

SIMULTANEOUS ROBUST SUBSPACE RECOVERY AND SEMI-STABILITY OF QUIVER REPRESENTATIONS

CALIN CHINDRIS AND DANIEL KLINE

ABSTRACT. We consider the problem of *simultaneously* finding lower-dimensional subspace structures in a given m -tuple $(\mathcal{X}^1, \dots, \mathcal{X}^m)$ of possibly corrupted, high-dimensional data sets all of the same size. We refer to this problem as *simultaneous robust subspace recovery* (SRSR) and provide a quiver invariant theoretic approach to it. We show that SRSR is a particular case of the more general problem of effectively deciding whether a quiver representation is semi-stable (in the sense of Geometric Invariant Theory) and, in case it is not, finding a subrepresentation certifying in an optimal way that the representation is not semi-stable. In this paper, we show that SRSR and the more general quiver semi-stability problem can be solved effectively.

CONTENTS

1. Introduction	1
2. Background on quiver invariant theory	4
3. Algorithm G	9
4. Shrunk subspaces and rank decreasing witnesses	11
5. Recovering lower-dimensional subspace structures	15
6. Semi-stability of quiver representations	17
References	22

1. INTRODUCTION

1.1. Formulation of the SRSR problem. Given a point cloud, classical robust subspace recovery asks for a lower-dimensional subspace containing a sufficiently large number of data points. In this paper, we take a quiver invariant theoretic approach to this problem. This allows us to work with multiple point clouds, all of the same cardinality, *simultaneously*. To decide what it means for a lower-dimensional subspace structure to contain enough points across a collection of point clouds, we use the Hilbert-Mumford numerical criterion for semi-stability of quiver representations [Kin94] as our guiding principle.

Let D, m, d_1, \dots, d_m be positive integers such that $D = \sum_{j=1}^m d_j$. Let $\mathcal{X} = \{X_1, \dots, X_n\} \subseteq \mathbb{R}^D$ be a data set such that the first d_1 coordinates of each one of the vectors X_i encode

Date: September 22, 2020.

2010 Mathematics Subject Classification. 16G20, 13A50, 14L24.

Key words and phrases. Capacity of completely positive operators, partitioned data sets, semi-stability of quiver representations, shrunk subspaces, simultaneous robust subspace recovery, Wong sequences.

measurements of a certain type-1, the next d_2 coordinates encode measurements of type-2, and so on, with the last d_m coordinates encoding measurements of type- m . We write

$$\mathcal{X}^j \subseteq \mathbb{R}^{d_j} \left\{ \begin{array}{c} X_1, \quad \cdots \quad X_i, \quad \cdots \quad X_n \\ \parallel \qquad \qquad \parallel \qquad \qquad \parallel \\ \begin{bmatrix} v_1^1 \\ \vdots \\ v_1^j \\ \vdots \\ v_1^m \end{bmatrix} \qquad \begin{bmatrix} v_i^1 \\ \vdots \\ v_i^j \\ \vdots \\ v_i^m \end{bmatrix} \qquad \begin{bmatrix} v_n^1 \\ \vdots \\ v_n^j \\ \vdots \\ v_n^m \end{bmatrix} \end{array} \right\} \text{ type-}j \text{ measurements}$$

with $v_i^j \in \mathbb{R}^{d_j}, i \in [n], j \in [m]$, and refer to \mathcal{X} as a *partitioned data set in \mathbb{R}^D of type (n, m, \underline{d})* where $\underline{d} = (d_1, \dots, d_m)$. It gives rise to the m -tuple of data sets $(\mathcal{X}^1, \dots, \mathcal{X}^m)$. Conversely, any tuple of data sets, all of them of the same size, defines a partitioned data set.

Let us now consider the m -tuple of data sets $\mathcal{X}^j := \{v_1^j, \dots, v_n^j\} \subseteq \mathbb{R}^{d_j}, j \in [m]$. Our goal in this paper is to perform *robust subspace recovery* on the data sets $\mathcal{X}^1, \dots, \mathcal{X}^m$ *simultaneously*. To explain what this means, let us consider an m -tuple of subspaces (T_1, \dots, T_m) with $T_j \subseteq \mathbb{R}^{d_j}, j \in [m]$, and set

$$\mathcal{I}_T := \{i \in [n] \mid v_i^j \in T_j, \forall j \in [m]\}.$$

Thus the initial data set $\mathcal{X} \subseteq \mathbb{R}^D$ contains precisely $|\mathcal{I}_T|$ vectors such that for each one of these vectors the first d_1 coordinates are from T_1 , the next d_2 coordinates are from T_2 , and so on. We are now ready to state what we mean by *simultaneous robust subspace recovery* (SRSR).

Problem. Given an m -tuple of data sets $\mathcal{X}^j := \{v_1^j, \dots, v_n^j\} \subseteq \mathbb{R}^{d_j}, j \in [m]$, is there an effective way to find, whenever possible, an m -tuple of subspaces (T_1, \dots, T_m) with $T_j \subseteq \mathbb{R}^{d_j}, j \in [m]$, such that


$$(1) \quad |\mathcal{I}_T| > \frac{\sum_{j=1}^m \dim T_j}{D} \cdot n?$$

By an *SRSR solution for \mathcal{X}* , we mean any m -tuple of subspaces (T_1, \dots, T_m) for which (1) holds. We say that \mathcal{X} admits a lower-dimensional subspace structure if there exists an SRSR solution for \mathcal{X} .

When $m = 1$, this is a central problem in *Robust Subspace Recovery* (RSR) (see for example [HM13], [LM18]). In [HM13], the authors argue that the RSR version ($m = 1$) of (1) is the optimal balance between robustness and efficiency. For arbitrary $m \geq 1$, we point out that (1) is intimately related to the Hilbert-Mumford numerical criterion for semi-stability applied to quiver representations (see Section 2 for more details).

1.2. The quiver approach. Our approach to SRSR is based on quiver invariant theory (see Section 2 for more details). We begin by viewing the m -tuple of data sets $\mathcal{X}^j \subseteq \mathbb{R}^{d_j}, j \in [m]$, as a representation $V_{\mathcal{X}}$ of the complete bipartite quiver $\mathcal{K}_{n,m}$ with n source vertices

and m sink vertices as follows:



The dimension vector of V_X , denoted by $\dim V_X \in \mathbb{N}^{nm}$, simply records the dimensions of the vector spaces attached to the vertices of $\mathcal{K}_{n,m}$, i.e.

$$\dim V_X(x_i) = 1, \forall i \in [n], \text{ and } \dim V_X(y_j) = d_j, \forall j \in [m].$$

By a *weight* of a quiver, we simply mean an assignment of integer numbers to the vertices of the quiver. In the SRSR setup, we are primarily interested in the weight $\sigma_0 \in \mathbb{Z}^{nm}$ defined by

$$(2) \quad \sigma_0(x_i) = D, \forall i \in [n], \text{ and } \sigma_0(y_j) = -n, \forall j \in [m].$$

One of the key concepts in quiver invariant theory is that of a semi-stable quiver representation with respect to a weight (see Definition 3). It turns out that the partitioned data set $\mathcal{X} = (\mathcal{X}^1, \dots, \mathcal{X}^m)$, viewed as the representation V_X of $\mathcal{K}_{n,m}$, is σ_0 -semi-stable if and only if

$$(3) \quad D \cdot |\mathcal{I}_T| - n \cdot \left(\sum_{j=1}^m \dim T_j \right) \leq 0,$$

for every m -tuple of subspaces (T_1, \dots, T_m) with $T_j \leq \mathbb{R}^{d_j}$, $j \in [m]$. (For more details, see Remark 5.) Thus Problem 1.1 can be rephrased as asking to decide whether the representation V_X is *not* σ_0 -semi-stable and, if that is the case, to find a subrepresentation of V_X for which (3) does *not* hold.

Our quiver approach combined with known algorithms, due to L. Gurvits [Gur04], and G. Ivanyos, Y. Qiao, and K.V. Subrahmanyam [IQS17], yields the following results.

Theorem 1. *Let $\mathcal{X}^j := \{v_1^j, \dots, v_n^j\} \subseteq \mathbb{R}^{d_j}$, $j \in [m]$, be an m -tuple of data sets. Then there exist a deterministic polynomial time algorithm and a probabilistic algorithm to decide whether $(\mathcal{X}^1, \dots, \mathcal{X}^m)$ has a lower-dimensional subspace structure and to compute a lower-dimensional subspace structure in case it exists.*

Moreover, if $(\mathcal{X}^1, \dots, \mathcal{X}^m)$ admits a lower-dimensional subspace structure then any output (T_1, \dots, T_m) of the deterministic or probabilistic algorithm is optimal in the following sense. If

(T'_1, \dots, T'_m) is any lower-dimensional subspace structure for $(\mathcal{X}^1, \dots, \mathcal{X}^m)$ then

$$(4) \quad |\mathcal{I}_T| - \frac{\sum_{j=1}^m \dim T_j}{D} \cdot n \geq |\mathcal{I}_{T'}| - \frac{\sum_{j=1}^m \dim T'_j}{D} \cdot n.$$

The probabilistic algorithm (Algorithm P) in the theorem above is a randomized version of the deterministic polynomial time algorithm from [IQS17]. We point out that Algorithm P, based on the Schwartz-Zippel Lemma [Sch80, Zip79], is easier to implement in practice.

In the general context of quiver representations, the *discrepancy* of a quiver datum (V, σ) , where V is a representation and σ is a non-zero weight of a quiver, is defined as

$$\text{disc}(V, \sigma) = \max\{\sigma \cdot \mathbf{dim} W \mid W \text{ is a subrepresentation of } V\}.$$

(We refer to Section 6 for the details behind our notations.) We have the following result, extending Theorem 1 to arbitrary quiver representations.

Theorem 2. *Let Q be a quiver without oriented cycles and let (V, σ) be a quiver datum (defined over \mathbb{Q}) such that $\sigma \cdot \mathbf{dim} V = 0$.*

- (i) *Then there exist deterministic polynomial time algorithms to check whether V is σ -semi-stable or not.*
- (ii) *Assume that Q is a bipartite quiver (not necessarily complete) and σ is nowhere zero. Then there exists a deterministic polynomial time algorithm that constructs a subrepresentation $W \leq V$ such that*

$$\text{disc}(V, \sigma) = \sigma \cdot \mathbf{dim} W.$$

In particular, if $\sigma \cdot \mathbf{dim} W = 0$ then V is σ -semi-stable. Otherwise, W is an optimal witness to V not being σ -semi-stable.

Acknowledgment. The authors would like to thank to Petros Valettas for many enlightening discussions on the paper, especially for suggesting the use of the Schwartz-Zippel Lemma in Proposition 24.

2. BACKGROUND ON QUIVER INVARIANT THEORY

Throughout, we work over the field \mathbb{R} of real numbers and denote by $\mathbb{N} = \{0, 1, \dots\}$. For a positive integer L , we denote by $[L] = \{1, \dots, L\}$.

2.1. Semi-stability of quiver representations. A quiver $Q = (Q_0, Q_1, t, h)$ consists of two finite sets Q_0 (vertices) and Q_1 (arrows) together with two maps $t : Q_1 \rightarrow Q_0$ (tail) and $h : Q_1 \rightarrow Q_0$ (head). We represent Q as a directed graph with set of vertices Q_0 and directed edges $a : ta \rightarrow ha$ for every $a \in Q_1$.

A representation of Q is a family $V = (V(x), V(a))_{x \in Q_0, a \in Q_1}$ where $V(x)$ is a finite-dimensional \mathbb{R} -vector space for every $x \in Q_0$, and $V(a) : V(ta) \rightarrow V(ha)$ is an \mathbb{R} -linear map for every $a \in Q_1$. A subrepresentation W of V , written as $W \leq V$, is a representation of Q such that $W(x) \leq V(x)$ for every $x \in Q_0$, and $V(a)(W(ta)) \leq W(ha)$ and $W(a)$ is the restriction of $V(a)$ to $W(ta)$ for every arrow $a \in Q_1$. Throughout, we assume that all of our quivers have no oriented cycles.

The dimension vector $\mathbf{dim} V \in \mathbb{N}^{Q_0}$ of a representation V is defined by $\mathbf{dim} V(x) = \dim_{\mathbb{R}} V(x)$ for all $x \in Q_0$. By a dimension vector of Q , we simply mean a \mathbb{N} -valued function

on the set of vertices Q_0 . For two vectors $\sigma, \mathbf{d} \in \mathbb{R}^{Q_0}$, we denote by $\sigma \cdot \mathbf{d}$ the scalar product of the two vectors, i.e. $\sigma \cdot \mathbf{d} = \sum_{x \in Q_0} \sigma(x) \mathbf{d}(x)$.


Definition 3. Let $\sigma \in \mathbb{Z}^{Q_0}$ be a weight of Q . A representation V of Q is said to be σ -semi-stable if

$$\sigma \cdot \mathbf{dim} V = 0 \text{ and } \sigma \cdot \mathbf{dim} W \leq 0$$

for all subrepresentations W of V .

We point out that, as shown by King [Kin94], these linear homogeneous inequalities come from the Hilbert-Mumford's numerical criterion for semi-stability applied to quiver representations.

Remark 4 (Classical RSR from quiver semi-stability). Let $\mathcal{X} = \{v_1, \dots, v_n\} \subseteq \mathbb{R}^D$ be a data set of n vectors in \mathbb{R}^D . We can view \mathcal{X} as the representation



Let us consider the weight $\sigma_0 = (D, \dots, D, -n)$ and note that $\sigma_0 \cdot \mathbf{dim} V_{\mathcal{X}} = 0$. For a subspace $T \leq \mathbb{R}^D$, let $\mathcal{I}_T = \{i \in [n] | v_i \in T\}$, and consider the subrepresentation $W_T \leq V_{\mathcal{X}}$ defined by

$$W_T(x_i) = \begin{cases} \mathbb{R} & i \in \mathcal{I}_T \\ 0 & i \notin \mathcal{I}_T \end{cases}, W_T(y_1) = T, \text{ and } W_T(a_i) = v_i, \forall i \in [n].$$

It turns out that these are all the subrepresentation of $V_{\mathcal{X}}$ that matters when checking whether $V_{\mathcal{X}}$ is σ_0 -semi-stable. Indeed it is easy to check that if W is a subrepresentation of $V_{\mathcal{X}}$ then

$$\sigma_0 \cdot \mathbf{dim} W \leq \sigma_0 \cdot \mathbf{dim} W_T,$$

where $T = W(y_1) \leq \mathbb{R}^D$. Thus we get that $V_{\mathcal{X}}$ is σ_0 -semi-stable if and only if

$$\sigma_0 \cdot \mathbf{dim} W_T \leq 0, \forall T \leq \mathbb{R}^D,$$

which is equivalent to

$$|\mathcal{I}_T| \leq \frac{\mathbf{dim} T}{D} \cdot n, \forall T \leq \mathbb{R}^D.$$

So \mathcal{X} has a lower-dimensional subspace structure precisely when $V_{\mathcal{X}}$ is not σ_0 -semi-stable as a representation of \mathcal{S}_n . By an *RSR solution for \mathcal{X}* , we mean any subspace $T \leq \mathbb{R}^D$ for which

$$|\mathcal{I}_T| > \frac{\mathbf{dim} T}{D} \cdot n.$$

□

Remark 5 (SRSR from quiver semi-stability). Let $\mathcal{X} = (\mathcal{X}^1, \dots, \mathcal{X}^m)$ be a partitioned data set with $\mathcal{X}^j = \{v_1^j, \dots, v_n^j\} \subseteq \mathbb{R}^{d_j}$, $j \in [m]$, and let $V_{\mathcal{X}}$ be the corresponding representation of the complete bipartite quiver $\mathcal{K}_{n,m}$. Recall that in this case the weight σ_0 of $\mathcal{K}_{n,m}$ is defined as

$$\sigma_0(x_i) = D, \forall i \in [n], \text{ and } \sigma(y_j) = -n, \forall j \in [m],$$

where $D = \sum_{j \in [m]} d_j$. Every m -tuple of subspaces (T_1, \dots, T_m) with $T_j \leq \mathbb{R}^{d_j}$, $j \in [m]$, gives rise to a subrepresentation W_T of $V_{\mathcal{X}}$ defined by

$$W_T(x_i) = \begin{cases} \mathbb{R} & \text{if } i \in \mathcal{I}_T \\ 0 & \text{if } i \notin \mathcal{I}_T \end{cases}, \quad W_T(y_j) = T_j, \text{ and } W_T(a_{ij}) = \begin{cases} v_i^j & \text{if } i \in \mathcal{I}_T \\ 0 & \text{if } i \notin \mathcal{I}_T \end{cases}$$

where a_{ij} denotes the arrow from x_i to y_j for all $i \in [n]$ and $j \in [m]$. These are all the subrepresentations of $V_{\mathcal{X}}$ that matter when checking the σ_0 -semi-stability of $V_{\mathcal{X}}$. Indeed, if W is a subrepresentation of $V_{\mathcal{X}}$ then

$$\sigma_0 \cdot \dim W \leq \sigma_0 \cdot W_T,$$

where $T = (W(y_1), \dots, W(y_m))$. Thus $V_{\mathcal{X}}$ is a σ_0 -semi-stable representation of $\mathcal{K}_{n,m}$ if and only if

$$(5) \quad \sigma_0 \cdot \dim W_T = D \cdot |\mathcal{I}_T| - n \cdot \left(\sum_{j=1}^m \dim T_j \right) \leq 0.$$

for every m -tuple of subspaces (T_1, \dots, T_m) with $T_j \leq \mathbb{R}^{d_j}$, $j \in [m]$. So \mathcal{X} has a (simultaneous) lower-dimensional subspace structure precisely when $V_{\mathcal{X}}$ is not σ_0 -semi-stable as a representation of $\mathcal{K}_{n,m}$. \square

Example 6. This is an example of a data set that has an RSR solution T whose projections do not yield an SRSR solution for the data set viewed as a partitioned data set. Consider the data set $\mathcal{X} \subseteq \mathbb{R}^5$ given by

$$\begin{array}{cccc} X_1 & X_2 & X_3 & X_4 \\ \parallel & \parallel & \parallel & \parallel \\ \begin{bmatrix} .1 \\ .2 \\ .1 \\ .2 \\ .1 \end{bmatrix} & \begin{bmatrix} .5 \\ .4 \\ .3 \\ .6 \\ .3 \end{bmatrix} & \begin{bmatrix} 1.3 \\ .9 \\ .3 \\ .5 \\ .1 \end{bmatrix} & \begin{bmatrix} .3 \\ .1 \\ .2 \\ .2 \\ .9 \end{bmatrix} \end{array}$$

Then $T = \text{Span}(X_1, X_2)$ is an RSR solution for \mathcal{X} since $|\mathcal{I}_T| = 2 > \frac{2}{5} \cdot 4$. Let us now view \mathcal{X} as a partitioned data set of type $(4, 2, (2, 3))$. Let $T_1 = \text{Span} \left(\begin{bmatrix} .1 \\ .2 \end{bmatrix}, \begin{bmatrix} .5 \\ .4 \end{bmatrix} \right) = \mathbb{R}^2$ and

$T_2 = \text{Span} \left(\begin{bmatrix} .1 \\ .2 \\ .1 \end{bmatrix}, \begin{bmatrix} .3 \\ .6 \\ .3 \end{bmatrix} \right) \cong \mathbb{R}$ be the projections of T onto \mathbb{R}^2 and \mathbb{R}^3 , respectively. Since

$|\mathcal{I}_{(T_1, T_2)}| = 2 \leq \frac{2+1}{5} \cdot 4$, the pair of subspaces (T_1, T_2) is not an SRSR solution for \mathcal{X} . In fact,

it is immediate to check that the partitioned data set \mathcal{X} has no SRSR solutions, i.e. $V_{\mathcal{X}}$ is σ_0 -semi-stable as a representation of $\mathcal{K}_{4,2}$. \square

Example 7. This is an example where each individual data set \mathcal{X}^j admits an RSR solution but the partitioned data set $\mathcal{X} = (\mathcal{X}^1, \dots, \mathcal{X}^m)$ does not have an SRSR solution. Consider the partitioned data set given by

$$\begin{array}{c} X_1 \\ \parallel \\ X_2 \\ \parallel \\ X_3 \\ \parallel \\ X_4 \\ \parallel \end{array} \left\{ \begin{array}{l} \mathcal{X}^1 \subseteq \mathbb{R}^2 \left\{ \begin{array}{l} \begin{bmatrix} 1 \\ .5 \end{bmatrix} \\ \begin{bmatrix} 2 \\ 1 \end{bmatrix} \\ \begin{bmatrix} 3 \\ 1.5 \end{bmatrix} \\ \begin{bmatrix} .75 \\ .25 \end{bmatrix} \end{array} \right. \\ \mathcal{X}^2 \subseteq \mathbb{R}^3 \left\{ \begin{array}{l} \begin{bmatrix} 3 \\ 1.25 \\ .75 \end{bmatrix} \\ \begin{bmatrix} 1 \\ .25 \\ .25 \end{bmatrix} \\ \begin{bmatrix} 1 \\ .5 \\ 2 \end{bmatrix} \\ \begin{bmatrix} .75 \\ 3.5 \\ 1.5 \\ 1.75 \end{bmatrix} \end{array} \right. \end{array} \right.$$

Then $T_1 = \text{Span} \left(\begin{bmatrix} 1 \\ .5 \end{bmatrix} \right)$ is an RSR solution for \mathcal{X}^1 , and $T_2 = \text{Span} \left(\begin{bmatrix} 3 \\ 1.25 \\ .75 \end{bmatrix}, \begin{bmatrix} 1 \\ .5 \\ 2 \end{bmatrix} \right)$ is a RSR solution for \mathcal{X}^2 . On the other hand, it is easy to check that the partitioned data set $\mathcal{X} = (\mathcal{X}^1, \mathcal{X}^2)$ does not have any SRSR solutions, i.e. $V_{\mathcal{X}}$ is σ_0 -semi-stable as a representation of $\mathcal{K}_{4,2}$. \square

Example 8. This is an example of a partitioned data set $\mathcal{X} = (\mathcal{X}^1, \mathcal{X}^2, \mathcal{X}^3)$ that has an SRSR solution while the data set \mathcal{X}^3 does not have an RSR solution. Consider the partitioned data set given by

$$\begin{array}{c} X_1 \\ \parallel \\ X_2 \\ \parallel \\ X_3 \\ \parallel \\ X_4 \\ \parallel \\ X_5 \\ \parallel \end{array} \left\{ \begin{array}{l} \mathcal{X}^1 \subseteq \mathbb{R}^3 \left\{ \begin{array}{l} \begin{bmatrix} 1 \\ .35 \\ .15 \end{bmatrix} \\ \begin{bmatrix} 2 \\ 3 \\ .15 \end{bmatrix} \\ \begin{bmatrix} 2 \\ 0.3 \\ .12 \end{bmatrix} \\ \begin{bmatrix} 5 \\ 1.75 \\ .75 \end{bmatrix} \\ \begin{bmatrix} 1 \\ 1 \\ .15 \end{bmatrix} \end{array} \right. \\ \mathcal{X}^2 \subseteq \mathbb{R}^3 \left\{ \begin{array}{l} \begin{bmatrix} 2 \\ 2 \\ 2 \end{bmatrix} \\ \begin{bmatrix} 1.3 \\ 1.3 \\ 1.3 \end{bmatrix} \\ \begin{bmatrix} 1 \\ 2.1 \\ 3 \end{bmatrix} \\ \begin{bmatrix} .8 \\ .8 \\ .8 \end{bmatrix} \\ \begin{bmatrix} 1 \\ 2 \\ 0.27 \end{bmatrix} \end{array} \right. \\ \mathcal{X}^3 \subseteq \mathbb{R}^2 \left\{ \begin{array}{l} \begin{bmatrix} 1 \\ .5 \end{bmatrix} \\ \begin{bmatrix} 2 \\ 1.1 \end{bmatrix} \\ \begin{bmatrix} 5 \\ 2.4 \end{bmatrix} \\ \begin{bmatrix} .5 \\ .25 \end{bmatrix} \\ \begin{bmatrix} 3 \\ 1.4 \end{bmatrix} \end{array} \right. \end{array} \right.$$

Let (T_1, T_2, T_3) be the tuple of subspaces defined by

$$T_1 = \left\langle \begin{bmatrix} 1 \\ .35 \\ .15 \end{bmatrix} \right\rangle \leq \mathbb{R}^3, T_2 = \left\langle \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right\rangle \leq \mathbb{R}^3, T_3 = \left\langle \begin{bmatrix} 1 \\ .5 \end{bmatrix} \right\rangle \leq \mathbb{R}^2.$$

Since $\mathcal{I}_T = \{1, 4\}$, we have

$$|\mathcal{I}_T| > \frac{1+1+1}{3+3+2} \cdot 5,$$

thus (T_1, T_2, T_3) is an SRSR solution for \mathcal{X} . On the other hand, it is immediate to check that the data set \mathcal{X}^3 does not have an RSR solution. \square

Example 9. In this example, we show that there exist partitioned data sets that have several SRSR solutions. We also illustrate the optimal SRSR solution (in the sense of (4)) recovered by Algorithms G and P. Consider the partitioned data set of type $(6, 2, (2, 3))$ given by

$$\begin{array}{cccccc} X_1 & X_2 & X_3 & X_4 & X_5 & X_6 \\ \parallel & \parallel & \parallel & \parallel & \parallel & \parallel \\ \begin{bmatrix} 1 \\ 2 \\ 4 \\ 5 \\ 7 \end{bmatrix} & \begin{bmatrix} 2 \\ 4 \\ 1 \\ 1 \\ 1 \end{bmatrix} & \begin{bmatrix} .1 \\ .2 \\ 3 \\ 1 \\ 3 \end{bmatrix} & \begin{bmatrix} .5 \\ 1 \\ 2 \\ 2 \\ 2 \end{bmatrix} & \begin{bmatrix} 3 \\ 6 \\ .3 \\ .1 \\ .3 \end{bmatrix} & \begin{bmatrix} 1 \\ 1 \\ 7 \\ 2 \\ 5 \end{bmatrix} \end{array}$$

This data set has exactly two SRSR solutions. First, consider the tuple of subspaces (S_1, S_2) where

$$S_1 = \left\langle \begin{bmatrix} 1 \\ 2 \end{bmatrix} \right\rangle \text{ and } S_2 = \mathbb{R}^3.$$

Then

$$|\mathcal{I}_{(S_1, S_2)}| = |\{1, 2, 3, 4, 5\}| > \frac{1+3}{2+3} \cdot 6 = \frac{24}{5},$$

so (S_1, S_2) is an SRSR solution.

Next, consider the tuple of subspaces (U_1, U_2) where

$$U_1 = \left\langle \begin{bmatrix} 1 \\ 2 \end{bmatrix} \right\rangle \text{ and } U_2 = \left\langle \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 3 \\ 1 \\ 3 \end{bmatrix} \right\rangle.$$

Then

$$|\mathcal{I}_{(U_1, U_2)}| = |\{2, 3, 4, 5\}| > \frac{1+2}{2+3} \cdot 6 = \frac{18}{5},$$

so (U_1, U_2) is an SRSR solution.

According to (4), (U_1, U_2) is the optimal SRSR solution returned by Algorithms G and P since

$$|\mathcal{I}_{(U_1, U_2)}| - \frac{1+2}{2+3} \cdot 6 = \frac{2}{5} > |\mathcal{I}_{(S_1, S_2)}| - \frac{1+3}{2+3} \cdot 6 = \frac{1}{5}.$$

□

2.2. Semi-stability via capacity of quiver representations. To address the algorithmic aspects of the semi-stability of quiver representations, we recall the notion of a completely positive operator.

Definition 10. (i) Let $\mathcal{A} = \{A_1, \dots, A_\ell\}$ be a collection of real $N \times N$ matrices. The *completely positive operator (cpo)* associated to \mathcal{A} is the linear operator $T_{\mathcal{A}}: \mathbb{R}^{N \times N} \rightarrow \mathbb{R}^{N \times N}$ defined by

$$X \mapsto \sum_{i=1}^n A_i X A_i^T$$

The matrices A_1, \dots, A_ℓ are called the *Kraus operators* of $T_{\mathcal{A}}$. When no confusion arises, we simply write T instead of $T_{\mathcal{A}}$.

(ii) The *capacity* of a cpo T is defined as

$$\text{cap } T = \inf \{ \det(T(Y)) \mid Y \text{ is positive definite with } \det Y = 1 \}$$

Next we show how to construct a cpo from our representation $V_{\mathcal{X}}$ of the complete bipartite quiver $\mathcal{K}_{n,m}$. Let $M = mn$, $N = nD$, and set

$$\mathcal{S}_{\sigma_0} := \left\{ (i, j, q, r) \left| \begin{array}{l} i \in [n], j \in [m], \\ (j-1)n < q \leq jn, \\ (i-1)D < r \leq iD \end{array} \right. \right\}.$$

For each index $(i, j, q, r) \in \mathcal{S}_{\sigma_0}$, let $V_{q,r}^{i,j}$ be the $M \times N$ block matrix (of size $N \times N$) whose (q, r) -block-entry is $v_i^j \in \mathbb{R}^{d_j \times 1}$, and all other block entries are zero matrices of appropriate size. Note that $V_{q,r}^{i,j}$ is a certain blow-up of the vector v_i^j .

Definition 11. The cpo associated to the representation $V_{\mathcal{X}}$ is the cpo $T_{\mathcal{X}}$ associated to the collection of matrices

$$\mathcal{A}_{\mathcal{X}} := \{V_{q,r}^{i,j} \mid (i, j, q, r) \in \mathcal{S}_{\sigma_0}\}.$$

We also refer to $T_{\mathcal{X}}$ as the cpo associated to the partitioned data set $\mathcal{X} = (\mathcal{X}^1, \dots, \mathcal{X}^m)$.

Remark 12. We point out $\mathcal{A}_{\mathcal{X}}$ is a linearly independent set of matrices in $\mathbb{R}^{N \times N}$ with $|\mathcal{A}_{\mathcal{X}}| = n^2 m D$. \square

Theorem 13. (see [CD19, Theorem 1]) Let $V_{\mathcal{X}}$ be the representation of the complete bipartite quiver $\mathcal{K}_{n,m}$ defined by the partitioned data set $\mathcal{X} = (\mathcal{X}^1, \dots, \mathcal{X}^m)$. Then $\text{cap}(T_{\mathcal{X}}) = 0$ if and only if $V_{\mathcal{X}}$ is not σ_0 -semi-stable.

Remark 14. The construction above, along with Theorem 13 can be formulated for representations of arbitrary bipartite quivers. In the general setting, the cpo above is known as the *Brascamp-Leib operator*, and has deep connections to the celebrated Brascamp-Lieb inequality from Harmonic Analysis. \square

3. ALGORITHM G

According to Theorem 13, a partitioned data set $\mathcal{X} = (\mathcal{X}^1, \dots, \mathcal{X}^m)$ has a lower-dimensional subspace structure if and only if the capacity of $T_{\mathcal{X}}$ is zero. To check whether or not the capacity is zero, we will make use of an algorithm originally published in [Gur04], known as Algorithm G. In order to state Algorithm G, we first recall some definitions.

Definition 15. Let $T_{\mathcal{A}}$ be a cpo with Kraus operators $\mathcal{A} = \{A_1, \dots, A_{\ell}\} \subseteq \mathbb{R}^{N \times N}$.

(i) The *dual* of $T_{\mathcal{A}}$, denoted $T_{\mathcal{A}}^*$, is the linear operator defined by

$$X \mapsto \sum_{i=1}^{\ell} A_i^T X A_i$$

(ii) We call $T_{\mathcal{A}}$ (and $T_{\mathcal{A}}^*$) *doubly stochastic* if

$$T_{\mathcal{A}}(I) = T_{\mathcal{A}}^*(I) = I.$$

(iii) The *distance of $T_{\mathcal{A}}$ to double stochasticity* is defined by:

$$\text{ds}(T_{\mathcal{A}}) := \text{tr}((T_{\mathcal{A}}(I) - I)^2) + \text{tr}((T_{\mathcal{A}}^*(I) - I)^2)$$

The next definition gives two specific instances of a more general process known as *operator scaling*. It leads to an efficient algorithm for testing for the positivity of the capacity of a cpo.

Definition 16. Let T be a cpo with Kraus operators $\{A_1, \dots, A_\ell\} \subseteq \mathbb{R}^{N \times N}$ such that $T(I)$ and $T^*(I)$ are both invertible.

(i) The *right normalization of T* , T_R , is the cpo with Kraus operators

$$A_1 \cdot T^*(I)^{-1/2}, \dots, A_\ell \cdot T^*(I)^{-1/2}.$$

(ii) The *left normalization of T* , T_L , is the cpo with Kraus operators

$$T(I)^{-1/2}A_1, \dots, T(I)^{-1/2}A_\ell.$$

We point out that $T_R^*(I) = T_L(I) = I$.

We are now ready to state Algorithm G, adapted slightly to our context from [GGOW15, page 47]. We point out that their version is concerned with estimating capacity within a given multiplicative error, whereas we are simply concerned with checking whether or not the capacity is positive. It has been proved that Algorithm G runs in (deterministic) polynomial time.

Algorithm G

Input: A cpo T whose Kraus operators (of size $N \times N$) lie over \mathbb{Z} and M the maximum magnitude of each entry.

Output: The positivity of $\text{cap}(T)$.

- (1) If either $T(I)$ or $T^*(I)$ is singular then output $\text{cap}(T) = 0$. Otherwise, continue.
- (2) For each integer $j \geq 1$, perform (alternatively) right and left normalization on $T = T_0$ with T_j the operator after j -steps.
- (3) If $\text{ds}(T_j) \leq \frac{1}{4N^3}$ after j steps, output $\text{cap}(T) > 0$.
- (4) If $\text{ds}(T_j) > \frac{1}{4N^3}$ after $4N^3(1 + 10N^2 \log(MN))$ steps, output $\text{cap}(T) = 0$.

Even though the algorithm is stated over \mathbb{Z} , it is valid over \mathbb{Q} , since multiplying the Kraus operators by a (non-zero) scalar does not affect the positivity of the capacity. While we have worked over \mathbb{R} , almost all data will be rational points, so there is no loss in working over \mathbb{Q} . For proof of correctness, we refer the reader to [GGOW15]. See also [Gur04].

Remark 17. Algorithm G has been developed by Gurvits in [Gur04] to solve Edmond's Problem for certain classes of matrices. Edmond's Problem, originally posed in [Edm67], asks whether or not the span of a collection of (square) matrices contains a non-singular

matrix. In [CK20], we show how Edmond's problem can be expressed via orbit semi-groups of quiver representations. \square

4. SHRUNK SUBSPACES AND RANK DECREASING WITNESSES

Given a partitioned data set \mathcal{X} , using Theorem 13 and Algorithm G, we can check in deterministic polynomial time whether \mathcal{X} has a lower-dimensional subspace structure (equivalently, whether $\text{cap}(T_{\mathcal{X}}) = 0$). When this is the case, we still need to find a way to recover this lower-dimensional subspace structure. In order to do this, we consider shrunk subspaces and second Wong sequences which we explain below.

Definition 18. Let $\mathcal{A} = \{A_1, \dots, A_\ell\}$ an ℓ -tuple of $N \times N$ real matrices and $T_{\mathcal{A}}$ the cpo with Kraus operators A_1, \dots, A_ℓ . Let c be a non-negative integer. A subspace $U \leq \mathbb{R}^N$ is called a c -shrunk subspace for \mathcal{A} (or $T_{\mathcal{A}}$) if

$$\dim U - \dim \left(\sum_{i=1}^{\ell} A_i(U) \right) \geq c.$$

We say that U is a *shrunk subspace* for \mathcal{A} if U is a c -shrunk subspace for some $c > 0$.

We point out that shrunk subspaces have also been studied by D. Eisenbud and J. Harris in [EH88], though they refer to them as *compression spaces*.

Definition 19. A cpo T on $\mathbb{R}^{N \times N}$ is said to be *rank-decreasing* if there exists an $N \times N$ positive semi-definite matrix Y such that $\text{rank } T(Y) < \text{rank } Y$. In this case, we call Y a *rank-decreasing witness* for T .

Proposition 20. Let \mathcal{X} be a partitioned data set, $\mathcal{A}_{\mathcal{X}} = \{A_1, \dots, A_\ell\}$ the Kraus operators associated to \mathcal{X} , and $T_{\mathcal{X}}$ the corresponding cpo. Then the following statements are equivalent:

- (i) $V_{\mathcal{X}}$ is not σ_0 -semi-stable;
- (ii) $\text{cap}(T_{\mathcal{X}}) = 0$;
- (iii) $T_{\mathcal{X}}$ is rank-decreasing;
- (iv) $T_{\mathcal{X}}$ has a shrunk subspace.

Proof. By Theorem 13, we know that (i) and (ii) are equivalent. The equivalence of (ii) and (iii) follows from [Gur04, Lemma 4.5].

For (iii) implies (iv), let Y be a rank-decreasing witness for $T_{\mathcal{X}}$ of rank r . Then we can write it as

$$Y = \sum_{j=1}^r \lambda_j u_j \cdot u_j^T$$

where $\lambda_1, \dots, \lambda_r > 0$ and the vectors u_1, \dots, u_r form an orthonormal set. Setting $U := \text{Span}(u_1, \dots, u_r)$, we get that $\dim U = \text{rank}(Y)$, and

$$\begin{aligned} \dim \left(\sum_{i=1}^{\ell} A_i(U) \right) &= \dim \text{Span} (A_i u_j \mid i \in [\ell], j \in [r]) \\ &= \text{rank} \left(\sum_{i,j} A_i u_j (A_i u_j)^T \right) = \text{rank } T_{\mathcal{X}}(Y). \end{aligned}$$

Thus

$$\dim U - \dim \left(\sum_{i=1}^{\ell} A_i(U) \right) = \text{rank}(Y) - \text{rank}(T_{\mathcal{X}}(Y)) > 0,$$

i.e. U is a shrunk subspace for $T_{\mathcal{X}}$.

For (iv) implies (iii), if U' is a shrunk subspace for $T_{\mathcal{X}}$, let Y' be the matrix of orthogonal projection onto U' . Write $Y' = Q \cdot Q^T$ where Q is a matrix whose columns form an orthonormal basis for U' . Then $\text{rank}(Y') = \dim U'$, and since $T_{\mathcal{X}}(Y') = \sum_{i=1}^{\ell} (A_i Q)(A_i Q)^T$, it follows that

$$\text{Im}(T_{\mathcal{X}}(Y')) \leq \sum_{i=1}^{\ell} \text{Im}(A_i Q) = \sum_{i=1}^{\ell} A_i(U')$$

which gives

$$\text{rank}(T_{\mathcal{X}}(Y')) \leq \dim \left(\sum_{i=1}^{\ell} A_i(U') \right) < \dim U' = \text{rank}(Y').$$

Thus

$$\text{rank}(Y') - \text{rank}(T_{\mathcal{X}}(Y')) \geq \dim U' - \dim \left(\sum_{i=1}^{\ell} A_i(U') \right) > 0,$$

i.e. Y' is a rank-decreasing witness for $T_{\mathcal{X}}$. □

We will use rank-decreasing witnesses for $T_{\mathcal{X}}$ to recover lower-dimensional subspace structures in \mathcal{X} (see Lemma 26). The proof above shows that finding a rank-decreasing witness for $T_{\mathcal{X}}$ is equivalent to finding a shrunk-subspace. In order to compute shrunk-subspaces, we recall the notion of *second Wong sequences*.

Definition 21. Let $\mathcal{A} = \{A_1, \dots, A_{\ell}\} \subset \mathbb{R}^{N \times N}$. For $B \in \mathbb{R}^{N \times N}$, the *second Wong sequence* associated to (\mathcal{A}, B) is the sequence of subspaces

$$\begin{aligned} W_0 &= 0 \\ W_1 &= \mathcal{A}(B^{-1}(W_0)) \\ W_2 &= \mathcal{A}(B^{-1}(W_1)) \\ &\vdots \\ W_k &= \mathcal{A}(B^{-1}(W_{k-1})) \end{aligned}$$

where for each j , $\mathcal{A}(B^{-1}(W_j)) = \langle A_1(B^{-1}(W_j)), \dots, A_{\ell}(B^{-1}(W_j)) \rangle$ and $B^{-1}(W_j)$ is the preimage of W_j under B .

It is not difficult to see that the W_i lie in a chain, i.e.

$$W_0 \leq W_1 \leq W_2 \leq \dots$$

Since we are working over finite-dimensional vector spaces, the second Wong sequence must converge, i.e. there exists a positive integer $L \leq N$ such that

$$W_0 \leq W_1 \leq \dots \leq W_{L-1} \leq W_L = W_{L+1} = \dots$$

W_L is called the *limit* of the associated second Wong sequence and is denoted by W^* .

The next result shows how shrunk subspaces can be computed from second Wong sequences.

Theorem 22. [IKQS15, Theorem 1 & Lemma 9] *Let $\mathcal{A} = \{A_1, \dots, A_\ell\} \subset \mathbb{R}^{N \times N}$ and $B \in \langle \mathcal{A} \rangle$. Then the following statements are equivalent:*

- (1) $W^* \leq \text{Im } B$;
- (2) $B^{-1}(W^*)$ is a $\text{corank}(B)$ -shrunk subspace.

If the operators A_1, \dots, A_ℓ are assumed to be of rank one then (1) and (2) are further equivalent to:

- (3) B has maximal rank among the matrices in $\langle \mathcal{A} \rangle$.

In fact, under the assumption that A_1, \dots, A_ℓ are of rank one, there exists a deterministic polynomial time algorithm to find a maximal rank matrix B in $\langle \mathcal{A} \rangle$.

Remark 23. We point out that in general (1) or (2) in Theorem 22 implies that the matrix B is of maximal rank.

In the context of SRSR, since the Kraus operators for $T_{\mathcal{X}}$ are all of rank one, Theorem 22 tells us that computing a shrunk subspace comes down to constructing a matrix $B \in \langle \mathcal{A}_{\mathcal{X}} \rangle$ of maximal rank. Furthermore it also provides a deterministic polynomial time algorithm for finding such a maximal rank matrix. However, from a practical standpoint, it is rather difficult to implement. In what follows, we are going to describe an algorithm that constructs a shrunk subspace with high probability.

Proposition 24. *Let $\mathcal{A}_{\mathcal{X}} = \{A_1, \dots, A_\ell\}$ be the set of Kraus operators for $T_{\mathcal{X}}$ where $\ell = n^2 m D$ and let $S \subseteq \mathbb{R}$ be any finite set of size at least $\frac{2n^2 D^2}{\epsilon}$ with $\epsilon > 0$.*

If $\alpha_1, \dots, \alpha_\ell$ are chosen independently and uniformly at random from S and $B := \alpha_1 A_1 + \dots + \alpha_\ell A_\ell \in \langle \mathcal{A}_{\mathcal{X}} \rangle$ then

$$\mathbb{P}(B^{-1}(W^*) \text{ is a } \text{corank}(B)\text{-shrunk-subspace for } \mathcal{A}_{\mathcal{X}}) \geq 1 - \epsilon.$$

Proof. Let \mathcal{V} be the subspace of $\mathbb{R}^{N \times N}$ spanned by $\mathcal{A}_{\mathcal{X}}$. Since $\mathcal{A}_{\mathcal{X}}$ is a linearly independent set of matrices, it is a basis for \mathcal{V} , and thus we can identify \mathcal{V} with \mathbb{R}^ℓ .

Let us denote by r the maximal rank of the matrices in \mathcal{V} , and let $P \in \mathbb{R}[t_1, \dots, t_\ell]$ be the non-zero polynomial defined by

$$P(t_1, \dots, t_\ell) = \sum \det(A')^2$$

where the sum is over all $r \times r$ submatrices A' of $t_1 A_1 + \dots + t_\ell A_\ell$. Note that the total degree of this polynomial is at most $2r^2 \leq 2N^2 = 2n^2 D^2$.

Then, for any $(\alpha_1, \dots, \alpha_\ell) \in \mathbb{R}^\ell$, we have that

$$\alpha_1 A_1 + \dots + \alpha_\ell A_\ell \in \mathcal{V} \text{ is not of maximal rank} \iff P(\alpha_1, \dots, \alpha_\ell) = 0.$$

At this point we can use the Schwartz-Zippel Lemma [Sch80, Zip79] to conclude that

$$(6) \quad \mathbb{P}(B \text{ is not of maximal rank}) \leq \frac{\deg(P)}{|S|} \leq \epsilon$$

It now follows from Theorem 22 and (6) that

$$\mathbb{P}(B^{-1}(W^*) \text{ is a } \text{corank}(B)\text{-shrunk-subspace for } \mathcal{A}_{\mathcal{X}}) = 1 - \mathbb{P} \left(\begin{array}{c} B \text{ is not of} \\ \text{maximal} \\ \text{rank} \end{array} \right) \geq 1 - \epsilon,$$

and this finishes the proof. \square

We are thus led to the following simple, efficient probabilistic algorithm for computing shrunk-subspaces for \mathcal{A}_X . This algorithm returns an output with high probability by Proposition 24. In what follows, we denote by $\text{corank}(\mathcal{X})$ the corank of a matrix $B \in \langle \mathcal{A}_X \rangle$ of maximal rank.

Algorithm P

Input: Kraus operators $\mathcal{A}_X = \{A_1, \dots, A_\ell\}$ with $\ell = n^2 m D$, and a sample set $S \subseteq \mathbb{R}$ of size at least $\frac{2n^2 D^2}{\epsilon}$.

Output: $\text{corank}(\mathcal{X})$ -shrunk-subspace for \mathcal{A}_X

- (1) Choose $\alpha_1, \dots, \alpha_\ell$ independently and randomly distributed from S , and set $B := \alpha_1 A_1 + \dots + \alpha_\ell A_\ell$.
- (2) Compute W^* , the limit of the second Wong sequence associated to (\mathcal{A}_X, B) .
- (3) If $W^* \leq \text{Im } B$, output $B^{-1}(W^*)$ as a $\text{corank}(\mathcal{X})$ -shrunk subspace for \mathcal{A}_X . Otherwise, return to (1).

Remark 25. (1) We point out that Proposition 24 can be restated to say that for a randomly chosen matrix B from $\mathcal{V} := \langle \mathcal{A}_X \rangle$, the subspace $B^{-1}(W^*)$ is a $\text{corank}(B)$ -shrunk-subspace for \mathcal{A}_X with probability one. Indeed, as in the proof of Proposition 24, we have the non-zero polynomial $P \in \mathbb{R}[t_1, \dots, t_\ell]$ such that for any $(\alpha_1, \dots, \alpha_\ell) \in \mathbb{R}^\ell$,

$$\alpha_1 A_1 + \dots + \alpha_\ell A_\ell \in \mathcal{V} \text{ is not of maximal rank} \iff P(\alpha_1, \dots, \alpha_\ell) = 0.$$

So the subset of \mathcal{V} consisting of all those matrices which are not of maximal rank is identified with the zero set of P in \mathbb{R}^ℓ . This zero set, denoted by $\mathbb{V}(P)$, has Lebesgue measure zero in \mathbb{R}^ℓ .

Now, let $(\alpha_1, \dots, \alpha_\ell) \in \mathbb{R}^\ell$ be a coefficient vector drawn from any probability distribution absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^ℓ (e.g. Gaussian distribution). If $B := \alpha_1 A_1 + \dots + \alpha_\ell A_\ell \in \mathcal{V}$ then

$$\mathbb{P}(B \text{ is not of maximal rank}) = \mathbb{P}((\alpha_1, \dots, \alpha_\ell) \in \mathbb{V}(P)) = 0,$$

proving that a randomly chosen matrix B from $\langle \mathcal{A}_X \rangle$ has maximal rank with probability one. This fact combined with Theorem 22 yields the desired claim.

(2) Let \mathcal{V} be an arbitrary matrix space of dimension $\ell \geq 1$. Then, using the same arguments as above, one can prove the following statements.

- (a) There exists a random process (based on the Schwartz-Zippel Lemma) that selects a random matrix from \mathcal{V} of maximal rank with high probability (see the proof of Proposition 24).
- (b) There exists a random process (based on absolutely continuous probability distributions) that selects a random matrix from \mathcal{V} of maximal rank with probability one.

From an implementation standpoint, the random process in (a) is more practical than that in (b). □

5. RECOVERING LOWER-DIMENSIONAL SUBSPACE STRUCTURES

In this section, $\mathcal{X} = (\mathcal{X}^1, \dots, \mathcal{X}^m)$ is a partitioned data set with $\mathcal{X}^j = \{v_1^j, \dots, v_n^j\} \subseteq \mathbb{R}^{d_j}$, $j \in [m]$. Let $V_{\mathcal{X}}$ be the representation of the complete bipartite quiver $\mathcal{K}_{n,m}$ corresponding to \mathcal{X} . We denote by $T_{\mathcal{X}}$ the cpo associated to the representation $V_{\mathcal{X}}$. We will be working with the weight σ_0 of $\mathcal{K}_{n,m}$ defined by

$$\sigma_0(x_i) = D, \forall i \in [n], \text{ and } \sigma(y_j) = -n, \forall j \in [m],$$

where $D = \sum_{j \in [m]} d_j$.

We begin by explaining how to construct a subrepresentation of $V_{\mathcal{X}}$ from a given rank-decreasing witness for $T_{\mathcal{X}}$.

Lemma 26. *Let $V_{\mathcal{X}}$ be the representation of $\mathcal{K}_{n,m}$ associated to the partitioned data set \mathcal{X} such that $\text{cap}(T_{\mathcal{X}}) = 0$. Let $Y = (y_{q,r})_{q,r}$ be an $N \times N$ rank-decreasing witness for $T_{\mathcal{X}}$ where $N = nD$. Set*

$$\mathcal{I} = \{i \in [n] \mid y_{r,r} \neq 0 \text{ for some } r \in \mathcal{I}_i\},$$

where $\mathcal{I}_i = [(i-1)D + 1, iD]$ for every $i \in [n]$. Let W be the subrepresentation of $V_{\mathcal{X}}$ defined by

$$W(x_i) = \begin{cases} \mathbf{0} & i \notin \mathcal{I} \\ \mathbb{R} & i \in \mathcal{I} \end{cases} \text{ for } i \in [n], \text{ and } W(y_j) = \text{Span}(v_i^j \mid i \in \mathcal{I}) \text{ for } j \in [m].$$

Then

$$\sigma_0 \cdot \mathbf{dim} W \geq \text{rank}(Y) - \text{rank}(T_{\mathcal{X}}(Y)) > 0.$$

Proof. It is clear that

$$(7) \quad \text{rank}(Y) \leq \sum_{i \in [n]} \sum_{r \in \mathcal{I}_i} \text{rank}(y_{r,r}) \leq \sum_{i \in [n]} D \cdot \dim W(x_i).$$

On the other hand, $T_{\mathcal{X}}(Y)$ turns out to be an $(mn) \times (mn)$ block diagonal matrix whose (q, q) -block entry is

$$\sum_{i=1}^n v_i^j \cdot \left(\sum_{r \in \mathcal{I}_i} y_{r,r} \right) \cdot (v_i^j)^T$$

for every $q \in [(j-1)n+1, jn]$ and $j \in [m]$. Thus we get that

$$\begin{aligned}
(8) \quad \text{rank}(T_{\mathcal{X}}(Y)) &= \sum_{j=1}^m n \text{rank} \left(\sum_{i=1}^n \left(v_i^j \left(\sum_{r \in \mathcal{I}_i} y_{r,r} \right) (v_i^j)^T \right) \right) \\
&= \sum_{j=1}^m n \text{rank} \left(\sum_{i,r} \sqrt{y_{r,r}} v_i^j (\sqrt{y_{r,r}} v_i^j)^T \right) \\
&= \sum_{j=1}^m n \dim \text{Span} (\sqrt{y_{r,r}} v_i^j \mid i \in [n], r \in \mathcal{I}_i) \\
&= - \sum_{j=1}^m \sigma_0(y_j) \dim W(y_j)
\end{aligned}$$

From (7) and (8), it follows that

$$(9) \quad \sum_{i \in [n]} \sigma_0(x_i) \dim W(x_i) \geq \text{rank}(Y) > \text{rank}(T_{\mathcal{X}}(Y)) = - \sum_{j \in [m]} \sigma_0(y_j) \dim W(y_j),$$

which proves our claim. \square

We are now ready to prove our main result.

Proof of Theorem 1. Let \mathcal{X} be a partitioned data set, $V_{\mathcal{X}}$ the associated representation of $\mathcal{K}_{n,m}$, and $T_{\mathcal{X}}$ the associated cpo. Since $\mathcal{A}_{\mathcal{X}}$ consists of matrices of rank one, $\mathcal{A}_{\mathcal{X}}$ has the so-called Edmonds-Rado property (see [Lov89]), meaning that $\text{cap}(T_{\mathcal{X}}) > 0$ if and only if $\text{corank}(\mathcal{X}) = 0$. (Recall that $\text{corank}(\mathcal{X})$ denotes the corank of a matrix $B \in \langle \mathcal{A}_{\mathcal{X}} \rangle$ of maximal rank.) It follows from Theorem 13 and Proposition 20 that the following statements are equivalent:

- \mathcal{X} has an SRSR solution;
- $\text{cap}(T_{\mathcal{X}}) = 0$;
- $\text{corank}(\mathcal{X}) > 0$.

Thus, to determine whether or not \mathcal{X} has a lower-dimensional subspace structure, we can use either Algorithm G, which runs in deterministic polynomial time, or the probabilistic Algorithm P.

Assume now that \mathcal{X} has a lower-dimensional subspace structure. Then we can compute a $\text{corank}(\mathcal{X})$ -shrunk subspace U for $T_{\mathcal{X}}$ by the deterministic polynomial time algorithm in Theorem 22 or by the probabilistic Algorithm P. More precisely, the output U is of the form $B^{-1}(W^*)$ where $B \in \langle \mathcal{A}_{\mathcal{X}} \rangle$ is of maximal rank and W^* is the limit of the second Wong sequence associated to $(\mathcal{A}_{\mathcal{X}}, B)$.

Next, let $Y = Q \cdot Q^T$ where Q is a matrix whose columns form an orthonormal basis for U . Then, according to Proposition 20, Y is a rank-decreasing witness for $T_{\mathcal{X}}$. By Lemma 26, we can finally use Y to find a subspace $W \leq V_{\mathcal{X}}$ such that

$$(10) \quad \sigma_0 \cdot \dim W \geq \text{rank}(Y) - \text{rank}(T_{\mathcal{X}}(Y)) \geq \text{corank}(\mathcal{X}) > 0.$$

For each $j \in [m]$, define $T_j = W(y_j)$. Then (10) yields

$$|\mathcal{I}_T| > \frac{\sum_{j \in [m]} \dim T_j}{D} n,$$

and so (T_1, \dots, T_m) is a SRSR solution.

It remains to show that (T_1, \dots, T_m) is an optimal structure. So let (T'_1, \dots, T'_m) be any SRSR solution for \mathcal{X} , and let W' be the corresponding subrepresentation of $V_{\mathcal{X}}$ (see Remark 5). Then there exists a c' -shrunk subspace for $T_{\mathcal{X}}$ with $c' = \sigma_0 \cdot \dim W'$. Indeed, let Y' be the $N \times N$ diagonal matrix whose r th diagonal entry is one if $r \in [(i-1)D+1, iD]$ with $i \in \mathcal{I}_{T'}$, and zero otherwise. It is easy to check that Y' is a rank-decreasing witness for $T_{\mathcal{X}}$ with

$$\text{rank}(Y') - \text{rank}(T_{\mathcal{X}}(Y')) \geq \sigma_0 \cdot \dim(W').$$

From Proposition 20 it follows that if U' denotes the shrunk subspace corresponding to Y' then

$$(11) \quad \dim U' - \dim \left(\sum_{i=1}^l A_i(U') \right) = \text{rank}(Y') - \text{rank}(T_{\mathcal{X}}(Y')) \geq \sigma_0 \cdot \dim(W'),$$

where $\mathcal{A}_{\mathcal{X}} = \{A_1, \dots, A_l\}$, i.e. U' is a $\sigma_0 \cdot \dim(W')$ -shrunk subspace.

On the other hand, we know that

$$(12) \quad \text{corank}(\mathcal{A}_{\mathcal{X}}) \geq \max\{c \in \mathbb{N} \mid \text{there exists a } c\text{-shrunk subspace for } \mathcal{A}_{\mathcal{X}}\}$$

(This inequality holds for any collection of matrices \mathcal{A} with the maximum on the right hand side also known as the discrepancy of \mathcal{A} ; see for example [IKQS15].) Using (10), (11), and (12) we get that

$$\begin{aligned} \sigma_0 \cdot \dim W &\geq \sigma_0 \cdot \dim W' \\ \Downarrow \\ D \cdot |\mathcal{I}_T| - n \cdot \left(\sum_{j=1}^m \dim T_j \right) &\geq D \cdot |\mathcal{I}_{T'}| - n \cdot \left(\sum_{j=1}^m \dim T'_j \right) \\ \Downarrow \\ |\mathcal{I}_T| - \frac{\sum_{j=1}^m \dim T_j}{D} \cdot n &\geq |\mathcal{I}_{T'}| - \frac{\sum_{j=1}^m \dim T'_j}{D} \cdot n. \end{aligned}$$

This finishes the proof. □

6. SEMI-STABILITY OF QUIVER REPRESENTATIONS

In this section, we address the general problem of effectively deciding whether a quiver representation is semi-stable (with respect to a weight) and, if it is not, finding a subrepresentation that certifies that the representation is not semi-stable.

Let $Q = (Q_0, Q_1, t, h)$ be an arbitrary connected bipartite quiver (not necessarily complete). This means that Q_0 is the disjoint union of two subsets $Q_0^+ = \{x_1, \dots, x_n\}$ and $Q_0^- = \{y_1, \dots, y_m\}$, and all arrows in Q go from Q_0^+ to Q_0^- . (We do allow multiple arrows between the vertices of Q .)

Let $\mathbf{d} \in \mathbb{N}^{Q_0}$ be a dimension vector of Q . The representation space of \mathbf{d} -dimensional representations of Q is the vector space

$$\text{rep}(Q, \mathbf{d}) := \prod_{a \in Q_1} \mathbb{R}^{\mathbf{d}(ha) \times \mathbf{d}(ta)}.$$

It is clear that any $V \in \text{rep}(Q, \mathbf{d})$ gives rise to a representation of Q of dimension vector \mathbf{d} and vice versa.

Let $\sigma \in \mathbb{Z}^{Q_0}$ be a weight of Q such that σ is positive on Q_0^+ , negative of Q_0^- , and

$$(13) \quad \sigma \cdot \mathbf{d} = \sum_{z \in Q_0} \sigma(z) \mathbf{d}(z) = 0.$$

Define

$$\sigma_+(x_i) = \sigma(x_i), \forall i \in [n], \text{ and } \sigma_-(y_j) = -\sigma(y_j), \forall j \in [m].$$

Then (13) is equivalent to

$$N := \sum_{i=1}^n \sigma_+(x_i) \mathbf{d}(x_i) = \sum_{j=1}^m \sigma_-(y_j) \mathbf{d}(y_j).$$

Let $M := \sum_{j=1}^m \sigma_-(y_j)$ and $M' := \sum_{i=1}^n \sigma_+(x_i)$. For each $j \in [m]$ and $i \in [n]$, define

$$\mathcal{I}_j^- := \{q \in \mathbb{Z} \mid \sum_{k=1}^{j-1} \sigma_-(y_k) < q \leq \sum_{k=1}^j \sigma_-(y_k)\},$$

and

$$\mathcal{I}_i^+ := \{r \in \mathbb{Z} \mid \sum_{k=1}^{i-1} \sigma_+(x_k) < r \leq \sum_{k=1}^i \sigma_+(x_k)\}.$$

In what follows, we consider $M \times M'$ block matrices of size $N \times N$ such that for any two indices $q \in \mathcal{I}_j^-$ and $r \in \mathcal{I}_i^+$, the (q, r) -block-entry is a matrix of size $\mathbf{d}(y_j) \times \mathbf{d}(x_i)$. Set

$$\mathcal{S}_\sigma := \{(i, j, a, q, r) \mid i \in [n], j \in [m], a \in \mathcal{A}_{i,j}, q \in \mathcal{I}_j^-, r \in \mathcal{I}_i^+\},$$

where $\mathcal{A}_{i,j}$ denotes the set of all arrows in Q from x_i to y_j , $i \in [n]$ and $j \in [m]$.

Let $V \in \text{rep}(Q, \mathbf{d})$ be a \mathbf{d} -dimensional representation. For each index $(i, j, a, q, r) \in \mathcal{S}_\sigma$, let $V_{q,r}^{i,j,a}$ be the $M \times M'$ block matrix of size $N \times N$ whose (q, r) -block-entry is $V(a) \in \mathbb{R}^{\mathbf{d}(y_j) \times \mathbf{d}(x_i)}$, and all other entries are zero. We denote by $\mathcal{A}_{V,\sigma}$ the set of all matrices $V_{q,r}^{i,j,a}$ with $(i, j, a, q, r) \in \mathcal{S}_\sigma$.

The Brascamp-Lieb operator associated to the quiver datum (V, σ) , defined in [CD19], is the cpo $T_{V,\sigma}$ with Kraus operators $V_{q,r}^{i,j,a}$, $(i, j, a, q, r) \in \mathcal{S}_\sigma$, i.e.

$$T_{V,\sigma}(X) := \sum_{(i,j,a,q,r)} V_{q,r}^{i,j,a} \cdot X \cdot (V_{q,r}^{i,j,a})^T, \forall X \in \mathbb{R}^{N \times N}.$$

It has been proved in [CD19, Theorem 1] that

$$V \text{ is } \sigma\text{-semi-stable} \iff \text{cap}(T_{V,\sigma}) > 0.$$

Consequently, as indicated in [CD19, Section 1.2], one can use Algorithm G to test whether V is σ -semi-stable or not.

Remark 27. In the context of arbitrary quivers and dimension vectors, when it comes to Algorithm P, even if B is chosen to have maximal rank, it is not guaranteed that $W^* \leq \text{Im } B$. This is because for arbitrary dimension vectors, $\mathcal{A}_{V,\sigma}$ may contain matrices that are not of rank one. Thus Algorithm P might not produce an output.

The following lemma allows us to go back and forth between subrepresentations on the one hand, and rank-decreasing witnesses and c -shrunk subspaces on the other. It plays a key role in the proof of Theorem 2

Lemma 28. Let Q be a bipartite quiver and (V, σ) a quiver datum as above. Let $\mathcal{A}_{V,\sigma} = \{A_1, \dots, A_L\}$ be the set of Kraus operators associated to (V, σ) .

- (1) Let W be a subrepresentation of V . For each $i \in [n]$, after choosing an orthonormal basis $u_1^i, \dots, u_{\mathbf{f}(i)}^i$ for $W(x_i) \leq \mathbb{R}^{\mathbf{d}(x_i)}$, define

$$Y_r := \sum_{l=1}^{\mathbf{f}(i)} u_l^i \cdot (u_l^i)^T, \forall r \in \mathcal{I}_i^+.$$

Then $Y := \bigoplus_{i \in [n]} \bigoplus_{r \in \mathcal{I}_i^+} Y_r \in \mathbb{R}^{N \times N}$ is a positive semi-definite matrix such that

$$\text{rank}(Y) - \text{rank}(T_{V,\sigma}(Y)) \geq \sigma \cdot \mathbf{dim } W.$$

Moreover, if $U \leq \mathbb{R}^N$ is the subspace associated to Y (see Proposition 20) then

$$\text{dim } U - \text{dim} \left(\sum_{i=1}^L A_i(U) \right) \geq \sigma \cdot \mathbf{dim } W.$$

- (2) Let $Y \in \mathbb{R}^{N \times N}$ be a positive semi-definite matrix. Viewing Y as an $M' \times M'$ block diagonal matrix, denote its block diagonal entries by $Y_r \in \mathbb{R}^{\mathbf{d}(x_i) \times \mathbf{d}(x_i)}$ with $r \in \mathcal{I}_i^+$ and $i \in [n]$. For each such r and i , write

$$Y_r = \sum_{l=1}^{d_{i,r}} \lambda_l^{i,r} u_l^{i,r} \cdot (u_l^{i,r})^T,$$

where $d_{i,r}$ is the rank of Y_r , the $\lambda_l^{i,r}$ are positive scalars, and the $u_l^{i,r}$ form an orthonormal set of vectors in $\mathbb{R}^{\mathbf{d}(x_i)}$. Define

$$W(x_i) := \text{Span} \left(\sqrt{\lambda_l^{i,r}} u_l^{i,r} \mid r \in \mathcal{I}_i^+, l \in [d_{i,r}] \right), \forall i \in [n],$$

and

$$W(y_j) := \sum_{i \in [n]} \sum_{a \in \mathcal{A}_{i,j}} V(a)(W(x_i)), \forall j \in [m].$$

Then $W := (W(x_i), W(y_j))_{i \in [n], j \in [m]}$ is a subrepresentation of V such that

$$\sigma \cdot \mathbf{dim } W \geq \text{rank}(Y) - \text{rank}(T_{V,\sigma}(Y)).$$

- (3) Let $U \leq \mathbb{R}^N$ be c -shrunk subspace for $T_{V,\sigma}$ for some $c \in \mathbb{N}$, and let $Y \in \mathbb{R}^{N \times N}$ be the matrix of orthogonal projection onto U . If W is the subrepresentation of V associated in (2) to Y then

$$\sigma \cdot \mathbf{dim } W \geq c.$$

Proof. (1) We have that $T_{V,\sigma}(Y)$ is the $M \times M$ block diagonal matrix whose (q, q) -block-diagonal entry is

$$\sum_{i=1}^n \sum_{a \in \mathcal{A}_{i,j}} V(a) \left(\sum_{r \in \mathcal{I}_i^+} Y_r \right) V(a)^T,$$

for all $q \in \mathcal{I}_j^-$ and $j \in [m]$. Thus we get that

$$\begin{aligned} \text{rank}(T_{V,\sigma}(Y)) &= \sum_{j=1}^m \sigma_-(y_j) \text{rank} \left(\sum_{i=1}^n \sum_{a \in \mathcal{A}_{i,j}} \sum_{l,r} V(a) u_l^i (V(a) u_l^i)^T \right) \\ &= \sum_{j=1}^m \sigma_-(y_j) \dim \left(\sum_{i=1}^n \sum_{a \in \mathcal{A}_{i,j}} V(a)(W(x_i)) \right) \leq \sum_{j=1}^m \sigma_-(y_j) \dim W(y_j). \end{aligned}$$

From this it follows that

$$\text{rank}(Y) - \text{rank}(T_{V,\sigma}(Y)) \geq \sum_{i=1}^n \sigma_+(x_i) \dim W(x_i) - \sum_{j=1}^m \sigma_-(y_j) \dim W(y_j) = \sigma \cdot \mathbf{dim} W.$$

(2) We have

$$(14) \quad \text{rank}(Y) \leq \sum_{i=1}^n \sum_{r \in \mathcal{I}_i^+} \text{rank}(Y_r) \leq \sum_{i=1}^n \sigma_+(x_i) \dim W(x_i),$$

and

$$\begin{aligned} (15) \quad \text{rank}(T_{V,\sigma}(Y)) &= \sum_{j=1}^m \sigma_-(y_j) \text{rank} \left(\sum_{i=1}^n \sum_{a \in \mathcal{A}_{i,j}} V(a) \left(\sum_{r \in \mathcal{I}_i^+} Y_r \right) V(a)^T \right) \\ &= \sum_{j=1}^m \sigma_-(y_j) \text{rank} \left(\sum_{i,a} \sum_{l,r} V(a) (\sqrt{\lambda_l^{i,r}} u_l^{i,r}) \left(V(a) (\sqrt{\lambda_l^{i,r}} u_l^{i,r}) \right)^T \right) \\ &= \sum_{j=1}^m \sigma_-(y_j) \dim \left(\sum_{i=1}^n \sum_{a \in \mathcal{A}_{i,j}} V(a)(W(x_i)) \right) = \sum_{j=1}^m \sigma_-(y_j) \dim W(y_j). \end{aligned}$$

It now follows from (14) and (15) that

$$\sigma \cdot \mathbf{dim} W \geq \text{rank}(Y) - \text{rank}(T_{V,\sigma}(Y)).$$

(3) Using the exact same arguments as in the proof of (iv) \implies (iii) in Proposition 20, we get that

$$\text{rank}(Y) - \text{rank}(T_{V,\sigma}(Y)) \geq \dim U - \dim \left(\sum_{i=1}^L A_i(U) \right).$$

This combined with part (2) yields

$$\sigma \cdot \mathbf{dim} W \geq c.$$

□

Definition 29. Let (V, σ) be a quiver datum as above. We define the *discrepancy* of (V, σ) to be

$$\text{disc}(V, \sigma) := \max\{\sigma \cdot \mathbf{dim} W \mid W \leq V\}.$$

Remark 30. It is clear that V is σ -semi-table if and only if $\text{disc}(V, \sigma) = 0$.

Recall that the discrepancy of $\mathcal{A}_{V, \sigma}$ is defined as

$$\text{disc}(\mathcal{A}_{V, \sigma}) := \max\{c \in \mathbb{N} \mid \text{there exists a } c\text{-shrunk subspace for } \mathcal{A}_{V, \sigma}\}.$$

We mention that $N - \text{disc}(\mathcal{A}_{V, \sigma})$ is the so-called non-commutative rank of the matrix space $\langle \mathcal{A}_{V, \sigma} \rangle \leq \mathbb{R}^{N \times N}$ (see for example [IQS17]).

Corollary 31. Let Q be a bipartite quiver and (V, σ) a quiver datum as above. Then

$$\text{disc}(V, \sigma) = \text{disc}(\mathcal{A}_{V, \sigma}).$$

Proof. From Lemma 28(3) we immediately get that $\text{disc}(V, \sigma) \geq \text{disc}(\mathcal{A}_{V, \sigma})$ while the reverse inequality follows from Lemma 28(1). □

We are now ready to prove Theorem 2.

Proof of Theorem 2. Let $\{x_1, \dots, x_n\}$ and $\{y_1, \dots, y_m\}$ be the vertices of Q where σ is positive and negative, respectively. Let Q^σ be the bipartite quiver with partite sets $\{x_1, \dots, x_n\}$ and $\{y_1, \dots, y_m\}$, respectively; furthermore, for every oriented path p in Q from x_i to y_j , we draw an arrow a_p from x_i to y_j in Q^σ . We denote the restriction of \mathbf{d} (or σ) to Q^σ by the same symbol.

Note every representation V of Q defines a representation V^σ of Q^σ as follows

- $V^\sigma(x_i) = V(x_i)$, $V^\sigma(y_j) = V(y_j)$ for all $i \in [n]$, $j \in [m]$, and
- $V^\sigma(a_p) = V(p)$ for every arrow a_p in Q^σ .

Then V (or V^σ) is σ -semi-stable as a representation of Q (or Q^σ) if and only if the same statement holds over the field of complex numbers (see [HS17, Proposition 2.4]). On the other hand, it has been proved in [CK20, Theorem 10] that over \mathbb{C} , $V_{\mathbb{C}}$ is σ -semi-stable as a representation of Q if and only if $V_{\mathbb{C}}^\sigma$ is σ -semi-stable as a representation of Q^σ . We thus conclude that

$$(16) \quad V \text{ is } \sigma\text{-semi-stable} \iff V^\sigma \text{ is } \sigma\text{-semi-stable.}$$

(i) This part of the theorem follows from the second part via (16).

(ii) For this part, we assume that Q is bipartite. According to [IQS17, Theorem 1.5], there exists a polynomial time algorithm that constructs a $\text{disc}(\mathcal{A}_{V, \sigma})$ -shrunk subspace U for $\mathcal{A}_{V, \sigma}$. Applying Proposition 28 to the subspace U , we can construct a subrepresentation W of V such that

$$\sigma \cdot \mathbf{dim} W \geq \text{disc}(\mathcal{A}_{V, \sigma}).$$

It now follows from Corollary 31 that

$$\sigma \cdot \mathbf{dim} W = \text{disc}(V, \sigma).$$

In particular, if $\sigma \cdot \mathbf{dim} W = 0$ then V is σ -semi-stable. Otherwise, W is the desired witness. □

- Remark 32.** (1) One of the key advantages of the Ivanyos-Qiao-Subrahmanyam’s algorithm in [IQS17] over Algorithm G is not only that it outputs a c -shrunk subspace but it constructs such a shrunk subspace certifying that c is precisely $\text{disc}(\mathcal{A}_{V,\sigma})$. Furthermore, the algorithm is algebraic in nature, working over large enough fields.
- (2) Keeping in mind the construction of Q^σ in the proof of Theorem 2, while any subrepresentation of V gives rise to a subrepresentation of V^σ , it is not clear how to extend a given subrepresentation of V^σ to a subrepresentation of V . So, when it comes to producing a certificate for V not being semi-stable in Theorem 2(1), we can only do it in the form of a subrepresentation of V^σ . □

REFERENCES

- [CD19] C. Chindris and H. Derksen, *The capacity of quiver representations and Brascamp-Lieb constants*, arXiv e-prints (2019), arXiv:1905.04783.
- [CK20] C. Chindris and D. Kline, *Orbit Semigroups of Quiver Representations and Edmond’s Problem*.
- [Edm67] J. Edmonds, *Systems of distinct representatives and linear algebra*, J. Res. Nat. Bur. Standards Sect. B **71B** (1967), 241–245. MR 0229540
- [EH88] D. Eisenbud and J. Harris, *Vector Spaces of Matrices of Low Rank*, Advances in Mathematics (1988).
- [GGOW15] A. Garg, L. Gurvits, R. Oliveira, and A. Wigderson, *Operator scaling: theory and applications*, ArXiv e-prints (2015).
- [Gur04] L. Gurvits, *Classical Complexity and quantum entanglement*, Journal of Computer and System Sciences (2004).
- [HM13] Moritz Hardt and Ankur Moitra, *Algorithms and hardness for robust subspace recovery*, Proceedings of Machine Learning Research, vol. 30, 2013, pp. 354–375.
- [HS17] V. Hoskins and F. Schaffhauser, *Rational points of quiver moduli spaces*, ArXiv e-prints (2017).
- [IKQS15] G. Ivanyos, M. Karpinski, Y. Qiao, and M. Santha, *Generalized wong sequences and their applications to edmonds’ problems*, Journal of Computer and System Sciences **81** (2015), no. 7, 1373 – 1386.
- [IQS17] G. Ivanyos, Y. Qiao, and K V. Subrahmanyam, *Constructive Non-Commutative Rank Computation Is in Deterministic Polynomial Time*, 8th Innovations in Theoretical Computer Science Conference (ITCS 2017) (Dagstuhl, Germany) (Christos H. Papadimitriou, ed.), Leibniz International Proceedings in Informatics (LIPIcs), vol. 67, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017, pp. 55:1–55:19.
- [Kin94] A.D. King, *Moduli of representations of finite-dimensional algebras*, Quart. J. Math. Oxford Ser.(2) **45** (1994), no. 180, 515–530.
- [LM18] Gilad Lerman and Tyler Maunu, *An overview of robust subspace recovery*, Proceedings of the IEEE **106** (2018), 1380–1410.
- [Lov89] L. Lovász, *Singular spaces of matrices and their application in combinatorics*, Bol. Soc. Brasil. Mat. (N.S.) **20** (1989), no. 1, 87–99. MR 1129080
- [Sch80] J. T. Schwartz, *Fast probabilistic algorithms for verification of polynomial identities*, J. ACM **27** (1980), no. 4, 701–717.
- [Zip79] R. Zippel, *Probabilistic algorithms for sparse polynomials*, Symbolic and Algebraic Computation (Berlin, Heidelberg) (Edward W. Ng, ed.), Springer Berlin Heidelberg, 1979, pp. 216–226.

UNIVERSITY OF MISSOURI-COLUMBIA, MATHEMATICS DEPARTMENT, COLUMBIA, MO, USA
Email address, Calin Chindris: chindrisc@missouri.edu

COLLEGE OF THE OZARKS, MATHEMATICS DEPARTMENT, POINT LOOKOUT, MO, USA
Email address, Daniel Kline: dkline@cofo.edu