{B} = (B_1, \dots, B_n) \in \text{Lin}(U, V)^n$ be the n -tuple of operators corresponding to the action of first-order partial derivatives; that is, B_i corresponds to the operator ∂_{x_i} . The adjoint algebra is then

$$\text{Adj} = \{(D, E) : \partial_{x_i} \cdot D = E \cdot \partial_{x_i} \text{ for all } i \in [n]\} \subseteq \text{Lin}(U, U) \times \text{Lin}(V, V).$$

To show the correctness of Step 2, it is sufficient to show that under the non-degeneracy condition, the adjoint algebra has dimension s .

Proposition B.2.

$$\dim(\text{Adj}) = s.$$

Note that the above proposition along with Proposition 4.1 implies the uniqueness of the decomposition of U as a direct sum of s subspaces.

Corollary B.1. *Let $A \subseteq \mathbb{R}^n$ be set of size N , and $d \geq 2$ be an integer. Let*

$$A = A_1 \uplus A_2 \uplus \dots \uplus A_s$$

be a partition of A that it is non-degenerate (see Definition B.3) with respect to d .

Then, the decomposition

$$\langle A^{\otimes d} \rangle = \langle A_1^{\otimes d} \rangle \oplus \langle A_2^{\otimes d} \rangle \oplus \dots \oplus \langle A_s^{\otimes d} \rangle$$

of $\langle A^{\otimes d} \rangle$ under the action of $\mathfrak{d}^=1$ is unique.

We start by proving a few lemmas characterizing the structure of the adjoint algebra.

Lemma B.4. *Let $(D, E) \in \text{Adj}$ be any element in the adjoint algebra. Then, there exist field constants $c_1, c_2, \dots, c_N \in \mathbb{R}$, such that for all $i \in [N]$ we have*

$$D \cdot \ell_i^d = c_i \cdot \ell_i^d, \quad \text{and } E \cdot \ell_i^{d-1} = c_i \cdot \ell_i^{d-1}.$$

Proof. Without loss of generality, suppose that $i = 1$. Further, after making a suitable change of variables, we can assume that $\ell_1 = x_1$ (note that the space of partial derivatives is closed under change of variables).

Now, for any $j \in [N] \setminus \{1\}$, we must have by the definition of the adjoint algebra that

$$\partial_{x_j} \cdot D \cdot x_1^d = E \cdot \partial_{x_j} \cdot x_1^d = 0.$$

This means that $(D \cdot x_1^d) \in \mathbb{R}[\mathbf{x}]^{=d}$ is a homogeneous degree d polynomial depending only on the variable x_1 . Hence, there exists $c_1 \in \mathbb{R}$ such that $(D \cdot x_1^d) = c_1 \cdot x_1^d$. Consequently, we also have

$$E \cdot x_1^{d-1} = \frac{1}{d} \cdot E \cdot \partial_{x_1} \cdot x_1^d = \frac{1}{d} \cdot \partial_{x_1} \cdot D \cdot x_1^d = \frac{1}{d} \cdot \partial_{x_1} \cdot c_1 x_1^d = c_1 \cdot x_1^{d-1}. \quad \square$$

Lemma B.5. *Let $(D, E) \in \text{Adj}$ be any element in the adjoint algebra, and let the constants $c_1, c_2, \dots, c_N \in \mathbb{R}$ be as in Lemma B.4. Suppose that $I \subseteq [N]$ is a minimal set such that $\{\ell_i^{d-1} : i \in I\}$ are linearly dependent. Then, $c_i = c_{i'}$ for all $i, i' \in I$.*

Proof. Without loss of generality we can assume that $I = \{1, 2, \dots, r\}$. By the minimality of I , let $\alpha_1, \dots, \alpha_r \neq 0$ be such that

$$\alpha_1 \cdot \ell_1^{d-1} + \alpha_2 \cdot \ell_2^{d-1} + \dots + \alpha_r \cdot \ell_r^{d-1} = 0. \quad (10)$$

Then, we have

$$\begin{aligned} E \cdot (\alpha_1 \cdot \ell_1^{d-1} + \dots + \alpha_r \cdot \ell_r^{d-1}) &= 0 \\ \implies \alpha_1 \cdot c_1 \cdot \ell_1^{d-1} + \dots + \alpha_r \cdot c_r \cdot \ell_r^{d-1} &= 0 \\ \implies \alpha_2 \cdot (c_2 - c_1) \cdot \ell_2^{d-1} + \dots + \alpha_r \cdot (c_r - c_1) \cdot \ell_r^{d-1} &= 0 \quad (\text{using 10}). \end{aligned}$$

Since each $\alpha_i \neq 0$, the minimality of I implies that for each $i \in [r] \setminus \{1\}$, we have $c_i = c_1$. \square

Lemma B.6. *Let $(D, E) \in \text{Adj}$ be any element in the adjoint algebra, and let the constants $c_1, c_2, \dots, c_N \in \mathbb{R}$ be as in Lemma B.4. Fix any $j \in [s]$, and suppose that $\langle A_j^{\otimes d} \rangle$ is irreducible (see Condition 2 in Definition B.3) under the action of \mathcal{B} . Let $I_j = \{i \in [N] : \ell_i \in A_j\}$. Then, it holds that $c_i = c_{i'}$ for each $i, i' \in I$.*

Proof. Fix any $(D, E) \in \text{Adj}$ and let $c_1, c_2, \dots, c_N \in \mathbb{R}$ be as in Lemma B.4. Let $I_j = \{i \in [N] : \ell_i \in A_j\}$ and $C_j = \{c_i : i \in I_j\}$, and let $t = |C_j|$.

Suppose for the sake of contradiction that $t > 1$, and $C_j = \{\tilde{c}_1, \dots, \tilde{c}_t\}$. Let $U_j = \langle A_j^{\otimes d} \rangle$, $V_j = \langle A_j^{\otimes (d-1)} \rangle$. Now, for each $k \in [t]$ define $U_{j,k} = \langle \{\ell_i^d : i \in I_j, c_i = \tilde{c}_k\} \rangle$, $V_{j,k} = \langle \{\ell_i^{d-1} : i \in I_j, c_i = \tilde{c}_k\} \rangle$. Then, by definition \mathcal{B} maps each $U_{j,k}$ into $V_{j,k}$. Further, it holds that $V_j = V_{j,1} \oplus \dots \oplus V_{j,t}$ is a direct sum: if we concatenate the bases of $(V_{j,k})_{k \in [t]}$, the list must be linearly independent, or else by Lemma B.5, some of the \tilde{c} 's must be equal, which is false by definition. This also implies the weaker condition that $U_j = U_{j,1} \oplus \dots \oplus U_{j,t}$ is a direct sum, and this contradicts the irreducibility assumption. \square

Next, we prove our main proposition regarding the dimension of the adjoint algebra.

Proof of Proposition B.2. Since U, V admit a decomposition into s subspaces, the dimension of the adjoint algebra is at least s (recall the adjoint algebra always contains the scaling maps; see comment after Definition 4.3). Suppose for the sake of contradiction that $\dim(\text{Adj}) = t > s$.

By Lemma B.4, it holds that for each $(D, E) \in \text{Adj}$, there exist $c_1, \dots, c_N \in \mathbb{R}$ such that for each $i \in [N]$, we have $D \cdot \ell_i^d = c_i \cdot \ell_i^d$ and $E \cdot \ell_i^{d-1} = c_i \cdot \ell_i^{d-1}$. Now, since $\dim(\text{Adj}) = t$, there exists an element $(D, E) \in \text{Adj}$ such that at least t of c_1, \dots, c_N are distinct (for example, any generic element satisfies this). On the other hand, the irreducibility of each component $\langle A_j^{\otimes d} \rangle$ under the action of \mathcal{B} , along with Lemma B.6, implies that there can be at most s distinct elements among c_1, \dots, c_N . This is a contradiction. \square

B.2.2 Completing the Proof

Proof of Theorem 9. The above subsection shows that Step 2 correctly obtains the subspaces $\langle A_j^{\otimes d} \rangle$ for $j \in [s]$. Then, Step 3 correctly obtains the subspaces $\langle A_j \rangle$ for each $j \in [s]$ (see Section B.1.1).

Now, observe that the non-degeneracy condition $\langle A^{\otimes(d-1)} \rangle = \langle A_1^{\otimes(d-1)} \rangle \oplus \langle A_2^{\otimes(d-1)} \rangle \oplus \dots \oplus \langle A_s^{\otimes(d-1)} \rangle$ implies that for each $i \in [N]$, there is a unique $j \in [s]$ such that $\mathbf{a}_i \in A_j$. Hence, it holds that $A_j = \langle A_j \rangle \cap A$. This completes the proof of correctness.

Runtime: Our algorithm deals with the set A of size N , and works with elements in the space $\mathbb{R}[\mathbf{x}]^{=d}$. The efficiency of the RVSD algorithm, and the efficiency of the RRSTP algorithm implies that Algorithm 4 runs in time $\text{poly}(N, n^d)$. \square

B.3 Subspace Clustering: Robust Case

In the noisy version of the subspace clustering problem, we are given the set A approximately, and we wish to cluster the points such that each cluster is close to a low-dimensional subspace. More formally, the problem is described as follows:

Problem 7. We are given as input an integer $s \in \mathbb{N}$, and a set $\tilde{A} = \{\tilde{\mathbf{a}}_1, \tilde{\mathbf{a}}_2, \dots, \tilde{\mathbf{a}}_N\} \subseteq \mathbb{R}^n$, where each $\tilde{\mathbf{a}}_i$ is close to an (unknown) point $\mathbf{a}_i \in \mathbb{R}^n$, such that the resulting set of points $A = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_N\} \subseteq \mathbb{R}^n$ can be clustered using s low-dimensional subspaces, i.e.

$$A = A_1 \uplus A_2 \uplus \dots \uplus A_s,$$

where each A_j satisfies $\dim(A_j) \leq t$. Our goal is to *efficiently* find an s -tuple of subspaces $\tilde{\mathbf{W}} = (\tilde{W}_1, \tilde{W}_2, \dots, \tilde{W}_s)$ such that (upto reordering) for each $j \in [s]$, it holds that $\text{dist}(\tilde{W}_j, \langle A_j \rangle)$ is small.

Proceeding as in Section 1.3, we devise an algorithm (Algorithm 5) for this problem, via a reduction to robust vector space decomposition.

Algorithm 5 Subspace Clustering.

Input: $(\tilde{A}, d, s, m_d, m_{d-1})$ where $\tilde{A} \subseteq \mathbb{R}^n$ is a set of size N , and $d \geq 2, m_d, m_{d-1}$ are positive integers.

Assumptions: There is a set $A \subseteq \mathbb{R}^n$ of size N , such that

- For each point $\mathbf{a} \in A$, there a unique point $\tilde{\mathbf{a}} \in \tilde{A}$, such that $\|\tilde{\mathbf{a}} - \mathbf{a}\|_2 \leq \epsilon$.
- The set A admits a partition $A = A_1 \uplus A_2 \uplus \dots \uplus A_s$, where each $\langle A_i \rangle$ is of dimension at most t .

Output: $\tilde{\mathbf{W}} = (\tilde{W}_1, \tilde{W}_2, \dots, \tilde{W}_s)$ such that (upto reordering) for each $j \in [s]$, $\text{dist}(\tilde{W}_j, \langle A_j \rangle)$ is small.

- 1: Compute the matrices $M_{\tilde{A}, d}$ and $M_{\tilde{A}, d-1}$ as in Definition B.4. Let \tilde{U} (resp. \tilde{V}) be the subspace spanned by the left singular vectors of $M_{\tilde{A}, d}$ (resp. $M_{\tilde{A}, d-1}$), corresponding to the top m_d (resp. m_{d-1}) singular values.
 - 2: Let $W_1 = \mathbb{R}[\mathbf{x}]^{=d}, W_2 = \mathbb{R}[\mathbf{x}]^{=(d-1)}$, and let $\mathcal{B} = (B_1, \dots, B_n) \in \text{Lin}(W_1, W_2)^n$, where B_i corresponds to the operator $\partial_{x_i} : W_1 \rightarrow W_2$.
Run RVSD Algorithm on $(W_1, W_2, s, \tilde{U}, \tilde{V}, \mathcal{B})$ ³⁹ (see Algorithm 2; use projections as defined in Section 4.3), and let the output be $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_s)$.
 - 3: For each $j \in [s]$, run RRSTP (Algorithm 3) on \tilde{U}_j , and let the output be \tilde{W}_j .
 - 4: Output $\tilde{\mathbf{W}} = (\tilde{W}_1, \dots, \tilde{W}_s)$.
-

The above algorithm gets the following guarantees:

Theorem 10. Let $A = \{\mathbf{a}_1, \dots, \mathbf{a}_N\} \subset \mathbb{R}^n$ be a finite set of N points of unit norm, which can partitioned as $A = A_1 \uplus \dots \uplus A_s$, where each $\langle A_i \rangle$ is subspace of dimension at most t .

Let $d \geq 2$ be an integer, let $\mathbf{U} = (U_1, \dots, U_s)$ (resp. $\mathbf{V} = (V_1, \dots, V_s)$) be an s -tuple of subspaces with $U_j = \langle A_j^{\otimes d} \rangle$ (resp. $V_j = \langle A_j^{\otimes d-1} \rangle$) for each $j \in [s]$. Let $U = \langle \mathbf{U} \rangle$ (resp. $V = \langle \mathbf{V} \rangle$) have dimension m_d (resp. m_{d-1}).

Suppose that:

- $U = U_1 \oplus \dots \oplus U_s, V = V_1 \oplus \dots \oplus V_s$, and for each $j \in [s]$, it holds that $\dim(U_j) = \binom{\dim(\langle A_j \rangle) + d - 1}{d}, \dim(V_j) = \binom{\dim(\langle A_j \rangle) + d - 2}{d-1}$.
- σ_A is the minimum of $\sigma_{m_d}(M_{A, d})$ and $\sigma_{m_{d-1}}(M_{A, d-1})$, where $M_{A, d}$ (resp. $M_{A, d-1}$) is the matrix whose columns are the polynomials $(\mathbf{a}_i \cdot \mathbf{x})^d$ (resp. $(\mathbf{a}_i \cdot \mathbf{x})^{d-1}$) (see Definition B.4).
- $\kappa(\mathbf{U})$ denotes the condition number of the tuple of subspaces \mathbf{U} (see Section 2).
- $\sigma_{-(s+1)}(\mathfrak{A})$ is the $(s+1)^{\text{th}}$ smallest singular value of the adjoint algebra map (see Definition 4.2), corresponding to the action of $\mathcal{B} = (B_1, \dots, B_n)$ on \mathbf{U}, \mathbf{V} , where B_i corresponds to the operator ∂_{x_i} .

Let $\tilde{A} = \{\tilde{\mathbf{a}}_1, \tilde{\mathbf{a}}_2, \dots, \tilde{\mathbf{a}}_N\} \subseteq \mathbb{R}^n$ be a set of unit norm vectors such that $\|\mathbf{a}_i - \tilde{\mathbf{a}}_i\|_2 \leq \epsilon$ for each $i \in [N]$. Let $\delta > 0$. Then, Algorithm 5, on input $(\tilde{A}, d, s, m_d, m_{d-1})$, runs in time

³⁹See Remark 8 for the parameter $\tau \in (0, 1)$

$\text{poly}(N, n^d)$ and outputs $\tilde{\mathbf{W}} = (\tilde{W}_1, \dots, \tilde{W}_s)$ such that with probability at least $1 - \delta$, it holds (upto reordering) that for each $j \in [s]$,

$$\begin{aligned} \text{dist}(\tilde{W}_j, \langle A_j \rangle) &\leq O\left(t^2 \sqrt{Nd^2} \cdot \frac{s^2}{\delta} \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \kappa(\mathbf{U})^3 \cdot \frac{1}{\sigma_A} \cdot \frac{1}{\sigma_{-(s+1)}(\mathfrak{A})} \cdot \epsilon\right) \\ &= \text{poly}\left(t, N, d, s, \frac{1}{\delta}, \kappa(\mathbf{U}), \frac{1}{\sigma_A}, \frac{1}{\sigma_{-(s+1)}(\mathfrak{A})}\right) \cdot \epsilon. \end{aligned}$$

Remark 8. In Theorem 10, we assume that the algorithm gets as input numbers $m_d = \dim(\langle A^{\otimes d} \rangle)$, and $m_{d-1} = \dim(\langle A^{\otimes d-1} \rangle)$ as input. Further, we implicitly assume that the algorithm knows a parameter $\tau \in (0, 1)$ needed for the RVSD Algorithm (Algorithm 2) that lies in the correct range (see Theorem 5).

As mentioned in Remark 3, for the final algorithm, we can iterate over the parameter τ , and similarly over m_d, m_{d-1} , and stop when the output $\tilde{\mathbf{W}} = (\tilde{W}_1, \tilde{W}_2, \dots, \tilde{W}_s)$ gives a valid clustering of the points in the set \tilde{A} into s low-dimensional subspaces; that is:

1. For each $j \in [s]$, it holds $\dim(\tilde{W}_j) \leq t$.
2. For each $i \in [N]$, there exists $j \in [s]$ such that $\tilde{\mathbf{a}}_i$ is close to \tilde{W}_j .

The actual inputs to the algorithm in this case will be (\tilde{A}, d, s, t) . The blow-up in the run time due to this iteration is $O(m_d \cdot m_{d-1} \cdot \log_2(\kappa(\mathbf{U}) \cdot t)) = \text{poly}(s, t^d, \log_2(\kappa(\mathbf{U})))$. Assuming that we have good upper bounds on the condition number $\log_2(\kappa(\mathbf{U}))$, as is required for the error bounds of Theorem 10, we can safely ignore this in the theorem statement.

Similar to the other parameters above, we may also iterate over s , if we know that the input to the problem can be clustered into a relatively small number of subspaces (also see Remark 5).

The remainder of this section is devoted to the proof of Theorem 10. The runtime guarantees follow from the guarantees of each of the individual steps., and we shall omit the details for that.

We shall fix the set $\tilde{A} = \{\tilde{\mathbf{a}}_1, \tilde{\mathbf{a}}_2, \dots, \tilde{\mathbf{a}}_N\} \subseteq \mathbb{R}^n$. Let $A = \{\mathbf{a}_1, \dots, \mathbf{a}_N\} = A_1 \uplus \dots \uplus A_s$ be as above, and let $m_d = \dim(\langle A^{\otimes d} \rangle)$ and $m_{d-1} = \dim(\langle A^{\otimes d-1} \rangle)$.

B.3.1 Closeness of Subspaces

Definition B.4. Consider the space $\mathbb{R}[\mathbf{x}]^{=d}$ with the Bombieri inner product, and of dimension $m_d = \binom{n+d-1}{d}$ (see Section 2). We define $M_{A, d} \in \mathbb{R}^{m_d \times N}$ (resp $M_{\tilde{A}, d}$) to be the matrix whose i^{th} column is the polynomial $(\mathbf{a}_i \cdot \mathbf{x})^d \in \mathbb{R}[\mathbf{x}]^{=d}$ (resp. $(\tilde{\mathbf{a}}_i \cdot \mathbf{x})^d$) written with respect to the Bombieri basis.

Similarly, we also define $m_{d-1} = \binom{n+d-2}{d}$, and $M_{A, d-1} \in \mathbb{R}^{m_{d-1} \times N}$ (resp. $M_{\tilde{A}, d-1}$) with columns $(\mathbf{a}_i \cdot \mathbf{x})^{d-1}$ (resp. $(\tilde{\mathbf{a}}_i \cdot \mathbf{x})^{d-1}$).

Then, $U \stackrel{\text{def}}{=} \langle A^{\otimes d} \rangle \subseteq \mathbb{R}[\mathbf{x}]^{=d}$ corresponds to the column-span of the matrix $M_{A, d}$. Let \tilde{U} be the m_d -dimensional subspace of $M_{\tilde{A}, d}$ closest to U : this equals the vector space spanned by the left-singular vectors of $M_{\tilde{A}, d}$, corresponding the top m_d singular values.

Lemma B.7.

$$\text{dist}(U, \tilde{U}) \leq \frac{2\sqrt{Nd}\epsilon}{\sigma_{m_d}(M_{A,d})},$$

where $\sigma_{m_d}(M_{A,d})$ is the m_d^{th} largest singular value of $M_{A,d}$.

Proof. We have

$$\begin{aligned} \|M_{A,d} - M_{\tilde{A},d}\|_F^2 &= \sum_{i=1}^N \left\| (\mathbf{a}_i \cdot \mathbf{x})^d - (\tilde{\mathbf{a}}_i \cdot \mathbf{x})^d \right\|_B^2 \\ &= \sum_{i=1}^N \left\| (\mathbf{a}_i \cdot \mathbf{x} - \tilde{\mathbf{a}}_i \cdot \mathbf{x}) \left(\sum_{j=0}^{d-1} (\mathbf{a}_i \cdot \mathbf{x})^j (\tilde{\mathbf{a}}_i \cdot \mathbf{x})^{d-1-j} \right) \right\|_B^2 \\ &\leq \sum_{i=1}^N \|\mathbf{a}_i \cdot \mathbf{x} - \tilde{\mathbf{a}}_i \cdot \mathbf{x}\|_B^2 \left\| \left(\sum_{j=0}^{d-1} (\mathbf{a}_i \cdot \mathbf{x})^j (\tilde{\mathbf{a}}_i \cdot \mathbf{x})^{d-1-j} \right) \right\|_B^2 \\ &\leq N\epsilon^2 d^2 \end{aligned}$$

where we use the submultiplicativity and triangle-inequality of the Bombieri norm, and the fact that $\|(\mathbf{b} \cdot \mathbf{x})^k\|_B = \|\mathbf{b}\|_2^k$ for all $\mathbf{b} \in \mathbb{R}^n, k \in \mathbb{N}$. Now, by Corollary A.1, we get

$$\text{dist}(U, \tilde{U}) \leq \frac{2\|M_{A,d} - M_{\tilde{A},d}\|_F}{\sigma_{m_d}(M_{A,d})} \leq \frac{2\sqrt{Nd}\epsilon}{\sigma_{m_d}(M_{A,d})}. \quad \square$$

In a similar manner, let $V = \langle A^{\otimes d-1} \rangle$, and let \tilde{V} be space spanned by the left singular values of $M_{\tilde{A},d-1}$, corresponding to the top m_{d-1} singular values. Then, we get

Lemma B.8.

$$\text{dist}(V, \tilde{V}) \leq \frac{2\sqrt{Nd}\epsilon}{\sigma_{m_{d-1}}(M_{A,d-1})}.$$

B.3.2 Using RVSD and RRSTP

Let $\mathbf{U} = (U_1, \dots, U_s)$, where $U_j = \langle A_j^{\otimes d} \rangle$ for each $j \in [s]$. Also, let $(\tilde{U}_1, \dots, \tilde{U}_s)$, be the output of RVSD in Step 2 of Algorithm 5.

Let $\delta > 0$. Then, by Corollary 4.1, we get that with probability at least $1 - \delta$, (upto reordering) for each $j \in [s]$,

$$\text{dist}(\tilde{U}_j, U_j) \leq O\left(t^{3/2} \cdot \kappa(\mathbf{U})^3 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \frac{\epsilon_1 + \epsilon_2}{\delta} \cdot \frac{\|\mathcal{B}\|_2}{\sigma_{-(s+1)}(\mathfrak{A})}\right),$$

where we have that

1. $\sigma_{-(s+1)}(\mathfrak{A})$ is the $(s+1)^{\text{th}}$ smallest singular value of the relevant adjoint algebra map.
2. $\mathcal{B} = (B_1, \dots, B_n) \in \text{Lin}\left(\mathbb{R}[\mathbf{x}]^{\leq d}, \mathbb{R}[\mathbf{x}]^{\leq (d-1)}\right)^n$, with B_i being the map corresponding to the operation ∂_{x_i} . Note that by Definition 2.1 and Lemma 2.1, we have $\|\mathcal{B}\|_2 = d$.

3. By Lemma B.7 and Lemma B.8, $\text{dist}(\tilde{U}, U) \leq \epsilon_1 = \frac{2\sqrt{Nd}\epsilon}{\sigma_{m_d}(M_{A,d})}$ and $\text{dist}(\tilde{V}, V) \leq \epsilon_2 = \frac{2\sqrt{Nd}\epsilon}{\sigma_{m_{d-1}}(M_{A,d-1})}$.

Simplifying, we get

$$\text{dist}(\tilde{U}_j, U_j) \leq O\left(t^{3/2}\sqrt{Nd}^2 \cdot \frac{s^2}{\delta} \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \kappa(\mathbf{U})^3 \cdot \frac{1}{\min\{\sigma_{m_d}(M_{A,d}), \sigma_{m_{d-1}}(M_{A,d-1})\}} \cdot \frac{1}{\sigma_{-(s+1)}(\mathfrak{A})} \cdot \epsilon\right).$$

Now, for each $j \in [s]$, let $W_j = \langle A_j \rangle$. Then, assuming that $\dim(U_j) = \binom{\dim(W_j)+d-1}{d}$, we have by Proposition B.1, that the output \tilde{W}_j of the RRSTP algorithm satisfies

$$\text{dist}(\tilde{W}_j, W_j) \leq O\left(t^2\sqrt{Nd}^2 \cdot \frac{s^2}{\delta} \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \kappa(\mathbf{U})^3 \cdot \frac{1}{\min\{\sigma_{m_d}(M_{A,d}), \sigma_{m_{d-1}}(M_{A,d-1})\}} \cdot \frac{1}{\sigma_{-(s+1)}(\mathfrak{A})} \cdot \epsilon\right).$$

B.4 Singular Value Analysis for The Adjoint Algebra

Analysis of the singular values of the adjoint algebra operator can be tedious. However, for subspace clustering we can obtain substantial lower bounds for $\sigma_{-(s+1)}(\mathfrak{A})$ which reveal how the geometry of the original subspaces affect the robustness of the algorithm. We dedicate section G to this analysis. We proceed by defining a special inner product for linear maps on sums of subspaces. This allows an inductive approach to separately bound the contributions of the diagonal and off-diagonal blocks of \mathfrak{A} to its smallest non-zero singular value. In particular we get the following theorem.

Theorem 11 (Theorem 15 restated). *For subspaces U, V of dimension t , let f_d be defined as $f_d(U, V) = \frac{d}{t} \left[\sum_{k \in [t]} \sin^2 \theta_k + (d-1) \sin^2 \theta_t \right]$ for $d \geq 2$, where $\theta_1 \geq \theta_2 \geq \dots \geq \theta_t$ are the canonical angles between U, V . Then, if \mathfrak{A} represents the adjoint algebra map corresponding to the subspace clustering problem for subspaces $\langle A_1 \rangle, \langle A_2 \rangle, \dots, \langle A_s \rangle$ with parameter d , we have*

$$\sigma_{-(s+1)}^2(\mathfrak{A}) \geq \left[\frac{d}{\kappa(\mathbf{U}, \mathbf{V})^2} \right]^2 \cdot \min\{\sigma_{\text{diag}}, \sigma_{\text{off-diag}}\}$$

where the above quantities are defined as follows:

$$\sigma_{\text{diag}} \stackrel{\text{def}}{=} \sqrt{\frac{d}{t^* + d - 1}} \left(1 - \sqrt{1 - \frac{1}{d} \frac{t^*}{t^* + d - 1}} \right),$$

$$\sigma_{\text{off-diag}} \stackrel{\text{def}}{=} \min_{j \neq k} \sqrt{\frac{d}{t_k + d - 1}} \left(1 - \sqrt{1 - \frac{1}{d} \frac{t_k}{t_k + d - 1} \cdot f_d(\langle A_j \rangle, \langle A_k \rangle)} \right),$$

where $\kappa(\mathbf{U}, \mathbf{V}) = \max\{\sigma_1(\mathbf{U}), \sigma_1(\mathbf{V})\} / \min\{\sigma_{-1}(\mathbf{U}), \sigma_{-1}(\mathbf{V})\}$, and t_1, \dots, t_s are dimensions of $\langle A_1 \rangle, \dots, \langle A_s \rangle$ respectively, with $t^* = \max_{i \in [s]} t_i$.

Refer to section G for the proof of Theorem 11.

B.5 Smoothed Analysis of Subspace Clustering

We analyse our algorithm for subspace clustering in a smoothed setting. We first describe the input model for our problem. For simplicity, we assume that each of the subspaces has the same dimension.

1. We have a tuple of s hidden subspaces of \mathbb{R}^n , $\mathbf{W} = (W_1, W_2, \dots, W_s)$, each of dimension t . Let $P_1, P_2, \dots, P_s \in \mathbb{R}^{n \times t}$ be matrices with orthonormal columns, such that the column span of P_i is W_i . Each subspace W_i is perturbed by perturbing P_i by a random Gaussian matrix $G_i \sim \mathcal{N}(0, \rho^2/n)^{n \times t}$. Let $\hat{P}_i = P_i + G_i$, and $\hat{W}_1, \hat{W}_2, \dots, \hat{W}_s$ be the column spans of $\hat{P}_1, \hat{P}_2, \dots, \hat{P}_s$ respectively.
2. Sample sets of points A_1, A_2, \dots, A_s from $\hat{W}_1, \hat{W}_2, \dots, \hat{W}_s$ respectively, of unit norm. For each $i \in [s]$, perturb each point in A_i with respect to \hat{W}_i to get the set of points \hat{A}_i . Formally, this means perturbing points in A_i by $\hat{B}_i \cdot v$, where \hat{B}_i is an $n \times t$ matrix describing an orthonormal basis for \hat{W}_i and $v \sim \mathcal{N}(0, \rho^2/t)^t$, and normalizing. Let $\hat{A} = \hat{A}_1 \cup \hat{A}_2 \cup \dots \cup \hat{A}_s$.
3. For each $\mathbf{a} \in \hat{A}$, add noise (and normalize) to get point \mathbf{a}' such that $\|\mathbf{a} - \mathbf{a}'\|_2 \leq \epsilon$. We are given \hat{A}' , the set of noise-added points.

Given the set of points \hat{A}' , the goal is to recover subspaces $\tilde{\mathbf{W}} = (\tilde{W}_1, \tilde{W}_2, \dots, \tilde{W}_s)$ such that $\text{dist}(\hat{W}_j, \tilde{W}_j)$ is small for each $j \in [s]$.

Algorithm 5 gets the following guarantees in the above smoothed setting.

Theorem 12. *Let $\mathbf{W} = (W_1, \dots, W_s)$, \hat{A}' and $\hat{A} = \hat{A}_1 \uplus \dots \uplus \hat{A}_s$ be as generated above, with smoothing parameter $\rho \in (0, 1)$. Let $\hat{\mathbf{U}} = (\hat{U}_1, \dots, \hat{U}_s)$ and $\hat{\mathbf{V}} = (\hat{V}_1, \dots, \hat{V}_s)$ be s -tuples such that $\hat{U}_j = \langle \hat{A}_j^{\otimes d} \rangle$, $\hat{V}_j = \langle \hat{A}_j^{\otimes d-1} \rangle$ for each $j \in [s]$, for some $d \geq 2$. Further, for each $j \in [s]$, let $M_{\hat{A}_j, d} = C_{j, d} \cdot M_{\hat{B}_j, d}$ and $M_{\hat{A}_j, d-1} = C_{j, d-1} \cdot M_{\hat{B}_j, d-1}$ where $C_{j, d}$ (resp. $C_{j, d-1}$) is a matrix with columns as an orthonormal basis of \hat{U}_j (resp. \hat{V}_j), and $B_j \subset \mathbb{R}^t$. Let $\hat{\mathbf{U}} = \langle \hat{\mathbf{U}} \rangle$, $\hat{\mathbf{V}} = \langle \hat{\mathbf{V}} \rangle$, and $m_d = \dim(\hat{\mathbf{U}})$, $m_{d-1} = \dim(\hat{\mathbf{V}})$. Let $\delta > 0$ and $t > d$. Then, Algorithm 5 on input $(\hat{A}', d, s, m_d, m_{d-1})$ outputs $\tilde{\mathbf{W}} = (\tilde{W}_1, \dots, \tilde{W}_s)$ such that with probability $1 - \delta - s^2 \exp(-\Omega(\rho^2 n))$,*

$$\text{dist}(\tilde{W}_j, \hat{W}_j) \leq O \left(\sqrt{N} \cdot \frac{s^2}{\delta} \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \kappa^6 \cdot \frac{1}{\sigma} \cdot t^{3/4} \cdot d^{3/4} \cdot n^{1/2} \cdot \frac{\epsilon}{\rho} \right)$$

where

$$\kappa = \frac{\max \{ \sigma_1(\mathbf{U}), \sigma_1(\mathbf{V}) \}}{\min \{ \sigma_{-1}(\mathbf{U}), \sigma_{-1}(\mathbf{V}) \}}, \quad \sigma = \min_{j \in [s]} \min \left\{ \sigma_{-1}(M_{\hat{B}_j, d}), \sigma_{-1}(M_{\hat{B}_j, d-1}) \right\}, \quad \text{and } N = |\hat{A}'|.$$

The rest of this section is devoted to the proof of Theorem 12.

We first show that the perturbed subspaces generated by the above procedure are well separated, with high probability.

Lemma B.9. *Let V be a subspace of \mathbb{R}^n of dimension $t < n$. Let \hat{V} be the perturbed subspace generated from V by perturbing an orthonormal basis of V by vectors sampled from $\mathcal{N}(0, \rho^2/n)^n$. Then, if U is another subspace of dimension t , we have, with high probability,*

$$\text{dist}(U, \hat{V}) \geq \frac{1}{5} \cdot \frac{\rho}{\sqrt{1 + \rho^2}} \cdot \sqrt{1 - \frac{t}{n}}$$

Proof. Without loss of generality, we can assume that $U = \langle \mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_t \rangle$. Let $V = \langle \mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_t \rangle$ be an orthonormal basis for V , and $\hat{V} = \langle \hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2, \dots, \hat{\mathbf{v}}_t \rangle$, where $\hat{\mathbf{v}}_i = \mathbf{v}_i + \mathbf{b}_i$, and $\mathbf{b}_i \sim \mathcal{N}(0, \rho^2/n)^n$. We have

$$\begin{aligned} \text{dist}(U, \hat{V}) &= \|\text{Proj}_U - \text{Proj}_{\hat{V}}\|_2 \\ &\geq \frac{\|\text{Proj}_U \cdot \hat{\mathbf{v}}_1 - \text{Proj}_{\hat{V}} \cdot \hat{\mathbf{v}}_1\|_2}{\|\hat{\mathbf{v}}_1\|_2} \\ &= \frac{\|\text{Proj}_U \cdot \hat{\mathbf{v}}_1 - \hat{\mathbf{v}}_1\|_2}{\|\hat{\mathbf{v}}_1\|_2}. \end{aligned}$$

Since \mathbf{v}_1 has unit norm, we have that

$$\Pr \left[\|\hat{\mathbf{v}}_1\|_2^2 \geq 1 + 5\rho^2 \right] \leq \exp(-\Omega(\rho^2 n)).$$

Also, since $U = \langle \mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_t \rangle$, if $\hat{\mathbf{v}}_1 = \sum_{i \in [n]} \hat{v}_{1i} \mathbf{e}_i$, then $\text{Proj}_U \cdot \hat{\mathbf{v}}_1 = \sum_{i \in [t]} \hat{v}_{1i} \mathbf{e}_i$. Therefore,

$$\|\text{Proj}_U \cdot \hat{\mathbf{v}}_1 - \hat{\mathbf{v}}_1\|_2^2 = \sum_{i=t+1}^n \hat{v}_{1i}^2.$$

Thus, we have

$$\Pr \left[\|\text{Proj}_U \cdot \hat{\mathbf{v}}_1 - \hat{\mathbf{v}}_1\|_2^2 \leq \frac{2\rho^2}{5} \cdot \left(1 - \frac{t}{n}\right) \right] \leq \exp(-\Omega(\rho^2 n)).$$

Therefore, we have, with probability at least $1 - \exp(-\Omega(\rho^2 n))$,

$$\begin{aligned} \frac{\|\text{Proj}_U \cdot \hat{\mathbf{v}}_1 - \hat{\mathbf{v}}_1\|_2^2}{\|\hat{\mathbf{v}}_1\|_2^2} &\geq \frac{2\rho^2}{25(1 + \rho^2)} \cdot \left(1 - \frac{t}{n}\right) \\ &\geq \frac{1}{25} \cdot \frac{\rho^2}{1 + \rho^2} \cdot \left(1 - \frac{t}{n}\right). \end{aligned}$$

□

As a corollary, we get a lower bound on the smallest non-zero singular value of the adjoint algebra for subspace clustering, in the smoothed setting.

Corollary B.2. *Let the subspaces $\hat{W}_1 = \langle \hat{A}_1 \rangle, \dots, \hat{W}_s = \langle \hat{A}_s \rangle$ be generated as given above. Let $\mathbf{U} = \left(\langle \hat{A}_1^{\otimes d} \rangle, \dots, \langle \hat{A}_s^{\otimes d} \rangle \right)$, $\mathbf{V} = \left(\langle \hat{A}_1^{\otimes d-1} \rangle, \dots, \langle \hat{A}_s^{\otimes d-1} \rangle \right)$, for $d \geq 2$. Let \mathfrak{A} be the adjoint algebra map corresponding to action of the derivative maps on \mathbf{U}, \mathbf{V} . Then, we have, with probability at least $1 - s^2 \exp(-\Omega(\rho^2 n))$,*

$$\sigma_{-(s+1)}^2(\mathfrak{A}) \geq \frac{1}{50\kappa^4} \cdot \frac{d^{5/2}}{(t+d-1)^{3/2}} \cdot \frac{\rho^2}{1 + \rho^2} \cdot \frac{1}{n}.$$

where $\kappa = \kappa(\hat{\mathbf{U}}, \hat{\mathbf{V}})$.

Proof. By Theorem 11, we have

$$\sigma_{-(s+1)}^2(\mathfrak{Q}) \geq \frac{d^2}{\kappa^4} \cdot \min\{\sigma_{\text{diag}}, \sigma_{\text{off-diag}}\},$$

where

$$\begin{aligned} \sigma_{\text{diag}} &= \sqrt{\frac{d}{t+d-1}} \left(1 - \sqrt{1 - \frac{1}{d} \cdot \frac{t}{d-1+t}} \right) \\ &\geq \frac{1}{2} \cdot \frac{t}{d^{1/2} \cdot (t+d-1)^{3/2}} \end{aligned}$$

and

$$\begin{aligned} \sigma_{\text{off-diag}} &= \min_{i \neq j} \sqrt{\frac{d}{t+d-1}} \cdot \left(1 - \sqrt{1 - \frac{1}{d} \cdot \frac{t}{d-1+t} \cdot f_d(\hat{W}_i, \hat{W}_j)} \right) \\ &\geq \frac{1}{2} \cdot \frac{t}{d^{1/2} \cdot (t+d-1)^{3/2}} \cdot \min_{i \neq j} f_d(\hat{W}_i, \hat{W}_j). \end{aligned}$$

For subspaces U, V with canonical angles $\theta_1 \geq \theta_2 \geq \dots \geq \theta_t \geq 0$, we have

$$f_d(U, V) = \frac{d}{t} \left[\sum_{k \in [t]} \sin^2 \theta_k + (d-1) \sin^2 \theta_t \right] \geq \frac{d}{t} \cdot \sin^2(\theta_1) = \frac{d}{t} \cdot \text{dist}(U, V)^2.$$

Thus, by Lemma B.9, we have, with probability $1 - s^2 \exp(-\Omega(\rho^2 n))$,

$$f_d(\hat{W}_i, \hat{W}_j) \geq \frac{1}{25} \cdot \frac{d}{t} \cdot \frac{\rho^2}{1 + \rho^2} \cdot \left(1 - \frac{t}{n}\right).$$

for all $i, j \in [s]$ such that $i \neq j$. Since $t \leq n - 1$, we have the required expression. \square

Using the above, and the fact that $\mathbf{W}, \tilde{\mathbf{W}}$ will be close with high probability, we have our result.

Proof of Theorem 12. By Corollary B.2 and Theorem 10, if Algorithm 5 returns subspaces $\tilde{W} = (\tilde{W}_1, \dots, \tilde{W}_s)$, we have, with probability $1 - \delta - s^2 \exp(-\Omega(\rho^2 n))$,

$$\text{dist}(\tilde{W}_j, W_j) \leq O \left(\sqrt{N} \cdot \frac{s^2}{\delta} \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \kappa^5 \cdot \frac{1}{\sigma'} \cdot t^{3/4} d^{3/4} n^{1/2} \cdot \frac{\epsilon}{\rho} \right)$$

and $\sigma' = \min \left\{ \sigma_{m_d(M_{\hat{A},d})}, \sigma_{m_{d-1}(M_{\hat{A},d-1})} \right\}$.

Note that we can write $M_{\hat{A},d}$ as

$$M_{\hat{A},d} = [C_{1,d} \quad \dots \quad C_{s,d}] \cdot \begin{bmatrix} M_{\hat{B}_1,d} & & \\ & \ddots & \\ & & M_{\hat{B}_s,d} \end{bmatrix}.$$

We can similarly write an expression for $M_{\hat{A},d-1}$. Thus, we get that

$$\sigma' \geq \frac{1}{\kappa} \min_{j \in [s]} \min \left\{ \sigma_{-1}(M_{\hat{B}_j,d}), \sigma_{-1}(M_{\hat{B}_j,d-1}) \right\}.$$

This gives us the required result. \square

C Learning Mixtures of Gaussians

In this section we give the proofs of technical claims pertaining to section 1.4.

Lemma C.1. [[GKS20]] **Blessing of Dimensionality for sums of powers of quadratics.** *Let positive n, k, d, s be non-negative integers satisfying the following constraints.*

$$s \cdot \binom{n+k-1}{k} \ll \binom{n+2d-k-1}{2d-k}.$$

Then with high probability over the random choice of $p_1(\mathbf{x}), p_2(\mathbf{x}), \dots, p_s(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]$ we have

$$\langle \mathcal{L} \cdot p \rangle = \langle \mathbf{y}^k \cdot q_1^{d-k} \rangle \oplus \langle \mathbf{y}^k \cdot q_2^{d-k} \rangle \oplus \dots \oplus \langle \mathbf{y}^k \cdot q_s^{d-k} \rangle,$$

where for each $i \in [s]$, $q_i(\mathbf{y}) \in \mathbb{R}[\mathbf{y}]$ is a restriction of p_i , and $|\mathbf{y}| = n$.

Remark 9. *The constraints/bounds on the parameters for which the conclusion of the above lemma holds is likely suboptimal and the conclusion is likely to hold for a larger range.*

Lemma C.2. [[GKS20]] **Blessing of Dimensionality for shifted spaces of powers of independent quadratics.** *Let positive n, k, e, s be non-negative integers satisfying the following constraints.*

$$s \cdot \binom{n+k-1}{k} \ll \binom{n+2e+k-1}{2e+k}.$$

Then with high probability over the random choice of $q_1(\mathbf{y}), q_2(\mathbf{y}), \dots, q_s(\mathbf{y}) \in \mathbb{R}[\mathbf{y}]$, $|\mathbf{y}| = n$ we have

$$\langle \mathbf{y}^k \cdot q_1^e \rangle + \langle \mathbf{y}^k \cdot q_2^e \rangle + \dots + \langle \mathbf{y}^k \cdot q_s^e \rangle = \langle \mathbf{y}^k \cdot q_1^e \rangle \oplus \langle \mathbf{y}^k \cdot q_2^e \rangle \oplus \dots \oplus \langle \mathbf{y}^k \cdot q_s^e \rangle$$

Remark 10. *The constraints/bounds on the parameters for which the conclusion of the above lemma holds is likely suboptimal and the conclusion is likely to hold for a larger range.*

Lemma C.3. *The adjoint algebra A for the action of \mathcal{B} has dimension s . In other words for any $(D, E) \in A$ there exists $(\lambda_1, \lambda_2, \dots, \lambda_s) \in \mathbb{R}^s$ such that for all $i \in [s]$, $f(\mathbf{y}) \in \mathbb{R}[\mathbf{y}]^k, g(\mathbf{y}) \in \mathbb{R}[\mathbf{y}]^{(k+2)}$*

$$D \cdot (q_i(\mathbf{y})^e \cdot f(\mathbf{y})) = \lambda_i \cdot (q_i(\mathbf{y})^e \cdot f(\mathbf{y}))$$

and

$$E \cdot (q_i(\mathbf{y})^{e-1} \cdot g(\mathbf{y})) = \lambda_i \cdot (q_i(\mathbf{y})^{e-1} \cdot g(\mathbf{y})).$$

Corollary C.1. *The decomposition of U under the action of \mathcal{B} is unique.*

We need a couple of technical lemmas.

Lemma C.4. *Fix any (D, E) in A . Then there exist $R_1, R_2, \dots, R_s \in \langle q_1^{d-k}, q_2^{d-k}, \dots, q_s^{d-k} \rangle$ such that for all $f(\mathbf{y}) \in \mathbb{R}[\mathbf{y}]^k$ we have $D \cdot (q_i^e \cdot f(\mathbf{y})) = R_i \cdot f(\mathbf{y})$.*

Lemma C.5. *Let $q_1(\mathbf{y}), q_2(\mathbf{y}), \dots, q_s(\mathbf{y})$ be independently chosen random quadratic forms. For all $i \in [s]$, there exists $B_i \in \mathcal{B}$ and $\ell_i \in \mathbb{R}[\mathbf{y}]^k$ such that $B_i \cdot (q_i^e \cdot \ell_i^k) = B_i \cdot \ell_i = 0$ and moreover for all $j \in [s] \setminus \{i\}$, we have $B_i \cdot (q_j^e \cdot \ell_i^k) \neq 0$.*

Proof of Lemma C.3. Let (D, E) be any element of the adjoint algebra A . By lemma C.4, there exist $R_1, R_2, \dots, R_s \in \langle q_1^{d-k}, q_2^{d-k}, \dots, q_s^{d-k} \rangle$ such that for all $f(\mathbf{y}) \in \mathbb{R}[\mathbf{y}]^k$ we have $D \cdot (q_i^e \cdot f(\mathbf{y})) = R_i \cdot f(\mathbf{y})$. Let $R_i = c_{i1}q_1^e + c_{i2}q_2^e + \dots + c_{is}q_s^e$. For an arbitrary $i \in [s]$, let $B_i \in \mathcal{B}, \ell_i \in \mathbb{R}[\mathbf{y}]^k$ be as provided by lemma C.5. Now

$$\begin{aligned}
& B_i \cdot (q_i^e \cdot \ell_i^k) = 0 \\
\implies & E \cdot B_i \cdot (q_i^e \cdot \ell_i^k) = 0 \\
\implies & B_i \cdot D \cdot (q_i^e \cdot \ell_i^k) = 0 \\
& \implies B_i \cdot (R_i \cdot \ell_i^k) = 0 \\
\implies & B_i \cdot \left(\left(\sum_{j \in [s]} c_{ij} q_j^e \right) \cdot \ell_i^k \right) = 0 \\
& \implies \sum_{j \in [s]} c_{ij} \cdot B_i \cdot (q_j^e \cdot \ell_i^k) = 0 \\
\implies & \forall j \in [s] : c_{ij} \cdot (B_i \cdot (q_j^e \cdot \ell_i^k)) = 0 \quad (\text{as } V_j\text{'s form a direct sum}) \\
& \implies \forall j \in ([s] \setminus \{i\}) : c_{ij} = 0 \quad (\text{using the property of } B_i, \ell_i \text{ from lemma C.5})
\end{aligned}$$

Thus $R_i(\mathbf{y}) = c_{ii} \cdot q_i(\mathbf{y})^e$ and so that for all $f \in \mathbb{R}[\mathbf{y}]^k, g \in \mathbb{R}[\mathbf{y}]^{(k+2)}$ we have $D \cdot q_i^e \cdot f = c_{ii} \cdot q_i^e \cdot f$ and $E \cdot q_i^{e-1} \cdot g = c_{ii} \cdot q_i^{e-1} \cdot g$. In particular, the adjoint algebra for \mathcal{B} is s -dimensional. \square

We will now provide the proofs of lemmas C.4 and C.5. For this we will in turn need a couple more preliminary observations.

Existence of nice basis. Let us first observe that the space of homogeneous polynomials admits a basis consisting of *very simple* polynomials for pretty much any definition of what is a very simple polynomial. In particular let us work with the specific notions below of what a very simple polynomial is. If a polynomial $g(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^k$ is of the form $g(\mathbf{x}) = \ell(\mathbf{x})^k$ for some linear form $\ell(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^1$, let us call it as a power of linear form and in short refer to it as a homogeneous- $\Sigma \wedge$ polynomial. The following is a classical result.

Theorem 13. Ellison [Ell69]. *There exists a basis of $\mathbb{R}[\mathbf{x}]$ consisting of homogeneous- $\Sigma \wedge$ polynomials.*

We will need variants of the above theorem for some other notions of what constitutes a very simple polynomial. If a polynomial $g(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^k$ is of the form

$$g(\mathbf{x}) = \ell_1(\mathbf{x}) \cdot \ell_2(\mathbf{x}) \cdot \dots \cdot \ell_k(\mathbf{x}),$$

$\ell_1(\mathbf{x}), \ell_2(\mathbf{x}), \dots, \ell_k(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^1$ are coprime linear forms let us call it as a product of coprime linear forms and in short refer to it as a homogeneous- $\Pi^{\text{coprime}} \Sigma$ polynomial. For an integer $r \geq 2$ we further call such a polynomial as a homogeneous- $\Pi^{\text{coprime}, r\text{-span}} \Sigma$ polynomial if we also have $\dim(\langle \ell_1(\mathbf{x}), \ell_2(\mathbf{x}), \dots, \ell_k(\mathbf{x}) \rangle) = r$.

Claim C.1. *There exists a basis of $\mathbb{R}[\mathbf{x}]^k$ consisting of homogeneous- $\Pi^{\text{coprime}, 2\text{-span}} \Sigma$ polynomials.*

Proof. By induction on the number of variables $n = |\mathbf{x}|$.

Base case. For the base case of $n = 2$, let $\alpha_1, \alpha_2, \dots, \alpha_{k+1} \in \mathbb{R}$ be any set of $k + 1$ distinct field elements. Let $p(x_1, x_2) \stackrel{\text{def}}{=} (x_1 + \alpha_1 x_2) \cdot (x_1 + \alpha_2 x_2) \cdot \dots \cdot (x_1 + \alpha_{k+1} x_2)$. Let $p_i \in \mathbb{R}[\mathbf{x}]^{\leq k}$ be defined as $p_i(x_1, x_2) \stackrel{\text{def}}{=} \prod_{j \neq i} (x_1 + \alpha_j x_2) = \frac{p}{(x_1 + \alpha_i x_2)}$. We claim that the p_i 's ($i \in [k + 1]$) form the required basis of $\mathbb{R}[\mathbf{x}]^{\leq k}$. Each of them is clearly a product of coprime linear forms and so it suffices to show they are linearly independent. Suppose that

$$c_1 p_1 + c_2 p_2 + \dots + c_{k+1} p_{k+1} = 0.$$

Making the substitution $x_1 = -\alpha_1 x_2$ in the above identity we get that

$$\begin{aligned} c_1 \cdot (-\alpha_1 x_2 + \alpha_2 x_2) \cdot (-\alpha_1 x_2 + \alpha_3 x_2) \cdot \dots \cdot (-\alpha_1 x_2 + \alpha_{k+1} x_2) \\ + c_2 \cdot 0 + c_3 \cdot 0 + \dots + c_{k+1} \cdot 0 = 0 \end{aligned}$$

from which we infer that $c_1 = 0$. Similarly we can infer $c_i = 0$ for all $i \in [k + 1]$ implying that the p_i 's are linearly independent, as required.

Inductive step. Suffices to show that any monomial $m \in \mathbb{R}[\mathbf{x}]^{\leq k}$ can be expressed as a linear combination of homogeneous- $\Pi^{\text{coprime}, r\text{-span}} \Sigma$ polynomials. Suppose that $m = x_1^{e_1} \cdot x_2^{e_2} \cdot \dots \cdot x_n^{e_n}$ where $\sum_{i \in [n]} e_i = k$. By theorem 13 let $m_1 \stackrel{\text{def}}{=} x_2^{e_2} \cdot x_3^{e_3} \cdot \dots \cdot x_n^{e_n}$ admit a representation as a sum of powers of linear forms as

$$m_1 = \sum_i \ell_i(x_2, x_3, \dots, x_n)^r, \text{ where } r = e_2 + e_3 + \dots + e_n.$$

So it suffices to show that for any i , $(x_1^{e_1} \cdot \ell_i^r)$ can be expressed as a sum of homogeneous- $\Pi^{\text{coprime}, 2\text{-span}} \Sigma$ polynomials. By making a suitable change of variables we can assume that $\ell_i = x_2$ and then we can infer the previous statement from the base case above. \square

Now consider a linear form $\ell(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^{\leq 1}$ that divides a product of two polynomials $q(\mathbf{x})$ and $p(\mathbf{x})$. If q is coprime to ℓ we can infer that ℓ divides p . Now suppose that ℓ divides $\partial_1(q(\mathbf{x}) \cdot p(\mathbf{x}))$. We would still like to infer that ℓ divides p . This is not true in general but with some mild conditions (a slightly more general version of) it does hold.

Claim C.2. Let $q(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^{\leq d}$ be a homogeneous polynomial of degree d and $p(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^{\leq k}$ be a homogeneous polynomial of degree k . Let $\ell(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^{\leq 1}$ be a linear form. For $\mathbf{a} = (a_1, a_2, \dots, a_n) \in \mathbb{R}^n$, let $(\mathbf{a} \cdot \partial) \stackrel{\text{def}}{=} (a_1 \partial_1 + a_2 \partial_2 + \dots + a_n \partial_n) \in \langle \partial^{\leq 1} \rangle$ be a first order differential operator. Suppose that there exists positive constants $c_1, c_2 \in \mathbb{R}_{>0}$ such that

$$\ell \mid (c_1((\mathbf{a} \cdot \partial)q) \cdot p + c_2 q \cdot ((\mathbf{a} \cdot \partial)p)). \quad (11)$$

If

$$\gcd(q, (\mathbf{a} \cdot \partial)q) = 1 \pmod{\ell(\mathbf{x})}$$

then $\ell \mid p$.

Proof. By making a suitable change of variables if needed, we can assume without loss of generality that $\ell(\mathbf{x}) = x_1$ and $\mathbf{a} = (0, 1, 0, 0, \dots, 0)$ so that $(\mathbf{a} \cdot \partial)$ is simply ∂_2 ,

i.e. the derivative with respect to the variable x_2 . Let $\hat{p} \stackrel{\text{def}}{=} p(0, x_2, x_3, \dots, x_n)$ and $\hat{q} \stackrel{\text{def}}{=} q(0, x_2, x_3, \dots, x_n)$. The conclusion of the above claim can be restated as claiming that $\hat{p} = 0$. Suppose not. Then we can write \hat{p} as $\hat{p} = \hat{q}^r \cdot f$, for some polynomial integer $r \geq 0$ and some polynomial $f \in \mathbb{R}[x_2, x_3, \dots, x_n]$ which is not divisible by \hat{q} . Now equation (11) implies that

$$\begin{aligned} c_1 \cdot (\partial_2 \hat{q}) \cdot \hat{p} + c_2 \cdot \hat{q} \cdot (\partial_2 \hat{p}) &= 0 \\ \implies (c_1 + rc_2) \cdot (\partial_2 \hat{q}) \cdot f + c_2 \cdot \hat{q} \cdot (\partial_2 f) &= 0 \end{aligned}$$

from which we can infer that \hat{q} divides $(c_1 + rc_2) \cdot (\partial_2 \hat{q}) \cdot f$. But $(c_1 + rc_2)$ is positive and hence nonzero and by assumption $\gcd(\hat{q}, (\partial_2 \hat{q})) = 1$ so that \hat{q} must divide f , a contradiction. Thus we must have $\hat{p} = 0$, or equivalently that $\ell(\mathbf{x})$ divides p . \square

A nondegenerate quadratic form (specifically, one which has rank at least 5) satisfies the desired property above.

Corollary C.2. *Let $\ell(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^1$ be a linear form. For $\mathbf{a} = (a_1, a_2, \dots, a_n) \in \mathbb{R}^n$, let $(\mathbf{a} \cdot \partial) \stackrel{\text{def}}{=} (a_1 \partial_1 + a_2 \partial_2 + \dots + a_n \partial_n) \in \langle \partial^1 \rangle$ be a first order differential operator. If $q(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^2$ is a quadratic form of rank at least 5 and $\ell(\mathbf{a}) = 0$ then it must hold that*

$$\gcd(q, (\mathbf{a} \cdot \partial)q) = 1 \pmod{\ell(\mathbf{x})}.$$

Consequently, for any polynomial $p(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^k$, if $\ell(\mathbf{x})$ divides $(\mathbf{a} \cdot \partial)(q \cdot p)$ then $\ell(\mathbf{x})$ divides p .

Proof. By making a suitable change of variables we can assume without loss of generality that $\ell(\mathbf{x}) = x_1$ and that the differential operator $(\mathbf{a} \cdot \partial)$ is ∂_2 and that $q(\mathbf{x}) = x_1 \cdot p_1(\mathbf{x}) + x_2 \cdot p_2(x_3, x_4, \dots, x_n) + c_3 \cdot x_3^2 \dots + c_r \cdot x_r^2$, where $p_1(\mathbf{x}), p_2(\mathbf{x})$ are linear forms and r is the rank of q . Then we have:

$$\begin{aligned} \gcd(q, (\mathbf{a} \cdot \partial)q) \pmod{\ell(\mathbf{x})} &= \gcd(x_2 \cdot p_2 + c_3 \cdot x_3^2 \dots + c_r \cdot x_r^2, p_2) \\ &= \gcd(c_3 \cdot x_3^2 \dots + c_r \cdot x_r^2, p_2) \\ &= 1, \quad (\text{as } (c_3 \cdot x_3^2 \dots + c_r \cdot x_r^2) \text{ is irreducible for } r \geq 5) \end{aligned}$$

The second conclusion follows from claim C.2 above. \square

We are now ready to prove our main technical lemmas.

Proof of Lemma C.4. Suffices to show that there exist constants $c_{11}, c_{12}, \dots, c_{1s} \in \mathbb{R}$ such that $D \cdot q_1^e \cdot f = (c_{11}q_1^e + c_{12}q_2^e + \dots + c_{1s}q_s^e) \cdot f$ for all polynomials $f \in \mathbb{R}[\mathbf{x}]^k$. Using the existence of nice bases of $\mathbb{R}[\mathbf{x}]^k$ as provided by claim C.1, it suffices to prove the lemma for polynomials $f(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^k$ of the form

$$f(\mathbf{x}) = \ell_1 \cdot \ell_2 \cdot \dots \cdot \ell_k,$$

where the ℓ_i 's are coprime linear forms spanning only 2 dimensions. By making a suitable change of variables we can assume that each ℓ_i is a linear form over the variables x_1 and x_2 only. Let

$$\begin{aligned} D \cdot (q_1^e \cdot f) &= q_1^e \cdot p_1 + q_2^e \cdot p_2 + \dots + q_s^e \cdot p_s, \quad \text{and} \\ D \cdot (q_1^e \cdot x_4 \ell_2 \ell_3 \ell_k) &= q_1^e \cdot \hat{p}_1 + q_2^e \cdot \hat{p}_2 + \dots + q_s^e \cdot \hat{p}_s. \end{aligned}$$

where the p_i 's and \hat{p}_j 's are in $\mathbb{R}[\mathbf{x}]^{=k}$. Now since f is a polynomial over only x_1 and x_2 we have $\partial_3 f = 0$ which implies that

$$\begin{aligned}
& x_4 \cdot \partial_3 \cdot (q_1^e \cdot \ell_1 \ell_2 \dots \ell_k) &= & \ell_1 \cdot \partial_3 (q_1^e \cdot x_4 \ell_2 \ell_3 \cdot \ell_k) \\
& \implies E \cdot (x_4 \cdot \partial_3 \cdot (q_1^e \cdot \ell_1 \ell_2 \dots \ell_k)) &= & E \cdot (\ell_1 \cdot \partial_3 (q_1^e \cdot x_4 \ell_2 \ell_3 \cdot \ell_k)) \\
& \implies (x_4 \cdot \partial_3) (D \cdot (q_1^e \cdot \ell_1 \ell_2 \dots \ell_k)) &= & (\ell_1 \cdot \partial_3) (D \cdot (q_1^e \cdot x_4 \ell_2 \ell_3 \cdot \ell_k)) \\
& \implies \sum_{i \in [s]} ((x_4 \cdot \partial_3) \cdot (q_i^e p_i) - (\ell_1 \cdot \partial_3) \cdot (q_i^e \hat{p}_i)) &= & 0 \\
& \implies \forall i \in [s] : ((x_4 \cdot \partial_3) \cdot (q_i^e p_i) - (\ell_1 \cdot \partial_3) \cdot (q_i^e \hat{p}_i)) &= & 0 \\
& \text{(as } (x_4 \cdot \partial_3) \cdot (q_i^e p_i) \in V_i \text{ and } (\ell_1 \cdot \partial_3) \cdot (q_i^e \hat{p}_i) \in V_i \text{ and the } V_i\text{'s form a direct sum)} \\
& \implies \forall i \in [s] : \ell_1 \mid (x_4 \cdot \partial_3) \cdot (q_i^e p_i) \\
& \implies \forall i \in [s] : \ell_1 \mid (\partial_3) \cdot (q_i^e p_i) \\
& \implies \forall i \in [s] : \ell_1 \mid p_i.
\end{aligned}$$

Similarly, we can show for all $j \in [k]$ and $i \in [s]$ that $\ell_j \mid p_i$. The ℓ_j 's are coprime and hence $\forall i \in [s] : (\prod_{j \in [k]} \ell_j) \mid p_i$. But the p_i 's are of degree k and hence we must have $p_i = c_{1i} \prod_{j \in [k]} \ell_j$ for some $c_{1i} \in \mathbb{R}$. Moreover it also follows that the c_{1i} 's are in fact independent of the choice of f and so we must have that $D \cdot (q_1^e \cdot f) = (c_{11} q_1^e + c_{12} q_2^e + \dots + c_{1s} q_s^e) \cdot f$ for all polynomials $f \in \mathbb{R}[\mathbf{x}]^{=k}$. \square

Proof of Lemma C.5. Assume without loss of generality that $i = 1$. Now, quadratic forms correspond to symmetric matrices which can be diagonalized over \mathbb{R} by orthogonal matrices. So Suppose that $q_1(\mathbf{y}) = c_1 \ell_1(\mathbf{y})^2 + c_2 \ell_2(\mathbf{y})^2 + \dots + c_n \ell_n(\mathbf{y})^2$, where the ℓ_j 's are pairwise orthogonal linear forms. By making a suitable orthonormal change of variables we can assume without loss of generality that

$$q_1(\mathbf{y}) = c_1 y_1^2 + c_2 y_2^2 + \dots + c_n y_n^2.$$

Consider the operator $B_1 \stackrel{\text{def}}{=} (c_1 y_1 \partial_2 - c_2 y_2 \partial_1)$ and let us apply it to polynomials of the form $q_j(\mathbf{y})^e \cdot x_3^k$. We have

$$B_1 \cdot (q_j^e \cdot x_3^k) = e \cdot q_j^{e-1} \cdot x_3^k \cdot (c_1 \cdot y_1 \cdot (\partial_2 q_j) - c_2 \cdot y_2 \cdot (\partial_1 q_j))$$

Thus $B_1 \cdot (q_1^e \cdot x_3^k) = 0$ and $B_1 \cdot (q_j^e \cdot x_3^k) \neq 0$ unless $y_1 \mid (\partial_1 q_j)$ and $y_2 \mid (\partial_2 q_j)$. For $j \neq 1$, with probability 1, this latter condition does not hold when q_j is chosen randomly and independent of q_1 and so $B_1 \cdot (q_j^e \cdot x_3^k) \neq 0$. \square

D Learning Arithmetic Circuits in the Presence of Noise

In this section, we will consider the problem of learning arithmetic circuits in the presence of noise, and follow the sketch in Section 1.1. Our algorithm is given below as Algorithm 6 and it gets the guarantees given by Theorem 14. Throughout this section, we shall assume that all vector spaces are inner product spaces, and we will work with appropriate orthonormal bases for the vector spaces.

Algorithm 6 Reconstructing the children of addition gates in the presence of noise.

Input: $(\tilde{f}(\mathbf{x}), s, d_U, d_V)$, where $\tilde{f}(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^{=d}$ is a polynomial, and s, d_U, d_V are positive integers.

Assumptions: $f(\mathbf{x}) = T_1(\mathbf{x}) + T_2(\mathbf{x}) + \dots + T_s(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]^{=d}$ is a polynomial such that each $T_i(\mathbf{x})$ belongs to a circuit class \mathcal{C} that admits operators \mathcal{L}, \mathcal{B} as follows: There are vector spaces W_1, W_2 , and a collection \mathcal{L} of linear maps $L : \mathbb{R}[\mathbf{x}]^{=d} \rightarrow W_1$, and a collection \mathcal{B} of linear maps $B : W_1 \rightarrow W_2$ such that:

- Let $U = \langle \mathcal{L} \cdot f \rangle$, and $U_i = \langle \mathcal{L} \cdot T_i \rangle$ for each $i \in [s]$. Then, $U = U_1 \oplus \dots \oplus U_s$.
- Let $V = \langle \mathcal{B} \cdot U \rangle = \langle \mathcal{B} \cdot \mathcal{L} \cdot f \rangle$, and $V_i = \langle \mathcal{B} \cdot U_i \rangle = \langle \mathcal{B} \cdot \mathcal{L} \cdot T_i \rangle$ for each $i \in [s]$. Then, $V = V_1 \oplus \dots \oplus V_s$.
- The decomposition of (U, V) under \mathcal{B} is strongly unique i.e. $\dim(\text{Adj}_{U,V}(\mathcal{B})) = s$.

The given integer inputs d_U and d_V are the dimensions of U and V respectively. The polynomial $\tilde{f}(\mathbf{x})$ is such that $\tilde{f}(\mathbf{x}) = f(\mathbf{x}) + \eta(\mathbf{x})$, with $\|\eta\| \leq \epsilon$.

Output: $\tilde{T}_1, \tilde{T}_2, \dots, \tilde{T}_s \in \mathbb{R}[\mathbf{x}]^{=d}$ such that $\|T_i - \tilde{T}_i\|$ is "small" for each $i \in [s]$ (upto reordering).

- 1: Compute $\tilde{U} \subseteq W_1$ spanned by top d_U left-singular vectors of \tilde{M} , the matrix with columns $(L \cdot \tilde{f})_{L \in \mathcal{L}}$.
 - 2: Compute $\tilde{V} \subseteq W_2$ spanned by top d_V left-singular vectors of \tilde{N} , the matrix with columns $(B \cdot L \cdot \tilde{f})_{B \in \mathcal{B}, L \in \mathcal{L}}$.
 - 3: Run RVSD algorithm on $(W_1, W_2, s, \tilde{U}, \tilde{V}, \mathcal{B})$ ⁴⁰; let the output be $\tilde{U} = (\tilde{U}_1, \dots, \tilde{U}_s)$, where $\tilde{U} = \tilde{U}_1 \oplus \dots \oplus \tilde{U}_s$.
 - 4: Let $\hat{L} : \mathbb{R}[\mathbf{x}]^{=d} \rightarrow W_1^t$ be as in Definition 2.1, where $|\mathcal{L}| = t$.
For each $i \in [s]$, let $\tilde{P}_i : W_1 \rightarrow W_1$ be the map which is identity on \tilde{U}_i , zero on each \tilde{U}_j for $j \neq i$, and zero on \tilde{U}^\perp .
For each $i \in [s]$, let $\text{Id}_t \otimes \tilde{P}_i : W_1^t \rightarrow W_1^t$ be the map given by $\text{Id}_t \otimes \tilde{P}_i \cdot (\mathbf{w}_1, \dots, \mathbf{w}_t) = (\tilde{P}_i \cdot \mathbf{w}_1, \dots, \tilde{P}_i \cdot \mathbf{w}_t)$.
For each $i \in [s]$, compute $\tilde{T}_i = \hat{L}^\dagger \cdot (\text{Id}_t \otimes \tilde{P}_i) \cdot \hat{L} \cdot \tilde{f}$.
 - 5: Output $\tilde{T}_1, \tilde{T}_2, \dots, \tilde{T}_s$.
-

Theorem 14. Let $f(\mathbf{x}) = T_1(\mathbf{x}) + T_2(\mathbf{x}) + \dots + T_s(\mathbf{x})$ with $\|f\| = 1$ be a polynomial such that each $T_i \in \mathbb{R}[\mathbf{x}]^{=d}$ belongs to a circuit class \mathcal{C} that admits operators \mathcal{L}, \mathcal{B} as follows.

Let W_1, W_2 be vector spaces, let \mathcal{L} be a collection of linear maps $L : \mathbb{R}[\mathbf{x}]^{=d} \rightarrow W_1$, and let \mathcal{B} be a collection of linear maps $B : W_1 \rightarrow W_2$, such that:

- Let $U = \langle \mathcal{L} \cdot f \rangle$, and $U_i = \langle \mathcal{L} \cdot T_i \rangle$ for each $i \in [s]$. Then, $U = U_1 \oplus \dots \oplus U_s$.
- Let $V = \langle \mathcal{B} \cdot U \rangle = \langle \mathcal{B} \cdot \mathcal{L} \cdot f \rangle$, and $V_i = \langle \mathcal{B} \cdot U_i \rangle = \langle \mathcal{B} \cdot \mathcal{L} \cdot T_i \rangle$ for each $i \in [s]$. Then, $V = V_1 \oplus \dots \oplus V_s$.
- The decomposition of (U, V) under \mathcal{B} is strongly unique i.e. $\dim(\text{Adj}_{U,V}(\mathcal{B})) = s$.

Consider the following:

⁴⁰See Remark 11 for the parameter $\tau \in (0, 1)$

- Let $\mathbf{U} = (U_1, \dots, U_s)$, $\mathbf{V} = (V_1, \dots, V_s)$. Let $d_U = \dim(U)$, $d_V = \dim(V)$, and let $d^* = \max_{i \in [s]} \dim(U_i)$, $d_* = \min_{i \in [s]} \dim(U_i)$.
- Let M and N be matrices with columns $(L \cdot f)_{L \in \mathcal{L}}$ and $(B \cdot L \cdot f)_{B \in \mathcal{B}, L \in \mathcal{L}}$ respectively, and let $\sigma_{M,N} = \min \{ \sigma_{d_U}(M), \sigma_{d_V}(N) \}$.
- Let \mathfrak{A} be the adjoint algebra map corresponding to (U, V, \mathcal{B}) (see Definition 4.2; consider the restriction of maps in \mathcal{B} to $\text{Lin}(U, V)$).
- Let the collections \mathcal{L}, \mathcal{B} be normalized such that $\|\mathcal{L}\|_2 = 1$ and $\|\mathcal{B}\|_2 = 1$. Further, let \hat{L} (see Definition 2.1) be injective, and let $\kappa(\mathcal{L})$ be the condition number $\kappa(\hat{L})$.
- Let $\delta > 0$ be arbitrary, and let $\theta = 10^6 \cdot \sqrt{\frac{d^{*3}}{d_*}} \cdot \frac{s^{5/2}}{\delta} \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \kappa(\mathbf{U})^5 \cdot \kappa(\mathcal{L}) \cdot \frac{1}{\sigma_{M,N}} \cdot \frac{1}{\sigma_{-(s+1)}(\mathfrak{A})} \cdot \epsilon < 1$.

Suppose that $\tilde{f}(\mathbf{x}) = f(\mathbf{x}) + \eta(\mathbf{x})$ such that $\|\eta\| \leq \epsilon$. Then, Algorithm 6, on input (\tilde{f}, s, d_U, d_V) , runs in time $\text{poly}(n^d, \dim W_1, \dim W_2)$, and outputs $\tilde{T}_1, \tilde{T}_2, \dots, \tilde{T}_s$ such that with probability at least $1 - \delta$, it holds (upto reordering) that for each $i \in [s]$,

$$\|T_i - \tilde{T}_i\| \leq \theta.$$

Remark 11. We note that it is possible to iterate over the parameter $\tau \in (0, 1)$ (which is the input to the RVSD algorithm), and also the parameters d_U, d_V, s , assuming that there is a way to check the validity of the polynomials $\tilde{T}_1, \dots, \tilde{T}_s$ obtained (which is usually the case in applications). The reader is referred to Remark 3, Remark 5, and Remark 8 for more details.

In the remainder of this section, we shall prove Theorem 14. It is not hard to see that the runtime is $\text{poly}(n^d, \dim W_1, \dim W_2)$: we work with orthonormal bases for all the vector spaces throughout, and also assume (without loss of generality) that all operators in \mathcal{L} and \mathcal{B} are linearly independent. We will omit the details for this.

Throughout, we will be following the notation defined in in the statement of Theorem 14. Also, let the operators \hat{L}, \hat{B} be as defined in Definition 2.1.

D.1 Applying Robust Vector Space Decomposition

As defined in Algorithm 6, let \tilde{M} and \tilde{N} be the matrices whose columns are $(L \cdot \tilde{f})_{L \in \mathcal{L}}$ and $(B \cdot L \cdot \tilde{f})_{B \in \mathcal{B}, L \in \mathcal{L}}$ respectively. Let \tilde{U} (resp. \tilde{V}) be the vector space spanned by the top d_U (resp. d_V) left singular vectors of \tilde{M} (resp. \tilde{N}).

Lemma D.1.

$$\begin{aligned} \text{dist}(U, \tilde{U}) &\leq \frac{2\epsilon}{\sigma_{d_U}(M)}. \\ \text{dist}(V, \tilde{V}) &\leq \frac{2\epsilon}{\sigma_{d_V}(N)}. \end{aligned}$$

Proof. Observe that U (resp. V) is the column space of the matrix M (resp. N). Then, by Corollary A.1,

$$\text{dist}(U, \tilde{U}) \leq \frac{2 \left\| M - \tilde{M} \right\|_2}{\sigma_{d_U}(M)}, \quad \text{dist}(V, \tilde{V}) \leq \frac{2 \left\| N - \tilde{N} \right\|_2}{\sigma_{d_V}(N)}.$$

Further, we have

$$\|M - \tilde{M}\|_F^2 = \sum_{L \in \mathcal{L}} \|L \cdot (f - \tilde{f})\|^2 = \|\hat{L} \cdot \eta\|^2 \leq \|\mathcal{L}\|_2^2 \cdot \|\eta\|^2 \leq \|\mathcal{L}\|_2^2 \cdot \epsilon^2 = \epsilon^2.$$

Similarly,

$$\|N - \tilde{N}\|_F^2 = \sum_{B \in \mathcal{B}} \sum_{L \in \mathcal{L}} \|B \cdot L \cdot \eta\|^2 \leq \sum_{L \in \mathcal{L}} \|\mathcal{B}\|_2^2 \cdot \|L \cdot \eta\|^2 \leq \|\mathcal{B}\|_2^2 \cdot \|\mathcal{L}\|_2^2 \cdot \epsilon^2 = \epsilon^2. \quad \square$$

Now, suppose that the output of the RVSD algorithm is $\tilde{\mathbf{U}} = (\tilde{U}_1, \dots, \tilde{U}_s)$. Then, Lemma D.1, along with Corollary 4.1 gives the following: For any $\delta > 0$, with probability at least $1 - \delta$, we have (upto reordering) that for each $i \in [s]$,

$$\text{dist}(U_i, \tilde{U}_i) \leq 15000 \cdot \sqrt{\frac{d^{*3}}{d_*}} \cdot \frac{s^2}{\delta} \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \kappa(\mathbf{U})^3 \cdot \frac{1}{\sigma_{M,N}} \cdot \frac{1}{\sigma_{-(s+1)}(\mathfrak{A})} \cdot \epsilon \stackrel{\text{def}}{=} \gamma.$$

Further, assuming that $2\gamma\sqrt{s} \cdot \kappa(\mathbf{U}) < 1$, we know $\tilde{\mathbf{U}} = \tilde{U}_1 \oplus \dots \oplus \tilde{U}_s$; note that this assumption follows from the assumption that $\theta < 1$ in the theorem statement.

D.2 Recovering the Polynomials

It remains to analyze the final step of the algorithm. We shall assume that $\gamma\sqrt{s} \cdot \kappa(\mathbf{U}) < 1/4$ (which follows from the assumption that $\theta < 1$ in the theorem statement).

For each $i \in [s]$, let $P_i : W_1 \rightarrow W_1$ (resp. \tilde{P}_i) be the map which is identity on U_i (resp. \tilde{U}_i), zero on each U_j (resp. \tilde{U}_j) for $j \neq i$, and zero on U^\perp (resp. \tilde{U}^\perp); note that these exist since $U = U_1 \oplus \dots \oplus U_s$ and $\tilde{U} = \tilde{U}_1 \oplus \dots \oplus \tilde{U}_s$.

Suppose that $|\mathcal{L}| = t$ (that is, \mathcal{L} has t operators). For each $i \in [s]$, let $\text{Id}_t \otimes P_i : W_1^t \rightarrow W_1^t$ (resp. $\text{Id}_t \otimes \tilde{P}_i$) be the map given by $\text{Id}_t \otimes P_i \cdot (\mathbf{w}_1, \dots, \mathbf{w}_t) = (P_i \cdot \mathbf{w}_1, \dots, P_i \cdot \mathbf{w}_t)$ (resp. $\text{Id}_t \otimes \tilde{P}_i \cdot (\mathbf{w}_1, \dots, \mathbf{w}_t) = (\tilde{P}_i \cdot \mathbf{w}_1, \dots, \tilde{P}_i \cdot \mathbf{w}_t)$).

Lemma D.2. *For all $i \in [s]$, it holds that*

$$\|P_i\|_2 \leq \kappa(\mathbf{U}), \quad \|P_i - \tilde{P}_i\|_2 \leq 14\gamma\sqrt{s} \cdot \kappa(\mathbf{U})^2.$$

Proof. Without loss of generality, we assume $W = \mathbb{R}^n$ with the usual inner product. For each $i \in [s]$, let $d_i = \dim(U_i)$; by the canonical decomposition (Theorem 6), we find an orthonormal basis $\mathbf{u}_{i,1}, \dots, \mathbf{u}_{i,d_i} \in \mathbb{R}^n$ of U_i and $\tilde{\mathbf{u}}_{i,1}, \dots, \tilde{\mathbf{u}}_{i,d_i} \in \mathbb{R}^n$ of \tilde{U}_i . Let $M_U \in \mathbb{R}^{n \times d_U}$ (resp. $M_{\tilde{U}} \in \mathbb{R}^{n \times d_{\tilde{U}}}$) be the \mathbf{U} (resp. $\tilde{\mathbf{U}}$)-associated matrix with columns $(\mathbf{u}_{i,j})_{i \in [s], j \in [d_i]}$ (resp. $(\tilde{\mathbf{u}}_{i,j})_{i \in [s], j \in [d_i]}$). Then, by the properties of the canonical decomposition, we have $\|M_U - M_{\tilde{U}}\|_2 \leq 2\gamma\sqrt{s}$ (note that this is essentially the same as Claim E.1).

Now, let $\Lambda_i \in \mathbb{R}^{d_U \times d_U}$ be a diagonal matrix defined as follows: let the d_U diagonal elements be split into s groups, of sizes d_1, \dots, d_s respectively; define Λ_i to have all ones in the i^{th} group, and zero otherwise. Then, we can write $P_i = M_U \cdot \Lambda_i \cdot M_U^\dagger$ and $\tilde{P}_i = M_{\tilde{U}} \cdot \Lambda_i \cdot M_{\tilde{U}}^\dagger$. This gives us

1.

$$\|P_i\|_2 \leq \kappa(M_U) \cdot \|\Lambda_i\|_2 = \kappa(\mathbf{U}) \cdot 1.$$

2. We assumed that $2\gamma\sqrt{s} \leq \frac{1}{2\kappa(\mathbf{U})} \leq \frac{\sigma_d(M_U)}{2} < 1$. Then, $\|M_{\tilde{U}}\|_2 \leq \|M_U\|_2 + 2\gamma\sqrt{s} \leq 2\|M_U\|_2$. Further, by Corollary A.2, we get

$$\begin{aligned}
\|P_i - \tilde{P}_i\|_2 &\leq \left\| (M_U - M_{\tilde{U}}) \cdot \Lambda_i \cdot M_U^\dagger \right\|_2 + \left\| M_{\tilde{U}} \cdot \Lambda_i \cdot (M_U^\dagger - M_{\tilde{U}}^\dagger) \right\|_2 \\
&\leq \|M_U - M_{\tilde{U}}\|_2 \cdot \|\Lambda_i\|_2 \cdot \|M_U^\dagger\|_2 + \|M_{\tilde{U}}\|_2 \cdot \|\Lambda_i\|_2 \cdot \|M_U^\dagger - M_{\tilde{U}}^\dagger\|_2 \\
&\leq 2\gamma\sqrt{s} \cdot 1 \cdot \frac{1}{\sigma_d(M_U)} + 2\|M_U\|_2 \cdot 1 \cdot \frac{3 \cdot 2\gamma\sqrt{s}}{\sigma_d(M_U)^2} \\
&\leq \frac{2\gamma\sqrt{s}}{\sigma_d(M_U)} + \frac{12\gamma\sqrt{s} \cdot \kappa(\mathbf{U})}{\sigma_d(M_U)} \leq 14\gamma\sqrt{s} \cdot \kappa(\mathbf{U})^2. \quad \square
\end{aligned}$$

Proof of Theorem 14. Fix any $i \in [s]$. Observe that for each $L \in \mathcal{L}$, we have $L \cdot T_i = P_i \cdot L \cdot T_i = P_i \cdot (L \cdot T_1 + \dots + L \cdot T_s) = P_i \cdot L \cdot f$. Hence, $\hat{L} \cdot T_i = (\text{Id}_t \otimes P_i) \cdot \hat{L} \cdot f$, and since \hat{L} is injective,

$$T_i = \hat{L}^\dagger \cdot (\text{Id}_t \otimes P_i) \cdot \hat{L} \cdot f.$$

This gives us

$$\begin{aligned}
\|T_i - \tilde{T}_i\|_2 &= \left\| \hat{L}^\dagger \cdot (\text{Id}_t \otimes P_i) \cdot \hat{L} \cdot f - \hat{L}^\dagger \cdot (\text{Id}_t \otimes \tilde{P}_i) \cdot \hat{L} \cdot \tilde{f} \right\|_2 \\
&\leq \left\| \hat{L}^\dagger \right\|_2 \cdot \left\| (\text{Id}_t \otimes P_i) \cdot \hat{L} \cdot f - (\text{Id}_t \otimes \tilde{P}_i) \cdot \hat{L} \cdot \tilde{f} \right\|_2 \\
&\leq \left\| \hat{L}^\dagger \right\|_2 \cdot \left(\left\| (\text{Id}_t \otimes P_i) \cdot \hat{L} \cdot (f - \tilde{f}) \right\|_2 + \left\| (\text{Id}_t \otimes P_i - \text{Id}_t \otimes \tilde{P}_i) \cdot \hat{L} \cdot \tilde{f} \right\|_2 \right) \\
&\leq \kappa(\mathcal{L}) \cdot \left(\left\| \text{Id}_t \otimes P_i \right\|_2 \cdot \|f - \tilde{f}\|_2 + \left\| \text{Id}_t \otimes P_i - \text{Id}_t \otimes \tilde{P}_i \right\|_2 \cdot \|\tilde{f}\|_2 \right) \\
&= \kappa(\mathcal{L}) \cdot \left(\|P_i\|_2 \cdot \|f - \tilde{f}\|_2 + \|P_i - \tilde{P}_i\|_2 \cdot \|\tilde{f}\|_2 \right) \\
&\leq \kappa(\mathcal{L}) \cdot \left(\kappa(\mathbf{U}) \cdot \epsilon + 14\gamma\sqrt{s} \cdot \kappa(\mathbf{U})^2 \cdot 2 \right) \\
&\leq 30 \cdot \kappa(\mathcal{L}) \cdot \gamma\sqrt{s} \cdot \kappa(\mathbf{U})^2.
\end{aligned}$$

Here we used Lemma D.2, $\|\tilde{f}\| \leq \|f\| + \|\eta\| \leq 1 + \epsilon \leq 2$, and $\gamma > \epsilon$. Plugging in the value of γ , we get the desired result. \square

E Analysis of the RRSB Algorithm

In this section, we will analyze our algorithm for the Robust Recovery from Scaling Maps problem, and prove Theorem 4.

In the following subsections, first we fix some notation, and give a simple proof of Lemma 3.1. Then, we will show some general properties about random matrices in the space of scaling maps. We analyze the algorithm in Section E.5. Finally, we prove the direct sum property in Section E.6.

E.1 Notation

Let $\mathbf{U} = (U_1, \dots, U_s)$ be an independent s -tuple of subspaces of W , and $U = U_1 \oplus U_2 \oplus \dots \oplus U_s \subseteq W$. Let $\dim(W) = n$, $\dim(U_i) = d_i$, and $\dim(U) = d = d_1 + \dots + d_s$. Let

$d^* = \max_{i \in [s]} \dim(U_i)$ and $d_* = \min_{i \in [s]} \dim(U_i)$. Since the algorithm is invariant to an orthogonal basis change, for the sake of the analysis, we can assume $W = \mathbb{R}^n$ with the canonical inner product.

Let $M \in \mathbb{R}^{n \times d}$ be a \mathbf{U} -associated matrix whose first d_1 columns form an orthonormal basis of U_1 , the next d_2 columns form an orthonormal basis of U_2 , and so on. In particular, we have $\kappa(\mathbf{U}) = \kappa(M)$.

For $\lambda = (\lambda_1, \dots, \lambda_s)^\top \in \mathbb{R}^s$, define $\Lambda(\lambda) = \text{diag}(\underbrace{\lambda_1, \dots, \lambda_1}_{d_1 \text{ times}}, \dots, \underbrace{\lambda_s, \dots, \lambda_s}_{d_s \text{ times}}) \in \mathbb{R}^{d \times d}$.

For each $i \in [s]$, let $P_i = M \cdot \Lambda(\mathbf{e}_i) \cdot M^\dagger \in \mathbb{R}^{n \times n}$, where $\mathbf{e}_i \in \mathbb{R}^s$ is the vector whose i^{th} coordinate is 1 and all other coordinates are 0. That is, the matrix P_i corresponds to the linear map which is identity on U_i , and zero on each U_j for $j \neq i$, and zero on U^\perp . Then, the matrices (P_1, \dots, P_s) form a basis of the space $S = S(\mathbf{U}) \subseteq \mathbb{R}^{n \times n}$ of scaling maps (when extended to $\mathbb{R}^{n \times n}$ appropriately; see Definition 3.1 and Definition 3.2).

We define $\text{vec} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n^2}$ to be the linear operator that maps $n \times n$ matrices to their corresponding flattened out matrix in \mathbb{R}^{n^2} . Let $\hat{M} \in \mathbb{R}^{n^2 \times s}$ be a matrix whose columns are $\text{vec}(P_1), \dots, \text{vec}(P_s)$ (see Definition 3.2).

The algorithm has access to a subspace $\tilde{U} \subseteq \mathbb{R}^n$, and $\tilde{S} \subseteq \mathbb{R}^{n \times n}$ such that $\text{dist}(S, \tilde{S}) \leq \epsilon$, and $\text{dist}(U, \tilde{U}) < 1$. Assuming $\epsilon < 1$, we know that $\dim(\tilde{U}) = d$ and $\dim(\tilde{S}) = s$.

E.2 Condition Number Relations

Now that we have established some notation, we begin by giving a simple proof of Lemma 3.1

Recall that $\hat{M} \in \mathbb{R}^{n^2 \times s}$ is the matrix whose columns are $\text{vec}(P_1), \dots, \text{vec}(P_s)$, where for each $i \in [s]$, $P_i = M \cdot \Lambda(\mathbf{e}_i) \cdot M^\dagger \in \mathbb{R}^{d \times d}$. Then, we have for each $\lambda \in \mathbb{R}^s$,

$$\|\hat{M}\lambda\|_2 = \left\| M \cdot \Lambda(\lambda) \cdot M^\dagger \right\|_F \leq \kappa(M) \cdot \|\Lambda(\lambda)\|_F \leq \kappa(M) \cdot \sqrt{\max_{i \in [s]} d_i} \cdot \|\lambda\|_2,$$

$$\|\hat{M}\lambda\|_2 = \left\| M \cdot \Lambda(\lambda) \cdot M^\dagger \right\|_F \geq \sqrt{\sum_{i \in [s]} d_i \lambda_i^2} \geq \sqrt{\min_{i \in [s]} d_i} \cdot \|\lambda\|_2.$$

The inequality in the second line above follows from Proposition A.1.

In particular, we get that $\|\hat{M}\|_2 \leq \kappa(M) \cdot \sqrt{\max_{i \in [s]} d_i}$, and $\kappa(\hat{M}) \leq \kappa(M) \cdot \sqrt{\frac{\max_{i \in [s]} d_i}{\min_{j \in [s]} d_j}}$. \square

E.3 Canonical Decomposition and Random Sampling

Now, we start analyzing the algorithm. Observe that the algorithm, in the first step, choose a random matrix $\tilde{A} \in \tilde{S}$. Using the canonical decomposition, we show that such a random matrix can be coupled with a random matrix $A \in S$, in a natural way.

Let $k = s - \dim(S \cap \tilde{S})$. Then, by Theorem 6, we can find an orthonormal basis $E_1, \dots, E_s, F_1, \dots, F_k, H_1, \dots, H_{n^2 - (r+k)}$ of $\mathbb{R}^{n \times n}$, and $\frac{\pi}{2} \geq \theta_1 \geq \dots \geq \theta_k > 0$ such that:

1. E_1, \dots, E_s form an orthonormal basis of S .

2. For $i \in [k]$, let $\tilde{E}_i = \cos(\theta_i) \cdot E_i + \sin(\theta_i) \cdot F_i$, and for $i \in [s] \setminus [k]$, let $\tilde{E}_i = E_i$. Then, $\tilde{E}_1, \dots, \tilde{E}_s$ form an orthonormal basis of \tilde{S} .
3. $\sin(\theta_1) = \text{dist}(S, \tilde{S}) \leq \epsilon$.

Under the above decomposition, we can couple random elements of S and \tilde{S} , as follows:

Definition E.1. (Coupled Random Matrices) Let $\alpha = (\alpha_1, \dots, \alpha_s) \in \mathbb{R}^s$, with each $\alpha_i \in \mathcal{N}(0, 1)$ chosen independently. Let $\beta = (\beta_1, \dots, \beta_s) \in \mathbb{R}^s$, with $\beta_i = \alpha_i / \|\alpha\|_2$ for each $i \in [s]$. Define $A = \sum_{i=1}^s \beta_i E_i \in S$ and $\tilde{A} = \sum_{i=1}^s \beta_i \tilde{E}_i \in \tilde{S}$.

Note that these satisfy $\|A\|_F = \|\tilde{A}\|_F = 1$ almost surely.

Observe that the distribution of the matrix \tilde{A} is independent of the choice of the orthonormal basis for \tilde{S} , and so the algorithm, in the first step, samples \tilde{A} from the same distribution.

Lemma E.1.

$$\|A - \tilde{A}\|_F \leq 2\epsilon.$$

Proof. This follows from Theorem 6 after observing that the relevant norm on $\mathbb{R}^{n \times n}$ is the Frobenius norm. \square

E.4 Bounding the Spectral Gap

As we will see later, the correctness of the algorithm depends on the gap between the eigenvalues of \hat{A} being large, and in this subsection we shall establish this fact.

Let $E \in \mathbb{R}^{n^2 \times s}$ be the matrix whose columns are given by $\text{vec}(E_1), \dots, \text{vec}(E_s)$. Recall that $\hat{M} \in \mathbb{R}^{n^2 \times s}$ is a matrix whose columns are $\text{vec}(P_1), \dots, \text{vec}(P_s)$.

Let β be as in Definition E.1, and let $\lambda = (\lambda_1, \dots, \lambda_s)^\top$ be such that $E\beta = \text{vec}(A) = \text{vec}(M \cdot \Lambda(\lambda) \cdot M^\top) = \hat{M}\lambda$, or equivalently, $\lambda = \hat{M}^\top E\beta$. Define

$$\text{Gap}(\lambda) \stackrel{\text{def}}{=} \min \left\{ \min_{i,j \in [s], i \neq j} |\lambda_i - \lambda_j|, \min_{i \in [s]} |\lambda_i| \right\}.$$

Lemma E.2. For any $\delta > 0$, it holds that

$$\Pr \left[\text{Gap}(\lambda) \leq \frac{\delta}{6 \cdot \|\hat{M}\|_2 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}}} \right] \leq \delta.$$

Proof. For each $i \in [s]$, let $\mathbf{e}_i \in \mathbb{R}^s$ be the vector whose i^{th} coordinate is 1 and all other coordinates are 0. Then, for each $i, j \in [s], i \neq j$, it holds that

$$|\lambda_i| = \mathbf{e}_i^\top \hat{M}^\top E\beta, \quad |\lambda_i - \lambda_j| = (\mathbf{e}_i - \mathbf{e}_j)^\top \hat{M}^\top E\beta.$$

Now, observe that

1. For each $i, j \in [s], i \neq j$,

$$\|\mathbf{e}_i^\top \hat{M}^\top E\|_2 \geq \sigma_s(\hat{M}^\top E), \quad \|(\mathbf{e}_i - \mathbf{e}_j)^\top \hat{M}^\top E\|_2 \geq \sqrt{2} \sigma_s(\hat{M}^\top E) \geq \sigma_s(\hat{M}^\top E).$$

2. Observe that $E, \hat{M} \in \mathbb{R}^{n^2 \times s}$ are both rank s matrices, each with column space S . Hence, by Lemma A.1, we have

$$\sigma_s(\hat{M}^\dagger E) = \frac{1}{\|(\hat{M}^\dagger E)^{-1}\|_2} = \frac{1}{\|E^\top \hat{M}\|_2} \geq \frac{1}{\|E^\top\|_2 \cdot \|\hat{M}\|_2} = \frac{1}{\|\hat{M}\|_2}.$$

Then, the result follows from Lemma A.10, and a union bound over the $s(s+1)/2 \leq s^2$ inequalities. \square

From now on, we shall consider fixed $\delta > 0$, and fixed $\lambda_1, \dots, \lambda_s$, and assume that $\text{Gap}(\lambda) > \frac{\delta}{6 \cdot \|\hat{M}\|_2 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}}}$ (by Lemma E.2, this occurs with probability at least $1 - \delta$).

E.5 Analysis of the Algorithm

E.5.1 Action on The Space of Scaling Maps

We consider matrices $\hat{A}, \tilde{A} \in \mathbb{R}^{n^2 \times n^2}$ as follows: these are the matrices corresponding to the linear maps defined in Definition 3.3 and Definition 3.4 (and extending the maps to be zero on S^\perp, \tilde{S}^\perp respectively), under the usual flattening of matrices by the map vec . Formally, we have that for each $B \in \mathbb{R}^{n \times n}$,

$$\begin{aligned} \hat{A} \cdot \text{vec}(B) &= \text{vec}(\text{Proj}_S(A \cdot \text{Proj}_S(B))), \\ \tilde{A} \cdot \text{vec}(B) &= \text{vec}(\text{Proj}_{\tilde{S}}(\tilde{A} \cdot \text{Proj}_{\tilde{S}}(B))), \end{aligned}$$

where $\text{Proj}_S, \text{Proj}_{\tilde{S}} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$ are the orthogonal projections onto S, \tilde{S} respectively.

Lemma E.3.

$$\|\hat{A} - \tilde{A}\|_2 \leq 4\epsilon.$$

Proof. For each $B \in \mathbb{R}^{n \times n}$, we have

$$\begin{aligned} \|(\hat{A} - \tilde{A}) \cdot \text{vec}(B)\|_2 &= \left\| \text{Proj}_S(A \cdot \text{Proj}_S(B)) - \text{Proj}_{\tilde{S}}(\tilde{A} \cdot \text{Proj}_{\tilde{S}}(B)) \right\|_F \\ &\leq \left\| (\text{Proj}_S - \text{Proj}_{\tilde{S}})(A \cdot \text{Proj}_S(B)) \right\|_F + \left\| \text{Proj}_{\tilde{S}}((A - \tilde{A}) \cdot \text{Proj}_S(B)) \right\|_F \\ &\quad + \left\| \text{Proj}_{\tilde{S}}(\tilde{A} \cdot (\text{Proj}_S - \text{Proj}_{\tilde{S}})(B)) \right\|_F. \end{aligned}$$

Then, using that $\|\text{Proj}_S - \text{Proj}_{\tilde{S}}\|_2 = \text{dist}(S, \tilde{S}) \leq \epsilon$, and Lemma E.1, we have for $\|B\|_F = 1$ that

$$\|(\hat{A} - \tilde{A}) \cdot \text{vec}(B)\|_2 \leq \epsilon + 2\epsilon + \epsilon = 4\epsilon. \quad \square$$

E.5.2 Perturbation Bound

Let $\gamma \stackrel{\text{def}}{=} 300 \cdot \kappa(\hat{M}) \cdot \|\hat{M}\|_2^2 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \frac{\epsilon}{\delta}$ be such that $\gamma < 1$, or else Theorem 4 holds trivially. Let $\eta \stackrel{\text{def}}{=} 100 \cdot \kappa(\hat{M}) \cdot \|\hat{M}\|_2 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \frac{\epsilon}{\delta}$. Then, we have $\eta \leq \frac{\gamma}{3\|\hat{M}\|_2} < 1$, since $\|\hat{M}\|_2 \geq \|P_1\|_F \geq 1$.

Following Definition 3.3, we can write

$$\hat{A} = \hat{M} \cdot \text{diag}(\lambda_1, \dots, \lambda_s) \cdot \hat{M}^\dagger.$$

This shows that \hat{A} has exactly s distinct non-zero eigenvalues (each with multiplicity one), which are equal to $\lambda_1, \dots, \lambda_s$, and these have eigenvectors $\text{vec}(P_1), \dots, \text{vec}(P_s)$ respectively.

Lemma E.4. 1. The matrix $\hat{A} \in \mathbb{R}^{n^2 \times n^2}$ has s non-zero, distinct eigenvalues, each of which is real and occurs with multiplicity 1. The eigenvalue 0 occurs with multiplicity $n^2 - s$, and so all the eigenvalues (and hence eigenvectors) are real.

2. Let $\text{vec}(\tilde{P}_1), \dots, \text{vec}(\tilde{P}_s) \in \mathbb{R}^{n^2}$ be eigenvectors of \hat{A} corresponding to the distinct non-zero eigenvalues, with $\|\tilde{P}_i\|_F = 1$ for each $i \in [s]$. Then, (upto reordering) for each $i \in [s]$, there exists $\zeta_i \in \{-1, 1\}$ such that

$$\left\| \zeta_i \tilde{P}_i - \frac{P_i}{\|P_i\|_F} \right\|_F \leq \eta.$$

Proof. 1. Let $\theta > 0$ be any number such that $\sigma_s(\hat{M}) \leq \theta \leq \sigma_1(\hat{M})$. We can write

$$\hat{A} = \hat{M}' \cdot \text{diag}(\lambda_1, \dots, \lambda_s, \underbrace{0, \dots, 0}_{n^2 - s \text{ times}}) \cdot (\hat{M}')^{-1},$$

where $\hat{M}' \in \mathbb{R}^{n^2 \times n^2}$ is the matrix whose first s columns are the same as that of \hat{M} , and the remaining columns are an arbitrary orthogonal basis of S^\perp with each column having ℓ_2 -norm equal to $\theta > 0$. The singular values of \hat{M}' are precisely the singular values of \hat{M} , along with the value θ which occurs as a singular value $n^2 - s$ times. By the choice of θ , we have $\kappa(\hat{M}') = \kappa(\hat{M})$.

Now, since $\kappa(\hat{M}') \cdot (4\epsilon) < \text{Gap}(\lambda)/2$ (which is implied by $\eta < 1$), by Lemma E.3 and Lemma A.7 we know that the eigenvalues of \hat{A} can be written as $\tilde{\lambda}_1, \dots, \tilde{\lambda}_{n^2} \in \mathbb{C}$ such that

- For each $i \in [s]$, it holds that $\tilde{\lambda}_i \in \mathbb{R}$, and that $|\tilde{\lambda}_i - \lambda_i| < \text{Gap}(\lambda)/2$. In particular, the eigenvalues $(\tilde{\lambda}_i)_{i \in [s]}$ are distinct.
- For each $i \in [n^2] \setminus [s]$, it holds that $|\tilde{\lambda}_i| < \text{Gap}(\lambda)/2$. In fact, since \hat{A} is defined to be zero on \tilde{S}^\perp , we know that the eigenvalue 0 occurs with multiplicity at least $n^2 - s$. Hence, each such $\tilde{\lambda}_i = 0$.

The above implies the first part of the Lemma.

2. This follows directly from Lemma A.8:

$$\left\| \zeta_i \tilde{P}_i - \frac{P_i}{\|P_i\|_F} \right\|_F \leq \frac{4\kappa(\hat{M}') \cdot 4\epsilon}{\text{Gap}(\lambda)} = \frac{16\kappa(\hat{M}) \cdot \epsilon}{\text{Gap}(\lambda)} \leq \eta. \quad \square$$

E.5.3 Recovering the Components

Fix some $i \in [s]$, and let $d_i = \dim(U_i)$.

We know that $\text{rank}(P_i) = d_i$, and so by Lemma E.4 and Weyl's inequality (Lemma A.3), the first d_i singular values of \tilde{P}_i (equal to those of $\zeta_i \tilde{P}_i$) are at least $\frac{\sigma_{d_i}(P_i)}{\|P_i\|_F} - \eta \geq \frac{1}{\|\hat{M}\|_2} - \eta$, where as the remaining ones are at most η . Here, we used the following observations

1. $\|P_i\|_F = \|\text{vec}(P_i)\|_2 \leq \|\hat{M}\|_2$.
2. $\sigma_{d_i}(P_i) \geq 1$ since the map P_i equals the identity map on the d_i -dimensional space U_i .

Hence, if $\eta < \tau \leq \frac{1}{\|\hat{M}\|_2} - \eta$, then the algorithm sets \tilde{U}_i to be the span of the left singular vectors of \tilde{P}_i (or equivalently, of $\zeta_i \tilde{P}_i$), corresponding to the top d_i singular vectors. For instance, this happens when $\eta \leq \frac{1}{3} \cdot \frac{1}{\|\hat{M}\|_2}$ and $\frac{1}{3} \cdot \frac{1}{\|\hat{M}\|_2} < \tau \leq \frac{2}{3} \cdot \frac{1}{\|\hat{M}\|_2}$. Then, by Lemma E.4 and Wedin's theorem (see Corollary A.1), we get that:

$$\text{dist}(\tilde{U}_i, U_i) \leq \frac{2 \cdot \left\| \zeta_i \tilde{P}_i - \frac{P_i}{\|P_i\|_F} \right\|_F}{\sigma_{d_i} \left(\frac{P_i}{\|P_i\|_F} \right)} \leq \frac{2\eta \cdot \|P_i\|_F}{\sigma_{d_i}(P_i)} \leq 2\eta \cdot \|\hat{M}\|_2 \leq \gamma. \quad \square$$

E.5.4 Runtime Analysis

In this section, we analyze the runtime of the algorithm. Recall that $n = \dim(W)$, $d = \dim(U) = \dim(\tilde{U})$, $s = \dim(S) = \dim(\tilde{S})$, and let $d_i = \dim(\tilde{U}_i)$. The input is given as $N = dn + sd^2$ field elements, consisting of an orthonormal basis of $\tilde{U} \subseteq W$ and an orthonormal basis of $\tilde{S} \in \text{Lin}(\tilde{U}, \tilde{U})$. The time taken by each step of the algorithm is as follows:

1. The random map $\tilde{A} \in \tilde{S}$ can be computed as a $d \times d$ matrix in time $O(sd^2)$, by taking a random linear combination of the basis elements of \tilde{S} .
2. The map $\hat{A} \in \text{Lin}(\tilde{S}, \tilde{S})$ can be computed as an $s \times s$ matrix with respect to the orthonormal basis of \tilde{S} : If the basis is $\tilde{\mathbf{s}}_1, \dots, \tilde{\mathbf{s}}_s \in \mathbb{R}^{d \times d}$, for $i, j \in [s]$, the (i, j) th entry of this matrix equals $\left\langle \tilde{A} \cdot \tilde{\mathbf{s}}_i, \tilde{\mathbf{s}}_j \right\rangle_F$, and can be computed in time $O(d^\omega)$, where ω is the matrix multiplication constant. So the total time taken in this step is $O(s^2 d^\omega)$.
3. The eigen-decomposition of \hat{A} can be computed in time $O(s^3)$.
4. Computing the eigenvectors $\tilde{P}_1, \dots, \tilde{P}_s$ as matrices in $\mathbb{R}^{d \times d}$, by taking appropriate linear combinations of the \tilde{S} basis elements takes time $O(s \cdot sd^2)$.
5. For each $i \in [s]$, the singular-value decomposition of \tilde{P}_i can be computed in time $O(d^3)$, and then computing the basis vectors of \tilde{U}_i as vectors in \mathbb{R}^n takes time $O(d \cdot dn)$. Therefore, the total time to compute the basis elements for all the \tilde{U}_i 's is $O(sd^3 + sd^2n)$.

Finally, we get that the total runtime is $O(s^3 + s^2 d^\omega + sd^3 + sd^2n) = O(N^{5/3})$, using $s \leq d \leq n$.

E.6 Direct Sum Property

We give a short proof of Proposition 3.1. It suffices to show that $\tilde{U}_1, \dots, \tilde{U}_s$ form a direct sum. The sum then equals \tilde{U} by counting dimensions; note that $\dim(U) = \dim(\tilde{U})$ and $\dim(U_i) = \dim(\tilde{U}_i)$ for each $i \in [s]$.

For each $i \in [s]$, by the canonical decomposition (Theorem 6), we can find an orthonormal basis $\mathbf{u}_{i,1}, \dots, \mathbf{u}_{i,d_i} \in \mathbb{R}^n$ of U_i and $\tilde{\mathbf{u}}_{i,1}, \dots, \tilde{\mathbf{u}}_{i,d_i} \in \mathbb{R}^n$ of \tilde{U}_i such that for each $\alpha_1, \dots, \alpha_{d_i} \in \mathbb{R}$, we have

$$\left\| \sum_{j \in [d_i]} \alpha_j \mathbf{u}_{i,j} - \sum_{j \in [d_i]} \alpha_j \tilde{\mathbf{u}}_{i,j} \right\|_2 \leq 2\gamma \cdot \sqrt{\sum_{j \in [d_i]} \alpha_j^2}.$$

Let $N \in \mathbb{R}^{n \times d}$ (resp. \tilde{N}) be the \mathbf{U} (resp. $\tilde{\mathbf{U}}$)-associated matrix with columns $(\mathbf{u}_{i,j})_{i \in [s], j \in [d_i]}$ (resp. $(\tilde{\mathbf{u}}_{i,j})_{i \in [s], j \in [d_i]}$).

Claim E.1.

$$\left\| N - \tilde{N} \right\|_2 \leq 2\gamma\sqrt{s}.$$

Proof. Let $\boldsymbol{\alpha} \in \mathbb{R}^d$ be indexed by $i \in [s], j \in [d_i]$. Then, we have

$$\begin{aligned} \left\| (\tilde{N} - N) \cdot \boldsymbol{\alpha} \right\|_2 &= \left\| \sum_{i \in [s], j \in [d_i]} \alpha_{i,j} (\tilde{\mathbf{u}}_{i,j} - \mathbf{u}_{i,j}) \right\|_2 \\ &\leq \sum_{i \in [s]} \left\| \sum_{j \in [d_i]} \alpha_{i,j} (\tilde{\mathbf{u}}_{i,j} - \mathbf{u}_{i,j}) \right\|_2 \\ &\leq \sum_{i \in [s]} 2\gamma \cdot \sqrt{\sum_{j \in [d_i]} \alpha_{i,j}^2} \\ &\leq 2\gamma \cdot \sqrt{s} \cdot \sqrt{\sum_{i \in [s], j \in [d_i]} \alpha_{i,j}^2}. \end{aligned}$$

□

Now, using Claim E.1 along with Weyl's inequality (Lemma A.3), we get

$$\sigma_d(\tilde{N}) \geq \sigma_d(N) - 2\gamma\sqrt{s} \geq \frac{1}{\kappa(\mathbf{U})} - 2\gamma\sqrt{s} > 0,$$

under the assumption $2\gamma\sqrt{s} \cdot \kappa(\mathbf{U}) < 1$. This completes the proof.

F Analysis of the RVSD Algorithm

In this section, we will analyze Algorithm 2 and prove Theorem 5.

F.1 Notation

Let W_1, W_2 be two vector spaces with $\dim(W_1) = n_1, \dim(W_2) = n_2$. Without loss of generality (by an orthogonal transformation), we can assume $W_1 = \mathbb{R}^{n_1}$ and $W_2 = \mathbb{R}^{n_2}$ with the canonical inner products on the two spaces.

Let $\mathbf{U} = (U_1, \dots, U_s)$ and $\mathbf{V} = (V_1, \dots, V_s)$ be independent s -tuples of subspaces in W_1 and W_2 respectively, and let $U = U_1 \oplus U_2 \oplus \dots \oplus U_s \subseteq \mathbb{R}^{n_1}, V = V_1 \oplus V_2 \oplus \dots \oplus V_s \subseteq \mathbb{R}^{n_2}$ be such that $\dim(U) = d_1, \dim(V) = d_2$. Let $\mathcal{B} = (B_1, \dots, B_m) \in (\mathbb{R}^{n_2 \times n_1})^m$ be an m -tuple of operators, with each B_j being the zero map on U^\perp , and such that for each $i \in [s]$, it holds that $\langle \mathcal{B} \cdot U_i \rangle \subseteq V_i$.

The algorithm has access to a subspaces $\tilde{U} \subseteq \mathbb{R}^{n_1}$ and $\tilde{V} \subseteq \mathbb{R}^{n_2}$ such that $\text{dist}(U, \tilde{U}) \leq \epsilon_1$ and $\text{dist}(V, \tilde{V}) \leq \epsilon_2$; assuming $\epsilon_1, \epsilon_2 < 1$, we have $\dim(\tilde{U}) = d_1, \dim(\tilde{V}) = d_2$. We also know $\tilde{\mathcal{B}} = (\tilde{B}_1, \dots, \tilde{B}_m)$, with each \tilde{B}_j being the zero map on \tilde{U}^\perp , such that \mathcal{B} and $\tilde{\mathcal{B}}$ are ϵ -close: Let $\hat{B} \in \mathbb{R}^{m n_2 \times n_1}$ (resp. $\hat{\tilde{B}} \in \mathbb{R}^{m n_2 \times n_1}$) be matrices formed by stacking the rows of B_1, \dots, B_m (resp. $\tilde{B}_1, \dots, \tilde{B}_m$). Then, $\|\hat{\tilde{B}} - \hat{B}\|_2 \leq \epsilon \|\hat{B}\|_2$ (see Definition 2.1).

F.2 Perturbation Bound on the Adjoint Algebra

Let $\text{Proj}_U, \text{Proj}_{\tilde{U}} \in \mathbb{R}^{n_1 \times n_1}$ and $\text{Proj}_V, \text{Proj}_{\tilde{V}} \in \mathbb{R}^{n_2 \times n_2}$ denote the orthogonal projections onto $U, \tilde{U}, V, \tilde{V}$ respectively. These also give us the orthogonal projection maps on the spaces of linear maps, for example, $\text{Proj}_{\text{Lin}(U, U)} : \mathbb{R}^{n_1 \times n_1} \rightarrow \mathbb{R}^{n_1 \times n_1}$ is given by $\text{Proj}_{\text{Lin}(U, U)}(D) = \text{Proj}_U \cdot D \cdot \text{Proj}_U$, where we think of $\text{Lin}(U, U)$ as a subspace of $\mathbb{R}^{n_1 \times n_1}$ in the natural way.

Let $\mathfrak{A} : \mathbb{R}^{n_1 \times n_1} \times \mathbb{R}^{n_2 \times n_2} \rightarrow (\mathbb{R}^{n_2 \times n_1})^m$ be the map defined as in Definition 4.2, extended to be the zero map on $\text{Lin}(U, U)^\perp \times \text{Lin}(V, V)^\perp$, given by

$$\mathfrak{A}(D, E) = (B_i \cdot (\text{Proj}_U \cdot D \cdot \text{Proj}_U) - (\text{Proj}_V \cdot E \cdot \text{Proj}_V) \cdot B_i)_{i=1}^m$$

Letting $\mathfrak{K} = \ker(\mathfrak{A})$. Then, we have that $\text{Adj} = \mathfrak{K} \cap (\text{Lin}(U, U) \times \text{Lin}(V, V)) \subseteq \mathbb{R}^{n_1 \times n_1} \times \mathbb{R}^{n_2 \times n_2}$ (see Definition 4.3). Here, we think of $\text{Lin}(U, U)$ as a subspace of $\mathbb{R}^{n_1 \times n_1}$, by extending each map to be zero on U^\perp , and similarly we think of $\text{Lin}(V, V)$ as a subspace of $\mathbb{R}^{n_2 \times n_2}$.

Note that assuming $\dim(\text{Adj}) = s$, we have that $\dim(\mathfrak{K}) = (n_1^2 - d_1^2) + (n_2^2 - d_2^2) + s$. We shall use $\sigma_{-(s+1)}(\mathfrak{A})$ to denote $(\dim(\mathfrak{K}) + 1)^{\text{th}}$ smallest singular value of \mathfrak{A} (this is also the smallest non-zero singular value). Note that this equals the value $\sigma_{-(s+1)}(\mathfrak{A})$ as defined in Theorem 5.

In a similar manner, we define $\tilde{\mathfrak{A}} : \mathbb{R}^{n_1 \times n_1} \times \mathbb{R}^{n_2 \times n_2} \rightarrow (\mathbb{R}^{n_2 \times n_1})^m$ as in Definition 4.5, extended to be the zero map on $\text{Lin}(\tilde{U}, \tilde{U})^\perp \times \text{Lin}(\tilde{V}, \tilde{V})^\perp$, given by

$$\tilde{\mathfrak{A}}(D, E) = (\tilde{B}_i \cdot (\text{Proj}_{\tilde{U}} \cdot D \cdot \text{Proj}_{\tilde{U}}) - (\text{Proj}_{\tilde{V}} \cdot E \cdot \text{Proj}_{\tilde{V}}) \cdot \tilde{B}_i)_{i=1}^m$$

Further, by Definition 4.5, letting $\tilde{\mathfrak{K}}$ denote the space spanned by the right singular vectors of $\tilde{\mathfrak{A}}$, corresponding to the $(n_1^2 - d_1^2) + (n_2^2 - d_2^2) + s$ smallest singular values, we have $\tilde{\text{Adj}} = \tilde{\mathfrak{K}} \cap (\text{Lin}(\tilde{U}, \tilde{U}) \times \text{Lin}(\tilde{V}, \tilde{V}))$, where again we think of $\text{Lin}(\tilde{U}, \tilde{U})$ and $\text{Lin}(\tilde{V}, \tilde{V})$ as subspaces of $\mathbb{R}^{n_1 \times n_1}$ and $\mathbb{R}^{n_2 \times n_2}$ respectively.

Lemma F.1.

$$\|\mathfrak{A} - \tilde{\mathfrak{A}}\|_2 \leq 2(\epsilon + \epsilon_1 + \epsilon_2) \|\mathcal{B}\|_2.$$

Proof. Let $(D, E) \in \mathbb{R}^{n_1 \times n_1} \times \mathbb{R}^{n_2 \times n_2}$ be such that $\|D\|_F^2 + \|E\|_F^2 = 1$. Then, by triangle inequality, and sub-multiplicativity of Frobenius norm (see Proposition A.1), we have that

$$\begin{aligned} \|\tilde{\mathfrak{A}}(D, E) - \mathfrak{A}(D, E)\|_2 &\leq \left\| \hat{B} - \hat{B} \right\|_2 \cdot \|D\|_F \\ &+ \left\| \hat{B} - \hat{B} \right\|_2 \cdot \|E\|_F + 2 \|\hat{B}\|_2 \cdot \|\text{Proj}_{\tilde{U}} - \text{Proj}_U\|_2 \cdot \|D\|_F + 2 \|\hat{B}\|_2 \cdot \|\text{Proj}_{\tilde{V}} - \text{Proj}_V\|_2 \cdot \|E\|_F. \end{aligned}$$

The lemma then follows. \square

Lemma F.2.

$$\text{dist}(\tilde{\mathfrak{K}}, \mathfrak{K}) \leq \frac{4(\epsilon + \epsilon_1 + \epsilon_2) \|\mathcal{B}\|_2}{\sigma_{-(s+1)}(\mathfrak{A})}.$$

Proof. This follows from Lemma F.1 and Corollary A.1. \square

Lemma F.3.

$$\text{dist}(\widetilde{\text{Adj}}, \text{Adj}) \leq \frac{4(\epsilon + \epsilon_1 + \epsilon_2) \|\mathcal{B}\|_2}{\sigma_{-(s+1)}(\mathfrak{A})} + 2\epsilon_1 + 2\epsilon_2 \leq \frac{6(\epsilon + \epsilon_1 + \epsilon_2) \|\mathcal{B}\|_2}{\sigma_{-(s+1)}(\mathfrak{A})}.$$

Proof. We have

$$\begin{aligned} \text{dist}(\widetilde{\text{Adj}}, \text{Adj}) &= \text{dist} \left(\tilde{\mathfrak{K}} \cap \left(\text{Lin}(\tilde{U}, \tilde{U}) \times \text{Lin}(\tilde{V}, \tilde{V}) \right), \mathfrak{K} \cap \left(\text{Lin}(U, U) \times \text{Lin}(V, V) \right) \right) \\ &\leq \text{dist}(\tilde{\mathfrak{K}}, \mathfrak{K}) + \text{dist} \left(\text{Lin}(\tilde{U}, \tilde{U}) \times \text{Lin}(\tilde{V}, \tilde{V}), \text{Lin}(U, U) \times \text{Lin}(V, V) \right) \\ &\leq \text{dist}(\tilde{\mathfrak{K}}, \mathfrak{K}) + \text{dist} \left(\text{Lin}(\tilde{U}, \tilde{U}), \text{Lin}(U, U) \right) + \text{dist} \left(\text{Lin}(\tilde{V}, \tilde{V}), \text{Lin}(V, V) \right). \end{aligned}$$

The first term is bounded by Lemma F.2. For the second term, we observe that for any $D \in \mathbb{R}^{n_1 \times n_1}$ with $\|D\|_F = 1$, we have

$$\left\| \text{Proj}_{\text{Lin}(\tilde{U}, \tilde{U})}(D) - \text{Proj}_{\text{Lin}(U, U)}(D) \right\|_2 = \left\| \text{Proj}_{\tilde{U}} \cdot D \cdot \text{Proj}_{\tilde{U}} - \text{Proj}_U \cdot D \cdot \text{Proj}_U \right\|_F \leq 2\epsilon_1.$$

The third term is bounded similarly.

For the last inequality, it suffices to show that $\sigma_{-(s+1)}(\mathfrak{A}) \leq \|\mathcal{B}\|_2$: For any (D, E) , with $E = 0$, $\|D\|_F = 1$, and $D \in \text{Lin}(U_1, U_2) \subseteq \mathbb{R}^{n_1 \times n_1}$, we have $(D, E) \in \mathfrak{K}^\perp$ and so

$$\sigma_{-(s+1)}(\tilde{\mathfrak{A}}) \leq \|\mathfrak{A} \cdot (D, E)\|_2 = \|\hat{B} \cdot D\|_2 \leq \|\hat{B}\|_2 = \|\mathcal{B}\|_2.$$

Note that we assumed $s \geq 2$, which is without loss of generality, as in the $s = 1$ the RVSD problem is trivial. \square

This also gives the following:

Lemma F.4.

$$\text{dist}(\widetilde{\text{Adj}}_1, \text{Adj}_1) \leq \frac{6(\epsilon + \epsilon_1 + \epsilon_2) \|\mathcal{B}\|_2}{\sigma_{-(s+1)}(\mathfrak{A})}.$$

F.3 Applying RRSB and Recovering The Component Subspaces

By Theorem 4, we get the following: Let $\hat{M} : \mathbb{R}^s \rightarrow \text{Lin}(U, U)$ be as defined in Definition 3.2, and suppose that $\tau \in (0, 1)$ satisfies that $\frac{1}{3} \cdot \frac{1}{\|\hat{M}\|_2} < \tau \leq \frac{2}{3} \cdot \frac{1}{\|\hat{M}\|_2}$. Then, for any $\delta > 0$, with probability at least $1 - \delta$, we get (upto reordering) that for each $i \in [s]$,

$$\text{dist}(U_i, \tilde{U}_i) \leq 300 \cdot \kappa(\hat{M}) \cdot \|\hat{M}\|_2^2 \cdot s^2 \sqrt{s + \ln \frac{s^2}{\delta}} \cdot \frac{1}{\delta} \cdot \frac{6(\epsilon + \epsilon_1 + \epsilon_2) \|\mathcal{B}\|_2}{\sigma_{-(s+1)}(\mathfrak{A})}.$$

The direct sum property $\tilde{U} = \tilde{U}_1 \oplus \dots \oplus \tilde{U}_s$ follows by Proposition 3.1. \square

F.4 Runtime Analysis

In this section, we analyze the runtime of the Algorithm 2. Recall that $n_1 = \dim(W_1)$, $n_2 = \dim(W_2)$, $d_1 = \dim(U) = \dim(\tilde{U})$, $d_2 = \dim(V) = \dim(\tilde{V})$. The input is given as $N = n_1 d_1 + n_2 d_2 + m d_1 d_2$ field elements, consisting of an orthonormal basis of $\tilde{U} \subseteq W_1$ and $\tilde{V} \subseteq W_2$, and a description of m -tuple $\mathcal{B} = \text{Lin}(\tilde{U}, \tilde{V})^m$ as matrices with respect to the above basis of \tilde{U} and \tilde{V} . The time taken by each step of the algorithm is as follows.

1. The adjoint algebra map can be computed as a $(m d_1 d_2) \times (d_1^2 + d_2^2)$ matrix. Setting $K = m d_1 d_2 \cdot (d_1^2 + d_2^2)$, this can be done in time $O(K)$, as each entry is simply (± 1 times) some entry of some $B_j, j \in [m]$.
2. The singular value decomposition of $\tilde{\mathfrak{A}}$ can be computed in time $O(K^3)$, from which a basis of $\widetilde{\text{Adj}}$ can be obtained in time $O(s \cdot (d_1^2 + d_2^2))$, and further a basis of $\widetilde{\text{Adj}}_1$ can be obtained in time $O(s d_1^2)$.
3. Assuming we run the second algorithm for RRSB (see Algorithm 1, Theorem 4), the last step takes time $O(s^3 + s^2 d_1^\omega + s d_1^3 + s d_1^2 n_1)$.

Finally, we get that the total runtime is $O((m d_1 d_2 \cdot (d_1^2 + d_2^2))^3 + s d_1^2 n_1) = O(N^6)$, using $s \leq d_1 \leq n_1$ and $s \leq d_2 \leq n_2$.

G Singular Values of the Adjoint Algebra Operator for Subspace Clustering

Section B discusses the problem of Subspace Clustering and culminates in theorem 10 where the robustness of the proposed algorithm is quantified. The quantity $\sigma_{-(s+1)}(\mathfrak{A})$ appears in the bounds. In this section we will compute a lower bound for this quantity.

The high level procedure for getting the bound is as follows. Let $U = U_1 \oplus \dots \oplus U_s$ and $V = V_1 \oplus \dots \oplus V_s$. Since the operators \mathcal{B} map each U_i to V_i , the action of the adjoint algebra operator \mathfrak{A} on input matrices $(D, E) \in \text{Lin}(U, U) \times \text{Lin}(V, V)$ can be separated into independent actions of smaller adjoint algebra operators \mathfrak{A}_{jk} acting on matrices $(D_{jk}, E_{jk}) \in \text{Lin}(U_k, U_j) \times \text{Lin}(V_k, V_j)$, for $k, j \in [s]$. Now, choosing a slightly modified inner product on these smaller spaces ensures that the map $\mathfrak{A}_{jk}^T \mathfrak{A}_{jk}$ can be roughly expressed as $(I - \Psi_{jk})$. These new maps Ψ_{jk} depend on the scaled partial derivative operators whose transposes turn out to be a sum of shifts. Here, we make the crucial observation that the derivatives of shifts on a polynomial space of degree $d + 1$ can be converted to shifts of derivatives on a polynomial space of degree d . This leads to an intricate inductive argument to calculate the singular values of Ψ_{jk} by induction on the degree d of the homogeneous polynomial space U_k . These singular values naturally fetch the required singular values of \mathfrak{A} , besides also allowing us to meaningfully relate them to the geometry of the underlying subspaces.

We prove the following theorem in the coming sections.

Theorem 15. *Let \mathfrak{A} , \mathbf{U} and \mathbf{V} be as in theorem 10 for clustering the collection $A = \{\mathbf{a}_1, \dots, \mathbf{a}_N\}$ into subspaces $\langle A_1 \rangle, \dots, \langle A_s \rangle$. For subspaces $\langle A_j \rangle, \langle A_k \rangle$ with canonical angles $\theta_1 \geq \dots \geq \theta_t$ between them, define $f_{jk} = f_d(\langle A_j \rangle, \langle A_k \rangle) = \frac{d+1}{t} [\sum_{k=1}^t \sin^2 \theta_k + d \sin^2 \theta_t]$. Then,*

$$\sigma_{-(s+1)}^2(\mathfrak{A}) \geq \frac{(d+1)^2}{\kappa^4(\mathbf{U}, \mathbf{V})} \cdot \min\{\sigma_{\text{diag}}, \sigma_{\text{off-diag}}\}$$

where the above quantities are defined as follows:

$$\begin{aligned} \sigma_{\text{diag}} &\stackrel{\text{def}}{=} \frac{\sqrt{(d+1)(t^*+d)} - \sqrt{(d+1)(t^*+d)-t^*}}{t^*+d}, \\ \sigma_{\text{off-diag}} &\stackrel{\text{def}}{=} \min_{j \neq k} \frac{\sqrt{(d+1)(t_k+d)} - \sqrt{(d+1)(t_k+d)-t_k \cdot f_{jk}}}{t_k+d}. \end{aligned}$$

Here t_1, \dots, t_s are the dimensions of $\langle A \rangle_1, \dots, \langle A \rangle_s$ and $t^* = \max_{i \in [s]} t_i$.

Further, remark 13 provides further interpretation of the relation between the quantity f_d and $\sigma_{-(s+1)}(\mathfrak{A})$. Note that in the statement of Theorem 15 we assume that $\langle \mathbf{U} \rangle \subseteq \mathbb{R}[\mathbf{x}]^{=d+1}$ and $\langle \mathbf{V} \rangle \subseteq \mathbb{R}[\mathbf{x}]^{=d}$.

G.1 Adjoint Algebra Operators corresponding to Partial Derivatives on Tensorized Spaces

We will begin by estimating the singular values of the adjoint algebra operator corresponding to a simplified subspace clustering instance when $s = 1$ (refer to section 1.3).

This means that we study the relevant operators on homogeneous polynomial spaces like $\mathbb{R}[\mathbf{x}]^d$, instead of a direct sum of many such spaces. As mentioned in the section overview, we will eventually relate the operators on the summed spaces to simpler operators which we will study here. We will take the aid of a special inner product to ease our calculations.

G.1.1 Matrix Representations of Derivatives and Shifts

Let $W = [\mathbf{w}_1 \ \dots \ \mathbf{w}_m]$, where $\mathbf{w}_i \in \mathbb{R}^n$ be a matrix with orthonormal columns. Given the set of variables $\mathbf{x} = (x_1, \dots, x_n)$ and positive integer d define the following:

$$\begin{aligned} \mathbf{y} &\stackrel{\text{def}}{=} W^T \cdot \mathbf{x} \text{ a new set of variables,} \\ U &\stackrel{\text{def}}{=} \mathbb{R}[\mathbf{y}]^{d+1} \text{ of dimension } m^{\binom{d+1}{d}} = \binom{m+d}{d+1}, \\ V &\stackrel{\text{def}}{=} \mathbb{R}[\mathbf{y}]^d \text{ of dimension } m^{\binom{d}{d}} = \binom{m+d-1}{d}. \end{aligned} \quad (12)$$

For $i \in [n]$ we would like to compute the matrix representation of the scaled partial derivatives $\{L_i\} = \left\{ \frac{\partial_i}{d+1} \right\}$ on U with respect to x_i s.

Let $\{p_\alpha\}_\alpha$ and $\{q_\beta\}_\beta$ be the Bombieri basis of U and V respectively with respect to the variables y . Note here that

$$(L_i)_{\beta\alpha} = \langle q_\beta, L_i p_\alpha \rangle = \frac{1}{d+1} \sqrt{\frac{(d+1)!}{\alpha!}} \langle q_\beta, \partial_i \mathbf{y}^\alpha \rangle = \sqrt{\frac{d!}{\alpha!(d+1)}} \left\langle q_\beta, \sum_{j=1}^m \alpha_j \mathbf{y}^{\alpha-j} w_{ji} \right\rangle,$$

where w_{ji} is the i -th coordinate of \mathbf{w}_j . Now, note that given α and β there may exist a j for which $\alpha_k = \beta_k$ for $k \neq j$, and $\alpha_j = \beta_j + 1$. This condition is compactly written as $\alpha = \beta \cup j$ or $\beta = \alpha - j$. In this case, such a j is unique. If no such j exists, \mathbf{y}^β and $\mathbf{y}^{\alpha-j}$ are orthogonal. This gives the matrix representation of L_i as

$$(L_i)_{\beta\alpha} = \begin{cases} w_{ji} \sqrt{\frac{\alpha_j}{d+1}} & \text{if } \exists j \in [m] \text{ such that } \alpha = \beta \cup j, \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Using the above representation we can find the action of L_i^T on q_β as follows:

$$\begin{aligned} L_i^T q_\beta &= \sum_{\alpha} (L_i^T)_{\alpha\beta} \cdot p_\alpha = \sum_{j=1}^m w_{ji} \sqrt{\frac{\beta_j+1}{d+1}} \cdot p_{\beta \cup j} = \sum_{j=1}^m w_{ji} y_j \cdot q_\beta \\ \implies L_i^T q_\beta &= \left(\sum_{j=1}^m w_{ji} y_j \right) q_\beta. \end{aligned} \quad (14)$$

Now, to distinguish the scaled partial derivatives acting on U and V we denote them \bar{L}_i and \underline{L}_i respectively. Then for any $q \in V$ we have

$$\begin{aligned} (\bar{L}_i \bar{L}_j^T) q &= \frac{\partial_i}{d+1} \left(\sum_{k=1}^m w_{kj} y_k \cdot q \right) \\ &= \frac{1}{d+1} \left[\sum_{k=1}^m w_{kj} w_{ki} \cdot q + \sum_{k=1}^m w_{kj} y_k \cdot \partial_i q \right] \\ &= \left[\frac{r_{ij}}{d+1} + \frac{d}{d+1} \underline{L}_j^T \bar{L}_i \right] q, \end{aligned}$$

where $r_{ij} \stackrel{\text{def}}{=} \sum_{k=1}^m w_{kj} w_{ki}$. Further, for any monomial \mathbf{y}^α and the operator L_i (irrespective of the degree of its polynomial domain) we have.

$$\begin{aligned} \left(\sum_{i=1}^n L_i^T L_i \right) \mathbf{y}^\alpha &= \sum_{j=1}^m \sum_{i=1}^n \alpha_j w_{ji} \cdot L_i^T \mathbf{y}^{\alpha-j} \\ &= \frac{1}{d+1} \sum_{j,k=1}^m \alpha_j \mathbf{y}^{\alpha-j+k} \left(\sum_{i=1}^n w_{ji} w_{ki} \right) \\ &= \frac{1}{d+1} \sum_{j=1}^m \alpha_j \mathbf{y}^\alpha \\ &= \mathbf{y}^\alpha. \end{aligned}$$

Hence the above equations give the following:

$$\bar{L}_i \bar{L}_j^T = \frac{r_{ij}}{d+1} I + \frac{d}{d+1} L_j^T L_i \text{ and } \sum_{i=1}^n L_i^T L_i = I. \quad (15)$$

Subsequently, we use the above identities with appropriate dimensions to get

$$\sum_{i=1}^n \bar{L}_i \bar{L}_i^T = \frac{I}{d+1} \sum_{i=1}^n r_{ii} + \frac{d}{d+1} \sum_{i=1}^n L_i^T L_i = \frac{m+d}{d+1} I = \frac{m^{(d+1)}}{m^{(d)}} I. \quad (16)$$

The above relations will turn out to be crucial in the following sections for the singular value analysis of relevant operators.

G.1.2 Adjoint Algebra Operator in a Special Inner Product

Let x, y be vectors in an arbitrary l -dimensional vector space H . Then for positive numbers t_1, \dots, t_l , the bilinear map

$$\langle x, y \rangle_\tau \stackrel{\text{def}}{=} \sum_{i=1}^l \frac{x_i \cdot y_i}{t_i} \quad (17)$$

defines an inner product. The following lemma builds on this special inner product.

Lemma G.1. *Let $A \in \text{Lin}(H, H)$. Let $\tau_1(A), \dots, \tau_l(A)$ denote the singular values of A with respect to the special inner product as defined in equation 17. Then for all $i \in [l]$*

$$\frac{t_*}{t^*} \sigma_i^2(A) \leq \tau_i^2(A) \leq \frac{t^*}{t_*} \sigma_i^2(A).$$

where $t_* = \min_i t_i$ and $t^* = \max_i t_i$. The above inequality is tight.

Proof. Let us first prove the above inequality for the largest singular values $\sigma_1(A)$ and $\tau_1(A)$. For any vector $x \in H$ let $\|x\|_\tau$ denote its norm that arises from the special inner product. Then

$$\|x\|_\tau^2 = \sum_{i=1}^l \frac{x_i^2}{t_i}.$$

As all the numbers involved above are positive we immediately get the inequality

$$\frac{\|x\|^2}{t^*} \leq \|x\|_\tau^2 \leq \frac{\|x\|^2}{t_*}.$$

Then observe that

$$\|Ax\|_\tau^2 \leq \frac{1}{t_*} \|Ax\|^2 \leq \frac{1}{t_*} \sigma_1^2(A) \|x\|^2 \leq \frac{t^*}{t_*} \sigma_1^2(A) \|x\|_\tau^2.$$

Similarly we also get $\|Ax\|^2 \leq \frac{t^*}{t_*} \tau_1^2(A) \|x\|^2$. These two inequalities give us

$$\frac{t_*}{t^*} \sigma_1^2(A) \leq \tau_1^2(A) \leq \frac{t^*}{t_*} \sigma_1^2(A). \quad (18)$$

In order to prove the inequality for other singular values we rely on the following characterisation of the $(i + 1)$ -st singular value.

$$\sigma_{i+1}(A) = \min_{B : \text{rank}(B)=i} \sigma_1(A - B).$$

The above holds for all inner products. Pick any $i \geq 1$. Let B be the matrix of rank i that minimizes the above for $\tau_{i+1}(A)$. Then using equation 18 we get

$$\sigma_{i+1}^2(A) \leq \sigma_1^2(A - B) \leq \frac{t_*}{t^*} \tau_1^2(A - B) = \tau_{i+1}^2(A).$$

A similar strategy obtains the inequality $\tau_i^2(A) \leq \frac{t^*}{t_*} \sigma_i^2(A)$. This proves the required inequality.

To show tightness it is sufficient to show the tightness of equation 18. Without loss of generality assume that $t^* = t_1$ and $t_* = t_l$. Consider a matrix A with $\sigma_1(A) = 1$ which maps the canonical basis vector e_1 to e_l . Then

$$\frac{\|Ae_1\|_\tau^2}{\|e_1\|_\tau^2} = \frac{t_1}{t_l} \implies \tau_1^2(A) \geq \frac{t_1}{t_l} = \frac{t_1}{t_l} \sigma_1^2(A).$$

Similarly considering a matrix A with $\tau_1(A) = 1$ which maps the vector $\sqrt{t_l}e_1$ to $\sqrt{t_1}e_1$ we get

$$\frac{\|Ae_l\|^2}{\|e_l\|^2} = \frac{t_1}{t_l} \implies \sigma_1^2(A) \geq \frac{t_1}{t_l} = \frac{t_1}{t_l} \tau_1^2(A).$$

This completes the proof of the lemma. \square

Now, we will adopt the special inner product to the context of our polynomial spaces. Let $W = [\mathbf{w}_1 \dots \mathbf{w}_m]$ and $\Omega = [\boldsymbol{\omega}_1 \dots \boldsymbol{\omega}_t]$ be two matrices with orthonormal columns with $\mathbf{w}_i, \boldsymbol{\omega}_i \in \mathbb{R}^n$. The set of variables \mathbf{y} and \mathbf{z} are defined as $\mathbf{y} = W^T \cdot \mathbf{x}$ and $\mathbf{z} = \Omega^T \cdot \mathbf{x}$. Further,

$$\begin{aligned} U_1 &= \mathbb{R}[\mathbf{y}]^{=d+1}, & V_1 &= \mathbb{R}[\mathbf{y}]^{=d}, \\ U_2 &= \mathbb{R}[\mathbf{z}]^{=d+1}, & V_2 &= \mathbb{R}[\mathbf{z}]^{=d}. \end{aligned} \quad (19)$$

And, for $i \in [n]$, the collections $\{L_{i1}\}$ and $\{L_{i2}\}$ are scaled partial derivatives (with respect to \mathbf{x}) on the spaces U_1 and U_2 respectively. That is

$$L_{i1} = \frac{\partial_i}{d+1} \Big|_{U_1} \quad \text{and} \quad L_{i2} = \frac{\partial_i}{d+1} \Big|_{U_2}.$$

Imitating the definition 17 we define the special inner product $\langle \cdot, \cdot \rangle_\tau$ on a space of linear operators by scaling the standard inner product by the dimension of the domain. For example, for E_1, E_2 in $\text{Lin}(V_2, V_1)$, the inner product is given by

$$\langle E_1, E_2 \rangle_\tau = \frac{\langle E_1, E_2 \rangle}{\dim V_2} = \frac{\langle E_1, E_2 \rangle}{t^{(d)}}. \quad (20)$$

For this inner product on a sum of spaces of linear operators, we scale each component as above. For example, for the space $\text{Lin}(U_2, U_1) \times \text{Lin}(V_2, V_1)$, we have

$$\langle (D_1, E_1), (D_2, E_2) \rangle_\tau = \langle D_1, D_2 \rangle_\tau + \langle E_1, E_2 \rangle_\tau = \frac{\langle D_1, D_2 \rangle}{t^{(d+1)}} + \frac{\langle E_1, E_2 \rangle}{t^{(d)}}. \quad (21)$$

With the above inner product in place, we can define a new operator Φ which we will see is closely connected to the adjoint algebra operator. Let

$$\Phi : \text{Lin}(V_2, V_1) \rightarrow \text{Lin}(U_2, U_1) \quad \text{with} \quad \Phi(E) \stackrel{\text{def}}{=} \sum_{i=1}^n L_{i1}^T E L_{i2}. \quad (22)$$

We can calculate the adjoint of this map with respect to the inner product defined above as done below:

$$\begin{aligned} \text{Tr} \Phi^*(D)^T E &= \dim V_2 \cdot \langle \Phi^*(D), E \rangle_\tau \\ &= \dim V_2 \cdot \langle D, \Phi(E) \rangle_\tau \\ &= \frac{\dim V_2}{\dim U_2} \cdot \text{Tr} \sum_{i=1}^n \left(L_{i1} D L_{i2}^T \right)^T E. \end{aligned}$$

This yields that

$$\Phi^*(D) = \frac{t^{(d)}}{t^{(d+1)}} \sum_{i=1}^n L_{i1} D L_{i2}^T. \quad (23)$$

Now recall that the adjoint algebra map $\mathfrak{A} : \text{Lin}(U_2, U_1) \times \text{Lin}(V_2, V_1) \rightarrow \text{Lin}(U_2, V_1)^n$ is given by

$$\mathfrak{A}(D, E) = (L_{11}D - EL_{12}, \dots, L_{n1}D - EL_{n2}).$$

Quickly note here that when $W = \Omega$, that is, $L_{i1} = L_{i2}$ for all $i \in [n]$, we have $\mathfrak{A}(I, I) = 0$. Thus, the bottom singular value of \mathfrak{A} is 0 in this case. Going forward we analyze the other singular values by looking at the eigenvalues of $\mathfrak{A}^T \mathfrak{A}$. We note that

$$\begin{aligned} \|\mathfrak{A}(D, E)\|_F^2 &= \sum_{i=1}^n \|L_{i1}D - EL_{i2}\|_F^2 \\ &= \text{Tr} D^T \left[\sum_{i=1}^n L_{i1}^T L_{i1} \right] D + \text{Tr} E^T E \left[\sum_{i=1}^n L_{i2} L_{i2}^T \right] \\ &\quad - \text{Tr} D^T \left[\sum_{i=1}^n L_{i1}^T E L_{i2} \right] - \text{Tr} E^T \left[\sum_{i=1}^n L_{i1} D L_{i2}^T \right] \\ &= t^{(d+1)} \left\langle \begin{bmatrix} D \\ E \end{bmatrix}, \begin{bmatrix} \frac{I}{t^{(d+1)}} & 0 \\ 0 & \frac{I}{t^{(d)}} \end{bmatrix} \begin{bmatrix} I & -\Phi \\ -\Phi^* & I \end{bmatrix} \begin{bmatrix} D \\ E \end{bmatrix} \right\rangle, \end{aligned}$$

where the map Φ is as defined in definition 22 and its adjoint Φ^* is taken according to the special inner product as defined in 23. Renaming (D, E) to X and giving appropriate names to the matrices appearing in the above equality, we obtain

$$\langle X, \mathfrak{A}^T \mathfrak{A} X \rangle = t^{(d+1)} \langle X, Y \Psi X \rangle,$$

where Y is the diagonal matrix. Note here that the matrix $Y \Psi$ is Hermitian in the standard inner product. This gives us

$$\mathfrak{A}^T \mathfrak{A} = t^{(d+1)} Y \Psi.$$

Let $\tau_i(\Phi)$ and $\tau_{-i}(\Phi)$ be used to index the singular values of Φ in decreasing and increasing orders respectively. From here we obtain the crucial relation between $\sigma_{-i}(\mathfrak{A})$ and $\tau_i(\Phi)$ as follows:

$$\begin{aligned} \sigma_{-i}^2(\mathfrak{A}) &= \sigma_{-i}(\mathfrak{A}^T \mathfrak{A}) \\ &\geq t^{(d+1)} \sigma_{-1}(Y) \sigma_{-i}(\Psi) \\ &\geq \left(\frac{t^{(d)}}{t^{(d+1)}} \right)^{1/2} \tau_{-i}(\Psi) \\ \implies \sigma_{-i}^2(\mathfrak{A}) &\geq \sqrt{\frac{d+1}{t+d}} (1 - \tau_i(\Phi)). \end{aligned} \quad (24)$$

Here we have used lemma G.1 and the following facts. For any two matrices A, B we have $\sigma_i(AB) \geq \sigma_{-1}(A) \sigma_i(B)$. And the eigenvalues of the Hermitian matrix Ψ (with respect to the special inner product) is the set $\{1 \pm \tau_i(\Phi)\}$.

G.1.3 Singular Values of the Φ Operator

We immediately delve into finding $\tau_i(\Phi)$ following the definition of $\langle \cdot, \cdot \rangle_\tau$ in 20 and Φ in 22. Recall that U_1, V_1 and U_2, V_2 are homogeneous polynomial spaces with respect to the variables \mathbf{y} and \mathbf{z} defined by the matrices W and Ω respectively, as given in 19.

The map Φ is defined on $\text{Lin}(V_2, V_1)$ which are linear maps on spaces of homogeneous polynomials of degree d . To emphasize this dependence on the degree we call Φ as Φ_d , and subsequently, for different d s the domain and the co-domain of the map Φ_d changes. In this spirit, to distinguish the scaled derivatives of larger and smaller spaces we imitate the notation in 15 as follows:

$$\bar{L}_{i1} = \frac{\partial_i}{d+1} \Big|_{U_1}, \quad \bar{L}_{i2} = \frac{\partial_i}{d+1} \Big|_{U_2}, \quad \underline{L}_{i1} = \frac{\partial_i}{d} \Big|_{V_1}, \quad \underline{L}_{i2} = \frac{\partial_i}{d} \Big|_{V_2}.$$

Therefore the above scaled derivatives define different Φ maps as follows:

$$\Phi_d(E) = \sum_{i=1}^n \bar{L}_{i1}^T E \bar{L}_{i2} \quad \text{and} \quad \Phi_{d-1}(E) = \sum_{i=1}^n \underline{L}_{i1}^T E \underline{L}_{i2}.$$

These identifications allow an inductive approach to calculate the relevant singular values, inspired by the inductive argument in recent works on analyzing eigenvalues for random walks on simplicial complexes (e.g. [ALGV19]).

Lemma G.2. Let $W^T \Omega = \text{diag}(\cos \theta_1, \dots, \cos \theta_t)$ (appended with 0s if necessary) where $\theta_1, \dots, \theta_t$ are the canonical angles between $\langle W \rangle$ and $\langle \Omega \rangle$ according to remark 7. Let $f_d(\langle W \rangle, \langle \Omega \rangle) = \frac{d+1}{t} [\sum_{k=1}^t \sin^2 \theta_k + d \sin^2 \theta_{\min}]$ and $g(W, \Omega) = \frac{1}{t} \sum_{k=1}^t \cos^2 \theta_k$. Then the top $m \cdot t$ singular values of Φ_d (with respect to the inner product defined in 20) are given by

$$\tau_i^2(\Phi_d) \leq 1 - \frac{1}{d+1} \frac{t}{t+d} [g + f_d - \tau_i^2(\Phi_0)].$$

Further, if $W = \Omega$, then the above inequality is an equality with $g = 1$ and $f_d = 0$.

Proof. We proceed with induction on d . Define operators Φ_d^\wedge and Φ_d^\vee on $\text{Lin}(V_2, V_1)$ as follows:

$$\Phi_d^\vee = \Phi_{d-1} \Phi_{d-1}^* \text{ and } \Phi_d^\wedge = \Phi_d^* \Phi_d.$$

Then, using the definition of Φ^* from 23 we can write Φ_d^\vee and Φ_d^\wedge as follows:

$$\Phi_d^\vee(D) = \frac{d}{t+d-1} \sum_{i,j=1}^n \underline{L}_{i1}^T \underline{L}_{j1} D \underline{L}_{j2}^T \underline{L}_{i2} \text{ and } \Phi_d^\wedge(D) = \frac{d+1}{t+d} \sum_{i,j=1}^n \bar{L}_{i1} \bar{L}_{j1}^T D \bar{L}_{j2} \bar{L}_{i2}^T.$$

Recalling the relation for converting derivatives of shifts to shifts of derivatives from 15, we obtain

$$\begin{aligned} (d+1)(t+d)\Phi_d^\wedge(D) &= \sum_{i,j=1}^n \left(r_{ij} I + d \underline{L}_{j1}^T \underline{L}_{i1} \right) D \left(q_{ij} I + d \underline{L}_{i2}^T \underline{L}_{j2} \right) \\ &= \sum_{i,j=1}^n \left[(r_{ij} \cdot q_{ij}) D + d \cdot q_{ij} \left(\underline{L}_{j1}^T \underline{L}_{i1} \right) D + d \cdot r_{ij} \cdot D \left(\underline{L}_{i2}^T \underline{L}_{j2} \right) \right. \\ &\quad \left. + d^2 \left(\underline{L}_{j1}^T \underline{L}_{i1} D \underline{L}_{i2}^T \underline{L}_{j2} \right) \right] \\ &= (tg \cdot D + GD + DH) + d(t+d-1)\Phi_d^\vee(D), \end{aligned} \quad (25)$$

where the involved quantities are defined as follows:

$$r_{ij} = \sum_{k \in [m]} \omega_{ki} \omega_{kj}, \quad q_{ij} = \sum_{k \in [t]} \omega_{ki} \omega_{kj}, \quad G = d \sum_{i,j=1}^n q_{ij} \cdot \underline{L}_{j1}^T \underline{L}_{i1} \text{ and } H = d \sum_{i,j=1}^n r_{ij} \cdot \underline{L}_{i2}^T \underline{L}_{j2},$$

which further yields the following relation used above:

$$\sum_{i,j=1}^n r_{ij} q_{ij} = \sum_{k=1}^m \sum_{i,j=1}^n (\omega_{ki} \omega_{ki}) (\omega_{kj} \omega_{kj}) = \sum_{k=1}^m \langle \mathbf{w}_k, \mathbf{w}_k \rangle^2 = \sum_{k=1}^t \cos^2(\theta_k) = tg.$$

Now, we express the matrices G and H in terms of the canonical angles θ_k s. Use the matrix representation of \underline{L}_{j1} from 13 to obtain

$$\begin{aligned} \left(\sum_{j=1}^n \omega_{kj} \underline{L}_{j1} \right)_{\beta\alpha} &= \sum_{j=1}^n \omega_{kj} \omega_{lj} \sqrt{\frac{\alpha_l}{d}} \\ &= \begin{cases} \langle \omega_k, \mathbf{w}_l \rangle \sqrt{\frac{\alpha_l}{d}} & \text{only when } \exists l \text{ such that } \alpha = \beta \cup l, \text{ and } 0 \text{ otherwise,} \\ \cos \theta_k \sqrt{\frac{\alpha_k}{d}} & \text{only when } \alpha = \beta \cup k, \text{ and } 0 \text{ otherwise.} \end{cases} \end{aligned}$$

From here, observe that G can also be written as

$$G = d \sum_{i,j=1}^n q_{ij} \cdot L_{j1}^T L_{i1} = d \sum_{k=1}^m \left(\sum_{j=1}^n \omega_{kj} L_{j1} \right)^T \left(\sum_{j=1}^n \omega_{kj} L_{j1} \right).$$

Plugging the appropriate quantities calculated above, and calculating similarly for H , we get

$$G = \text{diag} \left\{ \sum_{k=1}^m \alpha_k \cos^2 \theta_k \right\}_{\alpha} \quad \text{and} \quad H = \text{diag} \left\{ \sum_{k=1}^t \gamma_k \cos^2 \theta_k \right\}_{\gamma}$$

where α s and β s are multi-indices for the basis of d -degree homogeneous polynomials on m and t variables respectively.

Now, computing the i -th singular value from the equation 25, we get the following recurrence relation:

$$\begin{aligned} (d+1)(t+d) \cdot \tau_i(\Phi_d^\wedge) &\leq tg + \|G\| + \|H\| + d(t+d-1) \cdot \tau_i(\Phi_d^\vee) \\ \implies (d+1)(t+d) \cdot \tau_i^2(\Phi_d) &\leq tg + 2d \cos^2 \theta_{\min} + d(t+d-1) \cdot \tau_i^2(\Phi_{d-1}). \end{aligned}$$

Solving the above recurrence yields the required result. In the case when $W = \Omega$, we have $\theta_k = 0$ for all k , implying $G = H = dI$ and $g = 1$. Now, as Φ_d^\wedge and Φ_{d-1}^\wedge differ by a scaled identity, they are both diagonalizable in the same basis, and hence we get the following exact relation:

$$(t+d)(d+1) \cdot \Phi_d^\wedge = (t+2d) \cdot I + (t+d-1) \cdot \Phi_d^\vee.$$

Computing the singular values from the above recurrence gives the required result for this case. This completes the proof of the lemma. \square

Base case. Note that Φ_0 is a map defined as $\Phi_0 : \mathbb{R} \rightarrow \mathbb{R}^{t \times m}$ such that

$$\begin{aligned} \Phi_0(c) \cdot p &= c \sum_{i=1}^n L_{i1}^T L_{i2} \cdot p \\ &= c \sum_{i=1}^n \sum_{k=1}^m \sum_{j=1}^t w_{ji} \omega_{ki} p_k y_j \\ &= c \sum_{j=1}^m \left[\sum_{k=1}^t \langle \mathbf{w}_j, \omega_k \rangle p_k \right] y_j \\ &= c \cdot W^T \Omega \cdot p \end{aligned}$$

Therefore, with respect to our scaled inner product, we obtain

$$\tau_1^2(\Phi_0) = \frac{1}{t} \left\| W^T \Omega \right\|_F^2 = \frac{1}{t} \sum_{k \in [t]} \cos^2 \theta_k = g. \quad (26)$$

Now, as Φ_0 is defined on \mathbb{R} it admits at most 1 non-zero singular value. This shows that $\tau_i(\Phi_0) = 0$ for all $i \geq 2$. Clubbing the base case with the above lemma gives us the following remarks.

Remark 12. We get the following singular values of the Φ operator in the particular cases stated below:

1. When $W = \Omega$, we have $\tau_1(\Phi_0) = 1$. This gives $\tau_1(\Phi) = 1$.
2. Also, when $W = \Omega$, plugging $\tau_2(\Phi_0) = 0$ gives $\tau_2^2(\Phi) = 1 - \frac{1}{d+1} \frac{t}{t+d}$.
3. When $W \neq \Omega$, we plug $\tau_1^2(\Phi_0) = g$. This gives us $\tau_1^2(\Phi) = 1 - \frac{1}{d+1} \frac{t}{t+d} \cdot f_d$.

G.2 Adjoint Algebra Operator for Subspace Clustering

Recall the setting of subspace clustering problem in section B where we have a set of N points $A = \{\mathbf{a}_1, \dots, \mathbf{a}_N\} \in \mathbb{R}^n$ clusterable to the subspaces $\langle A_1 \rangle, \dots, \langle A_s \rangle$. These subspaces $\langle A_i \rangle$ of dimension t_i are also expressed as $\langle W_i \rangle$, where the columns of W_i are orthonormal. Let \mathbf{U} and \mathbf{V} be such that

$$\mathbf{U} = (U_1, \dots, U_s) = \left(\left\langle A_1^{\otimes d+1} \right\rangle, \dots, \left\langle A_s^{\otimes d+1} \right\rangle \right)$$

$$\text{and } \mathbf{V} = (V_1, \dots, V_s) = \left(\left\langle A_1^{\otimes d} \right\rangle, \dots, \left\langle A_s^{\otimes d} \right\rangle \right).$$

The section 1.3 tells us that the tuples \mathbf{U} and \mathbf{V} along with the differential operators $\mathcal{B} = \{\partial_i, \dots, \partial_n\}$ form an instance of the vector space decomposition problem. Following the conventions set in the equations in 12, the spaces U_i and V_i have dimensions $t_i^{(d+1)}$ and $t_i^{(d)}$ respectively. Further, we know that there exist bases P and Q , which are \mathbf{U} -associated and \mathbf{V} -associated matrices (as defined in section 2), such that in these bases the scaled derivatives, which are the operators we have studied in the above section, are block diagonal operators. That is

$$\mathcal{L} = \{L_1, \dots, L_n\} \text{ with } L_i = \text{diag}(L_{i1}, \dots, L_{is}) \text{ where } L_{ij} = \frac{\partial_i}{d+1} \Big|_{U_j}.$$

Thus, the adjoint algebra operator $\mathfrak{A} = \mathfrak{A}(\mathcal{B})$ in 10 can be expressed as follows:

$$\begin{aligned} \mathfrak{A}(D, E) &= (d+1) \cdot \left(QL_i P^{-1} D - EQ L_i P^{-1} \right)_{i \in [n]} \\ &= (d+1) \cdot Q \cdot \left(L_i (P^{-1} D P) - (Q^{-1} E Q) L_i \right)_{i \in [n]} \cdot P^{-1}. \end{aligned}$$

The above establishes the following relationship between \mathfrak{A} and the adjoint algebra operator $\mathfrak{A}_{\mathcal{L}}$ with respect to the operator collection \mathcal{L} :

$$\mathfrak{A} = (d+1) \cdot \Gamma_2 \circ \mathfrak{A}_{\mathcal{L}} \circ \Gamma_1$$

where

$$\Gamma_1(D, E) = (P^{-1} D P, Q^{-1} E Q) \text{ and } \Gamma_2(X_1, \dots, X_n) = Q \cdot (X_1, \dots, X_n) \cdot P^{-1}.$$

This immediately yields that

$$\sigma_{-(s+1)}(\mathfrak{A}) \geq (d+1) \cdot \sigma_{-(s+1)}(\mathfrak{A}_{\mathcal{L}}) \cdot \sigma_{-1}(\Gamma_1) \cdot \sigma_{-1}(\Gamma_2). \quad (27)$$

One can compute $\sigma_{-1}(\Gamma_1)$ and $\sigma_{-1}(\Gamma_2)$ from the definition of Γ_1 and Γ_2 to obtain

$$\sigma_{-1}(\Gamma_1) \cdot \sigma_{-1}(\Gamma_2) \geq \frac{1}{\kappa^2(\mathbf{U}, \mathbf{V})} \quad \text{where} \quad \kappa(\mathbf{U}, \mathbf{V}) \stackrel{\text{def}}{=} \frac{\max\{\sigma_1(\mathbf{U}), \sigma_1(\mathbf{V})\}}{\min\{\sigma_{-1}(\mathbf{U}), \sigma_{-1}(\mathbf{V})\}}. \quad (28)$$

We now focus on computing the quantity $\sigma_{-(s+1)}(\mathfrak{A}_{\mathcal{L}})$. We exploit the block diagonal structure of the matrices in the collection \mathcal{L} . Treating D and E also as block matrices $[D_{jk}]$ and $[E_{jk}]$, we note that the Frobenius norm of $\mathfrak{A}_{\mathcal{L}}$ can be separated across the blocks as follows:

$$\|\mathfrak{A}_{\mathcal{L}}(D, E)\|_F^2 = \sum_{i=1}^n \|L_i D - E L_i\|_F^2 = \sum_{j,k=1}^s \sum_{i=1}^n \|L_{ij} D_{jk} - E_{jk} L_{ik}\|_F^2 \stackrel{\text{def}}{=} \sum_{j,k=1}^s \|\mathfrak{A}_{jk}(D_{jk}, E_{jk})\|_F^2,$$

where each $\mathfrak{A}_{jk} : \text{Lin}(U_k, U_j) \times \text{Lin}(V_k, V_j) \rightarrow \text{Lin}(U_k, V_j)^t$ is a **block adjoint algebra operator** defined as follows:

$$\mathfrak{A}_{jk}(D_{jk}, E_{jk}) = (L_{1j} D_{jk} - E_{jk} L_{1k}, \dots, L_{nj} D_{jk} - E_{jk} L_{nk}).$$

Therefore we note that the adjoint algebra map $\mathfrak{A}_{\mathcal{L}}$ has the following block structure:

$$\mathfrak{A}_{\mathcal{L}}^T \mathfrak{A}_{\mathcal{L}} = \bigoplus_{j,k=1}^s \mathfrak{A}_{jk}^T \mathfrak{A}_{jk}, \quad (29)$$

where the notation above implies that the map $\mathfrak{A}^T \mathfrak{A}$ acts separately as $\mathfrak{A}_{jk}^T \mathfrak{A}_{jk}$ on different independent subspaces of the domain. This immediately implies that the singular values of the adjoint algebra operator \mathfrak{A} are the singular values of the block adjoint algebra operators \mathfrak{A}_{jk} collected together.

Now recall that equation 24 gives

$$\sigma_{-i}^2(\mathfrak{A}_{jk}) \geq \sqrt{\frac{d+1}{i+d}} (1 - \tau_i(\Phi_{jk})),$$

where the Φ_{jk} maps are related to scaled derivatives on polynomial spaces on variables $W_j^T \mathbf{x}$ and $W_k^T \mathbf{x}$ respectively. Now, as $\sigma_{-1}(\mathfrak{A}_{jj}) = 0$ for all $j \in [s]$ with (I, I) in its null space, we conclude that $\sigma_{-(s+1)}(\mathfrak{A}_{\mathcal{L}})$ must be $\sigma_{-2}(\mathfrak{A}_{jj})$ for some j , or $\sigma_{-1}(\mathfrak{A}_{jk})$ for some $j \neq k$. Thus, combining all the components above with remark 12 we establish theorem 15.

Finally, we end with a remark on how the function f_d helps relate $\sigma_{-(s+1)}(\mathfrak{A})$ to the geometry of the underlying input subspaces.

Remark 13. Recall for given subspaces $\langle A_i \rangle$ and $\langle A_j \rangle$ we have

$$f_d(\langle A_i \rangle, \langle A_j \rangle) = \frac{d+1}{t} \left[\sum_{k=1}^t \sin^2 \theta_k + d \sin^2 \theta_{\min} \right]$$

where θ_k s are the canonical angles between $\langle A_i \rangle$ and $\langle A_j \rangle$.

1. For any pair of subspaces note that

$$0 \leq t \cdot f_d \leq (t+d)(d+1).$$

Further, $f_d = 0$ if and only if $\theta_k = 0$ for all k , implying $\langle A_i \rangle \subseteq \langle A_j \rangle$. Thus, excluding the degenerate cases when one of the subspaces is contained in another subspace, we always have $\sigma_{s+1}(\mathfrak{A}) > 0$, thereby ensuring that the Adjoint algebra has dimension exactly equal to s .

2. Observe that we can write $f_d \geq \frac{(d+1)(t+d)}{t} \sin^2 \theta_{\min}$. Substituting this in $\sigma_{\text{off-diag}}$, we get

$$\sigma_{\text{off-diag}} \geq 1 - \cos \theta_{\min},$$

where θ_{\min} is the smallest canonical angle between any pair of distinct subspaces. This is a good heuristic for the contribution of the off-diagonal blocks of \mathfrak{A} to its $(s+1)$ -st smallest singular value in the case when the subspaces have trivial pairwise intersection.

3. When $f_d > 1$ for all pairs of distinct subspaces $\langle A_i \rangle$ and $\langle A_j \rangle$, then σ_{diag} is smaller than each term involved in the minimum in $\sigma_{\text{off-diag}}$. This gives

$$\sigma_{-(s+1)}(\mathfrak{A})^2 \geq \frac{(d+1)^2}{\kappa^4(\mathbf{U}, \mathbf{V})} \cdot \sigma_{\text{diag}}.$$

This is a bound that doesn't depend on geometry of the subspaces. This also shows that for any non-degenerate instance, one can choose d large enough so that $\sigma_{-(s+1)}^2(\mathfrak{A})$ has the form as above.