Non-commutative Rank, Quivers, and Tensors

by

Alana Huszar

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)
in The University of Michigan
2022

Doctoral Committee:

Professor Harm Derksen, Co-Chair
Professor Mel Hochster, Co-Chair
Professor Kayvan Najarian
Professor Andrew Snowden

Alana Gayle Huszar

huszara@umich.edu

ORCID iD: 0000-0001-5616-415X

ACKNOWLEDGEMENTS

Thank you so much to my wonderful advisor, Harm Derksen, for all the help, direction, and support over the years – I am so lucky to have had the opportunity to work with a kind and stellar advisor. Thank you also to the other members of my committee, Mel Hochster (my secondary advisor), Kayvan Najarian (my cognate member), and Andrew Snowden (my reader, who gave very helpful comments) for taking the time to review my dissertation, I truly appreciate it.

Of course, I need to thank my peers who overlapped with me at Michigan and supported me through it all, from the first year, to my final semester. Another thank you goes to Hanna Bennett for mentorship and beyond. And last, a final thank you to my parents, sister, and partner for helping me outside the program during my nearly six long years away from home.

Much of this work was completed with support from the National Science Foundation, DMS 1256260.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF ABBREVIATIONS

$\mathcal{A}$ (Matrix space)

$\mathcal{A}^{\{d\}}$ (The dth tensor blow-up of $\mathcal{A}$)

$A(\underline{x})$ (Matrix with homogeneous linear entries)

χ_σ (Character defined by weight σ)

e_x (Trivial path at vertex x)

$\mathrm{ext}(\alpha, \beta), \mathrm{ext}(\alpha, W), \mathrm{ext}(V, \beta)$ (General ext)

$\underline{\mathrm{dim}}(V)$ (Dimension vector of V)

$\mathrm{Ext}_Q(V, W)$ (Quiver Representation Extensions of V by W.)

frk (Placeholder for any notion of rank for tensors)

f_W^V (Linear map with kernel and cokernel defining $\mathrm{Hom}_Q(V, W)$ and $\mathrm{Ext}_Q(V, W)$)

$GL(\alpha)$ (Product of $\mathrm{GL}(\alpha(x))$)

$GL(n)$ (General Linear Group of $n \times n$ matrices)

$\overline{G \cdot v}$ (Zariski closure of orbit of v)

h (Head of an arrow, $h : Q_1 \to Q_0$)

$\mathrm{hom}(\alpha, \beta), \mathrm{hom}(\alpha, W), \mathrm{hom}(V, \beta)$ (General hom)

Im (Image of a linear map)

$\mathrm{Hom}_Q(V, W)$ (Quiver Representation Morphisms from V to W)

k^* (Non-zero elements of k)

kQ (Path Algebra)

$k[V]$ (Polynomial functions on vector space V)

$k[V]^G$ (Invariant polynomial functions on vector space V)

$\langle \alpha, \beta \rangle$ (Euler/Ringel form on $\mathbb{R}^{Q_0}$)

$M(\alpha)$ (Product of $M(\alpha(x))$)

$M_{n,m}$ (Space of n by m matrices)

$\text{nc}_{\text{bu}}\text{-frk}(\mathcal{T})$ (Non-commutative rank of $\mathcal{T}$, generalizes blow-up definition)

$\text{nc}_{\text{cs}}\text{-frk}(\mathcal{T})$ (Non-commutative rank of $\mathcal{T}$, generalizes c-shrunk subspace definition)

$\text{nc-rk}_{\alpha}^{G}(\mathcal{T})$ (Non-commutative G-stable rank of $\mathcal{T}$)

$\text{ncext}(\alpha, W), \text{ncext}(V, \beta)$ (Non-commutative general ext)

$\text{nchom}(\alpha, W), \text{nchom}(V, \beta)$ (Non-commutative general hom)

$\text{nc-slrk}_{\alpha}(\mathcal{T})$ (α-weighted non-commutative slice rank of $\mathcal{T}$)

$\mathcal{N}$ (Nullcone)

$\Phi_i(T)$ (ith flattening of a tensor T)

P_x (Indecomposable projective representation with bases defined by paths from x)

Q (Quiver)

Q_0 (Quiver vertex set)

Q_1 (Quiver arrow set)

Rk (Tensor Rank)

rk (Matrix rank)

rk_{α}^{G} (G-stable rank)

$\underline{\text{rk}}$ (Border rank)

Rep_Q (Category of Quiver Representations on Q)

$\text{Rep}_Q(\alpha)$ (Quiver Representations on Q with dimension vector α)

$\text{SI}(Q, \alpha)_{\sigma}$ (Semi-Invariant polynomials of weight σ for $\text{Rep}_Q(\alpha)$)

σ (Weight vector in $\mathbb{Z}^{Q_0}$)

$\sigma(\alpha)$ (Inner product of σ and α, $\sum \sigma(x)\alpha(x)$)

slrk_{α} (α-weighted slice rank)

$\mathcal{T}$ (Tensor space)

t (Tail of an arrow, $t : Q_1 \to Q_0$)

ABSTRACT

Fortin and Reutenauer defined the non-commutative rank for a matrix with entries that are linear functions. We will generalize this non-commutative rank to both the representation theory of quivers, and to tensor spaces. In particular, we will relate non-commutative rank to King's criterion for σ-stability of quiver representations, general Hom and Ext spaces studied by Kac and Schofield, and several notions of tensor rank.

CHAPTER I

Introduction

One of the earliest things you might learn in a first linear algebra course is row reduction of matrices. Step by step you clear out entries to solve a system of equations, or find the rank of a matrix. For instance, we might start with the skew-symmetric matrix

$$\begin{bmatrix} 0 & x & y \\ -x & 0 & z \\ -y & -z & 0 \end{bmatrix},$$

with entries in $k[x, y, z]$ for some field k. Clearing the bottom left corner by adding to the third row $-yx^{-1}$ times the second row, we get

$$\begin{bmatrix} 0 & x & y \\ -x & 0 & z \\ 0 & -z & -yx^{-1}z \end{bmatrix},$$

and similarly clearing the $-z$ with zx^{-1} times the first row, gives us:

$$\begin{bmatrix} 0 & x & y \\ -x & 0 & z \\ 0 & 0 & zx^{-1}y - yx^{-1}z \end{bmatrix}.$$

Before revisiting the title of this thesis, we might then note the bottom right entry in the matrix is zero, and state this matrix has rank two. However, we will be concerned with the "non-commutative" rank of this matrix, the rank when the variables do not commute. Doing this makes things much harder, because it is difficult to determine whether an identity in the free skew field is non-zero. Take for example the expression from Hua's identity [Hua49],

$$(x + xy^{-1}x)^{-1} + (x + y)^{-1} - x^{-1}.$$

Although it is not obvious, we can show this expression is zero. In general, as non-commutative expressions can get more complicated, requiring more and more nested inverses, it is not easy to prove whether an expression is zero or not. So in order to definitively show the non-commutative rank of our matrix, we will rely on other tools from linear algebra, representation theory, and even multilinear algebra, turning to tensors. In Chapter III, in addition to providing these more tractable definitions of non-commutative rank, we will show that the non-commutative rank of this matrix captures the fact that while any 3×3 skew-symmetric matrix has (commutative) rank 2, they do not collectively take any subspace of k^3 to a subspace of smaller dimension. As we will see, this non-commutative rank problem has rich connections to quiver representations, invariant theory, and tensors.

The three key results of this thesis are two generalizations and one application of non-commutative rank. We begin in Chapter II with background on quiver representations and geometric invariant theory. In Chapter III, more discussion on non-commutative rank is provided, and in Section 3.3, we come to the first key result, generalizing algorithms for non-commutative rank to those measuring semi-stability of quiver representations. We conclude that there are both polynomial time deterministic and random algorithms for not only measuring the lack of semi-stability, but providing a witnessing subrepresentation of this measurement.

In Section 3.4, we have our second key result, a new proof for the limit behavior of generic ext and hom, coming from applying insights from non-commutative rank. This new proof leads to a brand new corollary, asserting these limits can be algorithmically determined, and can be replaced by a simple maximum.

In Chapter IV, we begin with background on different notions of tensor rank, with our third key result in Section 4.2, generalizing non-commutative matrix rank to new definitions for non-commutative tensor rank. Due to the nature of non-commutative rank, these new definitions are some of the first described for rank of a space of tensors.

CHAPTER II

Quivers

2.1 Quivers and Quiver Representations

A *quiver* is a finite directed graph. We will denote the vertex set by Q_0, and the arrow set by Q_1. Each arrow has a head, a vertex denoted by ha, and a tail, a vertex denoted by ta; we can consider t and h as functions from Q_1 to Q_0. The term quiver comes from Gabriel in [Gab72], who decided to call the tuple (Q_0, Q_1, h, t) a quiver, as "graph" already had too many different meanings.

Example 2.1.

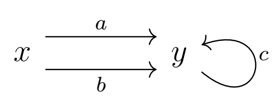

Here, we have a quiver with $Q_0 = \{x, y\}$ and $Q_1 = \{a, b, c\}$. Notice that we can have multiple arrows between vertices, as $ha = hb$ and $ta = tb$, and loops, as $hc = tc$.

In Chapter III, one quiver we will be particularly interested in is the *Generalized Kronecker Quiver*, a quiver on two vertices with m arrows, shown below.

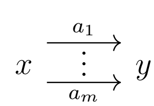

Fixing a field k, we will attach finite dimensional k-vector spaces to the vertices, and linear maps to the arrows so that the dimensions match up. This is a quiver representation, first described in [Gab72].

Definition 2.2. A *representation* V of a quiver Q is given by an assignment of a k-vector space, $V(x)$, to each vertex x in Q_0, and a linear map, $V(a) : V(ta) \to V(ha)$, to each arrow a in Q_1.

Example 2.3.

$$V : \ \mathbb{C}^2 \underset{\begin{bmatrix}0 & -1\end{bmatrix}}{\overset{\begin{bmatrix}1 & 2\end{bmatrix}}{\rightrightarrows}} \mathbb{C} \circlearrowleft 10$$

Here, we have a representation of the quiver in Example 2.1. Notice $V(x) = \mathbb{C}^2$, $V(y) = \mathbb{C}$, $V(a) = \begin{bmatrix}1 & 2\end{bmatrix}$, $V(b) = \begin{bmatrix}0 & -1\end{bmatrix}$, and $V(c) = 10$.

A *subrepresentation* of a quiver V is a representation W so that for all vertices x, $W(x)$ is a subspace of $V(x)$, and for all arrows a, $W(a) : W(ta) \to W(ha)$ is the restriction of $V(a)$ to $W(ta)$.

Example 2.4. A representation V of a quiver Q is called *trivial* if $V(x) = 0$ for all vertices x. The trivial representation of Q is a subrepresentation of all quiver representations of Q.

A *dimension vector* α for a quiver is a tuple in $\mathbb{N}^{Q_0}$. We denote the dimension vector of a quiver representation V by $\underline{\dim}(V)$, where $\underline{\dim}(V)(x) := \dim(V(x))$.

Example 2.5. In Example 2.3, $\underline{\dim}(V) = (2, 1)$. The trivial subrepresentation of V has dimension vector $(0, 0)$.

2.1.1 Categories of Quiver Representations

After fixing a field k and quiver Q, we can look at the category of representations, Rep_Q. To do that, we will need to define a morphism of quiver representations.

Definition 2.6. Let V and W be representations of a quiver Q. A *morphism* $\varphi : V \to W$ is given by a set of linear maps $\varphi(x) : V(x) \to W(x)$ defined for all vertices x so that for all arrows a, $\varphi(ha)V(a) = W(a)\varphi(ta)$. We denote the set of all morphisms from V to W by $\text{Hom}_Q(V, W)$.

We may also consider $\mathrm{Hom}_Q(V,W)$ as the kernel of the map:

$$f_W^V : \bigoplus_{x \in Q_0} \mathrm{Hom}\big(V(x), W(x)\big) \to \bigoplus_{a \in Q_1} \mathrm{Hom}\big(V(ta), W(ha)\big), \tag{2.1}$$

where $f_W^V(\varphi) = \big(\varphi(ha)V(a) - W(a)\varphi(ta) : a \in Q_1\big)$. The cokernel of f_W^V will be denoted by $\mathrm{Ext}_Q(V,W)$, and will be discussed more in Section 2.1.4.

Example 2.7. Let's find $\mathrm{Hom}_Q(V,W)$ for the V in Example 2.3 and the following representation W.

$$W:\ \mathbb{C}^2 \underset{\begin{bmatrix}0 & 1\end{bmatrix}}{\overset{\begin{bmatrix}1 & 0\end{bmatrix}}{\rightrightarrows}} \mathbb{C} \circlearrowleft 10$$

Let $\lambda : \mathbb{C} \to \mathbb{C}$ be the linear map from $V(y)$ to $W(y)$. Now, we need to define a linear map A from $V(x) \to W(x)$ so that the following diagrams commute.

$$\begin{array}{ccc} \mathbb{C}^2 & \xrightarrow{\begin{bmatrix}1 & 2\end{bmatrix}} & \mathbb{C} \\ {\scriptstyle A}\downarrow & & \downarrow{\scriptstyle \lambda} \\ \mathbb{C}^2 & \xrightarrow[\begin{bmatrix}1 & 0\end{bmatrix}]{} & \mathbb{C} \end{array} \qquad \begin{array}{ccc} \mathbb{C}^2 & \xrightarrow{\begin{bmatrix}0 & -1\end{bmatrix}} & \mathbb{C} \\ {\scriptstyle A}\downarrow & & \downarrow{\scriptstyle \lambda} \\ \mathbb{C}^2 & \xrightarrow[\begin{bmatrix}0 & 1\end{bmatrix}]{} & \mathbb{C} \end{array}$$

This gives the only option for $A = \left[\begin{smallmatrix}\lambda & 2\lambda \\ 0 & -\lambda\end{smallmatrix}\right]$, so $\mathrm{Hom}_Q(V,W)$ is a one-dimensional space determined completely by λ.

Two quiver representations, V and W, are isomorphic if there is a morphism $\varphi : V \to W$ and a morphism $\psi : W \to V$ so that $\psi \circ \varphi$ and $\varphi \circ \psi$ are both identity morphisms, or equivalently, if $\varphi(x)$ is bijective for all x. In order for two representations to be isomorphic, they must have the same dimension vector. We may restrict ourselves to quiver representations with a fixed dimension vector α, the set of quiver representations V on Q with $V(x) = k^{\alpha(x)}$. We call this vector space $\mathrm{Rep}_Q(\alpha)$, and can identify it with the product of matrix spaces $\prod_{a \in Q_1} M_{\alpha(ha),\alpha(ta)}$. On this product space, we have an action of

$$\mathrm{GL}(\alpha) := \prod_{x \in Q_0} \mathrm{GL}(\alpha(x)),$$

where the action of $(Y(x), x \in Q_0)$ takes $V(a)$ to $Y(ha)V(a)Y(ta)^{-1}$ for all arrows a. As we are acting on the vector space $\mathrm{Rep}_Q(\alpha)$ linearly, we can consider this action a group representation, which will be discussed in more detail starting in Section 2.2.

Example 2.8. Let's act on the same V (Example 2.3) by $(\left[\begin{smallmatrix} 0 & 1 \\ 2 & 3 \end{smallmatrix}\right], -3)$. We get the representation

$$\mathbb{C}^2 \underset{\begin{bmatrix} 2 & 0 \end{bmatrix}}{\overset{\begin{bmatrix} -1 & -1 \end{bmatrix}}{\rightrightarrows}} \mathbb{C} \circlearrowleft 10 .$$

Theorem 2.9 (See [DW17])**.** Two quiver representations V and W are isomorphic if and only if they are in the same $\mathrm{GL}(\alpha)$ orbit.

Proof. First, suppose $\varphi : V \to W$ is an isomorphism. Then, $\varphi(ha)V(a) = W(a)\varphi(ta)$ as we have a quiver morphism. But φ is invertible, giving us $\varphi(ha)V(a)\varphi(ta)^{-1} = W(a)$, and so defining $Y := (\varphi(x) \mid x \in Q_0)$ in $\mathrm{GL}(\alpha)$ gives $Y \cdot V = W$. In reverse, we now suppose $Y \cdot V = W$ for some $Y := (Y(x)|x \in Q_0)$ in $\mathrm{GL}(\alpha)$. This gives $Y(ha)V(a)Y(ta)^{-1} = W(a)$, so a map from V to W defined at each vertex by $Y(x)$ will be an isomorphism (with inverse similary defined by the $Y(x)^{-1}$). $\square$

Part of the discussion in the next sections will be on the different ways one might care about the orbits of this action. For instance, we might care about the isomorphism classes of a quiver Q.

Example 2.10. If $\alpha = (m, n)$, the isomorphism classes of the quiver $x \xrightarrow{a} y$ with dimension vector α are given by the rank of the linear map at a; from 0 to $\min(m, n)$.

2.1.2 Indecomposable Representations

Given two quiver representations, V and W, we can construct a quiver representation $V \oplus W$ by defining $(V \oplus W)(x) = V(x) \oplus W(x)$ at each vertex x, and

$$(V \oplus W)(a) = V(a) \oplus W(a) = \begin{pmatrix} V(a) & 0 \\ 0 & W(a) \end{pmatrix}$$

for all arrows a. For this reason, it is useful to try to look at the "smallest" quiver representations.

Definition 2.11. A non-trivial represenation V of Q is decomposable if $V \cong W_1 \oplus W_2$ for some non-trivial representations W_1 and W_2 of Q. A non-trivial representation is called *indecomposable* if it is not decomposable.

Theorem 2.12 (Krull-Remak-Schmidt [Kru24, Rem11, Sch29])**.** Every finite dimensional quiver representation is isomorphic to a direct sum of indecomposable representations. This decomposition is unique up to isomorphism and permutation of factors.

We will now find all indecomposable representations for the two quivers with only one arrow, following examples in [DW17].

Example 2.13. Let's find all isomorphism classes of indecomposable representations for the quiver $x \xrightarrow{a} y$ over $\mathbb{C}$. Let V be an indecomposable representation of our quiver. We begin by noting that if $V(x) = 0$ and $\dim V(y) \geq 2$, we can write $V(y) = W_1(y) \oplus W_2(y)$ for nonzero subspaces $W_1(y), W_2(Y)$ of $V(y)$. If we take $W_1(x) = W_2(x) = 0$ then W_1 and W_2 are nonzero subrepresentations of V with $V = W_1 \oplus W_2$. Contradiction. So $V(y)$ must be one-dimensional. Similarly, if $V(y) = 0$, $V(x)$ must be one-dimensional. Our first two isomorphism classes of indecomposables are then $0 \to \mathbb{C}$ and $\mathbb{C} \to 0$.

Now, we suppose the dimension at x and y are both positive. If $V(a)$ has kernel K, with complement K', we have $V \cong (K \to 0) \oplus (K' \to V(a))$, so $V(a)$ must be injective. Similarly, if $V(a)$ has cokernel C, with complement C', we have $V \cong (0 \to C) \oplus (V(x) \to C')$, so $V(a)$ must be surjective. Since $V(a)$ must then be an isomorphism, we can take any subspace $W_1(x)$ of $V(x)$, and its complement $W_2(x)$, and define $W_1(y) = V(a)W_1(x)$ and $W_2(y) = V(a)W_2(x)$. This gives $V = W_1 \oplus W_2$. So we must have $\dim(V(x)) = 1$ (as is $\dim(V(y))$, as we have an isomorphism), and our last isomorphism class is $\mathbb{C} \to \mathbb{C}$.

In the last example, we saw there were 3 isomorphism classes of indecomposable representations. Quivers with finitely many isomorphism classes of indecomposable representations are said to have *finite representation type*. As we will see in the next example, this is not always the case.

Example 2.14. Let's find all isomorphism classes of indecomposable representations over $\mathbb{C}$ for the quiver:

$$x \circlearrowleft a$$

Letting V be an indecomposable representation for this quiver, we see that for some basis, $V(a)$ will be in Jordan form, with Jordan blocks $J_1, \ldots, J_k$, of dimensions $d_1, \ldots, d_k$ respectively. We can decompose V into

$$\bigoplus_{i=1}^{k} \mathbb{C}^{d_i} \circlearrowleft J_i \ ,$$

so k must be 1, $V(a)$ must be a single $n \times n$ Jordan block, J, with eigenvalue λ. If we were able to decompose our quiver further, we would end up with a basis where J is a Jordan matrix with more than one block. So our isomorphism classes of indecomposables in this case are determined by n and λ.

Though the quiver in this example was not of finite representation type, it was *tame*, meaning after fixing the dimension vector (in our case n), the isomorphism classes are determined by finitely many one-parameter families (in our case, we had a single one-parameter family, given by λ). Quivers that are not of finite representation type nor tame are called *wild*. In [Gab72], Gabriel proved the quivers with finite representation type are exactly the quivers with underlying undirected graphs the Dynkin diagrams of type ADE. The extended Dynkin quivers of type ADE, are tame, work done on the description of their indecomposables was done in [Naz73] and [DF73]. The Dynkin diagrams and their extensions are shown in figure 2.1. For all other quivers, the representation type is *wild*, and finding indecomposables is much less tractable.

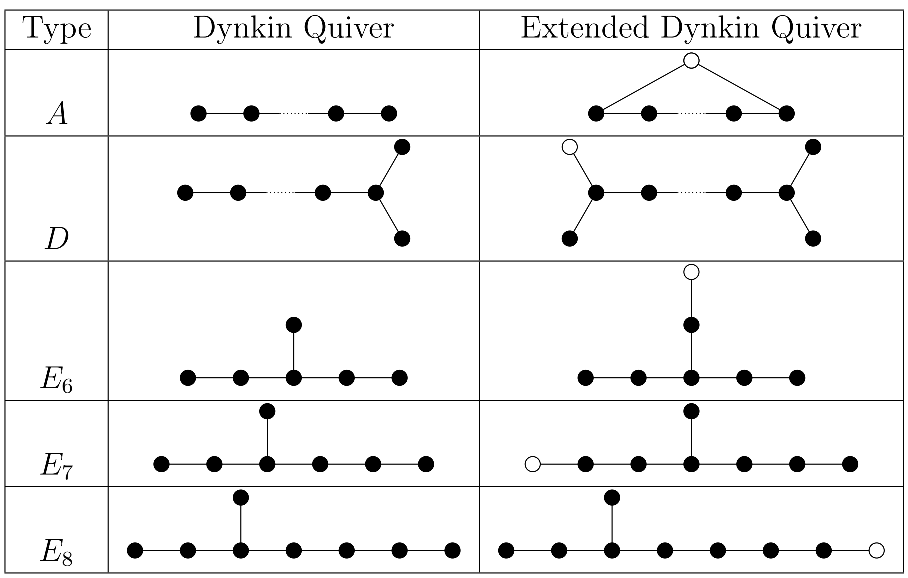

Figure 2.1: Dynkin Diagrams of Type ADE

2.1.3 Projective Representations

We will use a particular kind of indecomposable representation for constructions later in this thesis, indecomposable projective representations. To define these, we must first talk about the path algebra. A *path* in a quiver is a sequence of arrows, $p = a_\ell \dots a_1$ so that $ta_{i+1} = ha_i$ for $i = 1, \dots, \ell - 1$. The *length* of the path is ℓ, and the head and tail of the path are $hp = ha_\ell$ and $tp = ta_1$ respectively. For each vertex x, we have a trivial path, e_x, a path of length 0 with $he_x = te_x = x$. An *oriented cycle* in a quiver is a non-trivial path p with $tp = hp$. A quiver is *acylic* if it contains no oriented cycles. For the rest of this thesis, unless otherwise noted, all quivers are assumed to be acyclic.

Definition 2.15. Given a field k, the *path algebra* kQ is an k-algebra with basis labeled by all paths in Q. The multiplication in kQ is determined by concatenation $p \cdot q = pq$ if $hq = tp$, and $p \cdot q = 0$ if not. When concatenating with trivial paths, we have $pe_{tp} = p$ and $e_{hp}p = p$.

Example 2.16. Let's look at an example using the following quiver on 3 vertices,

$$Q: \; x \underset{b}{\overset{a}{\rightrightarrows}} y \xrightarrow{\;c\;} z.$$

In this case, kQ will have basis $e_x, e_y, e_z, a, b, c, ca, cb$. Multiplying the element $3e_y + c$ on the right by $a + e_z$ gives us $3a + ca + c$.

Proposition 2.17 (See [DW17], 1.5.4)**.** The category kQ-mod is equivalent to the category Rep_Q.

Although we will not go through the full proof of this equivalence of categories, we note that if V is a kQ-module, the corresponding quiver under the equivalence is defined by $V(x) = e_xV$ and $V(a) : e_taV \to e_haV$ as the restriction of the map $V \to V$ defined by left multiplication by a.

Definition 2.18. A representation P of V is projective if the functor $\operatorname{Hom}_Q(P, \_)$ is exact.

Proposition 2.19 (See [DW17])**.** A representation P of V is *projective* if and only if it is a direct summand of $(kQ)^r$ for some r.

Given a vertex x, in our quiver Q, define a representation by $P_x := kQe_x$, the quiver with basis at each vertex y given by paths from x to y. We will later see this representation is both projective and indecomposable. Let's look at P_x for our previous quiver.

Example 2.20. The representation P_x for the quiver

$$Q:\ x \overset{a}{\underset{b}{\rightrightarrows}} y \xrightarrow{c} z,$$

has a k-vector space at each vertex: $P_x(x)$ is spanned by e_x, $P_x(y)$ is spanned by a, b, and $P_x(z)$ is spanned by ca, cb. At the arrows, we have $P_x(a)$ the linear map sending e_x to a, $P_x(b)$ sending e_x to b, and $P_x(c)$ sending a and b to ca and cb respectively.

Proposition 2.21 (See [DW17])**.** For all vertices x of a quiver Q, the representation P_x is a projective kQ-module, and $kQ = \bigoplus_{x \in Q_0} P_x$.

Proof. Given u in kQ, we can consider it an element of $\bigoplus P_x$, as $u = u(\sum_x e_x) = \sum_x ue_x$. The sum is direct, as if $\sum_x a_x = 0$, for every vertex y we have $0 = (\sum a_x)e_y = a_y$. So each a_x is 0, and the sum is direct. In particular, each P_x is a projective representation. □

Proposition 2.22 (See [DW17])**.** Given a quiver Q, and a representation V, for all vertices x,

$$\mathrm{Hom}_Q(P_x, V) \cong V(x).$$

Proof. We see a map in $\mathrm{Hom}_Q(P_x, V)$ is completely determined by where in $V(x)$ we send e_x. So we may define $\psi : \mathrm{Hom}_Q(P_x, V) \to V(x)$ by $\varphi \mapsto \varphi(e_x)$. In the other direction, define $\Theta : V(x) \to \mathrm{Hom}_Q(P_x, V)$ by mapping v in $V(x)$ to the map φ_v, defined by sending v to e_x (and therefore taking any path p to $V(p)v$). our maps ψ and Θ are inverses, each giving us an isomorphism. □

Corollary 2.23 (See [DW17])**.** The indecomposable projective representations are exactly all the P_x, for x in Q_0.

Proof. From Proposition 2.22, we have $\operatorname{Hom}_Q(P_x, P_x) \cong ke_x$, which is one-dimensional. So all P_x must be indecomposable. We must have all indecomposable projective representations, as from Proposition 2.21, $(kQ)^r = \bigoplus_x (P_x)^r$, of which each projective representation must be a summand, and from Theorem 2.12, this factorization into indecomposables must be unique. □

2.1.4 Homological Algebra

Given a quiver representation V, a projective resolution is a sequence of projective representations P_i and maps f_i so that

$$\ldots \to P_2 \xrightarrow{f_2} P_1 \xrightarrow{f_1} P_0 \xrightarrow{f_0} V \to 0$$

is exact. We will see construction of such a sequence is always possible in 2.24, so there are enough projectives. For another quiver representation W, we apply $\operatorname{Hom}(\_, W)$, giving us

$$0 \to \operatorname{Hom}(P_0, W) \xrightarrow{\operatorname{Hom}(f_1,W)} \operatorname{Hom}(P_1, W) \xrightarrow{\operatorname{Hom}(f_2,W)} \operatorname{Hom}(P_2, W) \to \ldots,$$

computing $\operatorname{Ext}^i(V, W) := \ker(\operatorname{Hom}(f_i, W))/\operatorname{Im}(\operatorname{Hom}(f_{i-1}, W))$. The isomorphism class of each $\operatorname{Ext}^i(V, W)$ does not depend on the choice of projective resolution, see [Wei95] for background on homotopy equivalence of chain complexes. Note that from this construction, we get $\operatorname{Ext}^i(\_, W)$, the derived functors of $\operatorname{Hom}(\_, W)$.

Proposition 2.24 (See [DW17])**.** If Q is acyclic, given a quiver representation V, we have an exact sequence

$$0 \to \bigoplus_{a \in Q_1} V(ta) \otimes P_{ha} \xrightarrow{d^V} \bigoplus_{x \in Q_0} V(x) \otimes P_x \xrightarrow{f^V} V \to 0, \tag{2.2}$$

where $f^V(v \otimes p) = V(p)v$ for v in $V(x)$ and p in P_x, and $d^V(v \otimes p) = V(a)v \otimes p - v \otimes pa$

for v in $V(ta)$, and $p : ta \to ha$.

Proof. First, note that

$$f^V(d^V(v \otimes p)) = f^V(V(a)v \otimes p - v \otimes pa) = V(p)V(a)v - V(pa)v = 0.$$

Next, the image of f^V is all of V, as for any v in V, $f^V(v \otimes e_x) = v$. For the kernel of d^V, suppose $d^V\left(\sum_{a \in Q_1} v_a \otimes p_a\right) = 0$. As Q is acyclic, we can label vertices in Q_0 as $\{1, 2, \ldots, r\}$ so that $h(a) < t(a)$ for all a. Now, pick the maximal vertex y so that there's an arrow $ta = y$ with $v_a \otimes p_a \neq 0$, this is only possible (and always possible) when $\sum v_a \otimes p_a$ is non zero. For any arrow a' with $ha' = y$, we have $ta' > ha' = y$, and so $v_{ta'} \otimes p_{a'} = 0$. The component of $d^V\left(\sum v_a \otimes p_a\right)$ landing in $V(y) \otimes P_y$ is $\sum_{i=1}^{r} -v_{a_i} \otimes p_{a_i} a_i$ for all arrows $a_1, \ldots, a_r$ with tail y. By assumption, this must be 0, but the $p_i a_i$ are all independent, so there was no $v_a \otimes p_a \neq 0$, and the kernel of d^V is empty. Last, we have

$$\dim\left(\bigoplus_{x \in Q_0} V(x) \otimes P_x\right) - \dim\left(\bigoplus_{a \in Q_1} V(ta) \otimes P_x\right) =$$

$$= \sum_{x \in Q_0} \dim(V(x)) \left(\dim(P_x) - \sum_{\substack{a \in Q_1 \\ ta = x}} \dim(P_{ha})\right),$$

and $\dim(P_x) - \sum_{\substack{a \in Q_1 \\ ta = x}} \dim(P_{ha}) = 1$, as if $ta = x$, for all p in P_{ha}, pa is in P_x. Summing over all arrows a with $ta = x$, we geth the basis of paths starting at x for P_x, except for the empty path e_x. So, $\dim(P_x) - \sum \dim(P_{ha}) = 1$, and we get $\sum \dim(V(x)) = \dim(V)$, our sequence is exact. □

Corollary 2.25 (See [DW17])**.** For an acyclic quiver Q, given any two representations V and W, $\mathrm{Ext}^i(V, W) = 0$ for $i \geq 2$, and $\mathrm{Ext}^0(V, W) = \mathrm{Hom}_Q(V, W)$.

Proof. Apply $\mathrm{Hom}(\_, W)$ to the sequence (2.2). We get:

$$0 \to \mathrm{Hom}(V, W) \to \mathrm{Hom}(\bigoplus V(x) \otimes P_x, W) \xrightarrow{\mathrm{Hom}(d^V, W)} \mathrm{Hom}(\bigoplus V(ta) \otimes P_{ha}, W).$$

We have an isomorphism

$$\mathrm{Hom}(\bigoplus V(x) \otimes P_x, W) \cong \bigoplus \mathrm{Hom}(V(x), \mathrm{Hom}(P_x, W)) \cong \bigoplus \mathrm{Hom}(V(x), W(x))$$

by tensor-hom adjunction and Proposition 2.22, and similarly

$$\mathrm{Hom}(\bigoplus V(ta) \otimes P_{ha}, W) \cong \bigoplus \mathrm{Hom}(V(ta), \mathrm{Hom}(P_{ha}, W)) \cong \bigoplus \mathrm{Hom}(V(ta), W(ha)).$$

With these isomorphisms in mind, $\mathrm{Hom}(d^V, W) = d^V_W$ from 2.1, which has kernel $\mathrm{Hom}(V, W)$. The cokernel is $\mathrm{Ext}^1(V, W)$, which we denote by $\mathrm{Ext}(V, W)$, as all higher order Ext are 0. ☐

In the case of quivers with cycles, we may still define $\mathrm{Ext}_Q(V, W)$, either by the cokernel of d^V_W or by Yoneda extensions [Yon60], defined below.

Definition 2.26. An *extension* ξ of V by W is an exact sequence

$$\xi : 0 \to W \xrightarrow{i} E \xrightarrow{p} V \to 0.$$

Two extensions, ξ as above, and

$$\xi' : 0 \to W \xrightarrow{i'} E' \xrightarrow{p'} V \to 0$$

are *equivalent* if there is an isomorphism $f : E \to E'$ so that the diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & W & \stackrel{i}{\longrightarrow} & E & \stackrel{p}{\longrightarrow} & V & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle I_W} & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle I_V} & & \\ 0 & \longrightarrow & W & \stackrel{i'}{\longrightarrow} & E' & \stackrel{p'}{\longrightarrow} & V & \longrightarrow & 0 \end{array}$$

commutes.

Definition 2.27. Define the *Yoneda extension group*, $\mathrm{Ext}_Q(V, W)$, as the set of equivalence classes of extensions of V by W.

Given $[\xi]$ and $[\xi']$ in $\mathrm{Ext}_Q(V, W)$, with

$$\xi : 0 \to W \xrightarrow{i} E \xrightarrow{p} V \to 0, \text{ and } \xi' : 0 \to W \xrightarrow{i'} E' \xrightarrow{p'} V \to 0,$$

we get $[\xi] + [\xi'] = [\xi'']$, where $\xi'' : 0 \to W \xrightarrow{i''} E'' \xrightarrow{p''} V \to 0$ is defined as follows. Let $E'' = F/G$ with $F = \{(u, u') \in E \oplus E' | p(u) = p'(u')\}$, and $G = \{(i(w), -i'(w)) | w \in W\}$. Define $i'' : W \to F/G$ by $i''(w) = (i(w), 0) + G$, and p'' the morphism induced by the map $F \to V$, which contains G in its kernel.

2.2 Geometric Invariant Theory

Let's say we have a group G and a k-vector space V. We may want to act on V linearly instead of simply considering it as just a set. To do this, we define a *representation* of a group to be a homomorphism $\rho : G \to \mathrm{GL}(V)$. Then, our action of G is given by the representation, for g in G and v in V we get $g \cdot v = \rho(g)v$. The group G is called a *linear algebraic group* if it is an affine variety and its multiplication and inversion maps are morphisms of affine varieties. We point to the symmetric group for our first example.

Example 2.28. We can act on $V = \mathbb{C}^3$ by the symmetric group S_3 by sending each permutation to the corresponding permutation matrix in $\mathrm{GL}(3)$, giving us a representation.

Fix a basis for the vector space V, say $e_1, \ldots, e_n$, and let $x_1, \ldots, x_n$ be the corresponding coordinate functions. We define $k[V] := k[x_1, \ldots, x_n]$, the polynomial functions on V. We get an action of G on $k[V]$, given by

$$(g \cdot f)(a_1, \ldots, a_n) = f(g^{-1}(a_1, \ldots, a_n))$$

for g in G, f in $k[V]$, and $(a_1, \ldots, a_n)$ in V. Notice that we need the inverse in order to remain a left action of G.

Example 2.29. Again looking at S_3, we have $(123) \cdot (x_1x_2 + x_2x_3) = x_3x_1 + x_1x_2$.

In invariant theory, we are concerned with which polynomials in $k[V]$ are invariant under this action. We denote the set of invariant polynomials by $k[V]^G$. Constant polynomials and products and sums of invariant polynomials are invariant, so $k[V]^G$ is a ring. We will sometimes refer to invariant polynomials as simply "invariants".

The groups we will be concerned with are *reductive groups*, linear algebraic groups larger than one element whose largest connected normal unipotent subgroup is trivial (See [Hum75]). When working over characteristic 0, *linearly reductive groups*, groups whose rational representations are all completely reducible, are equivalent to reductive groups. Last, we have *geometrically reductive groups*, groups whose action on V is so that for every non-zero invariant vector v in V, there is an invariant homogeneous polynomial f of positive degree that does not vanish on v. Nagata and Miyata showed that all geometrically reductive groups are reductive [NM63]. The reverse was proven by Haboush, opening up work on Geometric Invariant Theory to all reductive groups.

Theorem 2.30 (Haboush's Theorem [Hab75])**.** Geometrically reductive groups are reductive over any characteristic.

Theorem 2.31 ([Hil90, Hil93])**.** If G is reductive, the invariant polynomial ring $k[V]^G$ is finitely generated.

One of the main problems in invariant theory is finding generating polynomials for $k[V]^G$. There's lots of work done in bounding the degree needed for a homogeneous set of generators. In characteristic 0, Noether proved for finite G, invariants with degree up to the order of G suffice to generate [Noe26], and Derksen and Kemper proved that if there is a uniform bound for a representation that only depends on G, G must be finite [DK02].

Example 2.32. The polynomial ring $\mathbb{C}[x_1, x_2, x_3]^{S_3}$ is generated by the elementary symmetric polynomials, $x_1 + x_2 + x_3$,$x_1x_2 + x_1x_3 + x_2x_3$, and $x_1x_2x_3$. Notice that in this case, we did not need to go all the way to Noether's bound, which would be $|S_3| = 6$; the largest degree in our generating set is 3.

Although the symmetric group is a great example, from this point foward, we will only be concerned with infinite groups. Luckily for us, the torus group $(k^*)^n$ along with $\mathrm{GL}(\alpha)$ and $\mathrm{SL}(n)$ over $\mathbb{C}$ are reductive.

Example 2.33. Now, let's look at orbits and invariants of the action of the reductive group $G = \left\{ \begin{bmatrix} t & 0 \\ 0 & t^{-1} \end{bmatrix} \middle| t \in \mathbb{C}^* \right\}$ on $\mathbb{C}^2$. Over the reals we will graph several orbits of this action, distinguished by color in figure 2.2 below.

Orbits: $\begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ -2.5 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \end{bmatrix}$

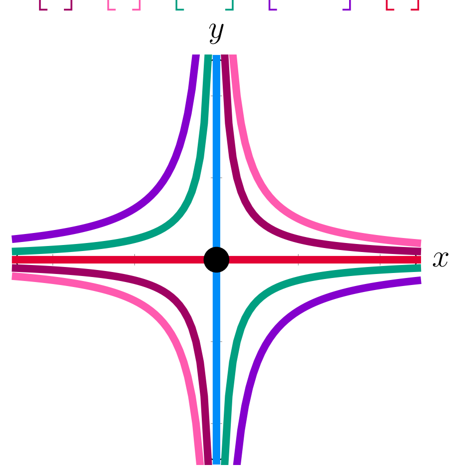

Figure 2.2: Orbits for Reductive G

The invariant polynomials for this action are generated by the polynomial xy. Notice that our generating invariant polynomial is constant on orbits. However, it can not distinguish orbits, as we see xy evaluates to 0 on both the x and the y axis orbits, along with the origin. What we are seeing is that the invariant xy distinguishes orbits up to Zariski closure; as the closure of the orbit for the x and y axis intersect at the origin, and no other orbits' closures intersect.

Example 2.34. Now let's look at when we do not have a reductive group, say

$$G = \left\{ \begin{bmatrix} 1 & u \\ 0 & 1 \end{bmatrix} \middle| u \in \mathbb{C} \right\},$$

acting on $\mathbb{C}^2$. We can quickly see G is not reductive, as it is not completely reducible - the first coordinate gives a subrepresentation W, but $\mathbb{C}^2/W$ is not a representation. Again over the reals we will graph several orbits of this action, distinguished by color in figure 2.3.

Orbits: $\begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \begin{bmatrix} 1 \\ -2 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \begin{bmatrix} 1 \\ 2.5 \end{bmatrix}, \begin{bmatrix} 1 \\ -0.5 \end{bmatrix}, \begin{bmatrix} x \\ 0 \end{bmatrix}$

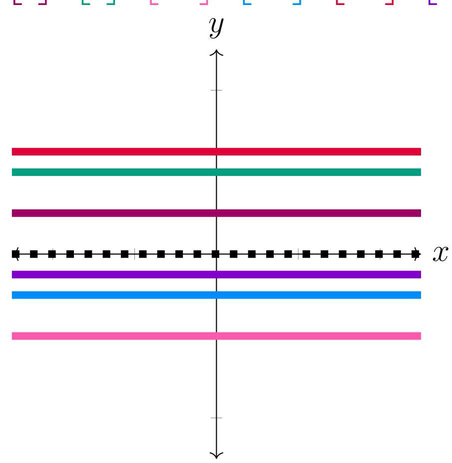

Figure 2.3: Orbits for Non-reductive G

Here, the invariant polynomials for this action are generated by the polynomial y. Notice that again, our generating invariant polynomial is constant on orbits. However, it can no longer distinguish orbit closures, as we see y evaluates to 0 on each of our closed orbits on the x-axis.

We have exemplified that invariant polynomials are constant on orbits, and when G is reductive, we can distinguish orbits up to Zariski closure, giving us the next theorem. From this point forward, we will only work with reductive groups G.

Theorem 2.35 (See [New78])**.** If G is a reductive group, given v, w in V, the intersection of orbits $\overline{G \cdot v} \cap \overline{G \cdot w}$ is empty if and only if there is an invariant f in $k[V]$ so that $f(v) \neq f(w)$.

Proof. Suppose u is in $\overline{G \cdot v}$ and $\overline{G \cdot w}$. For any invariant polynomial f, define $g(x) = f(x) - f(v)$. Notice g is also an invariant polynomial. We see g vanishes on $G \cdot v$ and at u. But u is in the Zariski closure of $G \cdot w$, so g must also vanish on $G \cdot w$. So, $f(w) - f(v) = 0$ for all invariant polynomials f. In the other direction, suppose the intersection of our two orbit closures is empty. We follow [New78, Lemma 3.3], and start with a polynomial h on V with $h(\overline{G \cdot v}) = 0$ and $h(\overline{G \cdot w}) = 1$. Now, consider the subspace of $k[V]$ spanned by $G \cdot h$, the orbit of h. This subspace has a basis, say $h_1, \ldots, h_m$, where $g \cdot h_i = \sum a_{ij}(g) h_j$ for all g in G. The assignment g to $a_{ij}(g)$ gives us a rational representation of G. Now, define $\psi : V \to k^m$ by sending the vector x to $(h_1(x), \ldots, h_m(x))$. Under ψ, the image of $\overline{G \cdot v}$ is zero, and of $\overline{G \cdot w}$ is a single, non-zero invariant point, w'. Now as G is reductive (and so equivalently, geometrically reductive), we can find an invariant polynomial $f' : k^m \to k$ with $f'(w') = 1$ and $f'(0) = 0$. Now, $f' \cdot \psi$ is an invariant polynomial in $k[V]$ separating the two orbit closures.

□

An important case is the orbits whose closures contain 0, the *nullcone*, defined as

$$\mathcal{N} := \{v \in V \mid 0 \in \overline{G \cdot v}\}.$$

On the nullcone, all homogeneous invariants of positive degree vanish. We call v in V *semi-stable* if v is not in the nullcone.

2.2.1 Semi-stability of Quiver Representations

Recall we have an action of

$$\mathrm{GL}(\alpha) := \prod_{x \in Q_0} GL(\alpha(x))$$

on $\mathrm{Rep}_Q(\alpha)$. We are interested in finding the nullcone of this action.

Example 2.36. Let's first look at the quiver $x \xrightarrow{a} y$, with dimension vector $\alpha = (2, 2)$. For any representation V, we have $V(a)$ a 2×2 matrix. Consider

$$(A, B) := \left(\begin{bmatrix} \varepsilon^{-1} & 0 \\ 0 & \varepsilon^{-1} \end{bmatrix}, \begin{bmatrix} \varepsilon & 0 \\ 0 & \varepsilon \end{bmatrix} \right)$$

for ε close to 0. No matter what $V(a)$ is, as ε approaches 0, we approach the trivial representation, so all representations V are in the nullcone, $\mathcal{N} = \mathrm{Rep}_Q(2, 2)$.

We can see that we may sometimes take advantage of the fact that $\mathrm{GL}(\alpha)$ is dense in $M(\alpha) = \prod_{x \in Q_0} M(\alpha(x))$ to show a representation is in the nullcone. In fact,

Proposition 2.37 (See [DW17], 9.7.5)**.** For all acyclic Q, and all α, all representations V in $\mathrm{Rep}_Q(\alpha)$ are in the nullcone.

For this reason, we will define a new notion of semi-stability that will enable us to get more information from acyclic quivers in the next section. Let's first look at semi-stability for cyclic quivers.

Example 2.38. Let's determine which representations of the 1-loop quiver Q (seen in example 2.14) with dimension vector $\alpha = 2$ are semi-stable. Let V be a representation of the quiver over $\mathbb{C}$, with $V(a) = X$, a 2×2 matrix. If we act on our quiver by B in $\mathrm{GL}(2)$, we get a new representation V', with $V'(a) = BXB^{-1}$. To find the representations in the nullcone of this action, we can find generating invariants, to see for which X, they all vanish. Notice that $\mathrm{Tr}(X)$ and $\det(X)$ are both polynomial in the entries of X, and are invariant under conjugation. Similarly, $\mathrm{Tr}(X^d)$ for any d is an invariant polynomial. Traces of powers generate $\mathbb{C}[\mathrm{End}(\mathbb{C}^n)]^{\mathrm{GL}(n)}$ (See [KP96] 2.4), and from the Cayley-Hamilton theorem, all $\mathrm{Tr}(X^d)$ for $d \geq 2$ are dependent on $\det(X)$ and $\mathrm{Tr}(X)$. This gives us our invariant ring $\mathbb{C}[\mathrm{Rep}_Q(2)]^{\mathrm{GL}(2)} = \mathbb{C}[\mathrm{Tr}(X), \det(X)]$. We can now find which representations are in the nullcone by setting all of our invariants equal to 0. Letting $X = \begin{bmatrix} x & y \\ z & w \end{bmatrix}$, we see that w must equal $-x$, and then $-x^2 - yz = 0$. So our nullcone is made up of the two subsets,

$$\left\{ \begin{bmatrix} x & y \\ \frac{x^2}{y} & -x \end{bmatrix} \middle| x, y \in \mathbb{C}, y \neq 0 \right\}, \text{ and } \left\{ \begin{bmatrix} 0 & y \\ 0 & 0 \end{bmatrix} \middle| y \in \mathbb{C} \right\}.$$

In general for cyclic quivers, traces of cyclic paths will always be invariant polynomials. Not only are these traces always invariant, they will also always generate the entire invariant ring, stated in the following theorem. We introduce the notation $V(p) = V(a_\ell) \dots V(a_1)$, where p is a path $a_\ell \dots a_1$.

Proposition 2.39 (Le Bruyn-Procesi, [LBP90]). For k with characteristic 0, the invariant ring $k[\mathrm{Rep}_\alpha(Q)]$ is generated by the traces $\mathrm{Tr}(V(p))$, for all cyclic paths p.

Bounds on generator degree for quiver representation invariants over fields including those of positive characteristic can be found in [DM17a].

2.2.2 σ-semi-stability of Quiver Representations

As we have seen, there are no semi-stable representations in $\mathrm{Rep}_Q(\alpha)$ when Q is acyclic. As we saw in the previous section, we could take advantage of the density of $\mathrm{GL}(\alpha)$ in $M(\alpha)$ to show any representation is in the nullcone. To combat this, we will introduce a weight σ, and instead look at what we will call σ-semi-stability. Note that in this section, we will assume all quivers are acyclic. For a weight σ in $\mathbb{Z}^{Q_0}$, we define the 1-dimensional representation χ_σ, a *character* with action of $\mathrm{GL}(\alpha)$ given by multiplication by

$$\chi_\sigma(Y(x), x \in Q_0) = \prod_{x \in Q_0} \det(Y(x))^{\sigma(x)}.$$

Definition 2.40. A representation W is σ-semi-stable if $(W, 1)$ is semi-stable in $\mathrm{Rep}_\alpha(Q) \oplus \chi_\sigma$.

The following theorem is stated in the context of quiver representations, but can be applied for any affine G-variety.

Proposition 2.41 (Hilbert-Mumford Criterion)**.** Given a represenation V, a in k, and a closed orbit $G \cdot (W, b)$ in $\mathrm{Rep}_Q(\alpha) \oplus \chi_\sigma$ with $G \cdot (W, b)$ contained in $\overline{G \cdot (V, a)}$, there exists a one-parameter subgroup $\lambda : k^* \to \mathrm{GL}(\alpha)$ so that $\lim_{t \to 0} \lambda(t) \cdot (V, a) = (W, b)$.

For a dimension vector, $\alpha = \underline{\dim} V$, we denote by $\sigma(\alpha)$ the sum $\sum_{x \in Q_0} \alpha(x)\sigma(x)$. In order to have σ-semi-stability, we must have $\sigma(\alpha) = 0$.

Lemma 2.42 (See [DW17])**.** If V in $\mathrm{Rep}_Q(\alpha)$ is σ-semi-stable, then $\sigma(\alpha) = 0$.

Proof. Assume $\sigma(\alpha)$ is non-zero. Define the one parameter subgroup λ by

$$\lambda(t) = (t \cdot I_\alpha(x) | x \in Q_0),$$

where $I_\alpha(x)$ is the identity map at vertex x. On χ_σ, $\lambda(t)$ acts by multiplication by $t^{\sigma(\alpha)}$. As $\sigma(\alpha)$ is non-zero, then either as t gets arbitrarily large or arbitrarily close to 0, $t^{\sigma(\alpha)} = 0$. This means that $(V, 0)$ is in the orbit closure of $(V, 1)$. But $(V, 0)$ is in the nullcone, as V is ayclic, so V is σ-semi-stable. □

Theorem 2.43 (King's Criterion [Kin94])**.** A representation V in $\mathrm{Rep}_Q(\alpha)$ is σ-semi-stable if and only if $\sigma(\alpha) = 0$ and $\sigma(\underline{\dim} W) \leq 0$ for all subrepresentations W of V.

Please note that this theorem is known as King's Criterion (for quivers), and should be distinguished from what is known as King's College Criteria (for livers)

[OAHW89]. Theorem 2.43 will inform definitions in the next chapter. For now, let's return to Example 2.36 now that we have tools to determine σ-semi-stability.

Example 2.44. Let $\alpha = (2,2)$, $\sigma = (1,-1)$, and $Q = x \xrightarrow{a} y$. Let V be a representation in $\mathrm{Rep}_Q(\alpha)$. Notice if $V(a)$ does not have full rank, we have a subpresentation W with $W(x) = V(a)$, and $W(y) = \mathrm{Im}(V(a))$, with $\dim W(y) < \alpha(y)$, giving us $\dim W(x) - \dim W(y) > 0$. We can also see representations with $V(a)$ less than full rank are not σ-semi-stable as follows: For some basis of $V(x)$ and $V(y)$, we have $V(a) = \begin{bmatrix} x & y \\ 0 & 0 \end{bmatrix}$ or $\begin{bmatrix} x & 0 \\ w & 0 \end{bmatrix}$. We will treat the former case, and note that the later case is similar. Let $\lambda(t)$ be given by the tuple $\left(\begin{bmatrix} t & 0 \\ 0 & 1 \end{bmatrix}, I_2 \right)$. We have

$$\lim_{t \to 0} (\lambda(t) V(a) \lambda(t)^{-1}, 1) = \lim_{t \to 0} \left(\begin{bmatrix} tx & ty \\ 0 & 0 \end{bmatrix}, t^2 \right) = (0,0),$$

so V is in the nullcone.

On the other hand, V is σ-semi-stable if $V(a)$ has full rank, as the only subrepresentations of V are then $0 \to \mathbb{C}^2$ and $0 \to 0$. So the nullcone in this case is made of matrices that do not have full rank.

In the next section, we will see that σ-semi-stability can characterize matrix spaces. Before moving on, we mention that in this case, instead of being concerned with invariants, we would instead look at *semi-invariants.* The space of semi-invariants of weight σ is

$$\mathrm{SI}(Q,\alpha)_\sigma = \big\{ f \in k[\mathrm{Rep}_Q(\alpha)] \big| \forall A \in \mathrm{GL}(\alpha), A \cdot f = \chi_\sigma(A) \cdot f \big\}.$$

We can find the σ-semi-stable points by looking at representations V that do not vanish on all homogeneous non-constant semi-invariants of weight σ.

Example 2.45. Returning to Example 2.44, let $\alpha = (2,2)$, $\sigma = (1,-1)$, and $Q = x \xrightarrow{a} y$. Let V be a representation in $\mathrm{Rep}_Q(\alpha)$, with $V(a) = X$. The action of

(A, B) in $\mathrm{GL}(\alpha)$ takes $(V, 1)$ to $(V', \det(A)\det(B))$, where $V'(a) = BXA^{-1}$. Letting $f(X) = \det(X)$, we see that $(A, B) \cdot f = \det(A)\det(B)f = \chi_\sigma(A, B)f$, and det is a semi-invariant of weight σ. The representations that vanish on det are exactly those without $V(a)$ of full rank. From Example 2.44, we already know these are the representations in the nullcone, so $\det(X)$ must generate the semi-invariants of weight σ.

Though we will not go into more detail on semi-invariants, complexity results in the next chapter rely on bounds of generators for the semi-invariant ring, proved in [DM17b].

CHAPTER III

Non-commutative Rank

3.1 History

Start with an $n \times n$ *linear matrix* $A(\underline{x})$, a matrix with each entry a homogeneous linear polynomial in $\mathbb{Z}[x_1, \ldots, x_m]$, where $\underline{x}$ stands for $x_1, \ldots, x_m$. In 1967, Edmonds asked what the rank of such a linear matrix is over $\mathbb{Q}(x_1, \ldots, x_m)$ [Edm67]. This problem is known as *Edmonds' Problem*, and the decision version of this question, namely, asking whether $A(\underline{x})$ has full rank or not, is called the *symbolic determinant identity testing problem* (SDIT). In this chapter, we consider $A(\underline{x}) = x_1A_1 + \ldots + x_mA_m$, with each A_i an $n \times n$ matrix with entries in any large field k with $|k| > 2n$.

Building on Edmonds' problem, we would like to determine the rank of $A(\underline{x})$ over the free skew field, first defined by Amitsur [Ami66], and described in more detail in [CR99]. This rank is called the *non-commutative rank* of $A(\underline{x})$, $\mathrm{ncrk}(A(\underline{x}))$, defined by Fortin and Reutenauer [FR04]. The question of finding $\mathrm{ncrk}(A(\underline{x}))$ is the *non-commutative Edmonds' problem*, and the relaxation in simply deciding whether $A(\underline{x})$ has full non-commutative rank is the *non-commutative full rank problem* (NCFullRank). Letting $\mathcal{A} = \mathrm{span}\{A_1, A_2, \ldots, A_m\}$, we alternatively denote the non-commutative rank of $A(\underline{x})$ by $\mathrm{ncrk}(\mathcal{A})$. Ivanyos, Qiao, and Subrahmanyam give equivalent formulations and history of NCFullRank in [IQS17]. We will be interested in the c-shrunk subspace, tensor blow-up, and particularly the nullcone formulations, which are discussed in

Section 3.2.

Lots of work from different angles has been done on this non-commutative rank. Cohn and Reutenauer proved NCFullRank was in PSPACE (can be solved using polynomial space) [CR99]. Fortin and Reutenauer connected non-commutative rank explicitly to c-shrunk subspaces [FR04]. Coming from studying non-commutative arithmetic circuits with divisions, Hrubeš and Wigderson proved that non-commutative rank was equivalent to rank for large enough tensor blow-ups [HW14]. Garg, Gurvitz, Oliveira, and Wigderson provide a polynomial time algorithm of non-commutative rank for fields of characteristic zero [GGOW16]. In [IKQS15], for certain matrix spaces, Karpinski, Ivanyos, Subrahmanyam, and Qiao use Wong sequences to calculate the non-commutative rank. Building on this using blow-ups, the latter three authors provide an algorithm for finding the non-commutative rank of any matrix space [IQS17]. Utilizing results on bounds from [DM17b], in [IQS18], they give a deterministic polynomial time algorithm. This algorithm returns a c-shrunk subspace certifying the non-commutative rank and works over any large enough field k. Additional results on the ratio between commutative rank and non-commutative rank are proved in [DM18b].

In this chapter we generalize this algorithm to one measuring semi-stability of quiver representations, giving us new results in this broader context. We also apply non-commutative rank to general ext and hom, studied by Kac, Schofield, and Crawley-Boevey [Kac82, Sch92, CB96]. Much of this Chapter is based on a paper on arXiv, for which I'd like to thank Chi-yu Cheng, Calin Chindris, Harm Derksen, Daniel Kline, and Visu Makam for comments.

3.2 Definitions

We will be concerned with the free skew field, made up of non-commuting polynomials, $k\langle x_1, \dots, x_m \rangle$, their inverses, and then enlarged to contain all sums, products, and inverses. The free skew field was first defined by Amitsur, [Ami66]. In the free skew field, there is no standardized way to express elements, and elements may need to be defined with nested inverses. For example, $(x + yz^{-1}w)^{-1}$ can not be written without a nested inverse [HW14].

Given a linear matrix, $A(\underline{x})$, with homogeneous linear polynomials in $k\langle x_1, \ldots, x_m\rangle$, the non-commutative analogue of the Edmonds' problem asks to determine the rank of $A(\underline{x})$ over the free skew field. We denote this rank by $\text{ncrk}(\mathcal{A})$. Similarly, the NCFullRank problem asks whether $A(\underline{x})$ has full rank over the free skew field. For example, we row reduce the following skew symmetric matrix over the free skew field to get:

$$T = \begin{bmatrix} 0 & x & y \\ -x & 0 & z \\ -y & -z & 0 \end{bmatrix} \sim \begin{bmatrix} 0 & x & y \\ -x & 0 & z \\ 0 & 0 & zx^{-1}y - yx^{-1}z \end{bmatrix}. \tag{3.1}$$

Unfortunately, by the nature of the free skew field, it is hard to determine polynomial identities — so it is not immediately clear if this linear matrix has non-commutative rank 2 or 3. For this reason, we explore additional equivalent formulations of non-commutative rank. For proof sketches on their equivalence, see [IQS17]. We can also define the non-commutative row and column rank of $A(\underline{x})$ by taking the row (respectively column) span over the free-skew field, leaving us with a finitely generated free module, of which we can take the rank.

We note here that many aspects of rank carry over to the non-commutative rank, for instance, the non-commutative row rank and column rank of $A(\underline{x})$ equal the $\text{ncrk}(A)$ and we must have a minor with full non-commutative rank equal to $\text{ncrk}(A)$. We must still be careful, as other aspects do not: naively finding the "determinant" of $A(\underline{x})$, and comparing it to zero will not tell us whether the non-commutative rank is full. In fact, even how to define a single determinant in this context is unclear [GR91].

3.2.1 Blow-ups

If $A(\underline{x}) = x_1A_1 + \ldots x_mA_m$, let $\mathcal{A} = \text{Span}\{A_1, \ldots, A_m\}$. The dth *tensor blow-up* of $\mathcal{A}$ is

$$\mathcal{A}^{\{d\}} := M_{d,d} \otimes \mathcal{A} \subseteq M_{dn,dn}.$$

The rank of a matrix space, $\operatorname{rk}\mathcal{A}$, is the maximal r so that there is a matrix with rank r in $\mathcal{A}$. When the base field k is large enough, d divides the rank of $\mathcal{A}^{\{d\}}$ [IQS17]. We have

$$\text{ncrk}(\mathcal{A}) = \lim_{d\to\infty} \frac{\operatorname{rk}\mathcal{A}^{\{d\}}}{d}.$$

The value of $(\operatorname{rk}\mathcal{A}^{\{d\}})/d$ is increasing as d increases, and is bounded by n. Using results of Derksen and Makam, if $\mathcal{A}$ has maximal non-commutative rank, then taking $d \geq n-1$ ensures $\operatorname{rk}\mathcal{A}^{\{d\}} = nd$ [DM17b]. If $\text{ncrk}(\mathcal{A}) = r < n$, then restricting to a full rank $r \times r$ submatrix of $A(\underline{x})$, we see that $\operatorname{rk}\mathcal{A}^{\{d\}} = nd$ for $d \geq r-1$. So we always have $\operatorname{rk}\mathcal{A}^{\{d\}} = nd$ for $d \geq n-1$.

Example 3.1. For our Example 3.1, take $d = 2$. We then look for 2×2 matrices D_1, D_2, D_3, so that $A_1 \otimes D_1 + A_2 \otimes D_2 + A_3 \otimes D_3$ has max rank. Letting

$$D_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \ D_2 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \ D_3 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix},$$

We find

$$\operatorname{rk}\left[\begin{array}{c|c|c} 0 & D_1 & D_2 \\ \hline -D_1 & 0 & D_3 \\ \hline -D_2 & -D_3 & 0 \end{array}\right] = 6,$$

which must be maximal, and so $\text{ncrk}(T) = 3$.

3.2.2 c-shrunk Subspaces

Definition 3.2. A subspace $U \subseteq k^n$ is a *c-shrunk subspace* of $\mathcal{A}$ if there exists a subspace $W \subseteq k^n$ with $\dim(W) \leq \dim(U) - c$, and for every A in $\mathcal{A}$, $A(U) \subseteq W$.

Example 3.3. Let $\mathcal{A} = \begin{bmatrix} x & 0 \\ y & 0 \end{bmatrix}$, and let e_1, e_2 be a basis for k^2. The entire space k^2 is 0-shrunk, as its image (under all matrices in $\mathcal{A}$) is all of k^2. The span of e_1 is not c-shrunk for any non-negative c, as its image is all of k^2, a space of higher dimension. The span of e_2 is in the kernel of any matrix in $\mathcal{A}$, and so it is a 1-shrunk subspace.

The NCFullRank problem is equivalent to determining whether $\mathcal{A}$ has no c-shrunk subspace for $c > 0$ [Coh95]. More generally we can define the non-commutative rank using c-shrunk subspaces [FR04].

Proposition 3.4 ([FR04])**.** Given a matrix space of $n \times n$ matrices $\mathcal{A}$, the non-commutative rank of $\mathcal{A}$, denoted $\text{ncrk}(\mathcal{A})$ is $n - \max\{c \mid \text{there is a } c\text{-shrunk subspace of } \mathcal{A}\}$.

Example 3.5. Returning to the previous example, for any non-zero a, b in k, the span of $ae_1 + be_2$ is not c-shrunk (it hits all of k^2), and the zero subspace is 0-shrunk. So we get $\text{ncrk}(\mathcal{A}) = 1$, as 1 is the max c so that $\mathcal{A}$ has a c-shrunk subspace.

Throughout the rest of this thesis, we let $c = n - \text{ncrk}(\mathcal{A})$, i.e. all c-shrunk subspaces discussed are so that c is maximal.

Lemma 3.6. Let $c = n - \text{ncrk}(\mathcal{A})$. If U_1, U_2 are c-shrunk subspaces of $\mathcal{A}$, then so are $U_1 \cap U_2$ and $U_1 + U_2$.

Proof. By assumption $\dim U_i - \dim \mathcal{A}(U_i) = c$. Let $U_3 = U_1 \cap U_2$, and $U_4 = U_1 + U_2$. We then have:

$$\begin{aligned} c + c &\geq (\dim U_3 - \dim \mathcal{A}(U_3)) + (\dim U_4 - \dim \mathcal{A}(U_4)) = \\ &= (\dim(U_1 \cap U_2) + \dim(U_1 + U_2)) - (\dim(\mathcal{A}(U_1) \cap \mathcal{A}(U_2)) + \dim(\mathcal{A}(U_1) + \mathcal{A}(U_2))) \geq \\ &\geq (\dim U_1 + \dim U_2) - (\dim \mathcal{A}(U_1) + \dim \mathcal{A}(U_2)) = \\ &= (\dim U_1 - \dim \mathcal{A}(U_1)) + (\dim U_2 - \dim \mathcal{A}(U_2)) = c + c. \end{aligned}$$

We conclude that $\dim U_3 - \dim \mathcal{A}(U_3) = \dim U_4 - \dim \mathcal{A}(U_4) = c$, as c is maximal. Therefore, U_3 and U_4 are c-shrunk subspaces. □

In particular, there is a unique c-shrunk subspace of the lowest dimension, namely, the intersection of all c-shrunk subspaces. Similarly, there is a unique largest c-shrunk subspace, the sum of all c-shrunk subspaces. A recent similar discussion can be found in [IMQ21]. In our skew-symmetric example (3.1), although any matrix in $\mathcal{A}$ has rank 2, the image of any subspace U of k^3 has an equal or larger dimension than U. In this case $c = n - \text{ncrk}(\mathcal{A}) = 3 - 3 = 0$, and the minimal c-shrunk subspace is the zero subspace.

3.2.3 Semi-stability of Generalized Kronecker Quiver

As promised, in this subsection we will define non-commutative rank using quivers and semi-stability. We begin by using quiver representations to determine whether a linear matrix has full non-commutative rank or not, the NCFullRank problem. For $A(\underline{x}) = x_1A_1 + \ldots, x_mA_m$, this determination is equivalent to determining whether the quiver representation W,

$$k^n \underset{A_m}{\overset{A_1}{\underset{\longrightarrow}{\overset{\longrightarrow}{\vdots}}}} k^n$$

is σ-semistable for $\sigma = (1, -1)$. In our skew-symmetric matrix example, (3.1), we would like to determine whether the above quiver with

$$A_1 = \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \; A_2 = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix}, \; A_3 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}$$

is $(1, -1)$-semi-stable.

We would also like to be able to relate quiver representations to the non-commutative rank, rather than just to NCFullRank. To do this, we need a way of measuring how far a representation V is from being σ-semistable. For this, we use Theorem 2.43. For a representation W, recall $\sigma(\underline{\dim}(W)) = \sum \dim(W(x))\sigma(x)$, summed over all vertices x.

Proposition 3.7. Given $\mathcal{A} = \text{Span}\{A_1, \ldots, A_m\}$, and c the maximum $\sigma(\underline{\dim}(W'))$ over all subrepresentations W' of the generalized Kronecker quiver with maps $A_1, \ldots, A_m$, the non-commutative rank of $\mathcal{A}$, $\text{ncrk}(\mathcal{A})$, is $n - c$.

Proof. If W is a representation of the generalized Kronecker quiver with m-arrows, and $\sigma = (1, -1)$, let W' be a subrepresentation with $c := \sigma(\underline{\dim}(W'))$ maximal. Then, since $\mathcal{A}(W'(x))$ is contained in $W'(y)$, $W'(x)$ is a c-shrunk subspace. On the other hand, if instead we start with a c-shrunk subspace U, we can define $W'(x) := U$, and $W'(y) := \mathcal{A}(U) = \sum_{i=1}^{m} A_i U$. This defines a subrepresentation W', where $\sigma(\underline{\dim}(W')) = c$. So for generalized Kronecker quivers, c-shrunk subspaces give us subrepresentations W' with $\sigma(\underline{\dim}(W'))$ maximal, and vice-versa. $\square$

So, the non-commutative rank of $\mathcal{A}$ is equal to the maximum of $\sigma(\underline{\dim}(W'))$ over all subrepresentations W' of the generalized Kronecker quiver with maps $A_1, \ldots, A_n$.

3.3 Non-commutative Rank and Semi-stability

In this Section, we will generalize results for the non-commutative rank, related to generalized Kronecker quivers, to arbitrary acylic quivers. We will let Q be an arbitrary acyclic quiver. Through Proposition 3.7, if a representation W is not σ-semi-stable, we can still measure its closeness to σ-semi-stability by finding a subrepresentation W' with $\sigma(\underline{\dim}(W'))$ maximal. In [CK21], this $\sigma(\underline{\dim}(W'))$ is called the *discrepancy* of (W, σ). Note that we are now no longer limited to the generalized Kronecker quiver — we can ask this question for any acyclic quiver representation W, for any σ. We call a subrepresentation W' which maximizes $c = \sigma(\underline{\dim}(W'))$ an optimal σ-witness. When σ is understood, we call this W' an *optimal witness*. We can generalize Lemma 3.6 for subrepresentations.

Proposition 3.8. If W_1, W_2 are optimal σ-witnesses of W, then so are $W_1 \cap W_2$ and $W_1 + W_2$. In particular, there is a minimal and maximal optimal σ-witness.

Proof. Let c be the discrepancy of (W, σ). Let $\sigma_+(x) = \max\{0, \sigma(x)\}$, and similarly, $\sigma_-(x) = -\min\{0, \sigma(x)\}$. For $i = 1, 2$, by assumption

$$\sigma(\underline{\dim}(W_i)) = \sum_{x \in Q_0} \big(\sigma_+(x) - \sigma_-(x)\big) \dim W_i(x) = c.$$

Let $W_3 = W_1 \cap W_2$, $W_4 = W_1 + W_2$. We then have:

$$\begin{aligned} c + c \geq \sum \big(\sigma_+(x) - \sigma_-(x)\big) \dim W_3(x) + \big(\sigma_+(x) - \sigma_-(x)\big) \dim W_4(x) = \\ = \sum \sigma_+(x) \dim W_3(x) - \sigma_-(x) \dim W_3(x) + \sigma_+(x) \dim W_4(x) - \sigma_-(x) \dim W_4(x) = \\ = \sigma_+(x)\big(\dim W_1(x) + \dim W_2(x)\big) - \sigma_-(x)\big(\dim W_1(x) + \dim W_2(x)\big) = c + c. \end{aligned}$$

We conclude that $\sigma(\underline{\dim}(W_3)) = \sigma(\underline{\dim}(W_4)) = c$, as c is maximal. Therefore, W_3 and W_4 are optimal σ-witnesses. We can find a minimal optimal σ-witness by taking the intersection of all optimal σ-witnesses, and similarly find a maximal optimal σ-witness by taking the sum of all optimal σ-witnesses. □

We would like to extend techniques in [IQS17] which find c-shrunk subspaces of matrix spaces to find an optimal σ-witness of a quiver representation. To do this, we reduce any acyclic quiver to the generalized Kronecker quiver. We use the construction described in [DM18a], but provide an altered set up, using presentations as in [DF15]. Let P_x be the indecomposable projective representation of Q with basis given by all paths starting at vertex x, as defined in Subsection 2.1.3.

Let P_x be the indecomposable projective representation corresponding to vertex x. So, $P_x(y) = e_y kQ e_x$, with basis given by paths from x to y. Let $\mathbf{P_1} := \bigoplus_{x \in Q_0} {P_x}^{\sigma_-(x)}$, and $\mathbf{P_0} := \bigoplus_{x \in Q_0} {P_x}^{\sigma_+(x)}$. Consider all possible morphisms φ between the quiver representations

$$\varphi : \mathbf{P_1} \to \mathbf{P_0}.$$

To our above set of morphisms, apply $\mathrm{Hom}_Q(\_, W)$ to get

$$A(\varphi) : \mathrm{Hom}_Q(\mathbf{P_0}, W) \to \mathrm{Hom}_Q(\mathbf{P_1}, W),$$

where $A(\varphi) := \mathrm{Hom}(\varphi, W)$. We can consider a subspace $\mathrm{Hom}(P_x{}^{\sigma_+(x)}, W')$ as $Z_+(x) \otimes W'(x)$ for some $Z_+(x) = k^{\sigma_+(x)}$. Notice $\mathrm{Hom}_Q(\mathbf{P_0}, W)$ is a right $\mathrm{End}(\mathbf{P_0})$-module by precomposition. Let x, y be so that both $\sigma_+(x)$ and $\sigma_+(y)$ are positive. Note $\mathrm{End}(\mathbf{P_0})$ contains $H = \prod_{x \in Q_0} \mathrm{GL}(\sigma_+(x))$, a reductive group, which acts on the $Z_+(x)$, leaving the $W(x)$ alone. So, an H-subrepresentation of $\bigoplus_{x \in Q_0} Z_+(x) \otimes W(x)$ must be of the form $\bigoplus_{x \in Q_0} Z_+(x) \otimes W'(x)$ for some subspaces $W'(x)$ of each $W(x)$. Our set of maps can also be considered between the spaces

$$A(\varphi) : \bigoplus_{x \in Q_0} W(x)^{\sigma_+(x)} \to \bigoplus_{x \in Q_0} W(x)^{\sigma_-(x)}.$$

Now, we have a matrix space $\mathcal{A}$ consisting of all $A(\varphi)$. This is the space of block matrices with blocks mapping $W(x)$ to $W(y)$ given by a linear combination of $W(p)$, where p is a path from x to y. For this new generalized Kronecker Quiver, we may run the algorithm in [IQS17] to get the minimal c-shrunk subspace of $\bigoplus_{x \in Q_0} W(x)^{\sigma_+(x)}$, U.

Lemma 3.9. The minimal c-shrunk subspace, $U \subseteq \mathrm{Hom}_Q(\mathbf{P_0}, W)$, is a right $\mathrm{End}(\mathbf{P_0})$ module, and $\sum_\varphi A(\varphi)U$ is a right $\mathrm{End}(\mathbf{P_1})$ module.

Proof. First, we prove that given any c-shrunk subspace, U, and invertible T in $\mathrm{End}(\mathbf{P_0})$, $U \cdot T$ is also c-shrunk. We have the image of $U \cdot T$:

$$\sum_\varphi A(\varphi)(U \cdot T) = \sum_\varphi A(T \cdot \varphi)U = \sum_\varphi A(\varphi)U.$$

Here the sum is taken over all morphisms φ as above. It follows that

$$\dim \sum_{\varphi} A(\varphi)(U \cdot T) = \dim \sum_{\varphi} A(\varphi)U.$$

As T is an automorphism, we also have $\dim U \cdot T = \dim U$, so $U \cdot T$ is c-shrunk. If U is the minimal c-shrunk subspace, $U \cdot T$ is also c-shrunk and of the same dimension, so $U \cdot T = U$. As $\mathrm{End}(\mathbf{P_0})$ is spanned by invertible elements, this shows that the minimal c-shrunk subspace U is a right $\mathrm{End}(\mathbf{P_0})$ module. Similarly, given S in $\mathrm{End}(\mathbf{P_1})$, we see that

$$\sum_{\varphi} A(\varphi)(U) \cdot S = \sum_{\varphi} A(\varphi \cdot S)U = \sum_{\varphi} A(\varphi)U.$$

□

Theorem 3.10. Given the minimal c-shrunk subspace for the set of linear maps

$$A(\varphi) : \bigoplus_{x \in Q_0} W(x)^{\,\sigma_+(x)} \to \bigoplus_{x \in Q_0} W(x)^{\,\sigma_-(x)},$$

we can construct a subrepresentation of W, W', so that $\sigma(\underline{\dim}(W'))$ is maximal. Furthermore, $\sigma(\underline{\dim}(W')) = c$.

Proof. Considered as a subspace of $\bigoplus Z(x) \otimes W(x)$, the minimal c-shrunk U is of the form $\bigoplus Z(x) \otimes W'(x)$, for some subspaces $W'(x)$ of $W(x)$. For y so that $\sigma_+(y) = 0$, define $W'(y) = \sum_{a:x\to y} W(a)W'(x)$. This ensures we have a subrepresentation. Note that $c \leq \sum \dim(W'(x)^{\sigma_+(x)}) - \sum \dim(W'(x)^{\sigma_-(x)})$, but c is maximal, so $\sigma(\underline{\dim}(W')) = c$. We note that the $W'(y)$ are similarly closed under the action of $\mathrm{End}(\mathbf{P_1})$.

If there were a subrepresentation W'' with $\sigma(\underline{\dim}(W''))$ less than c, Note that $U' = \bigoplus_{x \in Q_0} W''(x)^{\sigma_+(x)}$ is a shrunk subspace, with $\dim(U') - \dim(\mathcal{A}(U')) > c$, so c would not be maximal. □

3.3.1 Algorithms

After using this reduction of a quiver representation to a generalized Kronecker quiver, we can employ any previous algorithms or other techniques for finding a c-shrunk subspace. If we successfully find a c-shrunk subspace, U, that is not minimal, we can construct a c-shrunk subspace that is fixed under the action of $\mathrm{End}(\mathbf{P_0})$ by taking instead

$$\bigcap_{T \in \mathrm{Aut}(\mathbf{P_0})} U \cdot T.$$

Such a subspace will give a optimal σ witness.

Example 3.11. As we have seen, the space of the skew symmetric 3×3 matrices have non-commutative rank 3. The zero subspace is the minimal 0-shrunk subspace. However, all of k^3 is also 0-shrunk. This example is already reduced to generalized Kronecker form, with underlying quiver

$$x \; \substack{\longrightarrow \\ \longrightarrow \\ \longrightarrow} \; y$$

dimension vector $(3, 3)$ and weight and $\sigma = (1, -1)$. As $\sigma_+ = (1, 0)$, the representation $\mathbf{P_0}$ is given by simply P_x, the indecomposible projective representation with dimension vector $(1, 3)$. Any invertible morphism from P_x to itself is completely determined by a non-zero constant. So acting on k^3 by $\mathrm{Aut}(\mathbf{P_0})$ is then simply acting by non-zero constants, which will not change the space. We then see all of k^3 is also an acceptable c-shrunk subspace for building an optimal σ witness.

In [IKQS15], Wong sequences, originally defined by Kai-Tek Wong [Won74], are used in certain cases to find a c-shrunk subspace.

Algorithm 3.12. Second (generalized) Wong sequence [IKQS15].
Input: Matrix space of $n \times n$ matrices $\mathcal{A}$, spanned by $A_1, \ldots, A_m$, random matrix A in $\mathcal{A}$.

Output: Limit of Wong sequence, W^*

$W_0 = 0$;
for $i = 1$ **to** n **do**
    $W_i = \sum_{j=1}^{m} A_j(A^{-1}(W_{i-1}))$;
$W^* = W_n$

Note we can end the loop as soon as $W_i = W_{i-1}$. The limit of this sequence gives insights on the existence of c-shrunk subspaces of $\mathcal{A}$.

Proposition 3.13. [IKQS15] There is an $(n - \mathrm{rk}(A))$-shrunk subspace of $\mathcal{A}$ if and only if the limit of the Wong sequence, $W*$, is contained in the image of A. In this case, $A^{-1}(W^*)$ is an $(n - \mathrm{rk}(A))$-shrunk subspace of $\mathcal{A}$, and the non-commutative rank of $\mathcal{A}$ is A.

However, Algorithm 3.12 alone will not always find the non-commutative rank of $\mathcal{A}$, in the case W^* is not contained in the image of A, we need to search for a higher rank matrix and try the algorithm again. Even then, as we can see from our own Example 3.1, where there are no rank 3 matrices in our space of skew-symmetric matrices, this may be impossible. To combat this, in [IQS17], the algorithm is used in some blow-up of $\mathcal{A}$, where there must be some A with rank equal to $\mathrm{ncrk}(\mathcal{A})$ times the size of the blowup. When a blow-up must be used, the c-shrunk subspace of $\mathcal{A}$ can still be found by projecting.

We claim when Algorithm 3.12 returns a c-shrunk subspace, the subspace is the minimal c-shrunk subspace. Recall the minimal c-shrunk subspace, U, is the intersection of all c-shrunk subspaces, so $U \subseteq A^{-1}(W^*)$. The limit of the sequence W^* is the smallest subspace Z so that $\bigcup_{i=1}^{m} A_i^{-1}(Z)$ contains $A^{-1}(Z)$. So by minimality, U is returned when this sequence terminates with W^* contained in $\mathrm{Im}(A)$. In the case where the dth blow-up is invoked to find a c-shrunk subspace, the same sequence is used in the larger space, finding a $cd-$shrunk subspace. This by the same reasoning must be minimal, so when pulled back to a c-shrunk subspace in the original space, it must remain minimal.

Let $n = \min\{\sum \sigma_+(x) \dim W(x), \sum \sigma_-(y) \dim W(y)\}$. For sufficiently large fields, ($|k| > 2n$) there is a randomized algorithm to find a c-shrunk subspace [IQS17, Corollary 1.5]. This randomized algorithm is much simpler and typically must faster than the deterministic algorithm. In the context of representations, this algorithm immediately after reduction, blows up by a sufficiently large [IQS18, DM17b] factor, $d \geq n-1$. In this blow-up, randomly choose a matrix

$$A : \bigoplus_{x \in Q_0} W(x)^{\,d\sigma_+(x)} \to \bigoplus_{x \in Q_0} W(y)^{\,d\sigma_-(y)}, \tag{3.2}$$

where A is in $\mathcal{A}^{\{d\}} := M_{d,d} \otimes \mathcal{A}$. Through the Schwartz-Zippel-DeMillo-Lipton Lemma [Sch80, Zip79, DL78], which we will discuss more thoroughly in Section 3.3.2, if a field is large enough, evaluating a non-zero polynomial over that field at a randomly chosen point is likely to give a non-zero result. Taking the determinant of minors of a matrix in the blow-up, we are likely to have $\operatorname{rk} A = \operatorname{ncrk} \mathcal{A}^{\{d\}}$. Thus, running the Wong sequence on this A will result in the return of a $cd-$shrunk subspace [IKQS15, Lemma 9]. From this cd-shrunk subspace in the blow-up, we can find a c-shrunk subspace of $\bigoplus_{x \in Q_0} W(x)^{\,\sigma_+(x)}$, constructing a subrepresentation as above.

The deterministic Wong sequence algorithm for finding non-commutative rank, introduced in [IKQS15], uses a sequence of subspaces, testing its limit, W^* for evidence of a c-shrunk subspace. In the quiver representation context, we would like to instead use a sequence of subrepresentations. In this deterministic setting, we only need $|k| > n$.

To do this, we again start with a random matrix A in the blow-up, as above. Next, find a pseudo-inverse of A, a matrix B so that B's restriction to $\operatorname{Im}(A)$ is the inverse to A's restriction to a direct complement of $\ker(A)$. Note that B is a block matrix as well, with blocks mapping each $W(y)$ for $\sigma(y) < 0$ to each $W(x)$ with $\sigma(x) > 0$. Let I_x index the $|d\sigma(x)|$ copies of $W(x)$. Let $\pi_{x,i} : \bigoplus_{x \in Q_0} W(x)^{d\sigma_+(x)} \to W(x)$ be the projection to the ith copy of $W(x)$. Each projection can be thought of as coming from the action of $\mathrm{End}(\mathbf{P_0})$. Similarly, define this for vertices y with $\sigma_-(y) > 0$.

For each block, take the projection $\pi_{y,i} B \pi_{x,j}$. This gives a linear map from $W(x)$ to $W(y)$. Construct a new quiver representation, W^+, on a new quiver Q^+ by adding arrows $p : y \to x$ for each block in the pseudo-inverse, with each $W^+(p)$ defined as $\pi_{y,i} B \pi_{x,j}$.

Define a subspace at vertices x with $\sigma(x) > 0$ of W^+:

$$K(x) := \sum_{i \in I_x} \pi_{x,i} \ker(A).$$

For all other vertices, define $K(y) = 0$. Let W' be the smallest subrepresentation of W^+ containing each $K(x)$. Note that W' must also be a subrepresentation of our original W.

Proposition 3.14. For W' as defined above, $\bigoplus_{x \in Q_0} W'(x)^{d\sigma_+(x)}$ is cd-shrunk, with image (under $\mathcal{A}^{[d]}$) $\bigoplus_{x \in Q_0} W'(x)^{d\sigma_-(x)}$. Thus, W' is an optimal σ witness.

First, we claim that $\bigoplus_{x \in Q_0} W'(x)^{d\sigma_+(x)}$ is the minimal cd-shrunk subspace of $\mathcal{A}^{\{d\}}$. By construction, the Wong sequence algorithm returns the smallest subspace containing $\ker(A)$, and closed under $\mathcal{A}^{\{d\}}$ and our pseudo-inverse B. The $K(x)$ must remain inside the minimal shrunk subspace, as the projections come from $\mathrm{End}(\mathbf{P_0})$. Similarly, the new maps in W^+ come from the action of $\mathrm{End}(\mathbf{P_0}) \bigoplus \mathrm{End}(\mathbf{P_1})$, so in finding the smallest subrepresentation, we must still remain in the minimal shrunk subspace (at positive vertices). So in finding the minimal representation of W^+ that contains each $K(x)$, W', we get the smallest subspace $\bigoplus_{x \in Q_0} W'(x)^{d\sigma_+(x)}$ containing $\ker(A)$ and closed under $\mathcal{A}^{\{d\}}$ and B, i.e. the minimal cd-shrunk subspace.

Proposition 3.15. Given a quiver representation W, a weight vector σ, $|k| > n$, letting $n_x := \dim(W(x))$, and $N = \sum_{x \in Q_0} n_x$, there is an algorithm polynomial time in the n_x to find an optimal σ witness.

Recalling the above discussion, we first construct Q^+ and W^+. To do this, we choose a random matrix in the $d = \min\{\sum(n_x\sigma_+(x)), \sum(n_y\sigma_-(y))\} - 1$ blowup, A, and find its pseudo-inverse, B, which takes polynomial time ($\leq (dN)^3$). We then construct new linear maps for each of the $d^2\sigma_+\sigma_-$ blocks in B by composing B with projection maps. This composition is matrix multiplication, which can be done in polynomial time. Next we contruct $K(x)$ at each vertex x, which is the sum over the $d\sigma(x)$ projections of $\ker(A)$. We can find a basis for $\ker(A)$ itself in polynomial time using row reduction. Last, we use Algorithm 3.16 to loop through all our arrows N times, to find the optimal σ witness, W' from the $K(x)$. This algorithm will stabilize at the Nth loop or shorter, as each iteration of the outside loop will either raise the dimension of the current W', or will not (in which case, we are done, we have found the final W'). We can increase the dimension at most N times, so this must be a correct bound for the number of times to run the outer loop. Note we may terminate the outer loop as soon as the updates from the inner loop do not change W' at all.

Algorithm 3.16. Algorithm for finding W'.

Input: Quiver Q, Representation W of Q, subspaces $K(x) \subseteq W(x)$ for all vertices x.

Output: Smallest subrepresentation W' so that $K(x) \subseteq W'(x)$ for all vertices x.

$N = \sum_{x \in Q_0} \dim(W(x))$;

$W'(x) = K(x)$ for all x;

for $i = 1$ **to** N **do**

  for $a \in Q_1$ **do**

    $W'(ha) = W'(ha) + W(a)W'(ta)$;

Summarizing our algorithm, we start with a quiver Q, representation W, and weight vector σ. Our steps to find an optimal σ witness are as follows:

1. **Take a random matrix A in the blow-up of the generalized Kronecker reduction** (3.2). Here, we get a block matrix with d blocks for each of the $\sigma_+(x)$ copies of $W(x)$ mapping to each of the $\sigma_-(x)$ copies of $W(y)$. Inside each block is a random linear combination of $W(p)$, where p is a path from x to y in the quiver. The size of this matrix is $O((dN)^2)$.

2. **Get a pseudo-inverse of this matrix.** For an $i \times j$ or $j \times i$ matrix with $i >= j$, there's an algorithm to do this in $O(i^2 j)$ time. So here we take $O(d^3 N^3)$ time.

3. **Use projections to find maps for new quiver Q^+.** Here we are doing a matrix multiplication to project each of the blocks to a new linear map. A direct algorithm in this case would take $O(d^5 N^3)$ time, though we note there are faster algorithms for matrix multiplication.

4. **Use kernel of A to construct subspaces $K(x)$ at each vertex x.** Calculating the kernel of A using Gaussian elimination, we take $O((dN)^3))$ time. Then, we construct $K(x)$ for each x, which is a projection of the kernel - the kernel has max dimension $(\sum_{x \in Q_0} d\sigma_+(x) W(x))$, so we take $O((dN)^3)$ time.

   Now, we have a new quiver, $Q+$, with new arrows for each of the blocks in the pseudo-inverse, and a subspace $K(x)$ at each vertex. All that is left is to find the minimal subrepresentation that contains the $K(x)$s at each vertex.

5. **Build larger and larger representations of Q^+ containing $K(x)$ at each vertex using Algorithm 3.16.** This algorithm loops through N times for each of the $O(d^2)$ arrows to find the projection of current $W'(ta)$ onto $W'(ha)$ using $W^+(a)$, and accordingly, add it to $W'(ha)$. This multiplication is at most order $\max(\dim(W(x))^3$, and row reducing to find a new basis would also just be on the order at most $\max(\dim(W(x))^3$. So our loop overall is in $O(d^2 N^4)$.

As d is in $O(N)$, overall, our algorithm is in polynomial time in N. If our choice of $|\sigma|$ is very large (outside $O(n)$), the algorithm is polynomial in N and $|\sigma|$.

3.3.2 Additional Discussion and Examples

In this section, we will discuss the Schwartz-Zippel-DeMillo-Lipton lemma as promised, followed by a discussion on general objects, and examples of σ optimal witnesses.

Lemma 3.17 (Schwartz-Zippel-DeMillo-Lipton [Sch80, Zip79, DL78])**.** Given a non-zero polynomial f in $k[x_1, \ldots, x_n]$ of degree d, and a finite subset S of k, with $a_1, \ldots, a_n$ chosen independently and uniformly random from S,

$$\mathbb{P}(f(a_1, \ldots, a_n) = 0) \leq \frac{d}{|S|}.$$

Proof. We will prove this by induction on n. When $n = 1$, we have a degree d one variable polynomial, $f(x)$, which has at most d roots. So $\mathbb{P}(f(x) = 0) \leq \frac{d}{|S|}$. Assume now this holds for polynomials in up to n variables. Let $f(x_1, \ldots, x_n, x_{n+1})$ be a non-zero polynomial of degree d, and let $a_1, \ldots, a_n, a_{n+1}$ be our randomly chosen points from S. Factor out all instances of x_{n+1}, writing the polynomial as

$$f(x_1, \ldots, x_n, x_{n+1}) = \sum_{i=0}^{d} x_{n+1}^i f_{d-i}(x_1, \ldots, x_n).$$

Each f_{d-i} is a polynomial of degree $d-i$. Let f_k be of smallest degree while remaining non-zero (there must be at least one such f_k, else our original f is zero). Now, we have two cases, case $f_k(a_1, \ldots, a_n) = 0$. The probability of this happening is at most $\frac{|k|}{|S|}$ by induction. But we are concerned with this happening along with the the entire polynomial evaluating to 0, which is an even less likely event, so case one happens with a probability at most $\frac{|k|}{|S|}$. If we are in the other case, where $f_k(a_1, \ldots, a_n) \neq 0$, we need the probability of this happening along with the polynomial $f(a_1, \ldots, a_n, x_{n+1})$

evaluating to 0 at a_{n+1}. By our base case, the latter happens with probability at most $\frac{d-k}{|S|}$, and so similarly, the probability in this case overall is still at most $\frac{d-|k|}{|S|}$. Adding the probabilities in each disjoint case, we see the probability is bounded above by $\frac{d}{|S|}$. $\square$

The Schwartz-Zippel-DeMillo-Lipton lemma is useful in polynomial identity testing, as plugging in a value is much faster than multiplying out factored expressions. For example, it is much easier to plug in values into a determinant polynomial than it is to multiply out the polynomial. However, we can only use this lemma when we are okay with randomized algorithms, and cannot rely on it for deterministic algorithms.

We now slightly switch gears to define the term "general" which we will use in future examples and sections.

Definition 3.18. Given an algebraic variety X with the Zariski topology, a given property P is *general* if it holds on an open dense subset of X. We say a point x in X is *general* if it has property P.

Example 3.19. Letting k be an infinite field, A general $n \times n$ matrix over k has rank n. This is because all such matrices lie outside the closed set defined by $\det(\underline{x}) = 0$, and are therefore on an open set. As $k^{n\times n}$ is irreducible, any non-zero open subset is dense.

Example 3.20. We will now explicitly run through the algorithm to find an optimal σ witness for a simple quiver representation. Let Q be the following quiver:

$$x \underset{b}{\overset{a}{\rightrightarrows}} y \xrightarrow{c} z,$$

let $\sigma = (1, -1, 0)$, and let W be

$$\mathbb{C}^2 \underset{\begin{bmatrix}0 & 1\\ 0 & 0\end{bmatrix}}{\overset{\begin{bmatrix}1 & 0\\ 0 & 0\end{bmatrix}}{\rightrightarrows}} \mathbb{C}^2 \xrightarrow{\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix}} \mathbb{C}^2.$$

To find the optimal σ witness, we must first get a random matrix in the blow-up of the generalized Kronecker reduction. Here, our $d = 2 \cdot \sigma(x) - 1 = 1$, so a larger blow-up is not needed. Next, we construct the set of maps $\bigoplus W(x)^{\sigma_+(x)} \to \bigoplus W(y)^{\sigma_-(y)}$. In this case, we get the space of maps given by linear combinations of $W(a)$ and $W(b)$. Choosing a random map A from this , we will take $\begin{bmatrix}-1 & 10\\ 0 & 0\end{bmatrix}$, which has pseudo-inverse $P = \begin{bmatrix}-0.01 & 0\\ 0.099 & 0\end{bmatrix}$. Now, for each block in this pseudo-inverse, we add a map to Q, giving us Q^+. In this example, this adds only one arrow, from y to x:

$$x \underset{p}{\overset{a,\ b}{\rightrightarrows\!\!\!\leftarrow}} y \xrightarrow{\ c\ } z.$$

Our W^+ is then defined as W is for all shared vertices and arrows, along with $W^+(p) = P$. At vertex x, we have $K(x)$ the kernel of A, which is the span of $\begin{bmatrix}10\\ 1\end{bmatrix}$. As σ is only positive at x, $K(y)$ and $K(z)$ are both defined to be zero. We now want to find the smallest subrepresentation of W^+ that contains $K(x)$, $K(y)$, and $K(z)$. Denoting the span of a vector by $\langle v \rangle$, we go through Algorithm 3.16, starting with

$$\left\langle \begin{matrix}10\\ 1\end{matrix} \right\rangle \underset{P}{\overset{W(a),\ W(b)}{\rightrightarrows\!\!\!\leftarrow}} 0 \xrightarrow{\ W(c)\ } 0,$$

and looping through our arrows in the order p, a, b, c. Following P adds nothing to $K(x)$, following $W(a)$ and $W(b)$ will hit the first coordinate at y, and taking this first coordinate at y to z will hit the second coordinate, so running the inner loop of the algorithm once leaves us with:

$$\left\langle \begin{matrix} 10 \\ 1 \end{matrix} \right\rangle \;\underset{P}{\overset{W(a),\,W(b)}{\rightrightarrows\!\!\!\leftarrow}}\; \left\langle \begin{matrix} 1 \\ 0 \end{matrix} \right\rangle \xrightarrow{\quad W(c) \quad} \left\langle \begin{matrix} 0 \\ 1 \end{matrix} \right\rangle.$$

Running through the arrows again, we now see that following P will hit the span of $\left\langle \begin{matrix} 0.099 \\ -0.01 \end{matrix} \right\rangle$, bringing us to the whole space $\mathbb{C}^2$ at x. following $W(a)$ and $W(b)$ still only land us in the first coordinate, and so following $W(c)$ also changes nothing. After running the inner loop of the algorithm a second time, we have:

$$\mathbb{C}^2 \;\underset{P}{\overset{W(a),\,W(b)}{\rightrightarrows\!\!\!\leftarrow}}\; \left\langle \begin{matrix} 1 \\ 0 \end{matrix} \right\rangle \xrightarrow{\quad W(c) \quad} \left\langle \begin{matrix} 0 \\ 1 \end{matrix} \right\rangle.$$

At this point, looping through the arrows again will not increase the dimension at any vertex, and we are done. The quiver representation above without the arrow p is our optimal σ witness, giving us a discrepancy of $2(1) - 1(1) + 1(0) = 1$.

Example 3.21. Consider the quiver:

$$\begin{array}{ccccc} & & x_1 & & \\ & & \downarrow{\scriptstyle a_1} & & \\ x_2 & \xrightarrow{a_2} & y & \xleftarrow{a_3} & x_3 \end{array}$$

With representation V, given by $V(x_i) = \mathbb{C}^5$, $V(y) = \mathbb{C}^6$ and the $V(a_i)$ general. Let $\sigma(x_i) = 2$, and $\sigma(y) = -5$.

In this case, we take $d = 29$. Constructing our A gives a $30d \times 30d$ block matrix, computing its kernel gives a $3d$ dimensional space, which projects onto 3 dimensional spaces for each $V(x_i)$; the image of these projections in $V(y)$ is also 3 dimensional. As the only path from x_i to y is a_i, the additional maps in V^+ are simply pseudo-inverses of each $V(a_i)$, and will not increase the dimension of our subrepresentation. We then find $c = 2(3) + 2(3) + 2(3) - 5(3) = 3$, our quiver representation is not σ-semi-stable, and has an optimal σ witness of discrepancy 3.

3.4 Non-commutative General Ext and Hom

Recall we have defined $\mathrm{Hom}_Q(V, W)$ and $\mathrm{Ext}_Q(V, W)$ as the kernel and cokernel respectively of the map:

$$f_W^V : \bigoplus_{x \in Q_0} \mathrm{Hom}\big(V(x), W(x)\big) \to \bigoplus_{a \in Q_1} \mathrm{Hom}\big(V(ta), W(ha)\big),$$

where $f_W^V(\varphi) = \big(\varphi(ha)V(a) - W(a)\varphi(ta) : a \in Q_1\big)$.

Definition 3.22. Given dimension vectors α and β respectively, the *Euler form* or *Ringel form* on $\mathbb{R}^{Q_0}$ is

$$\langle \alpha, \beta \rangle = \sum_{x \in Q_0} \alpha(x)\beta(x) - \sum_{a \in Q_1} \alpha(ta)\beta(ha).$$

Proposition 3.23 (See [DW17])**.** Given quiver representations V and W with dimension vectors α and β respectively,

$$\langle \alpha, \beta \rangle = \dim \mathrm{Hom}_Q(V, W) - \dim \mathrm{Ext}_Q(V, W).$$

Proof. Extending f_W^V to an exact sequence by adding inclusion of $\mathrm{Hom}_Q(V, W)$ and projection of $\mathrm{Ext}_Q(V, W)$, the alternating sum of the dimension of spaces must be 0, and thus the difference in dimension between Hom and Ext is the same as the difference in dimension between $\bigoplus_{x \in Q_0} \mathrm{Hom}\big(V(x), W(x)\big)$ and $\bigoplus_{a \in Q_1} \mathrm{Hom}\big(V(ta), W(ha)\big)$. $\square$

In [Kac82], Kac studied the minimal dimension of $\mathrm{Hom}_Q(V, W)$ for representations

V and W of dimension vectors α and β. Define

$$Z^t(\alpha, \beta) := \{(V, W) \in \mathrm{Rep}_\alpha(Q) \times \mathrm{Rep}_\beta(Q) | \dim \mathrm{Hom}_Q(V, W) \geq t\}.$$

Each of these subsets of $\mathrm{Rep}_\alpha(Q) \times \mathrm{Rep}_\beta(Q)$ are closed. Take t the minimal positive value of $\dim \mathrm{Hom}_Q(V, W)$. Then, $Z^{t+1}(\alpha, \beta)$ is a proper closed subset. We call the pair (V, W) (α, β)-general if they are in the (open and dense) complement of $Z^{t+1}(\alpha, \beta)$. On this complement, $\dim \mathrm{Hom}_Q(V, W)$ is constant, as is $\dim \mathrm{Ext}_Q(V, W)$. In [Sch92], Schofield called these general hom and ext respectively:

$$\begin{aligned} \hom(\alpha, \beta) &= \dim(\mathrm{Hom}_Q(V, W)), \text{ and} \\ \mathrm{ext}(\alpha, \beta) &= \dim(\mathrm{Ext}_Q(V, W)). \end{aligned}$$

To calculate these, we may use the following theorem.

Theorem 3.24 (see [Sch92])**.** Let α and β dimension vectors. We have

$$\mathrm{ext}(\alpha, \beta) = \max_{\alpha'}\{-\langle \alpha', \beta\rangle\} = \max_{\beta'}\{-\langle \alpha, \beta'\rangle\},$$

where α' and β'' are so that for any (α, β)-general pair of representations (V, W), V has a subrepresentation of dimension α' and W has a factor representation of dimension β'.

Example 3.25. Let $\alpha = (1, 2)$, $\beta = (3, 3)$, and our quiver be

$$\bullet \;\begin{matrix} \xrightarrow{a} \\ \xrightarrow{b} \\ \xrightarrow{c} \end{matrix}\; \bullet \;.$$

For any V, let $V(a) = [x \quad x']$, $V(b) = [y \quad y']$, and $V(c) = [z \quad z']$. The set of maps

d_W^V are then of the form

$$\left[\begin{array}{c|c|c} -W(a) & xI & x'I \\ \hline -W(b) & yI & y'I \\ \hline -W(c) & zI & zI' \end{array}\right],$$

giving us $\hom(\alpha, \beta) = \operatorname{ext}(\alpha, \beta) = 0$, as these matrices have full rank in general. Using Theorem 3.24, we could see this by noting that in general, a representation W with dimension vector α only has subrepresentations with $W(x) = 0$ along with the entire representation itself. This means the only factor representation dimension vectors β' we need to go through are $(3, n)$ (for n between 0 and 3) and $(0, 0)$. Summarized by the following table, we see that the minimum negative Euler form also gives us 0.

β'	$-\langle \alpha, \beta' \rangle$
$(3, 0)$	-3
$(3, 1)$	-2
$(3, 2)$	-1
$(3, 3)$	0
$(0, 0)$	0

So we also see $\operatorname{ext}(\alpha, \beta)$ is 0 using the theorem, and get $\hom(\alpha, \beta) = 0$ by subtracting from $\langle \alpha, \beta \rangle$, which is also 0.

We can also look at the minimal dimension of the space of morphisms when one of the representations is fixed. Crawley-Boevey generalized the general hom in this way in [CB96], which we show along with the generalization of general ext. To do this, fix a representation W of Q, with dimension vector β. Define now

$$Z^t(\alpha, W) := \{V \in \operatorname{Rep}_\alpha(Q) | \dim \operatorname{Hom}_Q(V, W) \geq t\}.$$

Again, each of these subsets are closed, and we take t minimal so that $\dim \operatorname{Hom}_Q(V, W)$ is positive. We call V (α, W)-general if it is in the complement of $Z^{t+1}(\alpha, W)$.

Definition 3.26. Let V be an (α, W)-general representation. We define the (α, W)-

general hom and ext as

$$\begin{aligned}\hom(\alpha, W) &= \dim(\mathrm{Hom}_Q(V,W)), \text{ and} \\ \mathrm{ext}(\alpha, W) &= \dim(\mathrm{Ext}_Q(V,W)).\end{aligned}$$

Example 3.27. Returning to Example 3.25, let $\alpha = (1,2)$, but now fix W to have $W(a), W(b), W(c)$ given by

$$A_1 = \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \ A_2 = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix}, \ A_3 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}$$

respectively. For any V, again let $V(a) = [x \quad x']$, $V(b) = [y \quad y']$, and $V(c) = [z \quad z']$. The set of maps d_W^V for our fixed W are of the form

$$\left[\begin{array}{c|c|c} -A_1 & xI & x'I \\ \hline -A_2 & yI & yI' \\ \hline -A_3 & zI & zI' \end{array}\right],$$

giving us $\hom(\alpha, W) = \mathrm{ext}(\alpha, W) = 1$, as these matrices now have rank 8 in general. By fixing W to be the skew-symmetric matrices, we've increased the minimal dimension of $\mathrm{Hom}_Q(V,W)$, and $\mathrm{ext}(\alpha, W)$ is larger as well.

Lemma 3.28. We have

$$\mathrm{ext}(\alpha, W) \geq \max\{-\langle \alpha, \underline{\dim} W' \rangle \mid W' \text{ factor representation of } W\}.$$

Proof. If V is a general representation of dimension α, then applying $\mathrm{Hom}_Q(V, \_)$ to

$$0 \to W'' \to W \to W' \to 0$$

gives an exact sequence

$$\cdots \to \mathrm{Ext}_Q(V,W) \to \mathrm{Ext}_Q(V,W') \to 0,$$

so $\dim \mathrm{Ext}_Q(V,W') \leq \dim \mathrm{Ext}_Q(V,W)$ and $\mathrm{ext}(\alpha, W') \leq \mathrm{ext}(\alpha, W)$. We get

$$-\langle \alpha, \underline{\dim} W' \rangle = \mathrm{ext}(\alpha, W') - \hom(\alpha, W') \leq \mathrm{ext}(\alpha, W') \leq \mathrm{ext}(\alpha, W).$$

□

Definition 3.29. The non-commutative ext and hom are defined by the following limits of ext and hom:

$$\begin{aligned} \mathrm{ncext}(\alpha, W) &= \lim_{d\to\infty} \frac{\mathrm{ext}(d\alpha, W)}{d} \\ \mathrm{nchom}(\alpha, W) &= \lim_{d\to\infty} \frac{\hom(d\alpha, W)}{d}. \end{aligned}$$

Note that for every representation W of dimension β, we have $\mathrm{nchom}(\alpha, W) - \mathrm{ncext}(\alpha, W)$ equal to $\langle \alpha, \beta \rangle$. These limits were originally studied in [CB96], though we give them a name to highlight their connection to non-commutative rank, as seen in the next discussion and proposition.

We have a map

$$f_W^\alpha : \mathrm{Rep}_Q(\alpha) \longrightarrow \mathrm{Hom}\left(\bigoplus_{x\in Q_0} \mathrm{Hom}\left(k^{\alpha(x)}, W(x)\right), \bigoplus_{a\in Q_1} \mathrm{Hom}\left(k^{\alpha(ta)}, W(ha)\right)\right)$$

given by sending a representation V to the map $f_W^\alpha(V)$, which takes the set of $\varphi(x)$ from $\mathrm{Hom}(k^{\alpha(x)}, W(x))$ over all vertices x to the set of maps $\varphi(ha)V(a) - W(a)\varphi(ta)$ over all arrows a. Note that the kernel of each $f_W^\alpha(V)$ is $\mathrm{Hom}_Q(V,W)$, and and the cokernel is $\mathrm{Ext}_Q(V,W)$. From this point forward, we will refer to the image of f_W^α (the set of $f_W^\alpha(V)$ over all V), as simply f_W^α itself.

Next, we note that we can consider $\operatorname{Rep}_Q(d\alpha)$ as the blow-up of $\operatorname{Rep}_Q(\alpha)$ as follows. Each Z in $\operatorname{Rep}_Q(d\alpha)$ is so that $Z(x) \cong k^{\alpha(x)} \otimes U(x)$ for $U(x) \cong k^d$. At the arrows, we have $Z(a) \cong \sum V_i(a) \otimes U_i(a)$, a finite sum where each $U_i(a)$ is a $d \times d$ matrix over k, and each V_i is from $\operatorname{Rep}_Q(\alpha)$. Now, given a $\overline{V}$ in $\operatorname{Rep}_{d\alpha}$, we get a map:

$$\operatorname{Hom}\Big(\bigoplus_{x \in Q_0} \operatorname{Hom}\left(k^{d\alpha(x)}, W(x)\right) \xrightarrow{f_W^{d\alpha}(\overline{V})} \bigoplus_{a \in Q_1} \operatorname{Hom}\left(k^{d\alpha(ta)}, W(ha)\right)\Big).$$

Notice that we can find $\operatorname{ncrk}\left(f_W^{\alpha}\right)$ using $\operatorname{ncrk}\left(f_W^{d\alpha}\right)$ and dividing by d since $f_W^{d\alpha}$ is the dth blow-up of f_W^{α}.

Proposition 3.30. The rank and non-commutative rank of $f_W^{d\alpha}$ are equal if and only if $\operatorname{nchom}(\alpha, W) = \dfrac{\hom(d\alpha, W)}{d}$.

For a $(d\alpha, W)$-general $\overline{V}$ in $\operatorname{Rep}_{d\alpha}$, the kernel of $f_W^{d\alpha}(\overline{V})$ is of minimal dimension. So, $\operatorname{rk}(f_W^{d\alpha}) = \sum d\alpha(x)\beta(x) - \hom(d\alpha, W)$. We get

$$\frac{\operatorname{rk} f_W^{d\alpha}}{d} = \sum \alpha(x)\beta(x) - \frac{\hom(d\alpha, W)}{d},$$

showing that the d which maximizes the left-side (giving us the non-commutative rank), maximizes the right side (minimizing $\frac{\hom(d\alpha,W)}{d}$, giving us the non-commutative hom).

Corollary 3.31. Given dimension vector α, and a representation W of dimension β, the d in the limit of definition 3.29 can be chosen to be

$$\min\Big\{\sum_{x \in Q_0} \alpha(x)\beta(x) - 1, \sum_{a \in Q_1} \alpha(ta)\beta(ha) - 1\Big\}.$$

Proof. Recall the bound for non-commutative rank blow-ups from [DM17b] is $n-1$, where n is the dimension of both the domain and co-domain. We may not have a space of square matrices, so a large enough d will be found when we first reach either

$\sum_{x \in Q_0} \alpha(x)\beta(x) - 1$ or $\sum_{a \in Q_1} \alpha(ta)\beta(ha) - 1$. ☐

Theorem 3.32. We have

$$\text{ncext}(\alpha, W) = \max\{-\langle \alpha, \underline{\dim} W''\rangle \mid W'' \text{ factor representation of } W\}.$$

Proof. Choose d so that $\text{ncext}(d\alpha, W)$ equals $\frac{\text{ext}(d\alpha, W)}{d}$. Look at the set of maps:

$$\text{Hom}\left(\bigoplus_{x \in Q_0} \text{Hom}\left(k^{\alpha(x)}, W(x)\right) \xrightarrow{f_W^{d\alpha}(\overline{V})} \bigoplus_{a \in Q_1} \text{Hom}\left(k^{\alpha(ta)}, W(ha)\right)\right),$$

for all representations $\overline{V}$ in $\text{Rep}_{d\alpha}$. By Proposition 3.30, this set of maps has non-commutative rank equal to its rank. So we can find the minimal c-shrunk subspace, which:

1. has the form $\bigoplus_{x \in Q_0} \text{Hom}\left(k^{d\alpha(x)}, W'(x)\right)$, for some subrepresentation W' of W (from discussion in section 3.3), and

2. has image of the form $\bigoplus_{a \in Q_1} \text{Hom}\left(k^{d\alpha(ta)}, W'(ha)\right)$.

So we get $c = d \sum \alpha(x) \dim(W'(x)) - d \sum \alpha(ta) \dim(W'(ha)) = \langle d\alpha, \underline{\dim}(W')\rangle$, but c gives us the non-commutative rank, so also can be found by $\sum d\alpha(x)\beta(x) - \text{rk}(f_W^{d\alpha}) = \text{hom}(d\alpha, W)$. This leaves us with $\frac{\text{hom}(d\alpha, W)}{d} = \langle \alpha, \underline{\dim}(W')\rangle$ after dividing by d. As for non-commutative ext, we then get $\text{ncext}(\alpha, W) = \text{nchom}(\alpha, W) - \langle \alpha, \beta\rangle$, finally leaving us with $\text{ncext}(\alpha, W) = -\langle \alpha, \underline{\dim} W''\rangle$, for $W'' = W/W'$. ☐

Example 3.33. Returning to Example 3.27, we saw $\text{ext}(\alpha, W) = 1$. However, just as in Example 3.25, the only subrepresentations of W are with $W(x) = 0$ along with W itself. Looping through $-\langle \alpha, \underline{\dim} W''\rangle$ for these subrepresentations, we get a maximum value of 0. So although the general hom and ext for fixed W was larger than for general V and W, the non-commuative ext and hom are still 0.

We note that we can dually fix a representation V, and look at $\hom(V,\beta)$ and $\mathrm{ext}(V,\beta)$ to define $\mathrm{nchom}(V,\beta)$ and $\mathrm{ncext}(V,\beta)$.

Definition 3.34. The non-commutative ext and hom are defined by the following limits of ext and hom:

$$\begin{aligned}\mathrm{ncext}(V,\beta) &= \lim_{d\to\infty}\frac{\mathrm{ext}(V,d\beta)}{d}\\ \mathrm{nchom}(V,\beta) &= \lim_{d\to\infty}\frac{\hom(V,d\beta)}{d}\end{aligned}$$

Theorem 3.35. We have

$$\mathrm{ncext}(V,\beta) = \max\{-\langle \underline{\dim} V', \beta\rangle \mid V' \text{ subrepresentation of } V\}.$$

Proof. The proof follows from duality of Theorem 3.32. We note that this can also be seen by using Corollary 1 from [CB96], by subtracting $\langle \underline{\dim} V, \beta\rangle$. □

Corollary 3.36. For large enough $|k|$, there are both deterministic and randomized algorithms for calculating $\mathrm{ncext}(\alpha, W), \mathrm{nchom}(\alpha, W), \mathrm{ncext}(V,\beta)$, and $\mathrm{nchom}(V,\beta)$.

Proof. We can apply any of the algorithms used to find c-shrunk subspaces to the set of maps $f_W^{\alpha}(V)$ or $f_{\beta}^{V}(W)$ respectively, and use the dimension of the c-shrunk subspace to calculate the non-commutative ext and hom. □

CHAPTER IV

Tensors

4.1 Definitions

4.1.1 Rank

We now move onto tensors, a multilinear analog of matrices. Finding the rank of tensors is much more difficult than for matrices. Finding or bounding this rank, and approximating a tensor by low rank tensors has applications in signal processing, algebraic statistics, and more. For a nice overview, see [Lan12].

Let V, W be k-vector spaces. There is a pair (U, β) with U a k-vector space, β a bilinear map $\beta : V \times W \to U$, so that for every k-vector space Z and bilinear map $\gamma : V \times W \to Z$, there is a unique map $\widetilde{\gamma} : U \to Z$ so that the following diagram commutes:

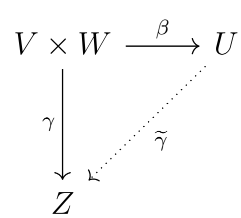

The pair (U, β) exists, as any bilinear map from $V \times W$ can be factored through the free vector space on $V \times W$ with the relations:

$$(av + v', w) = a(v, w) + a(v', w)$$
$$(v, aw + w') = a(v, w) + a(v, w')$$

for any v, v' in V, w, w' in W, and a in k, and this factorization will be unique. If we were to have another pair, (U', β') with this property, we get the diagram:

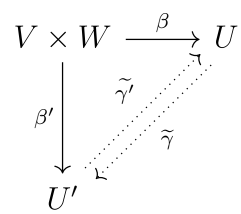

We have $\widetilde{\gamma} \circ \beta = \beta'$, so $\widetilde{\gamma'} \circ \widetilde{\gamma} \circ \beta = \widetilde{\gamma'} \circ \beta' = \beta$. We also have and $I_U \circ \beta = \beta$, so by uniqueness, $\widetilde{\gamma'} \circ \widetilde{\gamma} = I_U$. Similarly, $\widetilde{\gamma} \circ \widetilde{\gamma'} = I'_U$, and we have an isomorphism.

As pair (U, β) is unique up to isomorphism, we denote U by $V \otimes W$, and call it the *tensor product* of V and W. Elements of $V \otimes W$ are called tensors. Given a basis for V and W, $e_1, \ldots, e_n$ and $e'_1, \ldots, e'_{n'}$ respectively, $\{e_i \otimes e'_j | 1 \leq i \leq n, 1 \leq j \leq n'\}$ is a basis for their tensor product. We can define a tensor product of d vector spaces $V_1, V_2, \ldots, V_d$ by induction, getting $V = V_1 \otimes V_2 \otimes \ldots \otimes V_d$; tensoring is associative, and the order does not matter up to isomorphism. We will call d the order of a tensor in this space. The dimension of V is the product $\prod_{i=1}^{d} \dim(V_i)$.

Proposition 4.1. There is an isomorphism $\phi : V^* \otimes W \to \mathrm{Hom}(V, W)$.

Proof. Let $\phi(f \otimes w)(v) = f(v)w$. Let v_i be a basis for V with dual basis v_i^*, and let w_i be a basis for W. The image of $v_i^* \otimes w_j$ is the linear map sending v_i to w_j, and v_k to 0 for $k \neq i$. As these maps span $\mathrm{Hom}(V, W)$, ϕ must be surjective. As the dimensions of $V^* \otimes W$ and $\mathrm{Hom}(V, W)$ are both $\dim(W) \dim(V)$, this must be an isomorphism. □

Definition 4.2. A *simple tensor* in $V = V_1 \otimes V_2 \otimes \ldots \otimes V_d$ is a tensor of the form $v_1 \otimes v_2 \otimes \ldots \otimes v_d$ (where each v_j is a vector in V_j). The *Rank* of a tensor T in V is the minimum r so that T can be written as the sum of r simple tensors, i.e. the minimum r so that we may write

$$T = \sum_{i=1}^{r} v_1^{(i)} \otimes v_2^{(i)} \otimes \ldots \otimes v_d^{(i)},$$

with $v_j^{(i)}$ in V_j. We may also call a simple tensor a *Rank one tensor.*

As we will introduce several different notions of rank for tensors, we will capitalize Rank when we are using this definition to further distinguish it. From Proposition 4.1, we can see that the notion of rank for matrices is equivalent to the Rank for order 2 tensors.

Proposition 4.3. Given $\phi : V^* \otimes W \to \hom(V, W)$ from 4.1, and a rank one map, f in $\mathrm{Hom}(V, W)$, $\phi^{-1}(f)$ is a Rank one tensor.

Proof. A rank one matrix $A : V \to W$ can be decomposed as v^*w, for some non-zero vectors w in W, and v^* in V^*. However, notice this is exactly the image of $v^* \otimes w$, a Rank one tensor in $V^* \otimes W$. □

Decomposing a tensor into a form explicitly demonstrating its Rank is known by several names: Canonical Polyadic (CP) Decomposition [Hit27]; Canonical Decomposition (CANDECOMP) [CC70], and Parallel Factorization (PARAFAC) [Har70]. We caution the reader in that although one of these names is “canonical decomposition”, this decomposition is not unique.

Example 4.4. The Rank 2 tensor

$$\left(\begin{bmatrix}0\\1\end{bmatrix} \otimes \begin{bmatrix}0\\1\end{bmatrix}\right) + \left(\begin{bmatrix}1\\0\end{bmatrix} \otimes \begin{bmatrix}1\\0\end{bmatrix}\right) = \left(\begin{bmatrix}t\\1\end{bmatrix} \otimes \begin{bmatrix}0\\1\end{bmatrix}\right) + \left(\begin{bmatrix}1\\0\end{bmatrix} \otimes \begin{bmatrix}0\\-t\end{bmatrix}\right),$$

does not have a unique decomposition. As matrices, this is equivalent to:

$$\begin{bmatrix}1 & 0\\0 & 0\end{bmatrix} + \begin{bmatrix}0 & 0\\0 & 1\end{bmatrix} = \begin{bmatrix}1 & t\\0 & 0\end{bmatrix} + \begin{bmatrix}0 & -t\\0 & 1\end{bmatrix},$$

both sums of Rank one tensors give us the identity matrix.

The decomposition in some cases can be essentially unique - unique up to permutation of the simple tensors in the decomposition.

Definition 4.5. Given $v_1, v_2, \ldots, v_r$ in V, the Kruskal rank, K_V, of $S = \{v_1, \ldots, v_r\}$ is the largest k so that any k elements of S are linearly independent.

Theorem 4.6 (Kruskal's Theorem [Kru77])**.** Suppose $u_1, u_2, \ldots, u_r$ in U, $v_1, v_2, \ldots, v_r$ in V, and $w_1, w_2, \ldots, w_r$ in W have Kruskal ranks K_U, K_V, and K_W respectively, with $K_U, K_V, K_W \geq 1$ and $K_U + K_V + K_W \geq 2r + 2$. Then,

$$T = \sum_{i=1}^{r} u_i \otimes v_i \otimes w_i$$

has Rank r, and if

$$T = \sum_{i=1}^{r} a_i \otimes b_i \otimes c_i,$$

then $\{u_i \otimes v_i \otimes w_i\} = \{a_i \otimes b_i \otimes c_i\}$, i.e. the decomposition is essentially unique.

Example 4.7. Let $V = \mathbb{C}^4$ with basis e_1, e_2, e_3, e_4. Let $v_i = e_i$ for $1 \leq i \leq 4$, and let $v_5 = e_1 + e_2 + e_3 + e_4$. The Kruskal rank of the v_i is 4, and so the tensor

$$T = \sum_{i=1}^{5} v_i \otimes v_i \otimes v_i$$

satisfies the hypothesis of Kruskal's theorem, and has Rank 5.

Kruskal's theorem gives us one way of proving the tensor Rank for certain order 3 tensors, another theorem that allows us to do this follows.

Proposition 4.8 (See [Lan12])**.** Given a tensor T in $U \otimes V \otimes W$, $\mathrm{Rk}(T)$ is the smallest r so that the image of T_U, defined by the isomorphism in Propositon 4.1 as the map $U^* \to V \otimes W$, is contained in the span of r Rank one tensors.

Example 4.9. Consider the tensor

$$T = e_2 \otimes e_1 \otimes e_1 + e_1 \otimes e_2 \otimes e_1 + e_1 \otimes e_1 \otimes e_2,$$

in $U \otimes V \otimes W$, with $U \cong V \cong W \cong \mathbb{C}^2$. The isomorphism $U \otimes V \otimes W \cong \mathrm{Hom}(U^*, V \otimes W)$ sends T to T_U, the linear map taking e_1 and e_2 in U to $e_1 \otimes e_2 + e_2 \otimes e_1$ and $e_1 \otimes e_1$ in $V \otimes W$ respectively. As matrices, the image of T_U is the span of

$$\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \text{ and } \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix},$$

which requires the span of 3 rank 1 matrices to contain. Therefore T has tensor Rank 3, we can not write it as a sum of fewer Rank 1 tensors. We will return to this Rank 3 tensor several more times.

Although the rank of a matrix A over k does not change if we view A over a field extension of k, this is not the case for tensor Rank.

Example 4.10. Consider the tensor

$$T = e_1 \otimes e_1 \otimes e_1 - e_1 \otimes e_2 \otimes e_2 - e_2 \otimes e_1 \otimes e_2 - e_2 \otimes e_2 \otimes e_1.$$

Using Proposition 4.8 over $\mathbb{R}$, we get the image of T_U is contained in the span of

$$\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \text{ and } \begin{bmatrix} 0 & -1 \\ -1 & 0 \end{bmatrix}.$$

This subspace contains no rank 1 matrices, as $\begin{bmatrix} a & b \\ b & -a \end{bmatrix}$ will have rank 2 for all a and b not both non-zero. So $\mathrm{Im}(T_U)$ requires at least 3 rank 1 matrices to span. Three matrices that work for this are

$$\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \text{ and } \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix},$$

therefore T over $\mathbb{R}$ has Rank 3.

Instead now, consider this tensor over $\mathbb{C}$. Again using Proposition 4.8, we now can find the image of T_U is contained in the span of only 2 rank 1 matrices,

$$\begin{bmatrix} 1 & i \\ i & -1 \end{bmatrix} \text{ and } \begin{bmatrix} 1 & -i \\ -i & 1 \end{bmatrix}.$$

So $\mathrm{Rk}(T)$ is larger over $\mathbb{R}$ than over $\mathbb{C}$.

4.1.2 Border Rank

We will now define another notion of rank that is useful for algebraic geometric methods. As the set of tensors of Rank at most r is not a closed subset, there cannot be polynomial equations cutting out the tensors of Rank at most r. We instead define the border rank of a tensor, which allows for the existence of polynomial equations cutting out the tensors with border rank at most r. Border rank was first used in [BCRL79], and was first denoted as border rank in [BLR80].

Definition 4.11. The *border rank* of a tensor T, $\underline{\mathrm{rk}}(T)$, is the smallest r so that T is the limit of tensors of Rank at most r.

By definition, the border rank of a tensor is less than or equal to its Rank. Finding the equations for the tensors of border rank at most r in general remains an open problem. A summary of what is known on these equations can be found in [Lan12, Chapter 7].

Example 4.12. Let

$$T = e_2 \otimes e_1 \otimes e_1 + e_1 \otimes e_2 \otimes e_1 + e_1 \otimes e_1 \otimes e_2.$$

We have seen that this tensor has Rank 3, in Example 4.9. Define T_ε so that

$$\varepsilon T_\varepsilon = (e_1 \otimes \varepsilon e_2) \otimes (e_1 \otimes \varepsilon e_2) \otimes (e_1 \otimes \varepsilon e_2) - e_1 \otimes e_1 \otimes e_1.$$

Notice that T_ε is at most Rank 2. Taking the limit of T_ε as $\epsilon \to 0$ gives us T, so the border rank of T is at most 2, though T has Rank 3.

Note that for matrices, border rank is always equal to rank, as rank r matrices are a Zariski closed subset cut out by the $r \times r$ matrix minors.

4.1.3 Slice Rank

First introduced by Tao in [Tao16], the slice rank of a tensor is yet another notion of rank for tensors. Similar to tensor Rank, we will define slice rank by first defining slice rank one tensors.

Definition 4.13. A tensor in $V_1 \otimes V_2 \otimes \ldots \otimes V_d$ has *slice rank one* if it is contained in

$$V_1 \otimes \ldots \otimes V_{k-1} \otimes w \otimes V_{k+1} \otimes \ldots V_d$$

for some index k and w in V_k. A tensor T has *slice rank* $\operatorname{slrk} T = r$ if r is the smallest integer so that T can be written as the sum of r slice rank one tensors.

Example 4.14. Choosing the same tensor in $\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2$ from Example 4.9, let

$$T = e_2 \otimes e_1 \otimes e_1 + e_1 \otimes e_2 \otimes e_1 + e_1 \otimes e_1 \otimes e_2.$$

The tensor T is contained in

$$e_2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 + e_1 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2,$$

and has slice rank 2.

The slice rank was first denoted as such in [BCC$^+$17], and notes on it can be found in [ST16]. This notion of rank was introduced to tackle the cap set problem, which asks: What is the largest set in $\mathbb{F}_3^n$ containing no lines? It is conjectured that the maximum is c^n for some $c < 3$. Work on bounding this maximum from both above and below is an active area of research. A nice summary of the history and applications of this problem can be found in [Gro19]. We note now that slice rank does not have a "border" version.

Lemma 4.15 (See [Tao16])**.** The set of tensors T of slice rank less than or equal to r is Zariski closed.

Proposition 4.16. Given a tensor T, $\text{slrk}(T) \leq \underline{\text{rk}}(T)$.

Proof. By way of contradiction, suppose $s = \text{slrk}(T) > \underline{\text{rk}}(T) = b$. The slice rank of a tensor is less than or equal to its Rank. Note this means that if we can approach T by rank b tensors, we can approach it by slice rank less than or equal to b tensors. But by the previous lemma, the set of tensors with slice rank less than or equal to $s - 1$ is Zariski closed, this is impossible. □

For tensors in $V_1 \otimes \ldots \otimes V_d$, where the V_i are not all of the same dimension, it may be more interesting and useful to use a weighted version of the slice rank. Call a slice rank one tensor contained in $V_1 \otimes \ldots \otimes V_{k-1} \otimes w \otimes V_{k+1} \otimes \ldots V_d$ for some index k and w in V_k a slice rank one tensor *of slice* k.

Definition 4.17. Given a tensor T in $V_1 \otimes \ldots \otimes V_d$, and a d dimensional weight vector α of positive reals, the α-weighted slice rank of T, $\text{slrk}_\alpha(T)$ is the minimum $r = \sum_{k=1}^{d} \alpha(k) \dim(W_k)$ so T is contained in

$$\sum_{k=1}^{d} V_1 \otimes \ldots \otimes V_{k-1} \otimes W_k \otimes V_{k+1} \otimes \ldots V_d$$

where each W_k is a subspace of V_k.

The α-weighted slice rank With $\alpha = (1, 1, \ldots, 1)$ is equivalent to the slice rank. With the α-weighted slice rank, we can now give preference to slicing along certain vector spaces. In fact, we can even let the weights be infinity to exclude using certain orders for the slices (as long as at least one weight is still finite). If one is opposed to using infinite weights, the same result can be achieved using sufficiently large weights. A good choice of α is one so that all $\alpha(i)\dim(V_i)$ are equal.

Example 4.18. We now look a the tensor T from Example 4.9, but instead consider it in $\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^3$, and let $\alpha = (3, 3, 2)$. In Example 4.14 we saw T was contained in

$$e_2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^3 + e_1 \otimes \mathbb{C}^2 \otimes \mathbb{C}^3,$$

which would give us $\sum \alpha(k)\dim(W_k) = 6$. However, slicing instead along the third factor, we also have T contained in

$$\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes e_1 + \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes e_2,$$

showing the α-weighted slice rank of T is at most 4.

4.1.4 G-stable Rank

A tensor product is a vector space, so we still have tools from representation and invariant theory.

Example 4.19. We have an action of $G = \prod_{i=1}^d \mathrm{GL}(V_i)$ on $V_1 \otimes \ldots \otimes V_d$, given by: $(A_1, \ldots, A_n) \cdot v_1 \otimes \ldots \otimes v_d = A_1 v_1 \otimes \ldots \otimes A_d v_d$.

Also introduced with the cap set problem in mind, Derksen defined G-stable rank in [Der20], using the Hilbert-Mumford Criterion (see Proposition 2.41). Before defining the G-stable rank, we will first define some other tools we will need. The ith *flattening* of a tensor T in $V_1 \otimes \ldots \otimes V_d$, is the map

$$\Phi_i(T) : (V_1 \otimes \ldots \otimes \hat{V_i} \otimes \ldots \otimes V_d)^* \to V_d,$$

given by the isomorphism in Proposition 4.1. Two norms we will be working with are the Euclidean (or ℓ_2 norm), $||\cdot||$, defined on the tensors, and the spectral norm, $||\cdot||_\sigma$, defined on linear maps. Although defined over any field, we will give the definition of G-stable rank when working over $\mathbb{C}$.

Definition 4.20. Let $G = \prod_{i=1}^d \mathrm{GL}(V_i)$, and α, a d dimensional weight vector of positive reals. The G-stable rank of a tensor T in $V_1 \otimes \cdots \otimes V_d$ is

$$\mathrm{rk}_\alpha^G(T) := \sup_{g \in G} \min_i \frac{\alpha(i)||g \cdot T||^2}{||\Phi_i(g \cdot T)||_\sigma{}^2}.$$

When $\alpha = (1, 1, \ldots, 1)$, we denote the G-stable rank with rk^G. Similar to weighted slice rank, we may exclude the use of certain flattenings by setting weights equal to infinity (or sufficiently large). Unlike our other notions of rank, the G-stable rank does not have to be an integer.

Example 4.21. Let $\alpha = (1, 1, 1)$. Again returning to our tensor

$$T = e_2 \otimes e_1 \otimes e_1 + e_1 \otimes e_2 \otimes e_1 + e_1 \otimes e_1 \otimes e_2,$$

from Example 4.9, by symmetry, each flattening of T is

$$\Phi_i(T) = \left[\begin{array}{cc|cc} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 \end{array}\right],$$

which has largest singular value $\sqrt{2}$. This gives us

$$\mathrm{rk}^G(T) \geq \min_{1 \leq i \leq 3} \frac{||T||^2}{||\Phi_i(T)||_\sigma{}^2} = \frac{3}{2}.$$

It can also be shown that $\mathrm{rk}^G(T) \leq \frac{3}{2}$ using a more general definition for G-stable rank (see [Der20] Example 1.4). We have now seen that this tensor has Rank 3, slice rank 2, border rank 2, and G-stable rank $\frac{3}{2}$.

We will compare the slice rank and G-stable rank following [Der20].

Lemma 4.22. Given T in $V_1 \otimes \ldots \otimes V_d$ and a weight vector α in $\mathbb{R}^d_{>0}$,

$$\mathrm{rk}^G_\alpha(T) \geq \min_{1\leq i\leq d} \alpha(i).$$

Proof. The Euclidean norm of a tensor T is equivalent to the Euclidean form for any of its flattenings. For a linear map, the spectral norm is equal to its largest singular value, σ_1 and the Euclidean norm is equal to $\sqrt{\sigma_1^2 + \ldots + \sigma_n^2}$. This gives us

$$\begin{aligned} \frac{\alpha(i)||T||^2}{||\Phi_i(T)||_\sigma{}^2} &= \frac{\alpha(i)||\Phi_i(T)||^2}{||\Phi_i(T)||_\sigma{}^2} \\ &= \frac{\alpha(i)\sqrt{\sigma_1^2 + \ldots + \sigma_n^2}}{\sigma_1{}^2} \\ &\geq \alpha(i). \end{aligned}$$

So the G-stable rank is bounded below by the minimum component of the weight vector α. □

Lemma 4.23. A tensor of slice rank 1 has G-stable rank 1 for $\alpha = (1, 1, \ldots, 1)$.

Proof. Suppose T in $V_1 \otimes \ldots V_d$ has slice rank 1. Without loss of generality, let $T = v \otimes w$, with v in V_1 and w in $V_2 \otimes \ldots \otimes V_d$. Note that $g \cdot T$ is also slice rank one with concentration in the first order for all g. For any g, choose a basis of V so that $e_1 = gv$, the first flattening of $g \cdot T$ is the rank one map sending e_1 to gw, and all other e_i to zero. The only non-zero singular value for this map is $||gw||$, which is also the Euclidean norm for this map, giving us $\mathrm{rk}^G(T) \leq 1$, so $\mathrm{rk}^G(T)$ must be equal to 1. □

Lemma 4.24 (See [Der20])**.** For tensors T and S, $\mathrm{rk}^G_\alpha(T + S) \leq \mathrm{rk}^G_\alpha(T) + \mathrm{rk}^G_\alpha(S)$.

Proposition 4.25. Given a tensor T, $\mathrm{rk}^G_\alpha(T) \leq \mathrm{slrk}(T)$.

Proof. Suppose T has slice rank r, so we can write T as the sum of r slice rank 1 tensors, $T = T_1 + \ldots + T_r$. By Lemma 4.24,

$$\mathrm{rk}^G(T_1 + \ldots + T_r) \leq \sum_{i=1}^{r} \mathrm{rk}^G(T_i).$$

By Lemma 4.23, $\mathrm{rk}^G(T_i)$ is 1 for all i, so $\mathrm{rk}^G(T) \leq r$. □

In general, we have now shown:

$$\mathrm{rk}^G(T) \leq \mathrm{slrk}(T) \leq \underline{\mathrm{rk}}(T) \leq \mathrm{Rk}(T).$$

Similarly to slice rank, the G-stable rank has no border version.

Proposition 4.26 (See [Der20]). The set of tensors T with $\mathrm{rk}^G_\alpha(T) \leq r$ is Zariski Closed for all r.

We note that the role of α is similar to that of σ when we talk of σ-semi-stability for quivers. A good choice of α is one that makes all $\alpha(i)\dim(V_i)$ equal for tensors in $V_1 \otimes \cdots \otimes V_d$.

4.2 Non-commutative Tensor Rank

We have now seen four different notions of tensors, (tensor) Rank, border rank, slice rank, and G-stable rank. Many of these generalize matrix rank to higher order tensors. However, we have also seen a notion of rank for spaces of matrices - non-commutative rank. Although there has been lots of work on the rank of a single tensor, there is not much (if any) discussion or work on rank of a space of tensors. With this in mind, in this section, we explore how we can generalize rank to spaces of tensors, by providing completely new definitions and discussion in this section. We will bring our notions of tensors from the previous section to define rank on a space of tensors, analogous to non-commutative rank. Recall we define the rank (or commutative rank) of a linear span of matrices $\mathcal{A} = \text{Span}\{A_1, \ldots, A_m\}$ as $\text{rk}(\mathcal{A}) = \max\{r | A \in \mathcal{A}, \text{rk}(A) = r\}$. In Section 3.2, we saw four ways of defining non-commutative rank, by using either:

1. rank over the free skew field,
2. c-shrunk subspaces,
3. tensor blow-ups, or
4. semi-stability.

Can we generalize these definitions to instead, a linear subspace of tensors $\mathcal{T}$ in $V = V_1 \otimes \ldots \otimes V_d$? For the case of (commutative) rank of a tensor space $\mathcal{T}$, we can define it by the maximum rank among all tensors in $\mathcal{T}$. For non-commutative rank, we will explore generalizing each of these four definitions in the context of tensors.

4.2.1 Over the Free Skew Field: Failure to generalize

Given a matrix, $A(\underline{x})$, with homogeneous linear polynomials in $k\langle x_1, \ldots, x_m\rangle$, the non-commutative rank of $A(\underline{x})$, $\text{ncrk}(A(\underline{x}))$ is the rank of $A(\underline{x})$ over the free skew field.

Notice that for this definition, we do not think of our matrix space as a space, but instead as a matrix with homogeneous linear entries, e.g. $A = x_1 A_1 + \ldots + x_m A_m$.

Example 4.27. Recall the 3×3 skew-symmetric matrices. We may row reduce to try to determine the rank:

$$T = \begin{bmatrix} 0 & x & y \\ -x & 0 & z \\ -y & -z & 0 \end{bmatrix} \sim \begin{bmatrix} 0 & x & y \\ -x & 0 & z \\ 0 & 0 & zx^{-1}y - yx^{-1}z \end{bmatrix},$$

giving us non-commutative rank 3. As we'd like to generalize this to tensors, we will look at this example as a tensor. Following the notation that the tensor of two standard basis vector $e_i \otimes e_j$ is $[i, j]$ We have

$$x\big([2,1]-[1,2]\big)+y\big([3,1]-[1,3]\big)+z\big([3,2]-[2,3]\big).$$

We can see this is a tensor of (commutative) Rank at least 2, as when x is non-zero, it's equivalent to:

$$\begin{bmatrix}0 & x & y\end{bmatrix}\otimes\begin{bmatrix}1\\0\\\frac{-z}{x}\end{bmatrix}+\begin{bmatrix}-x & 0 & z\end{bmatrix}\otimes\begin{bmatrix}1\\0\\\frac{y}{x}\end{bmatrix}$$

In the case x is zero, we can similarly show this is (commutative) Rank 2.

In any case, we question what it would mean to manipulate this, or any tensor over the free skew field. Here, valid operations are unclear, so using this definition to generalize non-commutative rank is likely not worthwhile.

4.2.2 Generalizing c-shrunk Subspaces

Given a matrix space $\mathcal{A}$, recall a subspace $U \subseteq k^n$ is a *c-shrunk subspace* if there exists a subspace $W \subseteq k^n$ with $\dim(W) \leq \dim(U) - c$, and for every A in $\mathcal{A}$, $A(U) \subseteq W$. The non-commutative rank is $n - c$ where c is maximal so there is a c-shrunk subspace of $\mathcal{A}$.

To generalize this definition, we will look to slice rank, as there is not a canonical way to view an order d tensor as a linear map between two spaces.

Recall a non-zero tensor has slice rank 1 if it is contained in $V_1 \otimes \ldots \otimes V_{i-1} \otimes w \otimes V_{i+1} \otimes \ldots \otimes V_m$, for some index i and some w in V_i. The slice rank of any tensor T is the smallest r so that T is the sum of r slice rank 1 tensors. Now, we define non-commutative slice rank for a space of tensors $\mathcal{T}$.

Definition 4.28. A tensor space $\mathcal{T}$ has non-commutative slice rank r if r is the minimum so that $\mathcal{T}$ is contained in the span of r slice rank 1 tensors. Equivalently,

$$\text{nc-slrk}(\mathcal{T}) = \min\left\{\sum_{k=1}^{m} \dim(W_i) \middle| \mathcal{T} \subseteq \sum_{k=1}^{d} V_1 \otimes \ldots \otimes V_{k-1} \otimes W_k \otimes V_{k+1} \otimes \ldots \otimes V_m\right\},$$

where each W_k is a subspace of V_k.

Note we may also define a weighted version of the non-commutative slice rank building off of our Definition 4.17.

Definition 4.29. Given a tensor space $\mathcal{T}$ in $V_1 \otimes \ldots \otimes V_d$, and a d dimensional weight vector α of positive reals, the α-weighted non-commutative slice rank of $\mathcal{T}$, $\text{nc-slrk}_\alpha(\mathcal{T})$ is the minimum $r = \sum_{k=1}^{d} \alpha(k) \dim(W_k)$ so $\mathcal{T}$ is contained in

$$\sum_{k=1}^{d} V_1 \otimes \ldots \otimes V_{k-1} \otimes W_k \otimes V_{k+1} \otimes \ldots V_d$$

where each W_k is a subspace of V_k.

Equivalently, if $\mathcal{T}$ is spanned by $T_1, \ldots, T_m$, letting $\alpha' = (\alpha(1), \alpha(2), \ldots, \alpha(d), \infty)$, and $\{e_i\}$ a basis for k^m,

$$\text{nc-slrk}_\alpha(\mathcal{T}) = \text{slrk}_{\alpha'}\left(\sum_{j=1}^{m} T_j \otimes e_j\right),$$

i.e. we can bump the order up by one, and look at the slice rank of a single tensor defined by the spanning set of our original tensor space, simply excluding slicing in the newly used order $(d+1)$.

The non-commutative slice rank can be thought of as a generalization of non-commutative rank, as in the case of order two tensors (matrices) they are equivalent.

Proposition 4.30. Given $\mathcal{A}$, a matrix space in $\text{Hom}(\mathbb{C}^n, \mathbb{C}^n) \cong (\mathbb{C}^n)^* \otimes \mathbb{C}^n$, $\text{ncrk}(\mathcal{A}) = \text{nc-slrk}(\mathcal{A})$.

Proof. Let $\mathcal{A}$ be a matrix space, with U a c-shrunk subspace with c maximal. Let $\mathcal{A}(U)$ be the image of U under $\mathcal{A}$, i.e. the span of $A(U)$ for all A in $\mathcal{A}$, which must have dimension $\dim(U) - c$. Pick a basis so that $e_1, \ldots, e_k$ span U. Then, all matrices in $\mathcal{A}$ are of the form:

$$\left[\begin{array}{ccc:ccc} | & & | & | & & | \\ a_1 & \ldots & a_k & b_{k+1} & \ldots & b_n \\ | & & | & | & & | \\ \hdashline & 0 & & | & & | \end{array}\right],$$

where the submatrix made of the a_i is a $(k-c) \times k$ matrix. As tensors, each matrix in $\mathcal{A}$ can be written in the form

$$\sum_{i=1}^{k} e_i^* \otimes \begin{bmatrix} a_i \\ - \\ 0 \end{bmatrix} + \sum_{i=k+1}^{n} e_i^* \otimes b_i.$$

We see our tensor space $\mathcal{A}$ is contained in $(\mathbb{C}^n)^* \otimes \mathcal{A}(U) + e_{k+1}^* \otimes \mathbb{C}^n + \ldots + e_n^* \otimes \mathbb{C}^n$, and so $\text{ncslrk}(\mathcal{A}) \leq \dim(\mathcal{A}(U)) + n - k = n - c = \text{ncrk}(\mathcal{A})$. On the other hand, now suppose $\text{nc-slrk}(\mathcal{A}) = p + q$, where $\mathcal{A}$ is contained in $V^* \otimes \mathbb{C}^n + (\mathbb{C}^n)^* \otimes W$, with $\dim(V) = p$, $\dim(W) = q$. Change basis so that V^* is spanned by $e_1^*, \ldots, e_p^*$, and W is spanned by $e_1, \ldots, e_q$. Then, as a matrix space, $\mathcal{A}$ is contained in $\left[\begin{array}{c:c} * & * \\ \hdashline & 0 \end{array}\right]$, where the leftmost block is $n \times p$, and the top right block is size $q \times (n-p)$. We have $\text{span}\{e_{p+1}, \ldots, e_n\}$ shrinks by $(n-p) - q$ under this space, so $\text{ncrk}(\mathcal{A})$ is at most $n - ((n-p) - q) = p + q$, the slice rank we started with. So in the order 2 tensor case, $\text{nc-slrk}(\mathcal{A}) = \text{ncrk}(\mathcal{A})$. □

Example 4.31. Again returning to the skew symmetric example, most recently encountered in Example 4.27, consider the tensor

$$x\big([2,1]-[1,2]\big)+y\big([3,1]-[1,3]\big)+z\big([3,2]-[2,3]\big).$$

We've already seen when x is non-zero, it's equivalent to:

$$\begin{bmatrix}0 & x & y\end{bmatrix}\otimes\begin{bmatrix}1\\0\\\frac{-z}{x}\end{bmatrix}+\begin{bmatrix}-x & 0 & z\end{bmatrix}\otimes\begin{bmatrix}1\\0\\\frac{y}{x}\end{bmatrix},$$

and when x is 0, it is equivalent to

$$\begin{bmatrix}0 & 0 & 1\end{bmatrix}\otimes\begin{bmatrix}y\\z\\0\end{bmatrix}-\begin{bmatrix}y & z & 0\end{bmatrix}\otimes\begin{bmatrix}0\\0\\1\end{bmatrix}.$$

We see that although each tensor given a fixed x, y, and z has slice rank at most 2, we will need the span of 3 slice rank one tensors to contain the subspace of tensors obtained by varying x, y and z.

Slice rank is not unique in that there are other tensor ranks defined in terms of a rank 1 tensor. For example, our standard tensor Rank: A tensor has Rank 1 if it is of the form $v_1 \otimes \ldots \otimes v_m$, and a tensor has Rank r if it can be minimally written as the sum of r Rank 1 tensors. Take any notion of rank (your "favorite rank"), which we will denote in formulas with the placeholder frk to distinguish from the standard Rank, which we have further distinguished with capitalization. For any notion of rank, we can now similarly define what it means for a space of tensors to have non-commutative rank r by using the definition of rank 1 tensors.

Definition 4.32. Given a tensor space $\mathcal{T}$, and a tensor rank frk, the non-commutative rank of $\mathcal{T}$, $\text{nc}_{\text{cs}}\text{-frk}(\mathcal{T})$, is the minimum r so that $\mathcal{T}$ is contained in the span of subspaces $W_1, \ldots, W_r$ of V, where each W_i only contains tensors with frk=1.

If frk is slice rank, this is equivalent to Definition 4.28. As we will later define another non-commutative tensor rank, we distinguish this definition with cs, a nod to it being a generalization of the c-shrunk definition for non-commutative matrix rank.

4.2.3 Generalizing Blow-ups

Recalling the blow-up definition of non-commutative matrix rank, we start with a matrix space $\mathcal{A} = \text{Span}\{A_1, \ldots, A_m\}$. The dth *tensor blow-up* of $\mathcal{A}$ is

$$\mathcal{A}^{\{d\}} := M(d) \otimes \mathcal{A} \subseteq M(dn).$$

The rank of a matrix space, $\text{rk}\,\mathcal{A}$, is the max r so that there is a matrix with rank r in $\mathcal{A}$. We define the non-commutative rank of $\mathcal{A}$ as

$$\text{ncrk}(\mathcal{A}) = \max_{d \in \mathbb{N}} \frac{\text{rk}\,\mathcal{A}^{\{d\}}}{d}.$$

We can immediately define non-commutative versions of any notion of tensor rank by extending this definition. First, we will define the vertical tensor product, as defined in [Der16].

Definition 4.33. Given tensors $T = v_1 \otimes v_2 \otimes \ldots \otimes v_m$, and $S = w_1 \otimes w_2 \otimes \ldots \otimes w_m$, with T in $V = V_1 \otimes \ldots \otimes V_m$ and S in $W = W_1 \otimes \ldots \otimes W_d$, the *vertical tensor product* of T and S is

$$T \boxtimes S := (v_1 \otimes w_1) \otimes (v_2 \otimes w_2) \otimes \ldots \otimes (v_m \otimes w_m),$$

a tensor in $V \boxtimes W$, the *vertical tensor product of tensor spaces*, similarly defined as

$$V \boxtimes W := (V_1 \otimes W_1) \otimes (V_2 \otimes W_2) \otimes \ldots \otimes (V_m \otimes W_m).$$

Note that although $V \boxtimes W$ is isomorphic to $V \otimes W$, we are considering $V \boxtimes W$ as an order d tensor space, rather than an order $2d$ tensor space. We can now move on to define non-commutative rank for tensors by generalizing our definition using blow-ups.

Definition 4.34. Let $\mathcal{T}$ be a tensor space in $V_1 \otimes V_2 \otimes \ldots \otimes V_m$, and frk any notion of tensor rank. We define the non-commutative rank as

$$\text{nc}_{\text{bu}}\text{-frk}(\mathcal{T}) = \max_{d \in \mathbb{N}} \frac{\text{frk}\left(\mathcal{T} \boxtimes \bigotimes_{i=1}^{m} k^d\right)}{\text{frk}\left(\bigotimes_{i=1}^{m} k^d\right)}.$$

In the case of matrix rank, the generic rank is equal to the max rank of the whole space. That's not always true in the case of tensor ranks, so we have a choice in the denominator of this definition, to either choose the max rank among tensors in the whole space, or to choose the generic rank of tensors in the whole space. To get around this, we consider notions of rank that are semi-continuous, where these choices become equivalent. One way of doing this for non-semi-continuous ranks is to replace them with the border version of that rank.

4.2.4 Generalizing Semi-stability

Our last definition of non-commutative matrix rank utilized semi-stability. In [Der20], Derksen defines the G-stable rank, taking nods from Hilbert-Mumford criterion and semi-stability. To generalize the definition to tensor spaces, we will go up an order to build a single tensor.

Definition 4.35. Let $\mathcal{T}$ be a tensor space spanned by $T_1, \ldots, T_m$ in $V_1 \otimes V_2 \otimes \ldots \otimes V_d$, $G = \prod_{i=1}^{d} \text{GL}(V_i)$, and α, a d dimensional weight vector of positive reals. Letting $\alpha' = (\alpha(1), \ldots, \alpha(d), \infty)$, and $e_1, \ldots, e_m$ a basis for k^m, the non-commutative G-stable rank of $\mathcal{T}$ is

$$\text{nc-rk}_{\alpha}^{G}(\mathcal{T}) := \sup_{g \in G} \min_{i} \frac{\alpha'(i) || g \cdot \sum_{j=1}^{m} T_j \otimes e_j ||^2}{|| \Phi_i(g \cdot \sum_{j=1}^{m} T_j \otimes e_j) ||_{\sigma}{}^2}.$$

We similarly defined non-commutative slice rank of a tensor space in this way, looking at the minimum slice rank 1 tensors needed to give us $\sum_{j=1}^{m} T_j \otimes e_j$, restricting the slices to not be in the $(d+1)$st order. For this reason, the inequalities between slice rank and G-stable rank for tensors will hold for their non-commutative definitions for tensor spaces.

Example 4.36. Let $\alpha = (1, 1)$ and $\mathcal{T}$ be the skew-symmetric matrices (as a tensor space), which are spanned by the three tensors: $[2, 1] - [1, 2]$, $[3, 1] - [1, 3]$, and $[3, 2] - [2, 3]$. Bumping up to the next order to calculate the non-commutative G-stable rank, we get

$$T = [2, 1, 1] - [1, 2, 1] + [3, 1, 2] - [1, 3, 2] + [3, 2, 3] - [2, 3, 3],$$

a tensor with Euclidean norm $\sqrt{6}$. The first flattening of T is

$$\Phi_1(T) = \left[\begin{array}{ccc|ccc|ccc} 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & -1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \end{array}\right],$$

with largest singular value $\sqrt{2}$. Similarly, the second flattening has singular value $\sqrt{2}$, and so $\mathrm{rk}^G(\mathcal{T}) \geq 3$. However, as we saw in Example 4.31, $\mathcal{T}$ has non-commutative slice rank 3, and so this must be the non-commutative G-stable rank.

4.2.5 Inequalities

We've now defined non-commutative rank in two distinct ways, in terms of spanning subspaces (Definition 4.32) and blow-ups (Definition 4.34).

Proposition 4.37. Let $\mathcal{T}$ be a tensor space in $V_1 \otimes V_2 \otimes \ldots \otimes V_m$. Then, $\mathrm{nc_{cs}\text{-}frk}(\mathcal{T}) \leq \mathrm{nc_{bu}\text{-}frk}(\mathcal{T})$, where frk is chosen to be slice rank.

Proof. Let $\mathrm{nc_{cs}\text{-}frk}(\mathcal{T}) = r$, where $r = \sum \dim(W_i)$, and T is contained in the span of all $V_1 \otimes \ldots \otimes V_{i-1} \otimes W_i \otimes V_{i+1} \otimes \ldots \otimes V_m$. Then, every tensor in $\mathcal{T} \boxtimes \bigotimes_{i=1}^{m} k^d$ has slice rank at most $d \cdot \sum \dim(W_i)$. This can also be seen via the inequality in [BCC+17, Proposition 4.2]. The max slice rank of a tensor in $\bigotimes_{i=1}^m k^d$ is d, so the denominator of $\mathrm{nc}_{\mathrm{bu}}\text{-frk}(\mathcal{T}) = d$, and $\mathrm{nc}_{\mathrm{bu}}\text{-frk}(\mathcal{T}) \leq \mathrm{nc}_{\mathrm{cs}}\text{-frk}(\mathcal{T})$. □

BIBLIOGRAPHY

[Ami66] SA Amitsur. Rational identities and applications to algebra and geometry. *Journal of Algebra*, 3(3):304–359, 1966.

[BCC$^+$17] Jonah Blasiak, Thomas Church, Henry Cohn, Joshua A Grochow, Eric Naslund, William F Sawin, and Chris Umans. On cap sets and the group-theoretic approach to matrix multiplication. *Discrete Analysis*, page 1245, 2017.

[BCRL79] Dario Bini, Milvio Capovani, Francesco Romani, and Grazia Lotti. O(n2.7799) complexity for n $\times$ n approximate matrix multiplication. *Information Processing Letters*, 8(5):234–235, 1979.

[BD06] Matthias Bürgin and Jan Draisma. The Hilbert null-cone on tuples of matrices and bilinear forms. *Mathematische Zeitschrift*, 254(4):785–809, 2006.

[BLR80] Dario Bini, Grazia Lotti, and Francesco Romani. Approximate solutions for the bilinear form computational problem. *SIAM J. Comput.*, 9:692–697, 11 1980.

[CB96] William Crawley-Boevey. On homomorphisms from a fixed representation to a general representation of a quiver. *Transactions of the American Mathematical Society*, 348(5):1909–1919, 1996.

[CC70] J Douglas Carroll and Jih-Jie Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of "eckart-young" decomposition. *Psychometrika*, 35(3):283–319, 1970.

[CK20] Calin Chindris and Daniel Kline. Edmonds' problem and the membership problem for orbit semigroups of quiver representations. *arXiv preprint arXiv:2008.13648*, 2020.

[CK21] Calin Chindris and Daniel Kline. Simultaneous robust subspace recovery and semi-stability of quiver representations. *Journal of Algebra*, 577:210–236, 2021.

[Coh95] Paul M Cohn. Encyclopedia of mathematics and its applications, skew fields: Theory of general division rings, 1995.

[CR99] Paul Moritz Cohn and Christophe Reutenauer. On the construction of the free field. *International journal of Algebra and Computation*, 9(03n04):307–323, 1999.

[Der16] Harm Derksen. On the nuclear norm and the singular value decomposition of tensors. *Foundations of Computational Mathematics*, 16(3):779–811, 2016.

[Der20] Harm Derksen. The g-stable rank for tensors. *arXiv preprint arXiv:2002.08435*, 2020.

[DF73] Peter Donovan and Mary Ruth Freislich. The representation theory of finite graphs and associated algebras. Carleton Lecture Notes, 1973.

[DF15] Harm Derksen and Jiarui Fei. General presentations of algebras. *Advances in Mathematics*, 278:210–237, 2015.

[DK02] Harm Derksen and Gregor Kemper. On global degree bounds for invariants. *Invariant theory in all characteristics*, 35:37–41, 2002.

[DL78] Richard A. Demillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193 – 195, 1978.

[DM17a] Harm Derksen and Visu Makam. Generating invariant rings of quivers in arbitrary characteristic. *Journal of Algebra*, 489:435–445, 2017.

[DM17b] Harm Derksen and Visu Makam. Polynomial degree bounds for matrix semi-invariants. *Advances in Mathematics*, 310:44–63, 2017.

[DM18a] Harm Derksen and Visu Makam. Degree bounds for semi-invariant rings of quivers. *Journal of Pure and Applied Algebra*, 222(10):3282–3292, 2018.

[DM18b] Harm Derksen and Visu Makam. On non-commutative rank and tensor rank. *Linear and Multilinear Algebra*, 66(6):1069–1084, 2018.

[DW17] Harm Derksen and Jerzy Weyman. *An introduction to quiver representations*, volume 184. American Mathematical Soc., 2017.

[Edm67] Jack Edmonds. Systems of distinct representatives and linear algebra. *J. Res. Nat. Bur. Standards Sect. B*, 71(4):241–245, 1967.

[FR04] Marc Fortin and Christophe Reutenauer. Commutative/noncommutative rank of linear matrices and subspaces of matrices of low rank. *Séminaire Lotharingien de Combinatoire*, 52:B52f, 2004.

[Gab72] Peter Gabriel. Unzerlegbare Darstellungen I. *manuscripta mathematica*, 6(1):71–103, March 1972.

[GGOW16] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. A deterministic polynomial time algorithm for non-commutative rational identity testing. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 109–117. IEEE, 2016.

[GR91] Israel M Gel'fand and Vladimir S Retakh. Determinants of matrices over noncommutative rings. *Functional Analysis and Its Applications*, 25(2):91–102, 1991.

[Gro19] Joshua Grochow. New applications of the polynomial method: The cap set conjecture and beyond. *Bulletin of the American Mathematical Society*, 56(1):29–64, 2019.

[Gur04] Leonid Gurvits. Classical complexity and quantum entanglement. *Journal of Computer and System Sciences*, 69(3):448–484, 2004.

[Hab75] William J Haboush. Reductive groups are geometrically reductive. *Annals of Mathematics*, 102(1):67–83, 1975.

[Har70] Richard A Harshman. Foundations of the parafac procedure: Models and conditions for an" explanatory" multimodal factor analysis. 1970.

[Hil90] David Hilbert. Ueber die theorie der algebraischen formen. *Mathematische annalen*, 36(4):473–534, 1890.

[Hil93] David Hilbert. Über die vollen invariantensysteme. *Mathematische annalen*, 42(3):313–373, 1893.

[Hit27] Frank L Hitchcock. The expression of a tensor or a polyadic as a sum of products. *Journal of Mathematics and Physics*, 6(1-4):164–189, 1927.

[Hua49] Loo-Keng Hua. On the automorphisms of a sfield. *Proceedings of the National Academy of Sciences of the United States of America*, 35(7):386, 1949.

[Hum75] J.E. Humphreys. *Linear Algebraic Groups*. Graduate Texts in Mathematics. Springer, 1975.

[HW14] Pavel Hrubeš and Avi Wigderson. Non-commutative arithmetic circuits with division. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 49–66, 2014.

[IKQS15] Gábor Ivanyos, Marek Karpinski, Youming Qiao, and Miklos Santha. Generalized Wong sequences and their applications to Edmonds' problems. *Journal of Computer and System Sciences*, 81(7):1373–1386, 2015.

[IMQ21] Gábor Ivanyos, Tushant Mittal, and Youming Qiao. Symbolic determinant identity testing and non-commutative ranks of matrix lie algebras. *arXiv preprint arXiv:2109.06403*, 2021.

[IQS17] Gábor Ivanyos, Youming Qiao, and KV Subrahmanyam. Non-commutative Edmonds' problem and matrix semi-invariants. *Computational Complexity*, 26(3):717–763, 2017.

[IQS18] Gábor Ivanyos, Youming Qiao, and KV Subrahmanyam. Constructive non-commutative rank computation is in deterministic polynomial time. *Computational Complexity*, 27(4):561–593, 2018.

[Kac82] Victor G Kac. Infinite root systems, representations of graphs and invariant theory, ii. *Journal of algebra*, 78(1):141–162, 1982.

[Kin94] Alastair D King. Moduli of representations of finite dimensional algebras. *The Quarterly Journal of Mathematics*, 45(4):515–530, 1994.

[KP96] Hanspeter Kraft and Claudio Procesi. Classical invariant theory, a primer. Lecture Notes, 1996.

[Kru24] Wolfgang Krull. Algebraische theorie der ringe. ii. *Mathematische Annalen*, 91(1):1–46, 1924.

[Kru77] Joseph B Kruskal. Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear algebra and its applications*, 18(2):95–138, 1977.

[Lan12] Joseph M Landsberg. Tensors: geometry and applications. *Representation theory*, 381(402):3, 2012.

[LBP90] Lieven Le Bruyn and Claudio Procesi. Semisimple representations of quivers. *Transactions of the American Mathematical Society*, pages 585–598, 1990.

[Naz73] Liudmila A Nazarova. Representations of quivers of infinite type. *Mathematics of the USSR-Izvestiya*, 7(4):749, 1973.

[New78] PE Newstead. Lectures on Introduction to moduli problems and orbit spaces. Tata Institute of Fundamental Research, 1978.

[NM63] Masayoshi Nagata and Takehiko Miyata. Note on semi-reductive groups. *Journal of Mathematics of Kyoto University*, 3(3):379–382, 1963.

[Noe26] Emmy Noether. Der endlichkeitssatz der invarianten endlicher linearer gruppen der charakteristik p. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1926:28–35, 1926.

[OAHW89] John G O'Grady, Graeme JM Alexander, Karen M Hayllar, and Roger Williams. Early indicators of prognosis in fulminant hepatic failure. *Gastroenterology*, 97(2):439–445, 1989.

[Rem11] Robert Remak. Über die zerlegung der endlichen gruppen in direkte unzerlegbare faktoren. 1911.

[Sch29] Otto Schmidt. Über unendliche gruppen mit endlicher kette. *Mathematische Zeitschrift*, 29(1):34–41, 1929.

[Sch80] Jacob T Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM (JACM)*, 27(4):701–717, 1980.

[Sch92] Aidan Schofield. General representations of quivers. *Proceedings of the London Mathematical Society*, 3(1):46–64, 1992.

[ST16] Will Sawin and Terence Tao. Notes on the "slice rank" of tensors. Blog post, 2016. Available at `https://terrytao.wordpress.com/2016/08/24/notes-on-the-slice-rank-of-tensors`.

[Tao16] Terence Tao. A symmetric formulation of the croot-lev-pach-ellenberg-gijswijt capset bound. Blog post, 2016. Available at `https://terrytao.wordpress.com/2016/05/18/`.

[Wei95] Charles A Weibel. *An introduction to homological algebra*. Number 38. Cambridge university press, 1995.

[Won74] Kai-Tak Wong. The eigenvalue problem λtx+ sx. *Journal of differential Equations*, 16(2):270–280, 1974.

[Yon60] Nobuo Yoneda. On ext and exact sequences. *J. Fac. Sci. Univ. Tokyo Sect. I*, 8(507-576):1960, 1960.

[Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *International symposium on symbolic and algebraic manipulation*, pages 216–226. Springer, 1979.