Tensors: Entanglement, Geometry, and Combinatorics

by

Benjamin Lovitz

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Applied Mathematics (Quantum Information)

Waterloo, Ontario, Canada, 2022

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Gilad Gour Professor, Dept. of Math and Stats, University of Calgary
Supervisor(s):	William Slofstra Assistant Professor, Dept. of Pure Math, University of Waterloo
Internal Member:	Achim Kempf Professor, Dept. of Applied Math, University of Waterloo
Internal-External Member:	Rafael Oliveira Assistant Professor, Cheriton School of Computer Science, University of Waterloo
Other Member(s):	Matthew Satriano Associate Professor, Dept. of Pure Math, University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

The contents of Chapters 3 and 4 are based on the preprint [91], which was written in collaboration with Nathaniel Johnston. The contents of Chapter 5 are based on the preprint [95], which was written in collaboration with Vincent Steffan. The contents of Chapter 6 are based on the journal article [90], which is a single-author work. The contents of Chapter 7 are based on the preprint [94], which was written in collaboration with Fedor Petrov.

Abstract

In this thesis, we use algebraic-geometric and combinatorial techniques to study tensor decompositions, with applications in quantum information theory, algebraic complexity theory, and algebraic statistics. A *tensor* is a multilinear map. These objects naturally generalize matrices, and have many useful applications in math and science. They can be used to describe nearly any dataset, and any pure (or even, mixed) quantum state. A *decomposition* of a tensor v is an expression of v as a linear combination of "elementary tensors," which are defined according to the application. For a fixed choice of elementary tensors X, the *X-rank* of v is the minimum number of elementary tensors needed to span a space containing v. The *border X-rank* of v is the minimum number of elementary tensors needed to approximate v arbitrarily well. The most common choice of elementary tensors X are the product tensors, and we refer to the X-rank and border X-rank under this choice as simply the *rank* and *border rank*, respectively.

In quantum information theory, we think of a projective tensor $[v]$ (the projectivization of a tensor v) as a pure quantum state shared by multiple laboratories. Just as matrices are much better understood than multi-way tensors, bipartite entanglement is much better understood than multipartite entanglement. For example, while we know that the most useful state in the bipartite setting under the LOCC paradigm is the *canonical maximally entangled state*, along with its local unitary equivalents, this question becomes more difficult in the multi-party setting. As a natural recourse, we determine the "usefulness" of multi-way states for facilitating a particular task: local unambiguous state discrimination. We also study entangled subspaces, entanglement witnesses, and so-called *absolutely entangled sets* in the multi-way setting. A different set of elementary tensors that are relevant in quantum information are the so-called *stabilizer tensors*. The stabilizer rank of a pure quantum state $[v]$ represents the computational cost of classically simulating Clifford circuits applied to $[v]$ under the stabilizer formalism. We introduce new techniques from number theory and algebraic geometry for studying the stabilizer rank, and obtain simplified proofs of the best-known lower bounds on stabilizer rank up to a log factor.

In algebraic complexity theory, we think of a tensor v as a multilinear map. In this context, the rank of v is a useful barometer for the computational cost of implementing this multilinear map. We prove new lower bounds on tensor rank.

In algebraic statistics, we think of a tensor v as a probability vector for multiple (observable) random variables. Under an assumption of conditional independence, a decomposition of v into product tensors corresponds to a choice of latent (unobservable) random variable that gives rise to the observable random variables. If v has a unique tensor rank decomposition, this means that there is only one consistent choice of latent random variable (with the smallest number of outcomes) consistent with the probability vector v. We obtain a new sufficient condition for a given decomposition of v to be the unique rank decomposition of v, strengthening a theorem of Joseph Kruskal.

Acknowledgments

I am extremely grateful to my supervisor, William Slofstra, for agreeing to let me show him a tensor problem four years ago, and then taking me under his wing. Coming from physics, I had some catching up to do, and William showed me the way. This thesis is a product of taking his advice at nearly every turn.

Thanks to John Watrous for first taking me on as his PhD student. Even when my interests diverged from his purview, he continued to offer me tremendous time and energy. I am grateful for his encouragement to pursue my own research questions.

Thanks to Laura Frazee, Samanthi Sooriyabandara, Brian Ingalls, and Michael Waite for their administrative support throughout my PhD. Thanks to my committee members for reading this thesis and attending my defence.

Thanks to Tara, my love, for dancing with me; for sharing these years with me; for demonstrating such depth of compassion for others; for her strength, self-awareness and complete honesty. I am humbled by her love.

Thanks to Julian, my best friend and playmate, for his unending enthusiasm and limitless love; his kindness and care; his questions and storytelling.

Thanks to Ryan, my housemate and friend, for welcoming me into his home with kindness and warmth; for demonstrating empathy and thoughtfulness at every turn.

Thanks to my family. My father, Bob, for his good humour, and for always being there for me. My mother, Laura, for her kindness and care, and for showing me how to go on proper adventures. My brother, David, for keeping me grounded and playful.

Thanks to Connor, Dan, Jeff and Júlia for much needed emotional support and math gossip so many times over. Thanks to everyone who makes Kitchener my home. ♡

Table of Contents

Chapter 1

Introduction

Let $[m] = \{1, \dots, m\}$ when m is a positive integer, and let $[0] = \{\}$ be the empty set. For a vector space $\mathcal{V}$ over a field $\mathbb{F}$, let $\mathbb{P}(\mathcal{V})$ (or more briefly, $\mathbb{P}\mathcal{V}$) denote the set of one-dimensional linear subspaces of $\mathcal{V}$. Given a subset $X \subseteq \mathcal{V}$ for which $\operatorname{span}(X) = \mathcal{V}$, we define the *$X$-rank* of an element $v \in \mathcal{V}$ (or $[v] \in \mathbb{P}\mathcal{V}$), denoted $\operatorname{rank}_X(v)$ (or $\operatorname{rank}_X([v])$), to be the minimum number n for which there exist $x_1, \dots, x_n \in X$ such that $v \in \operatorname{span}\{x_1, \dots, x_n\}$. We refer to X as the set of *elementary vectors*. A decomposition of v into a linear combination of $\operatorname{rank}_X(v)$ elementary vectors is called an *X-rank decomposition* of v. An expression of v as a linear combination of elementary vectors (not necessarily of minimum number) is known simply as an *X-decomposition* of v. In this thesis, we concern ourselves with X-decompositions under several different (scientifically motivated) choices of X and $\mathcal{V}$.

Of particular interest to us are X-decompositions when $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ for vector spaces $\mathcal{V}_1, \dots, \mathcal{V}_m$ over $\mathbb{F}$, and X is the set of *product tensors* in $\mathcal{V}$, i.e., the set of non-zero vectors $z \in \mathcal{V}$ of the form $z = z_1 \otimes \cdots \otimes z_m$, with $z_j \in \mathcal{V}_j$ for all $j \in [m]$. In this case, we refer to vectors in $\mathcal{V}$ as *tensors* in order to emphasize the tensor product structure. In this case, we refer to the X-rank as the *tensor rank* (or simply the *rank*), refer to the X-border rank as simply the *border rank*, and use the shorthand $\operatorname{rank}(v) = \operatorname{rank}_X(v)$. Note that when $m = 2$, $\mathcal{V}_1 \otimes \mathcal{V}_2$ is naturally isomorphic to the vector space of linear maps from $\mathcal{V}_1^*$ to $\mathcal{V}_2$, and the rank of a tensor $v \in \mathcal{V}_1 \otimes \mathcal{V}_2$ is equal to the rank of the corresponding linear map. We refer to the spaces $\mathcal{V}_j$ that make up the space $\mathcal{V}$ as *subsystems*. We also study symmetric and alternating variants of tensor decompositions.

Also of interest to us are X-decompositions when $\mathcal{V} = (\mathbb{C}^2)^{\otimes m}$ and X is the set of *stabilizer tensors*, i.e. the set of tensors in the orbit of $\mathrm{e}_0^{\otimes m}$ under the action of the Clifford group, where $\{\mathrm{e}_0, \mathrm{e}_1\}$ denotes the standard basis of $\mathbb{C}^2$. We call the X-rank in this setting the *stabilizer rank*, and denote it by $\chi(\cdot)$.

The types of X-decompositions we have just described are examples of *tensor decompositions*: X-decompositions for which the underlying vector space $\mathcal{V}$ is endowed with a tensor product structure. In this thesis, we use algebraic-geometric and combinatorial techniques to study tensor decompositions. Our study of tensor decompositions is motivated by the following applications in quantum information theory, algebraic complexity theory, and algebraic statistics:

1. Quantum information theory: Elements of $\mathbb{P}(\mathbb{C}^{d_1} \otimes \cdots \otimes \mathbb{C}^{d_m})$ represent (pure) quantum states shared by m parties of local dimensions $d_1, \dots, d_m$. In this context, the tensor rank represents, to some extent, the amount of entanglement present in the state [36]. The stabilizer rank of a state in $\mathbb{P}((\mathbb{C}^2)^{\otimes m})$ measures the computational cost of classically simulating Clifford circuits applied to that state under the stabilizer formalism [28].

2. Algebraic complexity theory: The rank of a tensor in $\mathcal{V}_1 \otimes \mathcal{V}_2 \otimes \mathcal{V}_3$ corresponds to the multiplicative complexity of implementing the associated bilinear map $\mathcal{V}_1^* \otimes \mathcal{V}_2^* \to \mathcal{V}_3$. For the (well-studied) bilinear map which takes as input two matrices and outputs their product, more can be said: The (asymptotic) *total complexity* of this map is precisely quantified by the (asymptotic) tensor rank of the corresponding tensor, which is aptly known as the *matrix multiplication tensor* [32].

3. Algebraic statistics: In latent parameter models of statistical machine learning, tensors are used to represent the observed probabilities or expectation values of a collection of random variables. If a tensor has a unique tensor rank decomposition (up to trivialities), then one can uniquely infer the latent parameters of the model [4, 17].

In this thesis, we consider the following topics, with applications in each of the above three areas:

Chapter 3: Generic local state discrimination with pre-shared entanglement In the two-party setting, we know that the most useful state to share under the LOCC paradigm is the *canonical maximally entangled state* (along with its local unitary orbit), as this state can be locally converted into any other two-way state [98]. However, this question becomes more difficult in the multi-party setting. A natural sub-problem is to determine the most useful states for specific, fundamental tasks. We determine the most useful pre-shared entangled states for facilitating the task of local unambiguous state discrimination.

Chapter 4: Entangled subspaces and entanglement witnesses The border rank captures, to some extent, the amount of entanglement present in a state. We define an r*-entangled subspace* to be a linear subspace of multipartite space for which every non-zero element has border rank greater than r, which extends previous notions to higher, multipartite entanglement. We determine the maximum dimension of r-entangled subspaces, obtain explicit constructions of r-entangled subspaces of maximum dimension, and find applications of r-entangled subspaces to *entanglement witnesses*: quantum measurements that detect entanglement.

Chapter 5: New techniques for bounding stabilizer rank When X is the set of stabilizer states, the X-rank is a useful barometer for the classical simulation cost of quantum circuits under the stabilizer formalism. Despite its practical importance, few techniques are known for bounding this quantity. We introduce new techniques from number theory and algebraic geometry to bound the stabilizer rank.

Chapter 6: Decomposable correlation matrices We introduce and study the following question: Given a set of unit vectors $\{v_1, \dots, v_n\} \subseteq \mathbb{C}^n$ and a positive integer $d < n$, when does there exist a positive integer m and an isometry $V : \mathbb{C}^n \to (\mathbb{C}^d)^{\otimes m}$ for which $Vv_a \in (\mathbb{C}^d)^{\otimes m}$ is a product tensor for all $a \in \{1, \dots, m\}$? After motivating this question in quantum information theory, we identify properties of $\{v_1, \dots, v_n\}$ that either ensure or negate the existence of such a positive integer m and isometry V.

Chapter 7: Uniqueness of tensor rank decompositions Arguably the best-known uniqueness result for tensor decompositions, proven by Joseph Kruskal in 1977, states that a sum of product tensors constitutes a unique tensor rank decomposition if the so-called *k-ranks* of the product tensors are large. While several extensions of Kruskal's theorem are already present in the literature, all of these use Kruskal's original permutation lemma, and hence still cannot certify uniqueness when the k-ranks are below a certain threshold. We generalize Kruskal's theorem using a completely new (matroidal) proof technique. Our generalization contains many of these extensions, and can certify uniqueness below this theshold.

In the remainder of this introduction, we describe in greater detail the utility of tensor decompositions in quantum information theory, algebraic complexity theory, and algebraic statistics; and introduce more fully the topics covered in each chapter.

1.1 Tensors in quantum information theory

In this section, we introduce applications of tensor decompositions in quantum information theory, and present our work in this vein. For a complex vector space $\mathcal{V}$, let $\mathcal{S}(\mathcal{V})$ be the set of unit vectors in $\mathcal{V}$ (with respect to the Euclidean norm), and let $\mathrm{U}(\mathcal{V})$ be the set of unitary operators on $\mathcal{V}$. A *(pure) quantum state* is a unit vector modulo phase, i.e. an element of $\mathcal{S}(\mathcal{V})/\mathrm{U}(\mathbb{C})$. Under the canonical identification between this quotient and $\mathbb{P}(\mathcal{V})$, which sends a unit vector to its span, we also refer to elements of $\mathbb{P}(\mathcal{V})$ as quantum states. The quotient space $\mathcal{S}(\mathcal{V})/\mathrm{U}(\mathbb{C})$ is also naturally bijective to the set of rank-one projectors on $\mathcal{V}$, via the map which sends the equivalence class of $v \in \mathcal{S}(\mathcal{V})$ to vv^*. We therefore refer also to such rank-one projectors as pure quantum states. While our projective notation for pure quantum states may seem overly pedantic, we prefer it to other alternatives because it is mathematically precise, and because it allows us to identify certain, of-interest sets of quantum states with algebraic varieties. We will use such identifications in Chapters 3, 4, and 5 to study local state discrimination, entangled subspaces, and stabilizer rank, respectively. For brevity, we will refer to pure quantum states simply as "states," and we will refer to mixed quantum states (defined and treated only in Section 4.2) as "mixed states."

Let m be a positive integer, and let $\mathcal{V} = \bigotimes_{j=1}^m \mathbb{C}^{d_j}$ be an m-partite vector space. For a positive integer r, let $[\tau_{r,m}] = \left[\sum_{a\in[r]} \mathrm{e}_a^{\otimes m}\right] \in \mathbb{P}((\mathbb{C}^r)^{\otimes m})$ be the *tensor-rank-r GHZ state*. It is well-known that the set of states in $\mathbb{P}\mathcal{V}$ of tensor rank at most r is precisely the set of states in $\mathbb{P}\mathcal{V}$ obtainable from $[\tau_{r,m}]$ by SLOCC. In other words, the tensor rank of a state $[v] \in \mathbb{P}\mathcal{V}$ is equal to the size of the smallest GHZ state needed to obtain $[v]$ by SLOCC. Similarly, the border rank of $[v]$ is equal to the size of the smallest GHZ state needed to obtain arbitrarily good approximations to $[v]$ by SLOCC. In this way, the tensor rank and border rank provide loose barometers for the amount of entanglement present in a state [36]. In Section 1.1.3 we study a different notion of rank: the *stabilizer rank*, which is the X-rank when X is equal to the set of stabilizer states, which is motivated as the computational cost of classically simulating quantum circuits under the stabilizer formalism.

In the remainder of this section, we introduce our work on local state discrimination, entangled subspaces, stabilizer rank, and decomposable correlation matrices.

1.1.1 Chapter 3: Generic local state discrimination with pre-shared entanglement

It is practical to ask which entangled quantum states are the most useful for non-local quantum information processing. A natural non-local setting occurs when multiple spatially separated parties wish to jointly execute some task, but are physically constrained so that they can only perform quantum operations within their own laboratories, and classically communicate between each other. This is known as the *local operations and classical communication* (LOCC) paradigm. Assuming that the parties are allowed to use pre-shared entanglement to help them perform the task, what are the most useful entangled states for them to share?

It is well-known that for two parties of local dimensions r, the *canonical maximally entangled state* $[\tau_{r,2}] = \left[\frac{1}{\sqrt{r}} \sum_{a=1}^{r} \mathrm{e}_a \otimes \mathrm{e}_a\right]$, along with its local unitary equivalents, is the most useful state under the LOCC paradigm. (The vectors $\mathrm{e}_1, \dots, \mathrm{e}_r$ are the standard basis vectors for $\mathbb{C}^r$.) The state $[\tau_{r,2}]$ is the most useful for any task because it is majorized by every other state in the space, and hence, by Nielsen's theorem, it can be converted into every other state by LOCC [98]. For more than two parties, the results are not nearly so clear cut. Far from the existence of a single state that can be converted into any other, already in the three-qubit case an *infinite number* of pure states are required to reach every other state by LOCC, and in the four-qubit case a *full measure* set of states is required [46]. While LOCC convertibility cleanly determines the most useful entangled states under LOCC in the bipartite (two-party) case, it provides a rather impractical answer in the multipartite (more-than-two-party) case.

A natural alternative in the multipartite case is to determine the most useful states for specific, fundamental tasks. In Chapter 3, we consider the task of *quantum state discrimination*, in which one is handed a physical quantum state ρ, along with a classical description of n quantum states $(\rho_1, \dots, \rho_n)$, and promised that $\rho = \rho_a$ for some $a \in \{1, \dots, n\}$. The goal is to determine a by measuring ρ. While several variants of quantum state discrimination have been studied, we focus on *unambiguous state discrimination*. In this variant, one seeks a measurement that will either correctly output a (with non-zero probability), or else output "I don't know" (or more briefly, "?"). It is required that the measurement never erroneously outputs b for some $b \neq a$. Mathematically, an n-tuple of quantum states $(\rho_1, \dots, \rho_n)$ is *unambiguously discriminable* if there exists an $n+1$-outcome POVM measurement $\{M_1, \dots, M_n, M_?\}$ such that $\langle M_a, \rho_a \rangle > 0$ for all $a \in \{1, \dots, n\}$, and $\langle M_a, \rho_b \rangle = 0$ for all $a \neq b \in \{1, \dots, n\}$. It is not difficult to prove that a tuple of pure states is unambiguously discriminable if and only if the corresponding unit vectors are linearly independent. A similar statement can also be proven for mixed states.

We study *local* unambiguous state discrimination (LUSD), in which $m \geq 2$ parties wish to unambiguously discriminate an n-tuple of m-partite quantum states $(\rho_1, \dots, \rho_n)$ using LOCC. Suppose that the parties have access to a pre-shared (pure) entangled state $[w]$ that they can use to implement their measurement. How many states can be locally unambiguously discriminated with $[w]$? This number depends on the particular tuple of states to be discriminated. For example, n orthogonal product states can always be locally discriminated, but n identical states can never be. We therefore determine how many *generic* states can be discriminated. For any pure resource state $[w]$, we determine the maximum number $n_{\max}$ for which almost all $n_{\max}$-tuples of pure states $([v_1], \dots, [v_{n_{\max}}])$ are locally discriminable with $[w]$. We also prove a strong converse: for any $n > n_{\max}$, almost all n-tuples of pure states are *not* locally discriminable with $[w]$. We express $n_{\max}$ in terms of the dimension of a certain algebraic variety, which depends on $[w]$ and is known in several cases. The most useful states for generic LUSD are therefore the ones that maximize this dimension. Using algebraic group theoretic techniques, we prove that *almost all* resource states $[w]$ maximize this dimension, and hence are maximally useful for LUSD.

To conclude this section, we introduce our characterization of generic LUSD with pre-shared entanglement in more detail. For complex vector spaces $\mathcal{V}_1, \dots, \mathcal{V}_m$ and $\mathcal{W}_1, \dots, \mathcal{W}_m$, let $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ and $\mathcal{W} = \mathcal{W}_1 \otimes \cdots \otimes \mathcal{W}_m$. For a state $[w] \in \mathbb{P}\mathcal{W}$, the *SLOCC image of* $[w]$ *in* $\mathbb{P}\mathcal{V}$, denoted by $\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w])$, is the set of states in $\mathbb{P}\mathcal{V}$ obtainable from $[w]$ by stochastic local operations and classical communication (SLOCC). This set is also known as the *downward closure of* $[w]$ *with respect to SLOCC*, and can alternatively be characterized as the set of pure states in $\mathbb{P}\mathcal{V}$ that have a Tucker decomposition with core state $[w]$ [124, 104]. Letting $d = \dim(\overline{\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w])})$ be the Krull dimension of the closure of $\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w])$, we prove that almost all tuples of $n_{\max} := d + 1$ pure states are locally discriminable with $[w]$, and that for any $n > n_{\max}$, almost all n-tuples of pure states are *not* locally discriminable with $[w]$.

1.1.2 Chapter 4: Entangled subspaces and entanglement witnesses

An *entangled subspace* is a linear subspace of multipartite space for which every nonzero element exceeds some specified level entanglement. Entangled subspaces are connected to unextendible product bases, and have found applications, for example, in quantum error correction [59, 109] and quantum tomography [69]. An *entanglement witness* is a quantum measurement that can be used to detect entanglement. Entanglement witnesses have found many experimental applications [12, 23, 22].

As a generalization of previous notions of entangled subspaces, we define an *r-entangled subspace* to be a projective linear subspace that avoids $\overline{\mathrm{Im}([\tau_{r,m}])}$ (the set of pure states of border rank at most r). In bipartite space, the tensor rank is equal to the Schmidt rank, and $\mathrm{Im}([\tau_{r,2}])$ is already closed.

We use algebraic-geometric methods to generalize several known results on entangled subspaces to higher multipartite entanglement. We determine the maximum dimension of an r-entangled subspace in multipartite space, in terms of the dimension of $\overline{\mathrm{Im}([\tau_{r,m}])}$. Similarly to the task of LUSD, *almost all* subspaces of this maximum dimension are r-entangled. Despite the abundance of entangled subspaces of maximum dimension, it appears difficult to explicitly write down such a subspace. (This is analogous to the situation in classical error correcting codes, where it has long been known that almost all codes have optimal rate, but until recently we there were no explicit constructions of such codes.) We explicitly construct multipartite 2-entangled subspaces of maximal dimension, and extend our results to bosonic and fermionic quantum systems.

It is known that, under various notions of entanglement, the maximum dimension of an entangled subspace is precisely the maximum number of negative eigenvalues of an entanglement witness [8, 76, 77]. The number of negative eigenvalues quantifies "how good" the witness is at detecting entanglement. We prove that this connection between subspaces and negative eigenvalues holds under a much more general notion of witness, including multipartite r-entanglement witnesses.

1.1.3 Chapter 5: New techniques for bounding stabilizer rank

It is of great practical importance to determine the classical simulation cost of quantum computations. Indeed, lower bounds on the simulation cost indicate quantum speedups, while upper bounds can help us to understand the limitations of quantum computation. The stabilizer rank and approximate stabilizer rank are useful barometers for determining this cost under the stabilizer formalism. In this work, we introduce new techniques for bounding these quantities.

Recall that the *Clifford group* is the projective unitary group generated by tensor products of the so-called *Clifford gates* $\{X, S, CNOT\}$. A *stabilizer state* is a quantum state in the orbit of a computational basis state under the action of the Clifford group. For a (pure) quantum state $[v]$, we define its *stabilizer rank*, denoted $\chi([v])$, to be the smallest integer r for which $[v]$ can be written as a superposition of r stabilizer states. For a real number

$\delta > 0$, the δ-*approximate stabilizer rank*, $\chi_\delta([v])$, is defined as the minimum stabilizer rank over all quantum states that are δ-close to $[v]$.

The stabilizer rank $\chi([v])$ is a useful measure for the computational cost of simulating Clifford circuits with input state $[v]$. Indeed, by the Gottesman-Knill theorem, Clifford circuits with stabilizer state inputs can be classically simulated with cost polynomial in the number of qubits and Clifford gates [58]. Recent extensions of Gottesman-Knill, which handle an arbitrary input state $[v]$ and apply under various notions of classical simulation, reveal that Clifford circuits applied to $[v]$ can be simulated with computational cost polynomial in the number of qubits, Clifford gates, and $\chi([v])$ [26, 28, 27, 103]. The approximate stabilizer rank, $\chi_\delta([v])$, has similarly been shown to quantify the computational cost of classically approximating Clifford circuits with input state $[v]$ [27].

The T-*count* of a quantum state $[v]$ is the minimum number of T-gates needed to prepare $[v]$ using a circuit consisting only of Clifford+T-gates and post-selective computational basis measurements. It is known that if $[v]$ has T-count m, then $\chi([v]) \leq \chi([T^{\otimes m}])$, where $[T] = [\mathrm{e}_0 + e^{i\pi/4}\mathrm{e}_1]$ is the so-called T-*state*, and $\mathrm{e}_0, \mathrm{e}_1$ are the computational basis vectors in $\mathbb{C}^2$ [28]. Since the Clifford+T gate set is universal for quantum computation, it is of particular interest to determine $\chi([T^{\otimes m}])$ and $\chi_\delta([T^{\otimes m}])$.

Despite the importance of the stabilizer rank, few techniques are known for bounding this quantity. In this work, we introduce techniques from number theory and algebraic geometry for bounding the stabilizer rank. In particular, we:

1. Refine a theorem of Moulton on subset-sum representations of exponentially increasing sequences, and use this refinement to prove lower bounds on exact and approximate stabilizer rank. In particular, we:

   - Exhibit an explicit sequence of m-qubit product states of exponentially increasing stabilizer rank and constant δ-approximate stabilizer rank for any fixed $\delta > 0$.
   - Prove that for any non-stabilizer qubit state $[v]$, it holds that $\chi([\psi^{\otimes m}]) = \Omega(m/\log_2 m)$, and in particular, $\chi([T^{\otimes m}]) \geq \frac{m+1}{4\log_2(m+1)}$. Our asymptotic scaling matches the best-known lower bound $\chi([T^{\otimes m}]) \geq m/100$ up to a log factor [101].
   - Prove that for any non-stabilizer qubit state $[v]$, there exists a constant $\delta > 0$ for which $\chi_\delta([v^{\otimes m}]) \geq \sqrt{m}/(2\log_2 m)$ for all $m \geq 2$. In particular, our asymptotic scaling for the T-state matches the best-known lower bound $\chi_\delta([T^{\otimes m}]) = \Omega(\sqrt{m}/\log_2 m)$ [101].

2. Explicitly construct the first non-trivial examples of quantum states with multiplicative stabilizer rank under the tensor product.
3. Use algebraic-geometric techniques to prove bounds on, and useful reductions for, the *generic stabilizer rank*, which upper bounds $\chi([T^{\otimes m}])$.

In the remainder of this section, we expand on points 1, 2, and 3, and identify directions for future work. In Sections 5.1, 5.2, and 5.3 we prove the results introduced in points 1, 2, and 3, respectively.

Lower bounds on stabilizer rank and approximate stabilizer rank

In Section 5.1, we refine a number-theoretic theorem of Moulton to prove lower bounds on stabilizer rank and approximate stabilizer rank [97]. For integers $q \geq 2$ and $r \geq 1$, and tuples of non-zero complex numbers

$$\alpha = (\alpha_1, \dots, \alpha_q) \in \mathbb{C}^q$$
$$\beta = (\beta_1, \dots, \beta_r) \in \mathbb{C}^r,$$

we say that β is a *subset-sum representation* of α if for all $i \in [q]$ there exists a subset $R_i \subseteq [r]$ for which $\sum_{j \in R_i} \beta_j = \alpha_i$. We refer to the integer r as the *length* of the subset-sum representation $\beta \in \mathbb{C}^r$. For an integer $2 \leq p \leq q$, we say that $\alpha \in \mathbb{C}^q$ has an *exponentially increasing subsequence* of length p if there exists $i_1, \dots, i_p \in [q]$ for which

$$|\alpha_{i_{j+1}}| \geq 2|\alpha_{i_j}| \quad \text{for all} \quad j \in [p].$$

Moulton's theorem states that any subset-sum representation of a q-tuple containing the subsequence $(1, 2, 4, \dots, 2^{p-1})$ has length at least $p/\log_2 p$ [97]. We refine this result to prove that the same bound holds for any q-tuple that contains an exponentially increasing subsequence of length p (Theorem 21).

Since stabilizer states have coordinates in $\{0, \pm 1, \pm i\}$ in the computational basis (see Section 2.2), any decomposition of a state $[v]$ into a superposition of r stabilizer states can be converted into a length-$4r$-subset-sum representation of the coordinates of $[v]$. It follows that if the coordinates of $[v]$ contain an exponentially increasing subsequence of length p, then $\chi([v]) \geq p/(4\log_2 p)$ (Theorem 22). In particular, since $[T]$ is Clifford-equivalent to $[H] := [\mathrm{e}_0 + \frac{1}{\sqrt{2}-1}\mathrm{e}_1]$, and the coordinates of $[H^{\otimes m}]$ contain an exponentially increasing subsequence of length $m+1$, we obtain $\chi([T^{\otimes m}]) \geq \frac{m+1}{4\log_2(m+1)}$. By a similar argument, we prove that $\chi([v^{\otimes m}]) = \Omega(m/\log_2 m)$ for any non-stabilizer qubit state $[v]$ (Theorem 22).

We further use Theorem 21, along with standard concentration inequalities for the binomial distribution, to prove that for any non-stabilizer qubit state $[v]$ there exists a constant $\delta > 0$ for which it holds that $\chi_\delta([v^{\otimes m}]) \geq \sqrt{m}/(2\log_2 m)$ for all $m \geq 2$ (Theorem 25).

Very recently, Peleg et al. obtained similar results: They proved that $\chi([T^{\otimes m}]) \geq m/100$, and that there exists $\delta > 0$ for which $\chi_\delta([T^{\otimes m}]) = \Omega(\sqrt{m}/\log_2 m)$ [101]. Asymptotically, our bounds match theirs up to a log factor, and we suggest that our proof technique is much simpler. While both of our bounds follow quite quickly from our refinement of Moulton's theorem mentioned above, the two bounds of Peleg et al. use two different approaches from the analysis of boolean functions and complexity theory: For their lower bound on $\chi([T^{\otimes m}])$, they analyze directional derivatives of quadratic polynomials, and for their lower bound on $\chi_\delta([T^{\otimes m}])$, they use Razborov-Smolensky low-degree polynomial approximations and correlation bounds against the majority function [105, 113, 114]. It is interesting that the vastly different approaches of ours and Peleg et al. yield such similar results.

As a further application of our refinement of Moulton's theorem, we explicitly construct a sequence of n-qubit product states $[v^{\otimes m}]$ for which it holds that $\chi([v^{\otimes m}]) \geq \frac{2^m}{4m}$ and $\chi_\delta([v^{\otimes m}]) = \mathcal{O}(1)$ for any $\delta > 0$, simply by writing down a product state with exponentially increasing coordinate amplitudes. Note that the stabilizer rank of this sequence is very close to 2^m, the maximum possible of any m-qubit state. This result lies contrast to the situation for other notions of rank, in which it is a difficult open problem to explicitly construct sequences of states of near-maximal rank. For example, the maximum *border rank*, a relevant notion of rank in classical complexity theory, of a quantum state in three local spaces of (affine) dimensions d, is $\lceil \frac{d^3}{3d-2} \rceil$ for all $d \neq 3$, whereas the largest border rank of any known explicit sequence of states in this space is only linear in d (see [83] for the largest known border rank and [80] for a general introduction into the topic).

States with multiplicative stabilizer rank under the tensor product

It is a standard fact that the stabilizer rank is *sub-multiplicative* under the tensor product, i.e. $\chi([v \otimes v]) \leq \chi([v])^2$ for any quantum state $[v]$ [102, Section 2.1.3]. In [102, Section 4.4] it was remarked that there are no known examples of quantum states $[v]$ of stabilizer rank greater than one for which equality holds. In Section 5.2, we explicitly construct two-qubit states $[v]$ for which $\chi([v]) = 2$ and $\chi([v \otimes v]) = 4$. This is the smallest possible example of such a state, since for any single-qubit state $[u]$ it holds that $\chi([u \otimes u]) \leq 3$.

Generic stabilizer rank

For any positive integer m, all but finitely many qubit states $[v]$ maximize $\chi([v^{\otimes m}])$ (Fact 27). This motivates us to define the m-th *generic stabilizer rank*, denoted χ_m, to be the maximum stabilizer rank of any state of the form $[v^{\otimes m}]$. In Section 5.3, we prove new bounds on χ_m, along with some useful reductions for studying this quantity. In Proposition 28 we modestly improve the best-known upper bound on χ_m, recently obtained by Qassim et al., from $\mathcal{O}((m+1)2^{m/2})$ to $\mathcal{O}(2^{m/2})$ [103, Theorem 2]. In Propositions 29 and 30 we prove two useful reductions for studying χ_m, namely, that there must exist a single set of χ_m stabilizer states that span the symmetric subspace, and that it suffices to work over the real numbers. In Proposition 31 we introduce a technique for upper bounding χ_m when upper bounds on $\chi([v^{\otimes m}])$ are known for sufficiently many (linear in n) qubit states $[v]$. In Proposition 32 we compute an upper bound on the (finite) number of qubit states $[v]$ for which $\chi([v^{\otimes m}]) < \chi_m$.

Directions for future work

We believe that our work opens the door for new approaches on questions related to the stabilizer rank. Here, we present two of the most promising ones.

Lower bounds on stabilizer rank via the T-count As we mentioned in the second paragraph of the introduction, for any quantum state $[v]$ of T-count m, it holds that $\chi([T^{\otimes m}]) \geq \chi([v])$ [28]. We prove in Theorem 22 that if the coordinates of $[v]$ contain an exponentially increasing subsequence of length p, then $\chi([v]) \geq p/(4\log_2 p)$. Combining these two statements, if we can find a sequence of states $[v]$ with small T-counts and long exponentially increasing subsequences of coordinates, then we can obtain large lower bounds on $\chi([T^{\otimes m}])$. For example, if we can find a sequence of states $[\psi_m]$ with T-count $\mathcal{O}(2^{cm})$ for some $0 < c < 1$, whose coordinates contain an exponentially increasing subsequence of length $\Omega(2^m)$, then we would obtain a super-linear lower bound $\chi(T^{\otimes m}) = \Omega\left(\frac{m^{1/c}}{\log_2 m}\right)$.

Bounds on generic stabilizer rank In Section 5.3 we prove several reductions for studying the generic stabilizer rank χ_m, but only manage to modestly improve the best-known bounds on this quantity. We ask whether stronger bounds can be obtained from our reductions.

1.1.4 Chapter 6: On decomposable correlation matrices

For a field $\mathbb{F}$ and positive integers n and s, let $M^{\mathbb{F}}_{n\times s}$ denote the set of $n \times s$ matrices over $\mathbb{F}$, let $M^{\mathbb{F}}_n = M^{\mathbb{F}}_{n\times n}$, and let $\mathrm{Cor}\,(\mathbb{C}^n) \subset M^{\mathbb{C}}_n$ denote the set of *correlation matrices*: positive semidefinite matrices with diagonal entries all equal to one. Equivalently, a matrix $P \in M^{\mathbb{C}}_n$ is a correlation matrix if there exists a set of unit vectors $\{v_1, \dots, v_n\} \subset \mathcal{S}(\mathbb{C}^n)$ such that $P(a, b) = \langle v_a, v_b\rangle$ for all $a, b \in [n]$, in which case we say P is *generated* by $\{v_1, \dots, v_n\}$. We say a correlation matrix $P \in \mathrm{Cor}\,(\mathbb{C}^n)$ is *d-decomposable* if it can be written as the Schur product $\odot$ (also known as the Hadamard product, entrywise product, or pointwise product) of correlation matrices of rank $\leq d$, i.e.

$$P = R_1 \odot \cdots \odot R_m \tag{1.1}$$

for some positive integer m and correlation matrices $R_1, \dots, R_m \in \mathrm{Cor}\,(\mathbb{C}^n)$ with $\mathrm{rank}(R_i) \leq d$ for all $i \in [m]$. Equivalently, a correlation matrix is d-decomposable if it is generated by a set of unit product vectors, i.e. elements of

$$\mathcal{S}\mathrm{Prod}\left(\mathbb{C}^d : \cdots : \mathbb{C}^d\right) = \mathcal{S}((\mathbb{C}^d)^{\otimes m}) \cap \mathrm{Prod}\left(\mathbb{C}^d : \cdots : \mathbb{C}^d\right).$$

We denote the set of d-decomposable correlation matrices as $\mathrm{Cor}_d\,(\mathbb{C}^n)$.

Motivation for studying decomposable correlation matrices

The general topic of correlation matrices has received considerable interest in quantum information [50, 43, 25, 64, 129, 131]. This interest is due in part to Tsirelson's theorem [123], which reveals an intimate connection between correlation matrices and certain nonlocal correlations that can arise from bipartite quantum systems. Another motivation is the identification of correlation matrices with Schur channels, examples of which include physically relevant channels such as generalized dephasing channels, cloning channels, and the Unruh channel [25].

The particular topic of Schur products of correlation matrices has also been studied in quantum information [34, 71, 96]. Let $\{v_1, \dots, v_n\} \subset \mathbb{C}^s$ and $\{u_1, \dots, u_n\} \subset \mathbb{C}^t$ be unit vectors with $s \leq t$, and let $P, Q \in \mathrm{Cor}\,(\mathbb{C}^n)$ be the correlation matrices they generate respectively. It is well known that $P = Q$ if and only if there exists an isometry U such that $Uv_a = u_a$ for all $a \in [n]$. More generally, in [34, 71] it is proven that there exists a quantum channel Φ such that $\Phi(v_a v_a^*) = u_a u_a^*$ for all $a \in [n]$ if and only if there exists a correlation matrix $R \in \mathrm{Cor}\,(\mathbb{C}^n)$ such that $P = R \odot Q$. Moreover, it is not hard to show (using the Stinespring representation of Φ) that the rank of R is equal to the Choi rank of

Φ, which reveals a close relationship between this topic and d-decomposability. We note that [34, 71] also give results on transformations between sets of mixed states.

In [96] a very similar topic is studied. A characterization is found of what we call *CS-decomposable* correlation matrices: those that can be written as a Schur product of correlation matrices, each of which are generated by a set of coherent states (states of light produced by an ideal laser [21]). Equivalently, CS-decomposable correlation matrices are those that are generated by *multi-mode* coherent states: unit vectors that are tensor products of coherent states.

We now briefly summarize how the study of CS-decomposability is motivated in [96], as it will also motivate our study of r-decomposability. A pressing need in quantum information is to adapt or reinvent existing quantum protocols to be more experimentally realizable. Quantum fingerprinting and appointment scheduling are two examples of tasks for which this need has been recently addressed. The original protocols for these tasks use high dimensional entangled states that are difficult to prepare in a lab [30, 31], but both protocols have been adapted to use tensor products of coherent states and/or qubits (which are easier to produce experimentally), and simple quantum operations, while attaining similar figures of merit [7, 93, 122]. The experimental ease of producing tensor products of coherent states or low-dimensional unit vectors leads us to ask what other protocols can be adapted to use such states. This motivates the study of CS-decomposable and d-decomposable correlation matrices. By the discussion of the previous paragraph, protocols that use a fixed set of unit vectors can be adapted to use any other set of unit vectors that generate the same correlation matrix, simply by applying the corresponding isometry. Thus, if a protocol requires a set of unit vectors that generate a CS-decomposable (respectively, d-decomposable) correlation matrix, then the protocol can be adapted into a protocol that uses tensor products of coherent states (resp., d-dimensional unit vectors), which might be easier to implement than the original protocol. In this way, the study of CS-decomposable and d-decomposable correlation matrices could potentially give rise to more experimentally implementable protocols.

We further motivate our study of d-decomposability by the following entanglement detection scenario. Say we are given many copies of unknown pure states $v_1v_1^*, \dots, v_nv_n^*$, on which we are allowed to perform any of the measurements

$$\{v_1v_1^*, \mathbb{1} - v_1v_1^*\}, \dots, \{v_nv_n^*, \mathbb{1} - v_nv_n^*\},$$

and we wish to detect that for any partitioning of the space into subsystems of dimension $\leq d$, at least one of the states must be entangled. This scenario is similar to our d-decomposability question, as the only meaningful information to be gained from performing the allowed measurements is precisely the inner products $\langle v_a v_a^*, v_b v_b^* \rangle$ for $a, b \in [n]$. In Proposition 43 we find cases of this scenario in which one can indeed detect entanglement.

As a final motivation, decomposable correlation matrices are intimately related to so-called *absolutely entangled sets*, which have become a topic of recent interest in quantum information theory [33, 130, 86]. A correlation matrix P is called *absolutely entangled* if there does not exist any positive integer m and correlation matrices $R_1, \ldots, R_m$ with $\operatorname{rank}(R_i) \le \operatorname{rank}(P)$ for all $i \in [m]$ and $\prod_{i\in[m]} \operatorname{rank}(R_i) = \operatorname{rank}(P)$ for which

$$P = R_1 \odot \cdots \odot R_m.$$

Equivalently, $\{v_1, \ldots, v_n\} \subseteq \mathcal{S}(\mathbb{C}^n)$ is absolutely entangled if for every collection of positive integers $d_1, \ldots, d_m < \operatorname{rank}(P)$ for which $d_1 \cdots d_m = \operatorname{rank}(P)$, and every isometry

$$V : \mathbb{C}^n \to \mathbb{C}^{d_1} \otimes \cdots \otimes \mathbb{C}^{d_m},$$

there exists $a \in [n]$ for which $Vv_a \notin \operatorname{Prod}\left(\mathbb{C}^{d_1} : \cdots : \mathbb{C}^{d_m}\right)$. Note that non-decomposability is a stronger notion of entanglement than absolute entanglement: If P is not $(\operatorname{rank}(P)-1)$-decomposable, then it is absolutely entangled.

Results on decomposable correlation matrices

In Chapter 6, we obtain the following results on decomposable correlation matrices:

1. It is well known that $\operatorname{Cor}(\mathbb{C}^n)$ is a compact and convex set. To our knowledge, it is not known whether $\operatorname{Cor}_d(\mathbb{C}^n)$ is closed, and we leave this question unanswered. We show that $\operatorname{Cor}_d(\mathbb{C}^n)$ is not convex when $r \ge 1$ and $n \ge 2d+1$.

2. It is clear that $\operatorname{Cor}_d(\mathbb{C}^n) = \operatorname{Cor}(\mathbb{C}^n)$ for all $n \le d$. We prove that $\operatorname{Cor}_{n-1}(\mathbb{C}^n) = \operatorname{Cor}(\mathbb{C}^n)$ for all $n \ge 3$, but $\operatorname{Cor}_d(\mathbb{C}^n) \subsetneq \operatorname{Cor}(\mathbb{C}^n)$ for all $n \ge 2d+1$.

3. The previous point leaves open the question of whether the containment $\operatorname{Cor}_d(\mathbb{C}^n) \subseteq \operatorname{Cor}(\mathbb{C}^n)$ is strict for $n \in \{d+2, \ldots, 2d\}$, and in particular whether $\operatorname{Cor}_2(\mathbb{C}^4) \subseteq \operatorname{Cor}(\mathbb{C}^4)$ is strict. We reduce the latter to a simpler question of whether every element of a certain subset of $\operatorname{Cor}(\mathbb{C}^4)$ can be written as the Schur product of just two rank-two correlation matrices, which could make the problem more tractable for analytical or numerical approaches.

1.2 Tensors in algebraic complexity theory

In Section 1.1 we motivated the tensor rank in quantum information theory as the size of the smallest GHZ state needed to obtain a given pure state by SLOCC. In this section, we motivate the tensor rank in algebraic complexity theory, as the computational cost of implementing an associated multilinear map. Motivated by these applications, in Chapter 7 we will prove new lower bounds on tensor rank. We defer a formal introduction to Chapter 7 until Section 1.3, because the main results of that chapter concern uniqueness of tensor rank decompositions, which have stronger motivations in algebraic statistics.

Let $\mathbb{F}$ be a field, let $\mathcal{V}_1, \mathcal{V}_2, \mathcal{V}_3$ be $\mathbb{F}$-vector spaces, and let $\mathcal{V} = \mathcal{V}_1 \otimes \mathcal{V}_2 \otimes \mathcal{V}_3$. Recall the natural isomorphism between $\mathcal{V}_1 \otimes \mathcal{V}_2$ and $\mathrm{L}(\mathcal{V}_1^*, \mathcal{V}_2)$, the vector space of linear maps from $\mathcal{V}_1^*$ to $\mathcal{V}_2$, which sends $v_1 \otimes v_2$ to the map

$$\begin{aligned} \mathcal{V}_1^* &\to \mathcal{V}_2 \\ f &\mapsto f(v_1)v_2, \end{aligned}$$

and extends linearly. Likewise, there is a natural isomorphism between $\mathcal{V}_1 \otimes \mathcal{V}_2 \otimes \mathcal{V}_3$ and the vector space of bilinear maps from $\mathcal{V}_1^* \times \mathcal{V}_2^*$ to $\mathcal{V}_3$, which sends $v_1 \otimes v_2 \otimes v_3$ to the bilinear map

$$\begin{aligned} \mathcal{V}_1^* \times \mathcal{V}_2^* &\to \mathcal{V}_3 \\ (f,g) &\mapsto f(v_1)g(v_2)v_3. \end{aligned}$$

More generally, if a tensor $v \in \mathcal{V}$ has a decomposition $v = \sum_{a \in [n]} x_{a,1} \otimes x_{a,2} \otimes x_{a,3}$, then the corresponding bilinear map (which we also denote by v) is given by

$$\begin{aligned} v : \mathcal{V}_1^* \times \mathcal{V}_2^* &\to \mathcal{V}_3 \\ (f,g) &\mapsto \sum_{a \in [n]} f(x_{a,1})g(x_{a,2})x_{a,3}. \end{aligned}$$

We are interested in the cost of implementing the bilinear map v in terms of $\mathrm{rank}(v)$, using this observation. In the model of *multiplicative complexity*, we regard linear maps applied to the indeterminates f and g (e.g. $f \mapsto f(x_{a,1})$ and $g \mapsto g(x_{a,2})$) as free, and only charge for multiplications between indeterminate quantities (e.g. $f(x_{a,1})g(x_{a,2})$, which is regarded as a single multiplication of cost 1). Under this model, the expression above demonstrates that v can be implemented with multiplicative complexity at most r. In the other direction, Strassen proved that the multiplicative complexity of v is at least $\frac{1}{2}r$, so the tensor rank of a tensor v describes the multiplicative complexity of the associated bilinear map, up to a factor of two [117, 118, 120] (see also [32]).

In more precise terms, for an $\mathbb{F}$-algebra A, a finite subset $F \subseteq A$ and an *input set* $I \subseteq A$, the *multiplicative complexity of F with respect to I* is the minimum number of multiplications $*$ needed in a so-called *straight-line program* to obtain F on some input in $I^{\times p}$, for some $p \in \mathbb{N}$ [32]. A *straight-line program* is an algorithmic procedure that can manipulate the input using standard operations to arbitrary precision. Formally, it is a sequence of instructions, each of which can apply either a single standard binary operation $+, -, *, /$ to two elements in memory and store it (when $/$ is defined), store a scalar $\alpha \in \mathbb{F}$, or multiply an element in memory by a scalar. (Initially, the memory consists only of the input, but each operation appends an element of A to memory.) What we have dubbed the multiplicative complexity of v is precisely the multiplicative complexity of the set of bilinear forms $F = \{v_1, \ldots, v_{\dim(\mathcal{V}_1)}\}$ in the polynomial ring over the coordinate functions on $\mathcal{V}_1 \times \mathcal{V}_2$, with input set given by the standard coordinate functions on $\mathcal{V}_1 \times \mathcal{V}_2$; where

$$\begin{aligned} v_b : \mathcal{V}_1^* \times \mathcal{V}_2^* &\to \mathbb{F} \\ (f, g) &\mapsto \sum_{a \in [r]} f(x_{a,1}) g(x_{a,2}) e^b(x_{a,3}), \end{aligned}$$

for each $b \in [\dim(\mathcal{V}_3)]$, and $\{e^1, \ldots, e^{\dim(\mathcal{V})}\}$ is a standard basis for $\mathcal{V}_3^*$. The multiplicative complexity does not depend on any of the choices of bases, since linear maps are free. The *total complexity* is defined similarly to the multiplicative complexity, but instead of counting only the number of multiplications $*$ needed in a straight-line program, it counts all allowed operations ($+, -, *, /$, storing a scalar, or multiplying an element in memory by a scalar). The multiplicative complexity is a useful tool for lower bounding the total complexity, especially given that it can be elegantly described up to a factor of two by the tensor rank.

Perhaps the most well-studied example of a bilinear map in this context is the so-called *matrix multiplication map*

$$\mathrm{MaMu}_d : M_d^{\mathbb{F}} \times M_d^{\mathbb{F}} \to M_d^{\mathbb{F}},$$

which takes as input two $d \times d$ matrices and outputs their product (recall that $M_d^{\mathbb{F}}$ is the set of $d \times d$ matrices over $\mathbb{F}$). Determining the total complexity of this map is relevant in computational linear algebra, since the total complexities of problems such as matrix inversion, computation of the determinant, LUP-decomposition, and computation of the characteristic polynomial, are dominated asymptotically by this complexity. For MaMu_d, in the asymptotic limit, the connection between tensor rank and computational complexity has been shown to be even stronger than that observed above: Letting

$$\omega(\mathbb{F}) = \inf\{\alpha \in \mathbb{R} | \text{The total complexity of } \mathrm{MaMu}_d \text{ is } \mathcal{O}(d^\alpha)\},$$

then

$$\omega(\mathbb{F}) = \inf\{\alpha \in \mathbb{R} | \operatorname{rank}(\mathrm{MaMu}_d) = \mathcal{O}(d^\alpha)\}.$$

In short, in the limit of large d, $\operatorname{rank}(\mathrm{MaMu}_d)$ quantifies the total complexity (not just the multiplicative complexity) of matrix multiplication [117, 118, 120] (see also [32]).

1.3 Tensors in algebraic statistics

Tensors are useful tools for representing data. In this context, a decomposition of a tensor into a small number of product tensors allows one to both compress and interpret the data. If a tensor has a unique tensor rank decomposition, this corresponds to a unique interpretation of the data. For this reason, having efficient methods of certifying uniqueness of tensor decompositions is useful in statistical learning.

Many applications of uniqueness in statistical learning arise from the following observation: Let X, Y, Z, and L be finite random variables with the property that X, Y, Z are independent conditioned on L, i.e. for each x, y, z, and l in the sets of possible outcomes of X, Y, Z, and L, respectively, it holds that

$$\Pr(X = x, Y = y, Z = z | L = l) = \Pr(X = x | L = l) \Pr(Y = y | L = l) \Pr(Z = z | L = l). \tag{1.2}$$

Typically, X, Y, and Z are observable (i.e. measurable) random variables, and L is a *latent* random variable: A random variable that cannot be measured, but which we are nevertheless interested in knowing about. The goal of the scenario is to determine the full joint distribution $\Pr(X, Y, Z, L)$ given access only to the assumption (1.2) and the probability vector $\Pr(X, Y, Z)$. A common technique for obtaining the full distribution is through uniqueness of tensor decompositions. Note that

$$\begin{aligned}\Pr(X, Y, Z) &= \sum_l \Pr(L = l) \Pr(X, Y, Z | L = l)\\ &= \sum_l \Pr(L = l) \Pr(X | L = l) \otimes \Pr(Y | L = l) \otimes \Pr(Z | L = l),\end{aligned}$$

where the first line is obvious and the second line follows from (1.2). So a choice of latent random variable L corresponds to a decomposition of the tensor $\Pr(X, Y, Z)$. If $\Pr(X, Y, Z)$ happens to have a unique tensor rank decomposition, then there is a unique choice of latent parameter L with the minimum number of outcomes for which X, Y, and Z are conditionally independent. A similar observation holds even for continuous random variables [4]. This observation is used frequently in various applications of statistical machine learning, by finding or manufacturing conditional independence in a problem.

With these applications in mind, in Chapter 7 we present new sufficient conditions for a tensor decomposition to be the unique tensor rank decomposition.

1.3.1 Chapter 7: A generalization of Kruskal's theorem on tensor decomposition

Let $\mathbb{F}$ be a field, let $\mathcal{V}_1, \dots, \mathcal{V}_m$ be $\mathbb{F}$-vector spaces, and let $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$. Recall that the *tensor rank* (or *rank*) of a tensor $v \in \mathcal{V}$, denoted by $\operatorname{rank}(v)$, is the minimum number n for which v is the sum of n product tensors (any coefficients appearing in a linear combination of product tensors can be absorbed into the product tensors by linearity). An expression of v as a sum of product tensors (not necessarily of minimum number) is known simply as a *decomposition* of v. A decomposition of v

$$v = \sum_{a \in [n]} x_a \tag{1.3}$$

into a sum of product tensors $\{x_a : a \in [n]\}$ is said to be the *unique tensor rank decomposition* of v if for any decomposition

$$v = \sum_{a \in [r]} y_a \tag{1.4}$$

of v into the sum of $r \leq n$ product tensors $\{y_a : a \in [r]\}$, it holds that $r = n$ and $\{x_a : a \in [n]\} = \{y_a : a \in [n]\}$ as multisets. The decomposition (1.3) is said to be *unique in the j-th subsystem* if for any other decomposition (1.4), it holds that $r = n$ and there exists a permutation $\sigma \in S_n$ such that $x_{a,j} \in \operatorname{span}\{y_{\sigma(a),j}\}$ for all $a \in [n]$. Kruskal's theorem gives sufficient conditions for a given decomposition to constitute a unique tensor rank decomposition [79]. We refer to results of this kind as *uniqueness criteria*.

Uniqueness criteria have found scientific applications in signal processing and spectroscopy, among others mentioned previously in this introduction [84, 81, 42, 111]. In these circles, subsystems are also referred to as *factors* and *loadings*, and the tensor rank decomposition is also referred to as the *canonical decomposition (CANDECOMP), parallel factor (PARAFAC) model, canonical polyadic (CP) decomposition,* and *topographic components model.* Uniqueness of a tensor decomposition is also referred to as *specific identifiability*, and uniqueness criteria as *identifiability criteria.*

Kruskal's theorem, and a generalization

For a finite set S, let $|S|$ be the size of S. The *Kruskal-rank* (or *k-rank*) of a multiset of vectors $\{u_1, \dots, u_n\}$, denoted by $\text{k-rank}(u_1, \dots, u_n)$, is the largest number k for which $\dim \operatorname{span}\{u_a : a \in S\} = k$ for every subset $S \subseteq [n]$ of size $|S| = k$. Similarly, we call $\dim \operatorname{span}\{u_a : a \in [n]\}$ the *standard rank* (or *rank*) of $\{u_1, \dots, u_n\}$. Kruskal's theorem states that if a collection of product tensors $\{x_{a,1} \otimes \cdots \otimes x_{a,m} : a \in [n]\}$ has large enough k-ranks $k_j = \text{k-rank}(x_{1,j}, \dots, x_{n,j})$, then their sum constitutes a unique tensor rank decomposition. This theorem was originally proven for $m = 3$ subsystems over $\mathbb{R}$ [79], was later extended to more than three subsystems by Sidiropoulos and Bro [110], and then extended to an arbitrary field by Rhodes [106].

Theorem 1 (Kruskal's theorem)**.** *Let $n \geq 2$ and $m \geq 3$ be integers, let $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ be a vector space over a field $\mathbb{F}$, and let*

$$\{x_{a,1} \otimes \cdots \otimes x_{a,m} : a \in [n]\} \subseteq \mathcal{V} \setminus \{0\}$$

be a multiset of product tensors. For each $a \in [n]$, let $x_a = x_{a,1} \otimes \cdots \otimes x_{a,m}$. For each $j \in [m]$, let

$$k_j = \text{k-rank}(x_{1,j}, \dots, x_{n,j}).$$

If $2n \leq \sum_{j=1}^m (k_j - 1) + 1$, then $\sum_{a \in [n]} x_a$ constitutes a unique tensor rank decomposition.

In [49] it is shown that the inequality appearing in Kruskal's theorem cannot be weakened: there exist cases in which $2n = \sum_{j=1}^m (k_j - 1) + 2$ and the decomposition is not unique. While Kruskal's theorem gives sufficient conditions for uniqueness, necessary conditions are obtained in [78, 119, 87]. In [40] it is shown that Kruskal's theorem is *effective* over $\mathbb{R}$ or $\mathbb{C}$ in the sense that it certifies uniqueness on a dense open subset of the smallest semialgebraic set containing the set of rank n tensors. A robust form of Kruskal's theorem is proven in [17].

Our main result in this chapter is a "splitting theorem," which is not itself a uniqueness criterion, but implies a criterion that generalizes Kruskal's theorem. In our splitting theorem, the k-rank condition in Kruskal's theorem is relaxed to a standard rank condition. In turn, the conclusion is also relaxed to a statement describing the linear dependence of the product tensors. Before stating our splitting theorem, we first introduce the generalization of Kruskal's theorem it implies.

Theorem 2 (Generalization of Kruskal's theorem)**.** *Let $n \geq 2$ and $m \geq 3$ be integers, let $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ be a vector space over a field $\mathbb{F}$, and let*

$$\{x_{a,1} \otimes \cdots \otimes x_{a,m} : a \in [n]\} \subseteq \mathcal{V} \setminus \{0\}$$

be a multiset of product tensors. For each $a \in [n]$, let $x_a = x_{a,1} \otimes \cdots \otimes x_{a,m}$. For each subset $S \subseteq [n]$ and index $j \in [m]$, let

$$d_j^S = \dim \operatorname{span}\{x_{a,j} : a \in S\}.$$

If $2|S| \leq \sum_{j=1}^m (d_j^S - 1) + 1$ for every subset $S \subseteq [n]$ with $2 \leq |S| \leq n$, then $\sum_{a \in [n]} x_a$ constitutes a unique tensor rank decomposition.

Note that the computational cost of checking the conditions of our Theorem 2 is essentially the same as that of checking the conditions of Kruskal's theorem. In both cases, the quantities d_j^S must be computed for all $j \in [m]$ and $S \subseteq [n]$ with $2 \leq |S| \leq n$. To verify Kruskal's conditions, one uses these quantities to compute the Kruskal ranks, and then checks the single inequality $2n \leq \sum_{j=1}^m (k_j - 1) + 1$. To verify the conditions of our generalization, one checks a separate inequality $2|S| \leq \sum_{j=1}^m (d_j^S - 1) + 1$ for every S.

To see that Theorem 2 contains Kruskal's theorem, assume the conditions of Kruskal's theorem hold and note that for any subset $S \subseteq [n]$, the multiset of product tensors $\{x_a : a \in S\}$ satisfies $d_j^S \geq \min\{k_j, |S|\}$. Using this fact, it is easy to verify that $2|S| \leq \sum_{j=1}^m (d_j^S - 1) + 1$ for every subset $S \subseteq [n]$ with $2 \leq |S| \leq n$.

In Section 7.7 we compare Theorem 2 to the uniqueness criteria of Domanov, De Lathauwer, and Sørensen (DLS), which are the only known extensions of Kruskal's theorem that we are aware of [51, 52, 53, 116, 115]. All of these extensions rely on Kruskal's original permutation lemma, and as a result, still require the k-ranks to be above a certain threshold. Our generalization uses a completely new proof technique, can certify uniqueness below this threshold, and contains many of these extensions. The cited results of DLS contain many similar but incomparable criteria, which can be difficult to keep track of. For clarity and future reference, in Theorem 75 we synthesize these criteria into a single statement. Using insight gained from this synthesization and our generalization of Kruskal's theorem, we propose a conjectural uniqueness criterion that would contain and unify every uniqueness criteria of DLS into a single, elegant statement.

For $m \geq 4$, Kruskal's theorem can be "reshaped" by regarding multiple subsystems as a single subsystem. In Section 7.2 we present an analogous reshaping of Theorem 2, which has many more degrees of freedom to choose from than the reshaped Kruskal's theorem.

A splitting theorem for product tensors

We now state our splitting theorem, which we use in Section 7.2 to prove our generalization of Kruskal's theorem, and in Sections 7.4, 7.5, and 7.6 to obtain further results on tensor decompositions. We first require a definition.

Definition 3. Let $n \geq 2$ be an integer, and let $\mathcal{V}$ be a vector space over a field $\mathbb{F}$. We say that a multiset of non-zero vectors $\{v_1, \ldots, v_n\} \subseteq \mathcal{V} \setminus \{0\}$ *splits*, or is *disconnected*, if there exists a subset $S \subseteq \{v_1, \ldots, v_n\}$ with $1 \leq |S| \leq n-1$ for which

$$\operatorname{span}\{v_1, \ldots, v_n\} = \operatorname{span}(S) \oplus \operatorname{span}(S^c),$$

where $S^c := \{v_1, \ldots, v_n\} \setminus S$. In this case, we say that S *separates* $\{v_1, \ldots, v_n\}$. If $\{v_1, \ldots, v_n\}$ does not split, then we say it is *connected*.

Note that $\{v_1, \ldots, v_n\}$ splits if and only if it is disconnected as a matroid [99]. We now state our main result.

Theorem 4 (Splitting theorem). *Let $n \geq 2$ and $m \geq 2$ be integers, let $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ be a vector space over a field $\mathbb{F}$, let*

$$E = \{x_{a,1} \otimes \cdots \otimes x_{a,m} : a \in [n]\} \subseteq \mathcal{V} \setminus \{0\}$$

be a multiset of product tensors, and for each $j \in [m]$, let

$$d_j = \dim \operatorname{span}\{x_{a,j} : a \in [n]\}.$$

If $\dim \operatorname{span}(E) \leq \sum_{j=1}^{m}(d_j - 1)$, then E splits.

In Section 7.3 we use Derksen's result [49] to prove that the inequality appearing in Theorem 4 cannot be weakened.

We now give a rough sketch of how our splitting theorem implies Theorem 2, which we formalize in Section 7.2. First, a direct consequence of Theorem 4 is that E splits whenever $n \leq \sum_{j=1}^{m}(d_j - 1) + 1$ (see Corollary 49). To prove Theorem 2, let $\{x_a : a \in [n]\}$ be a multiset of product tensors satisfying the assumptions of Theorem 2, and let $\{y_a : a \in [r]\}$ be a multiset of $r \leq n$ product tensors for which $\sum_{a\in[n]} x_a = \sum_{a\in[r]} y_a$. Consider the multiset of $[n+r]$ product tensors

$$E = \{x_a : a \in [n]\} \cup \{-y_a : a \in [r]\}.$$

Since $2n \leq \sum_{j=1}^{m}(d_j^{[n]} - 1) + 1$, E splits. Since $\Sigma(E) = 0$, it follows that $\Sigma(S) = \Sigma(S^c) = 0$ for any separator S of E. Now, continue applying the splitting theorem to S and S^c, until every multiset has size 2, and contains one element each of $\{x_a : a \in [n]\}$ and $\{-y_a : a \in [r]\}$.

Further applications of the splitting theorem to tensor decompositions

In Sections 7.4, 7.5, and 7.6 we use the splitting theorem to prove further uniqueness results and sharp lower bounds on tensor rank. In Section 7.4 we prove a general statement that interpolates between our generalization of Kruskal's theorem and a natural offshoot of our splitting theorem (mentioned above), obtaining uniqueness results for weaker notions of uniqueness. In Section 7.5 we prove sharp lower bounds on tensor rank and *Waring rank*, a notion of rank for symmetric tensors. In Sections 7.4 and 7.6 we obtain uniqueness results for *non-rank* decompositions, a novel concept introduced in this work. We close this introduction by reviewing these results in more detail.

It is known that if a multiset of product tensors $\{x_a : a \in [n]\}$ satisfies

$$n + r \leq \sum_{j=1}^{m}(k_j - 1) + 1 \tag{1.5}$$

for $r = 0$, then it is linearly independent, and if it satisfies (1.5) for $r = 1$, then the only product tensors in $\text{span}\{x_a : a \in [n]\}$ are scalar multiples of $x_1, \dots, x_n$ [63]. When $r = n$, it holds that $\sum_{a\in[n]} x_a$ constitutes a unique tensor rank decomposition, by Kruskal's theorem. It is natural to ask what happens for $r \in \{0, 1, \dots, n\}$. In Section 7.4.1 we use our splitting theorem to prove that when the inequality (1.5) holds, the only rank $\leq r$ tensors in $\text{span}\{x_a : a \in [n]\}$ are those that can be written (uniquely) as a linear combination of $\leq r$ elements of $\{x_a : a \in [n]\}$, which interpolates between Kruskal's theorem for $r = n$, and the results of [63] for $r \in \{0, 1\}$. We generalize our interpolating statement in a similar manner to our generalization of Kruskal's theorem (Theorem 54). We also interpolate to weaker notions of uniqueness, which are explained further at the end of this introduction. We remark that the $m = 2, r = 0$ case of a result in this section was proven by Pierpaola Santarsiero in unpublished work, using a different proof technique.

The interpolating statement described in the previous paragraph immediately implies the following lower bound on tensor rank:

$$\operatorname{rank}\left[\sum_{a\in[n]} x_a\right] \geq \min\left\{n, \sum_{j=1}^{m}(k_j - 1) + 2 - n\right\}.$$

In Section 7.5 we use our splitting theorem to improve this bound. Namely, provided that the k-ranks are sufficiently balanced, we prove that two of the k-ranks k_i, k_j appearing in this bound can be replaced by standard ranks d_i, d_j, improving this bound when the ranks and k-ranks are not equal. Our improved bound specializes to Sylvester's matrix rank inequality when $m = 2$ [70]. In Section 7.5.1 we prove that our improved bound is sharp in a wide parameter regime.

In Section 7.6 we use our splitting theorem to prove uniqueness results for *non-Waring rank* decompositions of symmetric tensors. (Our terminology for symmetric tensor decompositions is analogous to that of general tensor decompositions, and we refer the reader to Chapter 2 for a formal introduction.) In particular, we prove a condition on a symmetric decomposition $v = \sum_{a\in[n]} \alpha_a v_a^{\otimes m}$ for which any other symmetric decomposition must contain at least $r_{\min}$ terms, where $r_{\min}$ depends on the rank and k-rank of $\{v_a : a \in [n]\}$. For $r_{\min} \leq n$, this gives a Waring rank lower bound that is contained in our lower bound described in the previous paragraph. For $r_{\min} = n + 1$, this gives a uniqueness result for symmetric tensors that is contained in Theorem 2, but is stronger than Kruskal's theorem in a wide parameter regime. Our main contribution in this section is the case $r_{\min} > n + 1$, which produces an even stronger statement than uniqueness: There are no symmetric decompositions of v into a linear combination of fewer than $r_{\min}$ terms, aside from $v = \sum_{a\in[n]} \alpha_a v_a^{\otimes m}$ (up to trivialities). This is an example of what we call a uniqueness result for *non-rank* decompositions of a tensor.

In Section 7.4.2 we prove further uniqueness results for non-rank decompositions of (possibly non-symmetric) tensors. In particular, we give conditions on a multiset of product tensors $\{x_a : a \in [n]\}$ for which whenever $\sum_{a\in[n]} x_a = \sum_{a\in[r]} y_a$ for some $r > n$ and multiset of product tensors $\{y_a : a \in [r]\}$, there exist subsets $R \subseteq [n]$, $Q \subseteq [r]$ such that $|Q| = |R| = q$ for some fixed positive integer q, and $\{x_a : a \in Q\} = \{y_a : a \in R\}$. In contrast to our non-rank uniqueness results of Section 7.6, which apply only to symmetric decompositions of symmetric tensors, the results of this subsection apply to arbitrary tensor decompositions.

In Section 7.6.2 we identify two potential applications of our uniqueness results for non-rank decompositions: First, they allow us to define a natural hierarchy of tensors in terms of "how unique" their decompositions are. Second, any uniqueness result for non-rank decompositions can be turned around to produce a result in the more standard setting, in which one starts with a decomposition into n terms, and wants to control the possible decompositions into fewer than n terms.

From the proof sketch of our generalization of Kruskal's theorem that appears at the end of the previous subsection, it is easy to surmise that if $\sum_{a\in[n]} x_a = \sum_{a\in[r]} y_a$, and $2n \leq \sum_{j=1}^{m}(d_j^{[n]} - 1) + 1$, then there exist non-trvial subsets $Q \subseteq [n]$ and $R \subseteq [r]$ for which $\sum_{a\in Q} x_a = \sum_{a\in R} y_a$. This conclusion can be viewed as an extremely weakened form of uniqueness, and it is natural to ask what statements can be made for notions of uniqueness in between the standard one and this weakened one. We answer this question in Sections 7.4.1 and 7.4.2.

We say that a set of non-zero vectors forms a *circuit* if it is linearly dependent and any proper subset is linearly independent. As a special case of our splitting theorem, in Corollary 60 we obtain an upper bound on the number of subsystems $j \in [m]$ for which a circuit of product tensors can have $d_j \geq 2$. This improves recent bounds obtained in [10, 9], and is sharp.

Chapter 2

Preliminaries

In this chapter, we review several necessary preliminaries for this thesis. In this preamble, we review some basic objects in affine and projective space, linear operators, and inner products, along with some basic set notation. We then introduce the Segre, Veronese, and Grassmannian varieties, which over $\mathbb{C}$ correspond to the unentangled quantum states of distinguishable, bosonic, and fermionic systems, respectively. We introduce the notion of secants to a variety, with particular emphasis on the Segre, Veronese, and Grassmannian varieties, and review the notions of X-rank and uniqueness of tensor decompositions introduced briefly in Chapter 1. Finally, we introduce the set of stabilizer states and its secants.

Let $[m] = \{1, \dots, m\}$ when m is a positive integer. For a non-zero vector space $\mathcal{V}$ over a field $\mathbb{F}$, let $\mathbb{P}(\mathcal{V})$ (or more briefly, $\mathbb{P}\mathcal{V}$) be the set of 1-dimensional linear subspaces of $\mathcal{V}$, and let $\mathbb{P}^D = \mathbb{P}(\mathbb{F}^{D+1})$. We abuse the $[\cdot]$ notation and also write $[v] \in \mathbb{P}\mathcal{V}$ for the span of a non-zero vector $v \in \mathcal{V}$. For a subset $X \subseteq \mathbb{P}\mathcal{V}$, let $\hat{X} \subseteq \mathcal{V}$ be the affine cone over X. Explicitly, $\hat{X} = \{x \in \mathcal{V} : [x] \in X\} \cup \{0\}$. In the other direction, for a subset $Z \subseteq \mathcal{V}$ that forms a cone (i.e. for all $z \in Z$ and $\alpha \in \mathbb{F}$ it holds that $\alpha z \in Z$), we let $\check{Z} = \{[z] \in \mathbb{P}\mathcal{V} : z \in Z \setminus \{0\}\}$ be the projectivization of Z.

For vector spaces $\mathcal{V}$ and $\mathcal{W}$ over a field $\mathbb{F}$, let $\mathrm{L}(\mathcal{V}, \mathcal{W})$ be the vector space of linear operators from $\mathcal{V}$ to $\mathcal{W}$, and let $\mathrm{L}(\mathcal{V}) = \mathrm{L}(\mathcal{V}, \mathcal{V})$. For a non-negative integer a, let e_a denote the standard basis vector with 1 in the a-th position and zeros elsewhere. When considering stabilizer states and stabilizer rank (e.g. in Chapter 5) we index the standard basis starting from $a = 0$, to more easily identify the index set with $\mathbb{F}_2$. In all other cases, we index the standard basis starting from $a = 1$.

For a complex vector space $\mathcal{V}$, let $\langle \cdot, \cdot \rangle : \mathcal{V} \times \mathcal{V} \to \mathbb{C}$ be the standard Euclidean inner product given by $\langle u, v\rangle = u^* v$, where $(\cdot)^*$ denotes the conjugate-transpose. Let $\|\cdot\| = \sqrt{\langle\cdot,\cdot\rangle}$ be the Euclidean norm, and define the set of *unit vectors* $\mathcal{S}(\mathcal{V})$ as the set of vectors $v \in \mathcal{V}$ that satisfy $\|v\| = 1$. For complex vector spaces $\mathcal{V}$ and $\mathcal{W}$, let $\mathrm{U}(\mathcal{V}, \mathcal{W}) \subset \mathrm{L}(\mathcal{V}, \mathcal{W})$ be the set of isometries from $\mathcal{V}$ to $\mathcal{W}$, i.e. the set of linear maps that preserve the Euclidean norm.

We write $S \cup T$ to denote the union of two sets S and T. If S and T happen to be disjoint, we often write $S \sqcup T$ instead to remind the reader of this fact. For a positive integer t, we say that a collection of subsets $S_1, \dots, S_t \subseteq T$ *partitions* T if $S_p \cap S_q = \{\}$ for all $p \neq q \in [t]$, and $S_1 \sqcup \cdots \sqcup S_t = T$.

2.1 The Segre, Veronese, and Grassmannian varieties (and their secants)

We will be particularly interested in the algebraic varieties known as the Segre, Veronese, and Grassmannian varieties, since over the complex numbers these correspond to the sets of unentangled states in a space of distinguishable, bosonic, and fermionic particles, respectively [60]. In this section, we briefly describe these varieties, as well as their r-th secants. We then review the notions of X-rank and uniqueness of tensor decompositions introduced briefly in Chapter 1. We refer the reader to [39, 81, 65, 14] for more in-depth treatments of these objects.

Let $\mathbb{F}$ be a field, let $\mathcal{V}_1, \dots, \mathcal{V}_m$, and $\mathcal{W}$ be non-zero $\mathbb{F}$-vector spaces, and let $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$. The Segre variety

$$Y = \mathrm{Seg}(\mathbb{P}\mathcal{V}_1 \times \cdots \times \mathbb{P}\mathcal{V}_m) \subseteq \mathbb{P}\mathcal{V} \tag{2.1}$$

is the image of the Segre embedding, and is equal to the set of projective tensors of the form $[x_1 \otimes \cdots \otimes x_m]$, where $x_j \in \mathcal{V}_j$ for each $j \in [m]$. If $\mathbb{F} = \mathbb{C}$, then the Segre variety corresponds to the unentangled (or, *product*) states in the space $\mathbb{P}(\bigotimes_{j=1}^m \mathcal{V}_j)$ of distinguishable particles.

For each permutation $\sigma \in S_m$, let $P_\sigma \in \mathrm{L}(\bigotimes^m \mathcal{W})$ be the linear map defined on product vectors as

$$P_\sigma(x_1 \otimes \cdots \otimes x_m) = x_{\sigma^{-1}(1)} \otimes \cdots \otimes x_{\sigma^{-1}(m)},$$

and extended linearly. Let $\mathbb{P}(S^m\mathcal{W}) \subseteq \mathbb{P}(\bigotimes^m \mathcal{W})$ denote the symmetric subspace, i.e, the set of projective tensors $[x] \in \mathbb{P}(\bigotimes^m \mathcal{W})$ such that $P_\sigma x = x$ for all $\sigma \in S_m$. If $\mathrm{Char}(\mathbb{F}) > m$

or $\mathrm{Char}(\mathbb{F}) = 0$, then $\mathbb{P}(S^m\mathcal{W})$ is spanned by projective tensors of the form $[x_1 \vee \cdots \vee x_m]$, where

$$[x_1 \vee \cdots \vee x_m] = \left[\sum_{\sigma \in S_m} x_{\sigma(1)} \otimes \cdots \otimes x_{\sigma(m)}\right].$$

If $\mathbb{F} = \mathbb{C}$, then the symmetric subspace represents a bosonic space of indistinguishable particles. (This space can also be thought of as the set of homogeneous polynomials of degree m in d variables.) The Veronese variety

$$\nu_m(\mathbb{P}\mathcal{W}) \subseteq \mathbb{P}(S^m\mathcal{W})$$

is the image of the m-th Veronese embedding, and is equal to the set of projective tensors in $\mathbb{P}(S^m\mathcal{W})$ of the form $[x^{\vee m}]$. If $\mathbb{F} = \mathbb{C}$, then the Veronese variety corresponds to the set of unentangled states in the bosonic space $\mathbb{P}(S^m\mathcal{W})$.

Let $\mathbb{P}(\bigwedge^m \mathcal{W}) \subseteq \mathbb{P}(\bigotimes^m \mathcal{W})$ denote the antisymmetric subspace, i.e. the set of projective tensors $[x] \in \mathbb{P}(\bigotimes^m \mathcal{W})$ such that $P_\sigma x = (-1)^{\mathrm{sgn}(\sigma)} x$ for all $\sigma \in S_m$. This space is spanned by the set of projective tensors of the form $[x_1 \wedge \cdots \wedge x_m]$, where

$$[x_1 \wedge \cdots \wedge x_m] = \left[\sum_{\sigma \in S_m} (-1)^{\mathrm{sgn}(\sigma)} x_{\sigma(1)} \otimes \cdots \otimes x_{\sigma(m)}\right].$$

If $\mathbb{F} = \mathbb{C}$, then $\mathbb{P}(\bigwedge^m \mathcal{W})$ represents a fermionic space of indistinguishable particles. The Grassmannian variety

$$\mathrm{Gr}(m-1, \mathbb{P}^{d-1}) \subseteq \mathbb{P}(\textstyle\bigwedge^m \mathcal{W})$$

is the set of projective tensors in $\mathbb{P}(\bigwedge^m \mathcal{W})$ of the form $[x_1 \wedge \cdots \wedge x_m]$. If $\mathbb{F} = \mathbb{C}$, then the Grassmannian variety corresponds to the set of unentangled states in the fermionic space $\mathbb{P}(\bigwedge^m \mathcal{W})$. The Grassmannian can also be viewed as the variety of projective $(m-1)$-planes in $\mathbb{P}\mathcal{W}$. We expound on this perspective in Section 3.1.1, as we will make frequent use of it.

2.1.1 Secants to the Segre, Veronese, and Grassmannian varieties

In this subsection, we define secant varieties, with particular emphasis on secants to the Segre, Veronese, and Grassmannian varieties. We then formally introduce the notions of tensor rank, border rank, Waring rank, and border Waring rank, as well as uniqueness of tensor rank decompositions and uniqueness of symmetric tensor rank decompositions.

For a vector space $\mathcal{V}$ and a non-degenerate projective variety $X \subseteq \mathbb{P}\mathcal{V}$, let

$$\Sigma_r^\circ(X) = \bigcup_{[x_1],\dots,[x_r]\in X} \operatorname{span}\{[x_1],\dots,[x_r]\} \subseteq \mathbb{P}\mathcal{V},$$

and let

$$\Sigma_r(X) = \overline{\Sigma_r^\circ} \subseteq \mathbb{P}\mathcal{V}$$

be the *r-th secant variety* to X, where the closure is taken with respect to the Zariski topology (over $\mathbb{C}$, the Zariski and Euclidean closures agree). Let $\hat{\Sigma}_r^\circ(X)$ and $\hat{\Sigma}_r(X)$ denote the affine cones over $\Sigma_r^\circ(X)$ and $\Sigma_r(X)$, respectively. For a point $[v] \in \mathbb{P}\mathcal{V}$, the *$X$-rank* of $[v]$, denoted $\operatorname{rank}_X([v])$, is the smallest integer r for which $[v]$ is contained in $\Sigma_r^\circ(X)$, and the *border X-rank* of $[v]$ is the smallest integer r for which $[v]$ is contained in $\Sigma_r(X)$. The X-rank and border X-rank of a vector $v \in \mathbb{F}^{D+1}$ are defined to be the X-rank and border X-rank of $[v]$.

The *tensor rank* (or simply, *rank*) of a point $[v] \in \mathbb{P}(\bigotimes_{j=1}^m \mathcal{V}_j)$, denoted $\operatorname{rank}([v])$, is defined to be the Y-rank of $[v]$, were Y denotes the Segre variety as in (2.1), and the *border rank* of $[v]$, denoted $\underline{\operatorname{rank}}([v])$, is the border Y-rank of $[v]$. When $\operatorname{Char}(\mathbb{F}) > m$ or $\operatorname{Char}(\mathbb{F}) = 0$, the *Waring rank* of a point $[v] \in \mathbb{P}(S^m\mathcal{W})$, denoted $\operatorname{WaringRank}([v])$, is defined to be the $\nu_m(\mathbb{P}\mathcal{W})$-rank of $[v]$, and the *border Waring rank* of $[v]$, denoted $\underline{\operatorname{WaringRank}}([v])$ is the border $\nu_m(\mathbb{P}\mathcal{W})$-rank of $[v]$.

Recall that we say a tensor $z \in \mathcal{V}$ is a *product tensor* if $z \neq 0$ and z is contained in the affine cone over the Segre variety, i.e. $z = z_1 \otimes \cdots \otimes z_m$ for some $z_1 \in \mathcal{V}_1, \dots, z_m \in \mathcal{V}_m$. Let $\operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ denote the set of (non-zero) product tensors in $\mathcal{V}$. We use symbols like a, b to index tensors, and symbols like i, j to index subsystems.

For a product tensor $z \in \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$, the vectors $z_j \in \mathcal{V}_j$ for which $z = z_1 \otimes \cdots \otimes z_m$ are uniquely defined up to scalar multiples $\alpha_1 z_1, \dots, \alpha_m z_m$ such that $\alpha_1 \cdots \alpha_m = 1$. For positive integers n and m, we frequently define multisets of product tensors

$$\{x_a : a \in [n]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$$

without explicitly defining corresponding vectors $\{x_{a,j}\}$ such that

$$x_a = x_{a,1} \otimes \cdots \otimes x_{a,m}$$

for all $a \in [n]$. In this case, we implicitly fix some such vectors, and refer to them without further introduction. We use the notation

$$x_{a,\hat{j}} = x_{a,1} \otimes \cdots \otimes x_{a,j-1} \otimes x_{a,j+1} \otimes \cdots \otimes x_{a,m},$$
$$\mathcal{V}_{\hat{j}} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_{j-1} \otimes \mathcal{V}_{j+1} \otimes \cdots \otimes \mathcal{V}_m,$$

so $x_{a,\hat{j}} \in \mathcal{V}_{\hat{j}}$. Recall that $\mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ is naturally isomorphic to $\mathrm{L}(\mathcal{V}_j^*, \mathcal{V}_{\hat{j}})$ for any $j \in [m]$, where $\mathcal{V}_j^*$ is the dual vector space to $\mathcal{V}_j$. The rank of a tensor in $\mathcal{V}_1 \otimes \mathcal{V}_2$ is equal to the rank of the corresponding linear operator in $\mathrm{L}(\mathcal{V}_1^*, \mathcal{V}_2)$. We denote the rank of a tensor $v \in \mathcal{V}$, viewed as an element of $\mathrm{L}(\mathcal{V}_j^*, \mathcal{V}_{\hat{j}})$, by $\mathrm{rank}_j(v)$. The *flattening rank* of v is defined as $\max\{\mathrm{rank}_1(v), \dots, \mathrm{rank}_m(v)\}$. Note that the tensor rank of v is lower bounded by the flattening rank of v.

With this notation, the tensor rank of an element $v \in \mathcal{V}$ is the smallest integer r for which there exist product tensors $\{x_a : a \in [r]\} \subseteq \mathrm{Prod}\,(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ such that $v = \sum_{a \in [r]} x_a$. An expression

$$v = \sum_{a \in [n]} x_a \tag{2.2}$$

of a tensor v as a sum of product tensors $\{x_a : a \in [n]\} \subseteq \mathrm{Prod}\,(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ (not necessarily of minimum number) is known simply as a *decomposition* of v. A decomposition (2.2) is said to be the *unique tensor rank decomposition* of v if for any decomposition

$$v = \sum_{a \in [r]} y_a \tag{2.3}$$

of v into the sum of $r \leq n$ product tensors $\{y_a : a \in [r]\} \subseteq \mathrm{Prod}\,(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$, it holds that $r = n$ and $\{x_a : a \in [n]\} = \{y_a : a \in [n]\}$ as multisets. The decomposition (2.2) is said to be *unique in the j-th subsystem* if for any other decomposition (2.3), it holds that $r = n$ and there exists a permutation $\sigma \in S_n$ such that $x_{a,j} \in \mathrm{span}\{y_{\sigma(a),j}\}$ for all $a \in [n]$. Kruskal's theorem gives sufficient conditions for a given decomposition to constitute a unique tensor rank decomposition [79]. We refer to results of this kind as *uniqueness criteria.*

We say that a tensor $v \in \mathcal{W}^{\otimes m}$ is *symmetric* if $v \in S^m(\mathcal{W})$. A *symmetric product tensor* is a non-zero element of $\hat{\nu}_m(\mathcal{W})$. Note that v is a symmetric product tensor if and only if $v = \alpha w^{\otimes m}$ for some $\alpha \in \mathbb{F}^\times$ and $w \in \mathcal{W} \setminus \{0\}$. Note that the Waring rank of a symmetric tensor $v \in S^m(\mathcal{W})$ is the smallest integer r for which v can be written as a sum of r symmetric product tensors.

An expression

$$v = \sum_{a \in [n]} \alpha_a v_a^{\otimes m}$$

of a symmetric tensor v as a sum of symmetric product tensors $\{\alpha_a v_a^{\otimes m} : a \in [n]\} \subseteq \hat{\nu}_m(\mathcal{W})$ (not necessarily of minimum number) is known simply as a *symmetric decomposition* of v.

A symmetric decomposition of v

$$v = \sum_{a \in [n]} \alpha_a v_a^{\otimes m} \tag{2.4}$$

is said to be the *unique Waring rank decomposition* of v if for any non-negative integer $r \leq n$, multiset of non-zero vectors $\{u_a : a \in [r]\} \subseteq \mathcal{W} \setminus \{0\}$, and non-zero scalars $\{\beta_a : a \in [r]\} \subseteq \mathbb{F}^\times$ for which

$$v = \sum_{a \in [r]} \beta_a u_a^{\otimes m}, \tag{2.5}$$

it holds that $r = n$ and

$$\{\alpha_a v_a^{\otimes m} : a \in [n]\} = \{\beta_a u_a^{\otimes m} : a \in [n]\}.$$

More generally, for a positive integer $\tilde{n} \geq n$, we say that the symmetric decomposition (2.4) is the *unique symmetric decomposition of v into at most $\tilde{n}$ terms* if for any $r \leq \tilde{n}$ and symmetric decomposition (2.5), either

$$\text{k-rank}(u_a : a \in [r]) = 1,$$

or $r = n$ and

$$\{\alpha_a v_a^{\otimes m} : a \in [n]\} = \{\beta_a u_a^{\otimes m} : a \in [n]\}.$$

Note that (2.4) is the unique Waring rank decomposition of v if and only if it is the unique symmetric decomposition of v into at most n terms. We refer to results that certify uniqueness of a symmetric decomposition into at most $\tilde{n} > n$ terms as *uniqueness results for non-Waring rank decompositions*. We present such results in Section 7.6.

Our assumption that $\text{Char}(\mathbb{F}) > m$ or $\text{Char}(\mathbb{F}) = 0$ in the symmetric case ensures that the symmetric subspace $S^m(\mathcal{W})$ is isomorphic to the space of homogeneous polynomials over $\mathbb{F}$ of degree m in d variables, and that every symmetric tensor has finite Waring rank (see e.g. [73, Appendix A] and [81, Section 2.6.4]).

2.1.2 Dimensions of secants to the Segre, Veronese, and Grassmannian varieties over $\mathbb{C}$

In this subsection, we comment on the dimensions of the secant varieties to the Segre, Veronese, and Grassmannian varieties over the field of complex numbers. It is a standard result that

$$\dim(\Sigma_r(X)) \leq \min\{D, r\dim(X) + r - 1\}.$$

If equality holds in this expression, then $\Sigma_r(X)$ is said to have the *expected dimension*, and otherwise it is said to be *defective*.

For the Segre variety $Y \subseteq \mathbb{P}(\bigotimes_{j=1}^m \mathbb{C}^{d_j})$, we have $\dim(Y) = \sum_{j=1}^m (d_j - 1)$ and $\dim(\mathbb{P}(\bigotimes_{j=1}^m \mathbb{C}^{d_j})) = d_1 \cdots d_m - 1$, so

$$\dim(\Sigma_r(Y)) \leq \min\left\{d_1 \cdots d_m - 1, r\sum_{j=1}^m (d_j - 1) + r - 1\right\}. \tag{2.6}$$

A conjecturally complete set of defective $\Sigma_r(Y)$ have been proposed in [2], which are nicely summarized in [14, Conjecture 6]. For example,

$$\dim(\Sigma_r(\mathrm{Seg}(\mathbb{P}^{d_1-1} \times \mathbb{P}^{d_2-1}))) = d_1 d_2 - (d_1 - \min\{d_1, r\})(d_2 - \min\{d_2, r\}) - 1, \tag{2.7}$$

so $\Sigma_r(\mathrm{Seg}(\mathbb{P}^{d_1-1} \times \mathbb{P}^{d_2-1}))$ is defective in many cases. Under the identification $\mathbb{C}^{d_1} \otimes \mathbb{C}^{d_2} \cong \mathrm{L}((\mathbb{C}^{d_1})^*, \mathbb{C}^{d_2})$, the variety $\Sigma_r(\mathrm{Seg}(\mathbb{P}^{d_1-1} \times \mathbb{P}^{d_2-1}))$ corresponds to the set of (projective) $d_2 \times d_1$ matrices of rank at most r.

For the Veronese variety $\nu_m(\mathbb{P}^{d-1}) \subseteq \mathbb{P}(S^m(\mathbb{C}^d))$, we have $\dim(\nu_m(\mathbb{P}^{d-1})) = d - 1$ and $\dim(\mathbb{P}(S^m(\mathbb{C}^d))) = \binom{d+m-1}{m-1} - 1$, so

$$\dim(\Sigma_r(\nu_m(\mathbb{P}^{d-1}))) \leq \min\left\{\binom{d-1+m}{m} - 1, rd - 1\right\}.$$

A complete set of defective Veronese secants are known [3]; see also Theorem 2 in [14]. In particular, $\Sigma_r(\nu_2(\mathbb{P}^{d-1}))$ is defective whenever $2 \leq r \leq d-1$, in which case

$$\dim(\Sigma_r(\nu_2(\mathbb{P}^{d-1}))) = \min\left\{\binom{d+1}{2} - 1, rd - \binom{r}{2} - 1\right\}.$$

Note that there is a typo in the expression of this dimension in Theorem 2 of [14]. Under the identification $\mathbb{C}^d \otimes \mathbb{C}^d \cong \mathrm{L}((\mathbb{C}^d)^*, \mathbb{C}^d)$, the variety $\Sigma_r(\nu_2(\mathbb{P}^{d-1}))$ corresponds to the set of (projective) symmetric $d \times d$ matrices of rank at most r.

For the Grassmannian variety $\mathrm{Gr}(m-1, \mathbb{P}^{d-1}) \subseteq \mathbb{P}(\bigwedge^m(\mathbb{C}^d))$, we have

$$\dim(\mathrm{Gr}(m-1, \mathbb{P}^{d-1})) = m(d-m),$$

and $\dim(\mathbb{P}(\bigwedge^m(\mathbb{C}^d))) = \binom{d}{m} - 1$, so

$$\dim(\Sigma_r(\mathrm{Gr}(m-1, \mathbb{P}^{d-1}))) \leq \min\left\{\binom{d}{m} - 1, rm(d-m) + r - 1\right\}.$$

As with the Segre variety, there are a conjecturally complete set of defective Grassmannian secants; see [14, Conjecture 7] and [13, 19, 15]. Similarly to the Segre and Veronese varieties, $\Sigma_r(\mathrm{Gr}(1, \mathbb{P}^{d-1}))$ is defective whenever $2 \leq r < \lfloor \frac{d}{2} \rfloor$, in which case

$$\dim(\Sigma_r(\mathrm{Gr}(1, \mathbb{P}^{d-1}))) = \binom{d}{2} - \binom{d-2r}{2} - 1 = 2r(d-r) - r - 1.$$

Under the identification $\mathbb{C}^d \otimes \mathbb{C}^d \cong \mathrm{L}((\mathbb{C}^d)^*, \mathbb{C}^d)$, the variety $\Sigma_r(\mathrm{Gr}(1, \mathbb{P}^{d-1}))$ corresponds to the set of projective antisymmetric $d \times d$ matrices of rank at most $2r$.

2.2 The set of stabilizer states (and its secants)

We restrict our attention to the complex numbers in this section. For a positive integer m, let $\mathcal{U}_m = (\mathbb{C}^2)^{\otimes m}$, and let $\mathrm{Pauli}_m \subseteq \mathrm{U}(\mathcal{U}_m)$ be the *Pauli group*: the group generated by all m-fold tensor products of elements of the set $\{X, Z, i\mathbb{1}_2\} \subseteq \mathrm{U}(\mathbb{C}^2)$, where

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

and $\mathbb{1}_k$ is the identity matrix in $\mathrm{U}(\mathbb{C}^k)$ for any positive integer k. The *Clifford group*, denoted $\mathrm{Cliff}_m \subseteq \mathrm{U}(\mathcal{U}_m)$, is the normalizer of the Pauli group in $\mathrm{U}(\mathcal{U}_m)$. The *projective Clifford group*, denoted $\mathcal{C}_m$, is defined as $\mathrm{Cliff}_m/\mathrm{U}(\mathbb{C})$, the Clifford group modulo phases. The projective Clifford group is generated by the (equivalence classes of) tensor products of elements of the set $\{H, S, CNOT\}$, where

$$H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \quad \text{and} \quad CNOT = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

A quantum state $[v] \in \mathbb{P}^{2^m-1}$ is called a *stabilizer state* if $[v] = U\left[\mathrm{e}_0^{\otimes m}\right]$ for some $U \in \mathcal{C}_m$. Let $\mathrm{Stab}_m = \mathcal{C}_m[\mathrm{e}_0^{\otimes m}]$ be the set of stabilizer states. It is well known that a state $[v] \in \mathbb{P}^{2^m-1}$ is a stabilizer state if and only if

$$[v] = \left[\sum_{x \in A} i^{l(x)} \cdot (-1)^{q(x)} \cdot \mathrm{e}_x\right] \tag{2.8}$$

for some affine linear subspace $A \subseteq \mathbb{F}_2^m$, linear form $l\colon \mathbb{F}_2^m \to \mathbb{F}_2$, and quadratic form $q\colon \mathbb{F}_2^m \to \mathbb{F}_2$ [47, 125]. Here and throughout, we define $\mathrm{e}_x := \mathrm{e}_{x_1} \otimes \cdots \otimes \mathrm{e}_{x_n}$ when $x \in \mathbb{F}_2^m$.

The *stabilizer rank* of a state $[v] \in \mathbb{P}^{2^m-1}$, denoted $\chi([v])$, is the Stab_m-rank of $[v]$, i.e. the smallest integer r for which $[v] \in \Sigma_r(\mathrm{Stab}_m)$. Note that $\Sigma_r^\circ(\mathrm{Stab}_m)$ is already closed, as it is a finite union of linear subspaces, so the stabilizer rank equals the border stabilizer rank. In concrete terms, the stabilizer rank of $[v]$ is the smallest integer r for which

$$v = \sum_{a \in [r]} \alpha_a \sigma_a$$

for some complex numbers $\alpha_a \in \mathbb{C}$ and stabilizer states $[\sigma_a] \in \mathrm{Stab}_m$.

We also have an approximate version of stabilizer rank. For a positive real number $\delta > 0$ and state $[v] \in \mathbb{P}^{2^m-1}$, we define the *$\delta$-approximate stabilizer rank* of $[v]$, denoted $\chi_\delta([v])$, as

$$\chi_\delta([v]) = \min\left\{\chi([u]) : \left\| \frac{u}{\|u\|} - \frac{v}{\|v\|} \right\| < \delta \right\}.$$

We say a quantum state $[v] \in \mathbb{P}^{2^m-1}$ is *real* if v is proportional to a state with only real coordinates in the computational basis. A quantum state $[v]$ is a real stabilizer state if and only if it can be written in the form (2.8) with $l = 0$. The set of real stabilizer states in $\mathbb{P}^{2^m-1}$, which we denote by $\mathrm{Stab}_m^{\mathbb{R}}$, is precisely the orbit of $[\mathrm{e}_0^{\otimes m}]$ under the group generated by the (equivalence classes of) H and $CNOT$ [47]. For a quantum state $[v] \in \mathbb{P}^{2^m-1}$, we define the *real stabilizer rank* of $[v]$, denoted $\chi^{\mathbb{R}}([v])$, to be the smallest integer r for which $[v]$ can be written as a (complex) superposition of r real stabilizer states.

We close this section by computing the number of stabilizer states and real stabilizer states, which we will use in Section 5.3 to upper bound the number of states of sub-generic rank. It is a standard fact that there are $\binom{m}{k}_2$ distinct k-dimensional linear subspaces of $\mathbb{F}_2^m$, where

$$\binom{m}{k}_2 = \prod_{i=0}^{k-1} \frac{2^{m-i}-1}{2^{k-i}-1}$$

is the *Gaussian binomial coefficient* (see e.g. [57]). Since there are 2^{m-k} distinct affine translations of a k-dimensional linear subspace of $\mathbb{F}_2^m$, it follows that there are $\binom{m}{k}_2 2^{m-k}$ distinct affine linear subspaces of dimension k. For each index $k \in [m]$ and each k-dimensional affine subspace $A \subseteq \mathbb{F}_2^m$, there are $2^{k(k+1)/2}$ distinct quadratic forms on A and 2^k distinct linear functions on A. It follows that

$$\begin{aligned} |\mathrm{Stab}_m| &= 2^m \sum_{k=1}^{n} \binom{m}{k}_2 2^{k(k+1)/2} \\ &= 2^m \prod_{k=1}^{m} (2^k + 1), \end{aligned}$$

where the second line follows from the Gaussian binomial theorem (see e.g. [57]). Similarly,

$$\left|\mathrm{Stab}_m^{\mathbb{R}}\right| = 2^m \sum_{k=1}^{m} \binom{m}{k}_2 .$$

The quantity $|\mathrm{Stab}_m|$ was previously computed in [62, Corollary 21] using a different proof technique.

Chapter 3

Generic local state discrimination with pre-shared entanglement

The contents of Chapters 3 and 4 are based on the preprint [91], which was written in collaboration with Nathaniel Johnston.

Walgate and Scott have determined the maximum number of generic pure quantum states that can be unambiguously discriminated by an LOCC measurement [6]. In this chapter, we determine this number in a more general setting in which the local parties have access to pre-shared entanglement in the form of a resource state. We find that, for an arbitrary pure resource state, this number is equal to the Krull dimension of (the closure of) the set of pure states obtainable from the resource state by SLOCC. Surprisingly, a generic resource state maximizes this number.

We work over $\mathbb{C}$ throughout this chapter. Recall the following definitions from Chapter 1. An *LOCC measurement* is a quantum measurement that can be implemented by local operations and classical communication (LOCC). We say an n-tuple of pure quantum states $([v_1], \dots, [v_n])$ is *locally (unambiguously) discriminable* if there exists an LOCC measurement with $n+1$ outcomes $\{1, \dots, n, ?\}$ that, when performed on any $[v_a]$, outputs either a or ?, with non-zero probability to output a. More generally, we say that $([v_1], \dots, [v_n])$ is *locally (unambiguously) discriminable with (a pure resource state)* $[w]$ if $([v_1], \dots, [v_n])$ can be locally discriminated using an LOCC measurement with pre-shared entanglement $[w]$.

Let $\mathcal{V} = \bigotimes_{j=1}^m \mathbb{C}^{d_j}$ and $\mathcal{W} = \bigotimes_{j=1}^m \mathbb{C}^{c_j}$ be m-partite vector spaces. Walgate and Scott determined that, a generic n-tuple of $\sum_{j=1}^m (d_j - 1) + 1$ pure states in $\mathcal{V}$ is locally discriminable, and this is the largest number for which this holds [6]. In this chapter, we extend Walgate and Scott's result to determine the number of generic pure states that can be locally discriminated with an arbitrary pure resource state $[w] \in \mathbb{P}\mathcal{W}$. For example, we prove that at most $r(\sum_{j=1}^{m}(d_j-1)+1)$ generic pure states are locally discriminable with the tensor-rank-r GHZ state $[\tau_{r,m}] = [\sum_{a=1}^{r} \mathrm{e}_a^{\otimes m}]$, where $\mathrm{e}_1, \dots, \mathrm{e}_r$ are standard basis vectors, and this bound is often achieved (Corollary 12). For $m = 2$, this bound cannot be achieved: precisely $d_1 d_2 - (d_1 - \min\{d_1, r\})(d_2 - \min\{d_2, r\})$ generic pure states are locally discriminable with $[\tau_{r,2}]$ (or any other Schmidt rank r state), and this is the largest number for which this holds (Corollary 14).

More generally, we characterize this number for an arbitrary pure resource state $[w] \in \mathbb{P}\mathcal{W}$ (Theorem 11). For an algebraic variety Y, we say that a property holds for a *generic* element of Y if there exists a Zariski-open-dense subset of Y on which that property holds. Recall that the *SLOCC image* of $[w]$ in $\mathbb{P}\mathcal{V}$, denoted by $\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])$, is the set of pure states obtainable from $[w]$ by stochastic local operations and classical communication (SLOCC). Letting $d = \dim(\overline{\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])})$ be the Krull dimension of the Zariski closure of $\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])$, we prove that a generic $(d+1)$-tuple of pure states is locally discriminable with $[w]$. We also prove a "strong converse:" A Zariski-open-dense set of $(d+2)$-tuples of pure states are *not* locally discriminable with $[w]$. To recover Walgate and Scott's result, observe that the SLOCC image of a trivial resource state is the set of unentangled (or, *product*) pure states. This set is already closed, and has dimension $d = \sum_{j=1}^{m}(d_j - 1)$.

Our characterization extends Walgate and Scott's result in two directions beyond the obvious addition of a resource state: First, our results hold under the algebraic-geometric notion of "generic" introduced above, which yields stronger statements than the measure-theoretic notion used by Walgate and Scott. Second, our "strong converse," mentioned above, is much stronger than the converse statement proven by Walgate and Scott (see the discussion following Corollary 13).

It is natural to quantify "how useful" a given resource state $[w]$ is for LUSD in the space $\mathcal{V}$ in terms of the number of generic pure states in $\mathbb{P}\mathcal{V}$ that can be locally discriminated with $[w]$. We use algebraic group theory to prove that, under this barometer, a generic resource state is maximally useful for LUSD whenever $c_j \leq d_j$ for all $j \in [m]$. This is quite surprising, because in many other contexts the most useful quantum states form a measure zero subset of the Hilbert space. For example, in bipartite space the maximally entangled states form a measure zero set, and in multipartite space the set of graph states (a class of states often regarded as the most useful states) has measure zero [68]. As one more example, it is known that most multipartite states are, in a sense, useless for measurement-based quantum computation [1].

We now sketch a proof of one direction (the "strong converse") of our characterization of generic LUSD with a resource state. We invoke an observation of Bandyopadhyay et al., that an n-tuple of pure states $([v_1], \dots, [v_n])$ is locally discriminable with $[w]$ if and only if there exist pure states $[u_1], \dots, [u_n] \in \operatorname{Im}_{\mathbb{P}\mathcal{V}}([w])$ for which $(u_a^\intercal v_b \neq 0 \iff a = b)$ [11]. In particular, $\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w]) \cap \operatorname{span}\{[v_1], \dots, [v_{n-1}]\}^\perp \neq \{\}$, because the intersection contains $[u_n]$. We also use a theorem in algebraic geometry, which states that for a quasiprojective variety X of dimension d, a generic (projective) linear subspace of codimension $d+1$ is disjoint from X, and a generic linear subspace of smaller codimension intersects X. Letting $X = \overline{\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w])}$, it follows from the algebraic geometry result that for a generic $(d+1)$-tuple of states $([v_1], \dots, [v_{d+1}])$, it holds that $X \cap \operatorname{span}\{[v_1], \dots, [v_{d+1}]\}^\perp = \{\}$, since $\operatorname{span}\{[v_1], \dots, [v_{d+1}]\}^\perp$ forms a generic projective linear subspace of codimension $d+1$. By the result of Bandyopadhyay et al., $([v_1], \dots, [v_{d+1}], [v_{d+2}])$ is not locally discriminable with $[w]$ for any pure state $[v_{d+2}]$. This proves that a generic $(d+2)$-tuple of pure states is not locally discriminable with $[w]$. We prove the other direction, that a generic $(d+1)$-tuple of pure states is locally discriminable with $[w]$, using a similar, but more complicated argument.

3.1 Preliminaries to Chapter 3

In this section, we review some background material for this chapter, including an alternative view of the Grassmannian variety as the set of projective planes of fixed dimension, a classical algebraic geometry result (Theorem 7) on the maximum dimension of a projective linear subspace disjoint from a variety, the SLOCC image, and local unambiguous state discrimination (LUSD). This background material is not needed in other chapters, aside from Theorem 7, which is used in Chapter 4 to determine the maximum dimension of entangled subspaces. As above, let $\mathcal{V} = \bigotimes_{j=1}^m \mathbb{C}^{d_j}$ and $\mathcal{W} = \bigotimes_{j=1}^m \mathbb{C}^{c_j}$ be m-partite vector spaces.

3.1.1 The set of projective $(n-1)$-planes as a projective variety

In this subsection, we recall the canonical bijection between the Grassmannian variety and the set of projective planes of a fixed dimension, which endows the latter set with the structure of a projective variety. We then review several facts about the Grassmannian that we will use to prove our characterization of generic LUSD with a resource state (Theorem 11).

Recall the canonical bijection between the set of projective $(n-1)$-planes and $\operatorname{Gr}(n-1, \mathbb{P}\mathcal{V})$, which identifies $[v_1 \wedge \cdots \wedge v_n]$ with $\operatorname{span}\{[v_1], \dots, [v_n]\}$ for any linearly independent set

$\{[v_1], \ldots, [v_n]\} \subseteq \mathbb{P}\mathcal{V}$. This identification endows the set of projective $n-1$ planes with the structure of a projective variety. In a slight abuse of notation, we will use $\mathrm{Gr}(n-1, \mathbb{P}\mathcal{V})$ to refer to both the projective variety of decomposable elements of $\mathbb{P}(\bigwedge^n \mathcal{V})$ and the projective variety of projective $(n-1)$-planes in $\mathbb{P}\mathcal{V}$.

It is important to keep in mind the canonical bijection between the set of projective $(n-1)$-planes in $\mathbb{P}\mathcal{V}$ (namely, $\mathrm{Gr}(n-1, \mathbb{P}\mathcal{V})$), and the set of n-dimensional linear subspaces of $\mathcal{V}$ (typically denoted $\mathrm{Gr}(n, \mathcal{V})$). While $\mathrm{Gr}(n, \mathcal{V})$ is perhaps more standard, we prefer $\mathrm{Gr}(n-1, \mathbb{P}\mathcal{V})$, as we would like the elements of each subspace to be states. Note that we have replaced the symbol m with the symbol n in this section, to match later notation in which the Grassmannian is viewed as a space of projective $(n-1)$-planes.

In accordance with our definition of *generic*, we say that a property holds for a generic projective $(n-1)$-plane if there exists a Zariski-open-dense subset of $\mathrm{Gr}(n-1, \mathbb{P}\mathcal{V})$ on which the property holds. We say a property holds for a generic n-tuple of states in $\mathbb{P}\mathcal{V}$ (or when $n = 1$, simply a generic state in $\mathbb{P}\mathcal{V}$) if there exists a Zariski-open-dense subset $U \subseteq \mathbb{P}(V)^{\times n}$ such that the property holds for every $([v_1], \ldots, [v_n]) \in U$. (Here, $\mathbb{P}(V)^{\times n}$ is viewed as a projective variety via the Segre embedding.)

In the remainder of this subsection, we review two facts that we will use to prove our characterization of generic LUSD with a resource state. First, a generic n-tuple of states spans a generic projective $(n-1)$-plane, and vice versa (Fact 5). Second, the bijection $\mathrm{Gr}(n-1, \mathbb{P}^d) \cong \mathrm{Gr}(d-n, \mathbb{P}^d)$, which sends a subspace to its orthogonal complement, defines an isomorphism of projective varieties (Fact 6). We defer the proofs of these facts to Appendix A.1.

Fact 5. A generic n-tuple of states spans a generic projective $(n-1)$-plane, and vice versa. In more details, let n be a positive integer, let $\mathcal{V}$ be a $\mathbb{C}$-vector space, and let

$$\tilde{\pi} : \mathbb{P}(\mathcal{V})^{\times n} \dashrightarrow \mathrm{Gr}(n-1, \mathbb{P}\mathcal{V}) \tag{3.1}$$

be the rational map defined by $\tilde{\pi}([v_1], \ldots, [v_n]) = [v_1 \wedge \cdots \wedge v_n]$. Then a subset $U \subseteq \mathrm{Gr}(n-1, \mathbb{P}\mathcal{V})$ is open-dense if and only if $\tilde{\pi}^{-1}(U) \subseteq \mathbb{P}(\mathcal{V})^{\times n}$ is open-dense.

Fact 6. The bijection $\mathrm{Gr}(n-1, \mathbb{P}^d) \cong \mathrm{Gr}(d-n, \mathbb{P}^d)$, which sends a subspace to its orthogonal complement with respect to some non-degenerate bilinear form $\langle \cdot, \cdot \rangle$, defines an isomorphism of projective varieties.

3.1.2 Projective linear subspaces disjoint from a variety

The following algebraic-geometric result will be instrumental in proving our characterization of generic LUSD with a resource state. We will also use it to determine the maximum dimension of an entangled subspace in Chapter 4.

Theorem 7. *Let $X \subseteq \mathbb{P}^D$ be a quasiprojective variety. Then*

$$\begin{aligned}\dim(X) = &\textit{The smallest non-negative integer } d \textit{ for which a generic projective } D-d-1 \textit{ plane}\\ &\quad\textit{is disjoint from } X\\ = &\textit{The largest non-negative integer } d \textit{ for which a generic projective } D-d \textit{ plane}\\ &\quad\textit{intersects } X.\end{aligned}$$

If X is projective, then

$$\begin{aligned}\dim(X) = &\textit{The largest non-negative integer } d \textit{ for which every projective } D-d \textit{ plane}\\ &\quad\textit{intersects } X.\end{aligned}$$

This characterization is taken to be the definition of the dimension of an irreducible variety in [65, Definition 11.2], and is shown to be equivalent to other standard notions of dimension, e.g. the Krull dimension. The above extension to the reducible case is straightforward: Let $X = X_1 \cup \cdots \cup X_t$ be the irreducible decomposition of X, let $d = \dim(X)$, and let $d_i = \dim(X_i)$ for each $i \in [t]$. Then $d = \max_{i \in [t]} d_i$ by definition. Since a general $D-d-1$ plane is disjoint from each X_i, it follows that a general $D-d-1$ plane is disjoint from X. Furthermore, d is the smallest non-negative integer that satisfies this property, because for any $\tilde{d} < d$ it holds that $\tilde{d} < d_i$ for some $i \in [t]$, and hence a general $D - \tilde{d} - 1$ plane intersects X_i (and thus, X). The refined statement when X is projective is handled similarly.

3.1.3 The SLOCC image

LOCC channels are quantum channels (completely positive, trace preserving maps) that can be implemented by local operations and classical communication (LOCC). *SLOCC maps* are completely positive, trace non-increasing maps that can be implemented with non-zero probability by an LOCC channel. In other words, SLOCC maps represent LOCC channels with postselection [127].

For a state $[w] \in \mathbb{P}\mathcal{W}$, recall that the *SLOCC image* of $[w]$ in $\mathbb{P}\mathcal{V}$, denoted $\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w]) \subseteq \mathbb{P}\mathcal{V}$, is the set of states in $\mathbb{P}\mathcal{V}$ obtainable from $[w]$ by SLOCC. A related notion is the *SLOCC orbit* of a state $[v] \in \mathbb{P}\mathcal{V}$, denoted $\mathcal{O}_{[v]} \subseteq \mathbb{P}\mathcal{V}$, which is the set of states in $\mathrm{Im}_{\mathbb{P}\mathcal{V}}([v])$ that can be converted back to $[v]$ by SLOCC, i.e.

$$\mathcal{O}_{[v]} = \{[u] \in \mathrm{Im}_{\mathbb{P}\mathcal{V}}([v]) : [v] \in \mathrm{Im}_{\mathbb{P}\mathcal{V}}([u])\}.$$

In this subsection, we observe several properties of $\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w])$ and $\mathcal{O}_{[v]}$ that we will use in Section 3.2. We can describe these sets mathematically as

$$\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w]) = \{[(A_1 \otimes \cdots \otimes A_m)w] : A_i \in \mathrm{L}(\mathbb{C}^{c_i}, \mathbb{C}^{d_i}) \quad \text{for all} \quad i \in [m] \\ \text{and} \quad (A_1 \otimes \cdots \otimes A_m)w \neq 0\},$$

and

$$\mathcal{O}_{[v]} = \{[(A_1 \otimes \cdots \otimes A_m)v] : A_i \in \mathrm{GL}(\mathbb{C}^{d_i}) \quad \text{for all} \quad i \in [m]\},$$

see [54]. True to its name, $\mathcal{O}_{[v]}$ is the orbit of $[v]$ under the standard action of the product (projective) general linear group.

If $\mathcal{V} = \mathbb{C}^{d_1} \otimes \mathbb{C}^{d_2}$ and $\mathcal{W} = \mathbb{C}^{c_1} \otimes \mathbb{C}^{c_2}$ are bipartite spaces, then $\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w])$ is the set of states in $\mathbb{P}\mathcal{V}$ of Schmidt rank less than or equal to the Schmidt rank of $[w]$, and $\mathcal{O}_{[v]}$ is the set of states of Schmidt rank equal to the Schmidt rank of $[v]$. In multipartite space, consider the tensor-rank-r GHZ state

$$[\tau_{r,m}] = \left[\sum_{a=1}^{r} \mathrm{e}_a^{\otimes m}\right] \subseteq \mathbb{P}((\mathbb{C}^r)^{\otimes m}). \tag{3.2}$$

It is straightforward to verify that $\operatorname{Im}_{\mathbb{P}\mathcal{V}}([\tau_{r,m}])$ is the set of states in $\mathbb{P}\mathcal{V}$ of tensor rank at most r, and hence $\overline{\operatorname{Im}_{\mathbb{P}\mathcal{V}}([\tau_{r,m}])} = \sigma_r(Y)$, where Y is the Segre variety of $\mathbb{P}\mathcal{V}$ defined in (2.1). It is clear that $\mathcal{O}_{[\tau_{r,m}]}$ is the set of states of the form $[\sum_{a=1}^{r} x_{a,1} \otimes \cdots \otimes x_{a,m}]$, where $\{x_{1,j}, \ldots, x_{r,j}\} \subseteq \mathbb{C}^r$ is linearly independent for all $j \in [m]$.

We conclude this subsection by proving that $\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w])$ and $\mathcal{O}_{[v]}$ are both irreducible and constructible (Fact 8), and that in many cases a generic state $[w] \in \mathbb{P}\mathcal{W}$ maximizes $\dim(\overline{\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w])})$ (Fact 9). We will use Facts 8 and 9, respectively, to prove our characterization of generic LUSD with a resource state, and to prove that a generic resource state is maximally useful for LUSD (see Theorem 11 and the subsequent discussion). We defer the proofs of these facts to Appendix A.1.

Fact 8. Let $\mathcal{V} = \bigotimes_{j=1}^{m} \mathbb{C}^{d_j}$ and $\mathcal{W} = \bigotimes_{j=1}^{m} \mathbb{C}^{c_j}$ be vector spaces, and let $[w] \in \mathbb{P}\mathcal{W}$ and $[v] \in \mathbb{P}\mathcal{W}$ be states. Then the sets $\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w])$ and $\mathcal{O}_{[v]}$ are both irreducible and constructible in the Zariski topology.

Since both $\operatorname{Im}_{\mathbb{P}\mathcal{V}}([w])$ and $\mathcal{O}_{[v]}$ are constructible, each contains an open-dense subset of its closure [5, Lemma 2.1]. In fact, $\mathcal{O}_{[v]}$ is itself an open-dense subset of its closure (i.e. $\mathcal{O}_{[v]}$ is *locally closed*) [72, Proposition 8.3]. Observe that $\overline{\mathcal{O}_{[v]}} = \overline{\operatorname{Im}_{\mathbb{P}\mathcal{V}}([v])}$.

Fact 9. Let $\mathcal{V} = \bigotimes_{j=1}^{m} \mathbb{C}^{d_j}$ and $\mathcal{W} = \bigotimes_{j=1}^{m} \mathbb{C}^{c_j}$ be vector spaces with $c_j \leq d_j$ for all $j \in [m]$. Then $\dim(\overline{\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])})$ is maximized for a generic state $[w] \in \mathbb{P}\mathcal{W}$.

It would be nice to know if the condition that $c_j \leq d_j$ for all $j \in [m]$ is necessary for Fact 9 to hold.

3.1.4 Local unambiguous state discrimination (LUSD)

We conclude this section by reviewing unambiguous state discrimination (USD), and its local counterpart, local unambiguous state discrimination (LUSD). We use these notions in Section 3.2 to characterize generic LUSD with a resource state.

An n-tuple of states $([v_1], \dots, [v_n]) \in \mathbb{P}(\mathcal{V})^{\times n}$ is *(unambiguously) discriminable* if there exists a quantum measurement with $n+1$ outcomes $\{1, \dots, n, ?\}$ that, when performed on any $[v_a]$, outputs either a or ?, with non-zero probability to output a. Mathematically, this is equivalent to the existence positive semidefinite operators $M_1, \dots, M_n, M_? \in \mathrm{Pos}(\mathcal{V})$ for which $M_1 + \cdots + M_n + M_? = \mathbb{1}$ and $(\langle v_b, M_a v_b \rangle \neq 0 \iff a = b)$. Note that $([v_1], \dots, [v_n])$ is discriminable if and only if it is linearly independent.

We say that an n-tuple of states

$$([v_1], \dots, [v_n]) \in \mathbb{P}(\mathcal{V})^{\times n}$$

is *locally (unambiguously) discriminable* if it is discriminable by a measurement implementable by an LOCC channel, with local subsystems $\mathbb{P}\mathbb{C}^{d_j}$. We say that an n-tuple of states $([v_1], \dots, [v_n])$ is *locally (unambiguously) discriminable with (resource state)* $[w] \in \mathbb{P}\mathcal{W}$ if $([v_1 \otimes w], \dots, [v_n \otimes w])$ is locally discriminable, where the local subsystem $j \in [m]$ is now the composite system $\mathbb{P}(\mathbb{C}^{d_j} \otimes \mathbb{C}^{c_j})$. This is equivalent to $([v_1], \dots, [v_n])$ being locally discriminable via an LOCC measurement with pre-shared entanglement $[w]$ [11].

3.2 Generic local state discrimination with pre-shared entanglement

In this section, we characterize the maximum number of generic pure states that can be locally discriminated with a fixed resource state $[w]$, and observe that a generic resource state $[w]$ maximizes this number. To prove this characterization, we require the following mathematical description of LUSD with a resource state.

Theorem 10 ([11]). *Let $\mathcal{V} = \bigotimes_{j=1}^m \mathbb{C}^{d_j}$ and $\mathcal{W} = \bigotimes_{j=1}^m \mathbb{C}^{c_j}$. An n-tuple of states $([v_1], \dots, [v_n]) \in \mathbb{P}(\mathcal{V})^{\times n}$ is locally discriminable with resource state $[w] \in \mathbb{P}\mathcal{W}$ if and only if there exist states*

$$[u_1], \dots, [u_n] \in \mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])$$

for which ($u_a^\intercal v_b \neq 0 \iff a = b$), where the transpose can equivalently be taken with respect to any product basis of $\mathcal{V}$.

Note that this statement indeed does not depend on the choice of product basis over which the transpose is taken, since any two product bases are related by a product change of basis $A = A_1 \otimes \cdots \otimes A_m \in \mathrm{GL}(\mathcal{V})$, and

$$[u_a] \in \mathrm{Im}_{\mathbb{P}\mathcal{V}}([w]) \iff [Au_a] \in \mathrm{Im}_{\mathbb{P}\mathcal{V}}([w]).$$

Also note that, if the Hermitian inner product is preferred, an alternative (equivalent) statement is that there exist $[u_1], \dots, [u_n] \in \mathrm{Im}_{\mathbb{P}\mathcal{V}}([\overline{w}])$ such that ($\langle u_a, v_b \rangle \neq 0 \iff a = b$), where $\overline{w}$ denotes the complex conjugate of w with respect to any product basis of $\mathcal{W}$.

The SLOCC image of a trivial (i.e. non-existent) resource state is simply the Segre variety Y of product states, defined in (2.1). This case of Theorem 10, first proven by Chefles, was used to characterize generic LUSD in [6], and we use the above generalization of Bandyopadhyay et al. to characterize generic LUSD with a resource state.

Theorem 11. *Let n be a positive integer, let $\mathcal{V} = \bigotimes_{j=1}^m \mathbb{C}^{d_j}$ and $\mathcal{W} = \bigotimes_{j=1}^m \mathbb{C}^{c_j}$ be vector spaces, let $[w] \in \mathbb{P}\mathcal{W}$ be a quantum state, let $X = \overline{\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])}$ be the closure of the SLOCC image of $[w]$, and let $d = \dim(X)$. If $n \leq d+1$, then a generic n-tuple of states in $\mathbb{P}\mathcal{V}$ is locally discriminable with $[w]$. If $n > d+1$, then a generic n-tuple of states in $\mathbb{P}\mathcal{V}$ is not locally discriminable with $[w]$.*

By Fact 9, if $c_j \leq d_j$ for all $j \in [m]$, then $\dim(\overline{\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])})$ is maximized for a generic resource state $[w] \in \mathbb{P}\mathcal{W}$. It follows from Theorem 11 that a generic resource state $[w]$ can be used to discriminate the maximum number of generic states in this setting. It would be nice to know if the condition that $c_j \leq d_j$ for all $j \in [m]$ can be dropped from Fact 9, as this would imply that a generic resource state can be used to discriminate the maximum number of generic states under any choice of $\mathcal{V}$ and $\mathcal{W}$.

We emphasize that the third sentence of Theorem 11 is much stronger than a simple converse to the second sentence: It asserts that an open-dense subset of $\mathbb{P}(\mathcal{V})^{\times n}$ is not locally discriminable with $[w]$.

We remark that similar characterizations of generic LUSD can be obtained in symmetric and antisymmetric space, but we omit stating such results explicitly.

Proof of Theorem 11. We will make use of the map $\tilde{\pi}$ defined in (3.1), as well as several facts observed in Section 3.1. Suppose first that $n > d+1$. By Theorem 7, there exists an open-dense subset $V \subseteq \mathrm{Gr}(d_1 \cdots d_m - n - 1, \mathbb{P}\mathcal{V})$ for which every element (i.e. subspace) in V is disjoint from X. By Facts 5 and 6, $\tilde{\pi}^{-1}(V^\perp) \subseteq \mathbb{P}(\mathcal{V})^{\times n-1}$ is open-dense, so

$$U := \mathbb{P}(\mathcal{V}) \times \tilde{\pi}^{-1}(V^\perp) \subseteq \mathbb{P}(\mathcal{V})^{\times n}$$

is open-dense. For any $([v_1], \dots, [v_n]) \in U$, it holds that

$$\mathrm{span}\{[v_2], \dots, [v_n]\}^\perp \cap X = \{\},$$

by definition of V. By Theorem 10, $([v_1], \dots, [v_n])$ is not locally discriminable with resource state $[w]$.

Conversely, suppose $n \leq d+1$. By Theorem 10, the desired result is equivalent to the existence of an open-dense subset of $\mathbb{P}(\mathcal{V})^{\times n}$ contained in $\bigcap_{a=1}^n S_a$, where

$$S_a = \{([v_1], \dots, [v_n]) : \text{there exists} \quad [u] \in \mathrm{Im}_{\mathbb{P}\mathcal{V}}([w]) \quad \text{such that} \quad (u^\intercal v_b \neq 0 \iff a = b)\}$$

for each $a \in [n]$. By Fact 8, $\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])$ is constructible. It follows that S_a is constructible, so it contains an open-dense subset of its closure [5, Lemma 2.1]. To complete the proof, it suffices to show that $\overline{S_a} = \mathbb{P}(\mathcal{V})^{\times n}$ all $a \in [n]$. We do so by constructing a subset of S_a that is dense in $\mathbb{P}(\mathcal{V})^{\times n}$.

We take $a = n$ to ease the notation (the other $a \in [n]$ follow by symmetry). Let $U \subseteq \mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])$ be an open-dense subset of X. Then U is an (irreducible) quasiprojective variety of dimension d. By Theorem 7 and the inequality $n \leq d+1$, there exists an open-dense subset $W \subseteq \mathrm{Gr}(d_1 \cdots d_m - n, \mathbb{P}\mathcal{V})$ for which every element (i.e. subspace) in W intersects U. Therefore, the set $Z = \tilde{\pi}^{-1}(W^\perp) \subseteq \mathbb{P}(\mathcal{V})^{\times n-1}$ is open-dense, and for each $v := ([v_1], \dots, [v_{n-1}]) \in Z$ there exists $[u_v] \in U$ with $u_v^\intercal v_b = 0$ for all $b \in [n-1]$. Let

$$T_v = \{[v_n] \in \mathbb{P}\mathcal{V} : u_v^\intercal v_n \neq 0\}.$$

The set

$$V_n = \bigcup_{v \in Z} \{v\} \times T_v$$

is clearly contained in S_n. To complete the proof, we show that $\overline{V_n} = \mathbb{P}(\mathcal{V})^{\times n}$. For any open-dense subset $S \subseteq \mathbb{P}(\mathcal{V})^{\times n}$, there exists $v \in Z$ for which the set $(\{v\} \times \mathbb{P}(\mathcal{V})) \cap S$ is open-dense inside $\{v\} \times \mathbb{P}(\mathcal{V})$. Since $T_v \subseteq \mathbb{P}\mathcal{V}$ is open-dense, it follows that $\{v\} \times T_v \subseteq \{v\} \times \mathbb{P}\mathcal{V}$ is open-dense, so

$$(\{v\} \times T_v) \cap S \subseteq \{v\} \times \mathbb{P}(\mathcal{V})$$

is open-dense. Thus, $V_n \cap S \neq \{\}$. Since V_n intersects every open-dense subset $S \subseteq \mathbb{P}(\mathcal{V})^{\times n}$, it follows that $\overline{V_n} = \mathbb{P}(\mathcal{V})^{\times n}$. This completes the proof. $\square$

Corollary 12. *Let n and r be positive integers, let $\mathcal{V} = \bigotimes_{j=1}^{m} \mathbb{C}^{d_j}$ be a vector space, let $Y \subseteq \mathbb{P}\mathcal{V}$ be the Segre variety defined in (2.1), let $X = \sigma_r(Y)$ be the r-th secant variety, and let $d = \dim(X)$. If $n \leq d+1$, then a generic n-tuple of states in $\mathbb{P}(\mathcal{V})$ is locally discriminable with the tensor-rank-r GHZ state $[\tau_{r,m}]$, defined in (3.2). If $n > d+1$, then a generic n-tuple of states in $\mathbb{P}(\mathcal{V})$ is not locally discriminable with $[\tau_{r,m}]$.*

Note that the standard upper bound on $\dim(\sigma_r(Y))$, reviewed in (2.6), yields the upper bound mentioned in the introduction on the number of generic states locally discriminable with $[\tau_{r,m}]$. The dimension of $\sigma_r(Y)$ is known in several cases (see Section 2.1), in particular, when $r = 1$ or $m = 2$. The next two corollaries follow from these known dimensions.

Corollary 13. *Let n be a positive integer, and let $\mathcal{V} = \bigotimes_{j=1}^{m} \mathbb{C}^{d_j}$ be a vector space. If $n \leq \sum_{j=1}^{m}(d_j - 1) + 1$, then a generic n-tuple of states in $\mathbb{P}(\mathcal{V})$ is locally discriminable. If $n > \sum_{j=1}^{m}(d_j - 1) + 1$, then a generic n-tuple of states in $\mathbb{P}(\mathcal{V})$ is not locally discriminable.*

Corollary 13 follows from the fact that the SLOCC image of a trivial (i.e. non-existent) resource state is simply the Segre variety, which has dimension $\sum_{j=1}^{m}(d_j - 1)$. Corollary 13 strengthens Theorem 4.3 in [6] in two ways: First, the notion of "generic" used by Walgate and Scott is measure-theoretic, which is weaker than our algebraic-geometric definition of "generic." Second, when $n > \sum_{j=1}^{m}(d_j - 1) + 1$, Walgate and Scott simply prove that the subset of $\mathbb{P}(\mathcal{V})^{\times n}$ consisting of locally discriminable n-tuples of states is not full measure, which is much weaker than our result that a full-measure (and in fact, open-dense) subset of $\mathbb{P}(\mathcal{V})^{\times n}$ is not locally discriminable.

The known dimensions of secant varieties in bipartite space allow us to quantify exactly how many generic states are locally discriminable with an arbitrary resource state:

Corollary 14. *Let n and r be positive integers, let $\mathcal{V} = \mathbb{C}^{d_1} \otimes \mathbb{C}^{d_2}$ and $\mathcal{W} = \mathbb{C}^{c_1} \otimes \mathbb{C}^{c_2}$ be vector spaces. If $n \leq d_1 d_2 - (d_1 - \min\{d_1, r\})(d_2 - \min\{d_2, r\})$, then a generic n-tuple of states in $\mathbb{P}(\mathcal{V})$ is locally discriminable with any Schmidt-rank-r resource state $[w] \in \mathbb{P}\mathcal{W}$. If $n > d_1 d_2 - (d_1 - \min\{d_1, r\})(d_2 - \min\{d_2, r\})$, then a generic n-tuple of states in $\mathbb{P}(\mathcal{V})$ is not locally discriminable with any Schmidt-rank-r resource state $[w] \in \mathbb{P}\mathcal{W}$.*

Chapter Acknowledgments I thank William Slofstra for his help on numerous occasions with the geometric arguments in this work. I thank Debbie Leung, Michał Oszmaniec and John Watrous for helpful discussions.

Chapter 4

Entangled subspaces and entanglement witnesses

The contents of Chapters 3 and 4 are based on the preprint [91], which was written in collaboration with Nathaniel Johnston.

In this chapter, we introduce r-entangled subspaces, which naturally generalize previously studied spaces to higher multipartite entanglement. We use algebraic-geometric methods to determine the maximum dimension of an r-entangled subspace, and present novel explicit constructions of such spaces. We obtain similar results for symmetric and antisymmetric r-entangled subspaces, which correspond to entangled subspaces of bosonic and fermionic systems, respectively. We apply these results to entanglement witnesses.

Let m be a positive integer, let $\mathcal{V} = \bigotimes_{j=1}^{m}(\mathbb{C}^{d_j})$ be an m-partite vector space, and let $Y = \mathrm{Seg}(\mathbb{P}^{d_1-1} \times \cdots \times \mathbb{P}^{d_m-1})$ be the set of product states in $\mathbb{P}\mathcal{V}$. In Chapter 1, we defined an *r-entangled subspace* to be a projective linear subspace that avoids $\Sigma_r(Y)$. In bipartite space, the tensor rank is equal to the Schmidt rank, and $\Sigma_r(Y)$ is already closed. Wallach determined the maximum dimension of a 1-entangled subspace [126], and Bhat explicitly constructed a 1-entangled subspace of maximum dimension [18] (Parthasarathy proved Wallach's result independently in [100]). Cubitt et al. proved analogous results for r-entangled subspaces of bipartite space [45]. By Theorem 7, if X is a projective variety, then by Theorem 7 the minimum codimension of a projective linear subspace disjoint from X is $\dim(X) + 1$ (and furthermore, almost all subspaces of this codimension avoid X). It follows, from a standard upper bound on $\dim(\Sigma_r(Y))$, that there always exists an

r-entangled subspace of dimension

$$d_1 \cdots d_m - r \sum_{j=1}^{m} (d_j - 1) - r - 1,$$

whenever this quantity is non-negative. Furthermore, this is often the maximum dimension of such a subspace. Using this bound, we explicitly construct 2-entangled subspaces of maximum dimension in tripartite space with local (affine) dimensions $d_1, d_2 \in \{2, 3\}$ and $d_3 = 2$ (i.e. qubits and qutrits); and in quadripartite space with local (affine) dimensions $d_1 = d_2 = d_3 = d_4 = 2$ (i.e. all qubits). To show that these subspaces are indeed 2-entangled, we solve an equivalent ideal membership problem using the Macaulay2 software package [61, 92]. While ideal membership problems are notoriously intractable in general, our positive results reveal that this may not be the case for verifying r-entangled subspaces (at least for small r).

We also define r-entangled subspaces of the symmetric and antisymmetric spaces, which correspond to bosonic and fermionic entangled subspaces, respectively [60]. We explicitly construct maximal symmetric and antisymmetric r-entangled subspaces of bipartite space for arbitrary r, and of multipartite space for $r = 1$, which matches the cases of standard r-entangled subspace constructions presented in [18, 45].

It is known that, under various notions of entanglement, the maximum dimension of an entangled subspace is precisely the maximum number of negative eigenvalues of an entanglement witness [8, 76, 77]. The number of negative eigenvalues quantifies "how good" the witness is at detecting entanglement. We prove that this connection between subspaces and negative eigenvalues holds under a much more general notion of witness, including multipartite r-entanglement witnesses.

There are other types of entangled subspaces that have been studied in previous works: *non-positive partial transpose subspaces*, for which every mixed state supported on that subspace has non-positive partial transpose [76, 77]; *genuinely entangled subspaces*, for which every element is non-product with respect to every bipartition [45, 48]; and subspaces of bipartite space with high entropy of entanglement [67]. Entangled subspaces are connected to unextendible product bases, and have found applications, for example, in quantum error correction [59, 109] and quantum tomography [69].

4.1 Entangled subspaces of maximum dimension

Recall that, for a vector space $\mathcal{V} = \bigotimes_{j=1}^{m} \mathbb{C}^{d_j}$, we define a *(standard) r-entangled (projective linear) subspace* of $\mathbb{P}\mathcal{V}$ to be a projective linear subspace disjoint from $\Sigma_r(Y)$, where $Y =$

$\mathrm{Seg}(\mathbb{P}^{d_1-1} \times \cdots \times \mathbb{P}^{d_m-1})$. In other words, an r-entangled subspace is one that does not contain any states of border rank at most r. Similarly, we define a *symmetric r-entangled (projective linear) subspace* of $\mathbb{P}(S^m\mathbb{C}^d)$ to be a projective linear subspace disjoint from $\Sigma_r(\nu_m(\mathbb{P}^{d-1}))$, and an *antisymmetric r-entangled (projective linear) subspace* of $\mathbb{P}(\bigwedge^m \mathbb{C}^d)$ to be a projective linear subspace disjoint from $\Sigma_r(\mathrm{Gr}(m-1, \mathbb{P}^{d-1}))$. These three types of r-entangled subspaces correspond to entangled subspaces in systems of distinguishable, bosonic, and fermionic particles, respectively [60].

In Appendix B.1 we use Theorem 7 to determine the maximum dimensions of these entangled subspaces. In particular, Corollaries 79, 80, and 81 establish the maximum possible dimension of standard r-entangled, symmetric r-entangled, and antisymmetric r-entangled subspaces, respectively, and show that a generic subspaces of these dimensions are r-entangled. Despite the abundance of entangled subspaces of maximum dimension, explicit constructions of them are only known for standard r-entangled subspaces in the $r = 1$ case [18] and the $m = 2$ case [45]. In this section, we match these results for symmetric and antisymmetric r-entangled subspaces by providing explicit constructions in these two cases. We also construct maximal 2-entangled subspaces of $\mathbb{P}(\mathbb{C}^{d_1} \otimes \mathbb{C}^{d_2} \otimes \mathbb{C}^2)$ for $d_1, d_2 \in \{2, 3\}$, and of $\mathbb{P}(\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2)$.

Many of these constructions (and indeed, many of the already-known constructions of similar subspaces from [45, 69, 35]) are based on *totally non-singular matrices*, which are matrices with the property that all of their minors (i.e., determinants of square submatrices) are non-zero. For example, every Vandermonde matrix

$$\begin{bmatrix} 1 & \alpha_1 & \alpha_1^2 & \cdots & \alpha_1^{d-1} \\ 1 & \alpha_2 & \alpha_2^2 & \cdots & \alpha_2^{d-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha_d & \alpha_d^2 & \cdots & \alpha_d^{d-1} \end{bmatrix} \tag{4.1}$$

with $\alpha_i \neq \alpha_j \neq 0$ for all $i \neq j$ is totally non-singular (a slightly stronger property of these matrices, called total positivity, was proved in the case when $0 < \alpha_1 < \cdots < \alpha_d$ in [55], but the same proof works for total non-singularity in general). In fact, total non-singularity is a generic phenomenon: the set of totally non-singular matrices is open-dense.

The following result, which was proved in [45, Lemma 9], provides the reason that totally non-singular matrices are of use to us.

Lemma 15. *Let M be an $n \times n$ totally non-singular matrix with $n \geq k$, and let $v \in \mathbb{C}^n$ be a linear combination of k of the columns of M. Then v contains at least $n-k+1$ non-zero entries.*

4.1.1 Maximal symmetric r-entangled subspaces of bipartite space

To construct a symmetric r-entangled subspace of $\mathbb{P}(S^2(\mathbb{C}^d))$ attaining the bound (B.3), we first note that the isomorphism $\mathbb{C}^d \otimes \mathbb{C}^d \cong \mathrm{L}((\mathbb{C}^d)^*, \mathbb{C}^d)$ shows that it is equivalent to construct a projective linear subspace of symmetric $d \times d$ matrices of rank greater than r of dimension

$$\binom{d-r+1}{2} - 1.$$

We construct such a subspace by placing columns of totally non-singular matrices along the super- and sub-diagonals of those symmetric matrices. More specifically, for each $0 \leq i \leq d-r-1$ and $1 \leq j \leq d-r-i$, let M_j^i be the matrix that has the j-th column of some $(d-i) \times (d-i)$ totally non-singular matrix along its i-th super-diagonal (where for each fixed i, the same totally non-singular matrix is used for all j). For example, if $d = 6$, $r = 2$, $i = 1$, and we choose the totally non-singular matrix to be the Vandermonde matrix (4.1) with $\alpha_j = j$ for $1 \leq j \leq d-i = 5$, then

$$M_1^1 = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad M_2^1 = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 & 0 & 5 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad M_3^1 = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 4 & 0 & 0 & 0 \\ 0 & 0 & 0 & 9 & 0 & 0 \\ 0 & 0 & 0 & 0 & 16 & 0 \\ 0 & 0 & 0 & 0 & 0 & 25 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

We claim that the following set of symmetric matrices is a basis of a symmetric r-entangled subspace of $\mathbb{P}(\mathbb{C}^d \vee \mathbb{C}^d)$ (once we convert the matrices back into states in the canonical way):

$$B = \Big\{ \Big[M_j^i + \big(M_j^i\big)^\mathsf{T} \Big] : 0 \leq i \leq d-r-1, 1 \leq j \leq d-r-i \Big\}.$$

The fact that the set B is linearly independent (and thus a basis of its span) follows immediately from Lemma 15: every non-zero linear combination of those basis matrices has, for each $0 \leq i \leq d-r-1$, at least $(d-i)-(d-r-i)+1 = r+1 \geq 1$ non-zero entries along its i-th super-diagonal, and thus does not equal the zero matrix.

In fact, this argument also shows why this subspace is r-entangled: every non-zero diagonal of a matrix $M \in \mathrm{span}(B)$ contains at least $r+1$ non-zero entries, so there is an $(r+1) \times (r+1)$ submatrix of M that is upper triangular with non-zero diagonal entries (and is thus invertible), so $\mathrm{rank}(M) \geq r+1$. Since the rank of a matrix corresponds to the symmetric tensor rank in $\mathbb{P}(S^2(\mathbb{C}^d))$, the result follows.

All that remains is to count the number of vectors in B:

$$\begin{aligned}|B| = \sum_{i=0}^{d-r-1} \big(d-r-i\big) &= (d-r)^2 - \sum_{i=0}^{d-r-1} i \\ &= (d-r)^2 - \frac{1}{2}(d-r)(d-r-1) = \binom{d-r+1}{2}.\end{aligned}$$

Since the projective dimension of the subspace is $|B|-1$, this completes the proof.

4.1.2 Maximal antisymmetric r-entangled subspaces of bipartite space

To construct an antisymmetric r-entangled subspace of $\mathbb{P}(\bigwedge^2(\mathbb{C}^d))$ attaining the bound (B.5), we note that the isomorphism $\mathbb{C}^d \otimes \mathbb{C}^d \cong \mathrm{L}((\mathbb{C}^d)^*, \mathbb{C}^d)$ shows that it is equivalent to construct a projective linear subspace of antisymmetric $d \times d$ matrices with rank greater than $2r$ (not r) of dimension

$$\binom{d-2r}{2} - 1.$$

The construction of this subspace is identical to the symmetric construction from Section 4.1.1, except we omit the M_j^i matrices with non-zero entries on the main diagonal (i.e., the ones with $i = 0$), and we subtract in the lower-triangular portion of each matrix instead of adding. That is, a basis of this subspace is

$$B = \Big\{ \Big[M_j^i - \big(M_j^i\big)^{\mathsf{T}}\Big] : 1 \leq i \leq d-2r-1, 1 \leq j \leq d-2r-i \Big\},$$

so the dimension of this subspace is

$$|B| - 1 = \binom{d-2r+1}{2} - (d-2r) - 1 = \binom{d-2r}{2} - 1,$$

as desired.

4.1.3 Maximal symmetric 1-entangled subspaces of multipartite space

By Equation (B.2) and the subsequent discussion, the maximum dimension of a symmetric 1-entangled subspace of $\mathbb{P}(S^m\mathbb{C}^d)$ is

$$\binom{m+d-1}{m} - d - 1.$$

In this section, we construct such a subspace.

Consider the subspace spanned by the linearly independent set

$$\{[e_{a_1} \vee \cdots \vee e_{a_m}] : (a_1, \ldots, a_m) \in [d]^{\times n} \setminus \Delta_d^n\},$$

where $\Delta_d^n = \{(a, \ldots, a) : a \in [d]\}$. This subspace clearly has the correct dimension, and if

$$\sum_{a \in [d]^{\times n} \setminus \Delta_d^n} \alpha_a (e_{a_1} \vee \cdots \vee e_{a_m}) = x^{\otimes m}$$

for some $x = \sum_{b=1}^d \beta_b e_b$, then for each $b \in [d]$ the coefficient of $e_b^{\otimes m}$ in the expansion of $x^{\otimes m}$ is zero, and hence $\beta_b = 0$. It follows that $x = 0$, a contradiction.

4.1.4 Maximal antisymmetric 1-entangled subspaces of multipartite space

By Equation (B.4) and the subsequent discussion, the maximum dimension of an antisymmetric 1-entangled subspace of $\mathbb{P}(\bigwedge^m \mathbb{C}^d)$ is

$$\binom{d}{m} - m(d-m) - 2,$$

whenever $d \geq m$ (otherwise, $\bigwedge^m \mathbb{C}^d = 0$). We construct a subspace that attains this bound in a somewhat similar manner to the non-positive partial transpose subspaces constructed in [76, 77]. Let

$$J = \left\{ \binom{m}{2} + m - 1, \binom{m}{2} + m, \ldots, dm - \binom{m}{2} - 1, dm - \binom{m}{2} \right\}.$$

For each $s \in J$, let

$$I_s = \left\{ (a_1, \dots, a_m) \in [d]^m : 1 \leq a_1 < \cdots < a_m \leq d \quad \text{and} \quad \sum_{j=1}^{m} a_j = s \right\}.$$

Let

$$\mathbb{P}\mathcal{W} = \{[v] \in \mathbb{P}(\bigwedge^m \mathbb{C}^d) : \textstyle\sum_{a \in I_s} v_a = 0\},$$

where v_a is the coefficient of $e_{a_1} \wedge \cdots \wedge e_{a_m}$ in the expansion of v with respect to the standard basis

$$\{[e_{a_1} \wedge \cdots \wedge e_{a_m}] : 1 \leq a_1 < \cdots < a_m \leq d\}$$

of $\mathbb{P}(\bigwedge^m \mathbb{C}^d)$. We first observe that $\mathbb{P}\mathcal{W}$ has the correct dimension. Note that $|J| = m(d-m)+1$, and one linear constraint is placed on $\mathcal{W}$ for each $s \in J$, so

$$\dim(\mathbb{P}\mathcal{W}) = \dim(\mathbb{P}(\bigwedge^m \mathbb{C}^d)) - |J| = \tbinom{d}{m} - m(d-m) - 2,$$

as desired.

To complete the proof, we need only show that $\mathbb{P}\mathcal{W}$ does not contain any state of the form $[x_1 \wedge \cdots \wedge x_m]$. Proving this property is quite technical, so we begin by proving it in the special case $m = 2$ as a warm-up.

Let $[v] \in \mathbb{P}\mathcal{W}$ be arbitrary, and let

$$t = \min\{s \in J : v_a \neq 0 \quad \text{for some} \quad a \in I_s\}.$$

Let $a, b \in I_t$ be any two multi-indices for which $a \neq b$ and $v_a, v_b \neq 0$. Since $a_1 + a_2 = b_1 + b_2 = t$, there exists a permutation $\sigma \in S_2$ for which

$$b_{\sigma(1)} < a_{\sigma(1)} < a_{\sigma(2)} < b_{\sigma(2)}.$$

Under the inclusion $\wedge^2(\mathbb{C}^d) \subseteq \mathcal{V}_1 \otimes \mathcal{V}_2$, where $\mathcal{V}_1, \mathcal{V}_2 \cong \mathbb{C}^d$, we can regard v as an element of $\mathrm{L}(\mathcal{V}_2^*, \mathcal{V}_1)$. Under this identification, consider the 4×4 submatrix of v corresponding to the column index $\{b_{\sigma(1)}, a_{\sigma(1)}, a_{\sigma(2)}, b_{\sigma(2)}\}$ and row index $\{b_{\sigma(2)}, a_{\sigma(2)}, a_{\sigma(1)}, b_{\sigma(1)}\}$. It is straightforward to verify that this matrix takes the form

$$\begin{array}{cc} \begin{array}{cccc} b_{\sigma(2)} & a_{\sigma(2)} & a_{\sigma(1)} & b_{\sigma(1)} \end{array} & \\ \begin{bmatrix} \pm v_b & 0 & 0 & 0 \\ * & \pm v_a & 0 & 0 \\ * & * & \pm v_a & 0 \\ * & * & * & \pm v_b \end{bmatrix} & \begin{array}{c} b_{\sigma(1)} \\ a_{\sigma(1)} \\ a_{\sigma(2)} \\ b_{\sigma(2)} \end{array} \end{array},$$

where an asterisk $(*)$ denotes an entry we don't care about, and a ± denotes a sign we don't care about. It follows that the rank of v under this identification is at least four. This proves that $[v] \notin \mathrm{Gr}(1, \mathbb{P}^{d-1})$, since any such state has rank two under this identification. This completes the proof that $\mathbb{P}\mathcal{W}$ is antisymmetric 1-entangled in the case $m = 2$.

We now proceed to prove that $\mathbb{P}\mathcal{W}$ is antisymmetric 1-entangled for arbitrary m. As before, let $[v] \in \mathbb{P}\mathcal{W}$ be arbitrary, let

$$t = \min\{s \in J : v_a \neq 0 \quad \text{for some} \quad a \in I_s\},$$

and let $a, b \in I_t$ be any two multi-indices for which $a \neq b$ and $v_a, v_b \neq 0$. Let $A = \{a_1, \dots, a_m\}$, $B = \{b_1, \dots, b_m\}$, and let $S = A \cap B$. Let $q = m - |S|$, and let $s_{q+1} < \dots < s_m \in [d]$ be such that $\{s_{q+1}, \dots, s_m\} = S$.

Suppose toward contradiction that $[v] = [x_1 \wedge \dots \wedge x_m]$ for some $x_1, \dots, x_m \in \mathbb{C}^d$. Writing $x_{j,a} \in \mathbb{C}$ for the a'th coordinate of x_j in the standard basis, we may assume that for any $j \in [m]$ and $i \in \{q+1, \dots, m\}$, it holds that $(x_{j,s_i} \neq 0 \iff i = j)$. This can be observed by considering the element of $\pi^{-1}[v]$ that is in reduced row echelon form with respect to the re-ordered standard basis in which the elements $e_{s_{q+1}}, \dots, e_{s_{q+m}}$ come first. In this nice form, $v_a \neq 0$ implies

$$w_{A \setminus S} \neq 0,$$

where $w = x_1 \wedge \dots \wedge x_q$, so $[w] \in \mathrm{Gr}(q-1, \mathbb{P}^{d-1})$. The notation $w_{A\setminus S}$ denotes the coefficient of w with respect to the standard basis element $e_{c_1} \wedge \dots \wedge e_{c_q}$, where $c_1 < \dots < c_q \in [d]$ are such that $\{c_1, \dots, c_q\} = A \setminus S$. Similarly, $v_b \neq 0$ implies $w_{B\setminus S} \neq 0$. Let $Q = (A \cup B) \setminus S$, and let $\tilde{c}_1 < \dots < \tilde{c}_{2q} \in [d]$ be such that $\{\tilde{c}_1, \dots, \tilde{c}_{2q}\} = Q$. Regarding w as an element of $\mathrm{L}(\bigotimes_{k=2}^{q} \mathcal{V}_k^*, \mathcal{V}_1)$, consider the submatrix of w with column index Q and row index $\{\tilde{c}_{\hat{1}}, \dots, \tilde{c}_{\widehat{2q}}\} \subseteq [d]^{\times q-1}$, where

$$\tilde{c}_{\hat{j}} = \begin{cases} (A \setminus S) \setminus \{\tilde{c}_j\}, & \tilde{c}_j \in A \setminus S \\ (B \setminus S) \setminus \{\tilde{c}_j\}, & \tilde{c}_j \in B \setminus S \end{cases} \quad \text{for each} \quad j \in [2q].$$

Similarly to before, we have implicitly identified the set $(A \setminus S) \setminus \{\tilde{c}_j\}$ with a strictly increasing $q-1$-tuple, and likewise for B. By the minimality of t, it follows that this submatrix is lower-triangular with each diagonal entry equal to $\pm w_{A\setminus S}$ or $\pm w_{B\setminus S}$, and hence has rank $2q$. This is a contradiction, as any element of $\mathrm{Gr}(q-1, \mathbb{P}^{d-1})$ has rank q when viewed in this way. This completes the proof that $\mathbb{P}\mathcal{W}$ is antisymmetric 1-entangled.

4.1.5 Some maximal 2-entangled subspaces of multipartite space

There are a handful of higher-rank multipartite cases where it is straightforward to construct a 2-entangled subspace of $\mathbb{P}(\bigotimes_{j=1}^{m} \mathbb{C}^{d_j})$ of maximum dimension, and we consider these well known.

For example, if $m = 3$ and $d_1 = d_2 = d_3 = 2$ then the maximum dimension of a 2-entangled subspace is zero, and an explicit example is simply given by

$$[e_2 \otimes e_1 \otimes e_1 + e_1 \otimes e_2 \otimes e_1 + e_1 \otimes e_1 \otimes e_2].$$

Similarly, if $m = 3$ and $d_1 = 3$, $d_2 = d_3 = 2$ then we see that 2-entangled subspaces cannot be any larger than 1-dimensional, since we can identify $\mathbb{C}^3 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \cong \mathbb{C}^3 \otimes \mathbb{C}^4$ in the natural way and then use the bipartite bound from [45]. Furthermore, we can construct a 2-entangled subspace of this dimension just by using the fact that a state has border rank ≥ 3 if and only if it has flattening rank ≥ 3 [82, Theorem 5.1], so the explicit construction of a 2-entangled subspace of $\mathbb{P}(\mathbb{C}^3 \otimes \mathbb{C}^4)$ of dimension 1 from [45] also works in $\mathbb{P}(\mathbb{C}^3 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2)$.

In general, however, explicit constructions of r-entangled subspaces in this multipartite higher-rank setting are rather ad-hoc. Our contribution here is to present explicit examples of 2-entangled subspaces of $\mathbb{P}(\mathbb{C}^3 \otimes \mathbb{C}^3 \otimes \mathbb{C}^2)$ and of $\mathbb{P}(\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2)$ (i.e., the smallest non-trivial and previously unknown cases) that are 5-dimensional, which is maximal by Equation (B.1).

In $\mathbb{P}(\mathbb{C}^3 \otimes \mathbb{C}^3 \otimes \mathbb{C}^2)$, we first let $\delta, \epsilon, \theta, \kappa \in \mathbb{C}^4$ and then consider the following set of 6 states:

$$\begin{aligned} B = \Big\{ & \big[\big(e_1 \otimes e_1 + e_2 \otimes e_2 + e_3 \otimes e_3\big) \otimes e_1\big], \\ & \big[\big(e_1 \otimes e_1 + e_2 \otimes e_2 + e_3 \otimes e_3\big) \otimes e_2\big], \\ & [\delta_1(e_1 \otimes e_2 \otimes e_1) + \delta_2(e_2 \otimes e_1 \otimes e_1) + \delta_3(e_3 \otimes e_1 \otimes e_2) + \delta_4(e_1 \otimes e_3 \otimes e_2)], \\ & [\epsilon_1(e_1 \otimes e_2 \otimes e_1) + \epsilon_2(e_2 \otimes e_1 \otimes e_1) + \epsilon_3(e_3 \otimes e_1 \otimes e_2) + \epsilon_4(e_1 \otimes e_3 \otimes e_2)], \\ & [\theta_1(e_1 \otimes e_2 \otimes e_1) + \theta_2(e_3 \otimes e_1 \otimes e_1) + \theta_3(e_3 \otimes e_2 \otimes e_2) + \theta_4(e_2 \otimes e_3 \otimes e_2)], \\ & [\kappa_1(e_1 \otimes e_2 \otimes e_1) + \kappa_2(e_3 \otimes e_1 \otimes e_1) + \kappa_3(e_3 \otimes e_2 \otimes e_2) + \kappa_4(e_2 \otimes e_3 \otimes e_2)] \Big\} \end{aligned} \tag{4.2}$$

It is straightforward to show that if $\{\delta, \epsilon, \theta, \kappa\}$ is linearly independent then so is the set (4.2), so it is a basis of its span. Under the choice $\delta = (0, 1, 1, 1)$, $\epsilon = (1, 1, 2, 0)$, $\theta = (1, 1, 1, 0)$, and $\kappa = (0, 2, 1, 1)$; we have verified using the Macaulay2 software package [61] that $\mathrm{span}(B)$ is indeed 2-entangled (i.e. all of its members have border rank ≥ 3). Alternatively, in Appendix B.2 we provide an explicit proof that $\mathrm{span}(B)$ is 2-entangled.

The details of this proof are somewhat more complicated than in the other constructions that we have considered (we elaborate on this point at the end of the appendix). We have developed a heuristic algorithm in MATLAB which suggests that most choices of $\{\delta, \epsilon, \theta, \kappa\}$ produce a 2-entangled subspace. Our Macaulay2 and MATLAB codes are available on github.com [92].

In $\mathbb{P}(\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2)$, we similarly let $\phi, \psi \in \mathbb{C}^4$ and then consider the following set of 6 states:

$$\begin{aligned}
\Big\{ &[e_1 \otimes e_1 \otimes e_1 \otimes e_2 + e_1 \otimes e_2 \otimes e_2 \otimes e_1 + e_2 \otimes e_2 \otimes e_1 \otimes e_2], \\
&[e_1 \otimes e_1 \otimes e_2 \otimes e_1 + e_2 \otimes e_1 \otimes e_1 \otimes e_2 + e_2 \otimes e_2 \otimes e_2 \otimes e_1], \\
&[e_1 \otimes e_2 \otimes e_1 \otimes e_1 + e_2 \otimes e_1 \otimes e_1 \otimes e_1 + e_2 \otimes e_1 \otimes e_2 \otimes e_2 + e_2 \otimes e_2 \otimes e_1 \otimes e_2], \\
&[e_1 \otimes e_1 \otimes e_2 \otimes e_1 + e_1 \otimes e_2 \otimes e_1 \otimes e_1 + e_1 \otimes e_2 \otimes e_2 \otimes e_2 + e_2 \otimes e_1 \otimes e_2 \otimes e_2], \\
&\big[\phi_1(e_1^{\otimes 4}) + \phi_2(e_1 \otimes e_1 \otimes e_2 \otimes e_2) + \phi_3(e_2 \otimes e_2 \otimes e_1 \otimes e_1) + \phi_4(e_2^{\otimes 4})\big], \\
&\big[\psi_1(e_1^{\otimes 4}) + \psi_2(e_1 \otimes e_1 \otimes e_2 \otimes e_2) + \psi_3(e_2 \otimes e_2 \otimes e_1 \otimes e_1) + \psi_4(e_2^{\otimes 4})\big] \Big\}
\end{aligned} \tag{4.3}$$

It is straightforward to show that if $\{\phi, \psi\}$ is linearly independent then so is the set (4.3), so it is a basis of its span. Furthermore, if $\phi = (0, 1, 1, 1)$ and $\psi = (1, 2, 1, 0)$ then we have verified via Macaulay2 that its span is 2-entangled. As before, our numerics suggest that most choices of $\{\phi, \psi\}$ produce a 2-entangled subspace [92].

4.2 Witnesses and multipartite Schmidt number

This is the only section in which we consider mixed states, so we introduce notation and related definitions for them here. For a vector space $\mathcal{V}$ (which we always take to be over $\mathbb{C}$), let $\mathrm{D}(\mathcal{V}) \subseteq \mathrm{L}(\mathcal{V})$ be the set of positive semidefinite operators on $\mathcal{V}$ with trace one, which we refer to as the set of *density operators*. We identify the set of rank-one density operators with the set of pure quantum states under the natural bijection $vv^* \mapsto [v]$. We call a density operator of rank greater than one a *mixed state*. Let $\mathrm{Herm}(\mathcal{V}) \subseteq \mathrm{L}(\mathcal{V})$ be the set of Hermitian operators on $\mathcal{V}$. For a positive semidefinite operator $P \in \mathrm{Pos}(\mathcal{V})$, we define the *support* of P to be the image of P, denoted $\mathrm{Im}(P)$. For a subset $Z \subseteq \mathcal{V}$ we say that ρ is *supported on* Z if $\Im(P) \subseteq Z$. We say that $P, Q \in \mathrm{Pos}(\mathcal{V})$ have *orthogonal support* if $\langle v, u \rangle = 0$ for all $v \in \mathrm{Im}(P)$ and $u \in \mathrm{Im}(Q)$.

The notion of an operator that "witnesses" a particular property of a (mixed) quantum state is an important one in quantum information theory—it provides a way of demonstrating that property via a single measurement, without the need to have complete information about the state.

Definition 16. Given a subset $Z \subseteq \mathcal{V}$, a *not-Z witness* is a Hermitian matrix $W \in \mathrm{Herm}(\mathcal{V})$ for which

1. $\mathrm{Tr}(W\rho) \geq 0$ for all density operators $\rho \in \mathrm{D}(\mathcal{V})$ supported on Z, and
2. there exists some density operator $\sigma \in \mathrm{D}(\mathcal{V})$ such that $\mathrm{Tr}(W\sigma) < 0$.

For example, if $\hat{Y} \subseteq \mathbb{C}^{d_1} \otimes \mathbb{C}^{d_2}$ is (the affine cone over) the set of product states $Y = \mathrm{Seg}(\mathbb{P}^{d_1} \times \mathbb{P}^{d_2})$, then not-$\hat{Y}$ witnesses are called *entanglement witnesses* (since any mixed state σ for which $\mathrm{Tr}(W\sigma) < 0$ is then guaranteed to be entangled). Slightly more generally, if $\hat{\Sigma}_r(Y) \subseteq \mathbb{C}^{d_1} \otimes \mathbb{C}^{d_2}$ is (the affine cone over) the set of states of tensor rank (i.e., Schmidt rank) at most r, then the not-$\hat{\Sigma}_r(Y)$ witnesses are called r-entanglement witnesses [108] (or r-*block positive* [112]), and any mixed state σ for which $\mathrm{Tr}(W\sigma) < 0$ is said to have *Schmidt number* greater than r [121].

Even more generally, if $\hat{Y} \subseteq \bigotimes_{j=1}^{m} \mathbb{C}^{d_j}$ is (the affine cone over) the set of product states $Y = \mathrm{Seg}(\mathbb{P}^{d_1} \times \cdots \times \mathbb{P}^{d_m})$, and $\hat{\Sigma}_r(Y) \subseteq \bigotimes_{j=1}^{m} \mathbb{C}^{d_j}$ is (the affine cone over) the set of states with tensor rank at most r, then we still refer to not-$\hat{\Sigma}_r(Y)$ witnesses as r-entanglement witnesses, and the witnessed (mixed) states with support not in $\hat{\Sigma}_r(Y)$ are said to have (multipartite) Schmidt number larger than r [37]. However, a hiccup that occurs in this multipartite case that does not in the bipartite case is that, since the set of states with (multipartite) Schmidt number $\leq r$ is not closed, it is not true that every mixed state with Schmidt number $> r$ can be detected by some r-entanglement witness. Indeed, only the mixed states that are outside of the *closure* of that set of states can be detected, so these r-entanglement witnesses are better thought of as witnesses for border rank, not tensor rank.

It is well known, at least in the bipartite case, that entangled subspaces can be used to construct entanglement witnesses with the maximum number of negative eigenvalues (see [76, 77], for example). We now show that the same is true in the multipartite case, and even for not-Z witnesses in general, as long as Z is a Euclidean closed cone.

Theorem 17. *Suppose $Z \subseteq \mathcal{V}$ is a Euclidean closed cone. The maximum number of negative eigenvalues that a not-Z witness can have is equal to the maximum dimension of a linear subspace $\mathcal{W} \subseteq \mathcal{V}$ for which $\mathcal{W} \cap Z = \{0\}$.*

Proof. Throughout this proof, we let n be the maximum dimension of a linear subspace of $\mathcal{V}$ that trivially intersects Z.

Suppose $W \in \mathrm{Herm}\,(\mathcal{V})$ is a not-Z witness with s negative eigenvalues. To see that $s \leq n$, write $W = W_+ - W_-$ where W_+, W_- are positive semidefinite with orthogonal support (i.e., they come from the spectral decomposition of W). Letting $\mathcal{W} = \mathrm{Im}(W_-)$, we have $\dim(\mathcal{W}) = s$. Furthermore, for all $v \in \mathcal{W} \setminus \{0\}$, it holds that

$$\mathrm{Tr}\left(Wvv^*\right) = \mathrm{Tr}\left(W_+vv^* - W_-vv^*\right) = -\,\mathrm{Tr}(W_-vv^*) < 0,$$

from which it follows that $\mathcal{W}$ has trivial intersection with Z, so $s \leq n$.

Conversely, to see that there is a not-Z witness with n negative eigenvalues, let P be the orthogonal projection onto some linear subspace $\mathcal{W} \subseteq \mathcal{V}$ of dimension n that trivially intersects Z. Then $\mathrm{Tr}(Pvv^*) < \langle v, v\rangle$ for all $v \in Z$, and the fact that Z is a closed cone implies that there exists a real constant $0 < \epsilon < 1$ such that $\mathrm{Tr}(Pvv^*) \leq \epsilon\langle v, v\rangle$ for all $v \in Z$. If we define the matrix $W = \mathbb{1} - \frac{1}{\epsilon}P$, then W has exactly $\dim(\mathcal{W}) = n$ negative eigenvalues. Furthermore, $\mathrm{Tr}(Wvv^*) \geq 0$ for all $v \in Z$, so W is a not-Z witness. □

Combining Theorem 17 with Theorem 7, we get that for any projective variety $X \subseteq \mathbb{P}^D$, the maximum number of negative eigenvalues of a not-$\hat{X}$ witness is $D - \dim(X)$. We close this section by applying this result to entanglement witnesses.

Corollary 18. *The maximum number of negative eigenvalues that an r-entanglement witness can have is exactly*

$$d_1 \cdots d_m - \dim(\Sigma_r(Y)) - 1,$$

where $Y = \mathrm{Seg}(\mathbb{P}^{d_1-1} \times \cdots \times \mathbb{P}^{d_m-1})$.

For example, if $P \in \mathrm{Pos}(\mathbb{C}^3 \otimes \mathbb{C}^3 \otimes \mathbb{C}^2)$ is the orthogonal projection onto the 6-dimensional subspace spanned by (the affine cone over) the set B in (4.2), then there exists some scalar $\alpha > 1$ such that $W = I - \alpha P$ is a 2-entanglement witness (here we have set $\alpha = 1/\epsilon$, where $\epsilon < 1$ is as in the proof of Theorem 17), and furthermore there is no 2-entanglement witness with more negative eigenvalues than this one (which has 6). Importantly, this gives (at least in principle) a *measurable* way of showing that a multipartite state has large Schmidt number: if measuring a state $\rho \in \mathrm{Pos}(\mathbb{C}^3 \otimes \mathbb{C}^3 \otimes \mathbb{C}^2)$ produces an expectation value $\mathrm{Tr}(W\rho)$ that is strictly negative, then ρ must have Schmidt number at least 3.

Unfortunately, finding an explicit value of $\alpha > 1$ that actually works to make W a 2-entanglement witness is a non-trivial task, which we have not been able to solve analytically.

Numerics performed in MATLAB, however, strongly suggest that the optimal choice of α is approximately 1.0113.

A similar construction with $P \in \mathrm{Pos}(\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2)$ being the orthogonal projection onto the 6-dimensional subspace spanned by (the affine cone over) the set B in (4.3) gives a 2-entanglement witness $W = I - \alpha P \in \mathrm{L}(\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2)$ with 6 negative eigenvalues (for some value of $\alpha > 1$).

We can analogously define a symmetric (or antisymmetric) r-entanglement witness to be a Hermitian operator $W \in \mathrm{Herm}\left(S^m\mathbb{C}^d\right)$ (or $W \in \mathrm{Herm}\left(\bigwedge^m \mathbb{C}^d\right)$) such that $\mathrm{Tr}(W\rho) \geq 0$ for all $\rho \in \mathrm{D}\left(S^m\mathbb{C}^d\right)$ (or $\rho \in \mathrm{D}\left(\bigwedge^m \mathbb{C}^d\right)$) with rank $\leq r$, and for which there exists some $\sigma \in \mathrm{D}\left(S^m\mathbb{C}^d\right)$ (or $\sigma \in \mathrm{D}\left(\bigwedge^m \mathbb{C}^d\right)$) such that $\mathrm{Tr}(W\sigma) < 0$.

Corollary 19. *The maximum number of negative eigenvalues that a symmetric r-entanglement witness can have is exactly*

$$\binom{d-1+m}{m} - \dim(\Sigma_r(\nu_m(\mathbb{P}^{d-1}))) - 1.$$

Corollary 20. *The maximum number of negative eigenvalues that an antisymmetric r-entanglement witness can have is exactly*

$$\binom{d-1+m}{m} - \dim(\Sigma_r(\mathrm{Gr}(m-1, \mathbb{P}^{d-1}))) - 1.$$

For example, in the bipartite $m = 2$ case, we see that the maximum number of negative eigenvalues that symmetric and antisymmetric r-entanglement witnesses can have is

$$\binom{d-r+1}{2} \quad \text{and} \quad \binom{d-2r}{2},$$

respectively.

Chapter Acknowledgments I thank William Slofstra for his help on numerous occasions with the geometric arguments in this work. I thank Debbie Leung, Michał Oszmaniec and John Watrous for helpful discussions.

Chapter 5

New techniques for bounding stabilizer rank

The contents of this chapter are based on the preprint [95], which was written in collaboration with Vincent Steffan.

As we reviewed in Chapter 1, the stabilizer rank is a useful quantifier for the classical simulation cost of quantum circuits under the stabilizer formalism. In this chapter, we present number-theoretic and algebraic-geometric techniques for bounding the stabilizer rank. First, we refine a number-theoretic theorem of Moulton to exhibit an explicit sequence of product states with exponential stabilizer rank but constant approximate stabilizer rank, and to provide alternate (and simplified) proofs of the best-known asymptotic lower bounds on stabilizer rank and approximate stabilizer rank, up to a log factor. Second, we find the first non-trivial examples of quantum states with multiplicative stabilizer rank under the tensor product. Third, we use algebraic-geometric techniques to prove new bounds on the generic stabilizer rank.

5.1 Lower bounds on stabilizer rank and approximate stabilizer rank

In this section, we refine a number-theoretic theorem of Moulton, and use this to prove lower bounds on the stabilizer rank and approximate stabilizer rank. Recall the definitions of exponentially increasing subsequences and subset-sum representations given in Section 1.1.3. Moulton proved that any subset-sum representation of a q-tuple containing the subsequence $(1, 2, 4, \ldots, 2^{p-1})$ has length at least $p/\log_2 p$ [97]. In Section 5.1.1, we refine this result to prove that the same bound holds for any q-tuple containing an exponentially increasing subsequence of length p (Theorem 21). In Section 5.1.2 we use our refinement to prove that any quantum state whose coordinates contain an exponentially increasing subsequence of length p has stabilizer rank at least $p/(4\log_2 p)$ (Theorem 22). We then use this result to explicitly construct a sequence of product states of exponential stabilizer rank, to prove that $\chi([T^{\otimes m}]) \geq \frac{m+1}{4\log_2(m+1)}$, and to prove that $\chi([v^{\otimes m}]) = \Omega(m/\log_2 m)$ for any non-stabilizer qubit state $[v]$. In Section 5.1.3 we use our refinement of Moulton's theorem to prove that, for any non-stabilizer qubit state $[v]$, there exists $\delta > 0$ for which $\chi_\delta([v^{\otimes m}]) \geq \sqrt{m}/(2\log_2 m)$ for all $m \in \mathbb{N}$.

5.1.1 A refinement of Moulton's theorem

Theorem 21 (Refinement of Theorem 1 in [97])**.** *Let $2 \leq p \leq q$ be integers, and let $\alpha \in \mathbb{C}^q$ be a q-tuple of non-zero complex numbers. If α contains an exponentially increasing subsequence of length p, then any subset-sum representation of α has length at least $p/\log_2(p)$.*

Proof. It suffices to consider the case $p = q$ and $2|\alpha_i| \leq |\alpha_{i+1}|$ for all $i \in [q-1]$. Let $\beta \in \mathbb{C}^r$ be a subset-sum representation of α. Then for each $i \in [q]$, there exists $c_i \in \{0,1\}^r$ such that $\alpha_i = \beta^\intercal c_i$ Suppose that, for some $u_1, \ldots, u_q, v_1, \ldots, v_q \in \{0,1\}$, we have

$$\sum_{i=1}^{q} u_i c_i = \sum_{i=1}^{q} v_i c_i.$$

Applying $\beta^\intercal$ to both sides gives

$$\sum_{i=1}^{q} u_i \alpha_i = \sum_{i=1}^{q} v_i \alpha_i.$$

It follows that $u_i = v_i$ for all $i \in [q]$. Indeed, it suffices to prove that $|\alpha_{i+1}| > |\alpha_1 + \cdots + \alpha_i|$ for all $i \in [q-1]$, which in turn can be easily verified by an inductive argument. By assumption, $|\alpha_2| > |\alpha_1|$, and by induction,

$$|\alpha_1 + \cdots + \alpha_i| \leq |\alpha_1 + \cdots + \alpha_{i-1}| + |\alpha_i| < 2|\alpha_i| \leq |\alpha_{i+1}|.$$

The remainder of the proof is identical to that of [97]. There are at most $2^q - 1$ choices of $u_1, \ldots, u_q \in \{0,1\}$, excluding the case $u_1 = \cdots = u_q = 1$. For each of these choices, the sum $\sum_{i=1}^{q} u_i c_i$ can take one of $q^r - 1$ possible choices in $\{0, 1, \ldots, q-1\}^{\times r}$ (the choice $(q-1, q-1, \ldots, q-1)^\intercal$ is excluded since the u_i are not all equal to 1). Since each choice of $u_1, \ldots, u_q$ yields a different vector, we must have $q^r - 1 \geq 2^q - 1$, i.e. $r \geq q/\log_2(q)$. □

We emphasize that excluding the case $u_1 = \cdots = u_q = 1$ in the last paragraph of the above proof is just a choice that ends up slightly improving the bound on r over what we would get if we performed a similar count with this case included: With this case included, the sum $\sum_{i=1}^{q} u_i c_i$ can take one of $(q+1)^r$ possible choices in $\{0, 1, \ldots, q\}^{\times r}$, which results in a bound of $(q+1)^r \geq 2^q$, i.e. $r \geq q/\log_2(q+1)$.

5.1.2 Lower bounds on stabilizer rank

In this subsection we use Theorem 21 to prove lower bounds on stabilizer rank.

Theorem 22. *Let $p \geq 2$ be an integer, and let $[v] \in \mathbb{P}^{2^m - 1}$ be a quantum state. If the coordinates of v contain an exponentially increasing subsequence of length p, then $\chi([v]) \geq p/(4\log_2 p)$.*

Proof. Let $r = \chi([v])$, let $x_1, \ldots, x_p \in \mathbb{F}_2^m$ be such that $|v_{x_i}| \leq 2|v_{x_{i+1}}|$ for all $i \in [p-1]$, and let $\alpha = (v_{x_1}, \ldots, v_{x_p}) \in \mathbb{C}^p$. Without loss of generality, there exist complex numbers $\{\beta_a : a \in [r]\} \subseteq \mathbb{C}$ and stabilizer states $\{[\sigma_a] : a \in [r]\} \subseteq \mathrm{Stab}_m$ such that for all $a \in [r]$, every coordinate of σ_a is an element of $\{0, \pm 1, \pm i\}$, and $v = \sum_{a=1}^{r} \beta_a \sigma_a$. Let

$$S = (\sigma_1, \ldots, \sigma_r) \in \{0, \pm 1, \pm i\}^{\{0,1\}^m \times r}$$

and

$$\beta = (\beta_1, \ldots, \beta_r) \in \mathbb{C}^r,$$

so that $S\beta = v$. In particular, there exists a $p \times r$ submatrix T of S for which $T\beta = \alpha$. Let $T_1, T_2, T_3, T_4 \in \{0,1\}^{p \times r}$ be such that

$$T = T_1 - T_2 + i(T_3 - T_4).$$

Then

$$(T_1, T_2, T_3, T_4)(\beta, -\beta, i\beta, -i\beta)^\intercal = T\beta = \alpha,$$

so $(\beta, -\beta, i\beta, -i\beta)$ is a subset-sum representation of α. It follows from Theorem 21 that $4r \geq p/(\log_2 p)$. This completes the proof. □

Theorem 22 allows us to construct an explicit sequence of product states with exponential stabilizer rank and constant approximate stabilizer rank.

Corollary 23. *For any $m \in \mathbb{N}$, let*

$$v_m = \bigotimes_{i=1}^{m} (\mathrm{e}_0 + 2^{2^{i-1}} \mathrm{e}_1) \in \mathcal{U}_m.$$

Then $\chi([v_m]) \geq \frac{2^m}{4m}$ and for any constant $\delta > 0$, $\chi_\delta([v_m]) = \mathcal{O}(1)$.

Proof. For each m, the coordinates of v_m form an exponentially increasing sequence of length 2^m. It follows from Theorem 22 that $\chi([v_m]) \geq \frac{2^m}{4m}$. The bound $\chi_\delta([v_m]) = \mathcal{O}(1)$ is obtained as follows. For each $i \in \{0, 1, \ldots, 2^m - 1\}$, let

$$c_i = \sum_{j=0}^{i} 4^j = \frac{4^{i+1} - 1}{3},$$

and observe that for any positive integer k, the tensor $u_{m,k} \in \mathcal{U}_m$ obtained by setting all but the k largest coordinates of v_m to zero satisfies

$$\begin{aligned}
\left\| \frac{v_m}{\|v_m\|} - \frac{u_{m,k}}{\|u_{m,k}\|} \right\|^2 &= \left\| \frac{1}{\sqrt{c_{2^m-1}}} \sum_{i=0}^{2^m-1} 2^i \mathrm{e}_i - \frac{1}{\sqrt{c_{2^m-1} - c_{2^m-k-1}}} \sum_{i=2^m-k}^{2^m-1} 2^i \mathrm{e}_i \right\|^2 \\
&= \frac{c_{2^m-k-1}}{c_{2^m-1}} + \left(\frac{1}{\sqrt{c_{2^m-1}}} - \frac{1}{\sqrt{c_{2^m-1} - c_{2^m-k-1}}} \right)^2 (c_{2^m-1} - c_{2^m-k-1}) \\
&= \frac{c_{2^m-k-1}}{c_{2^m-1}} + \left[\sqrt{1 - \frac{c_{2^m-k-1}}{c_{2^m-1}}} - 1 \right]^2 \\
&= \frac{4^{1-k} - 4^{1-2^m}}{1 - 4^{1-2^m}} + \left[\sqrt{\frac{1 - 4^{1-k}}{1 - 4^{1-2^m}}} - 1 \right]^2, \qquad (5.1)
\end{aligned}$$

where we have re-indexed the computational basis of $\mathcal{U}_m$ as $\mathrm{e}_0, \ldots, \mathrm{e}_{2^m-1}$ for clarity in this proof. The quantity (5.1) can clearly be set to less than any constant $\delta > 0$ by appropriate choice of $k = \mathcal{O}(1)$. Since $\chi([u_{m,k}]) \leq k$, this completes the proof. □

Theorem 22 also implies the following lower bound on $\chi([T^{\otimes m}])$, and more generally, on $\chi([v^{\otimes m}])$ for any non-stabilizer qubit state $[v]$.

Corollary 24. *For any state $[v] \in \mathbb{P}^1$ that is not a stabilizer state, $\chi([v^{\otimes m}]) = \Omega(m/\log_2 m)$. In particular,*

$$\chi(\left[T^{\otimes m}\right]) \geq \frac{m+1}{4\log_2(m+1)}.$$

Proof. Since $[v]$ is not a stabilizer state, there exists $\alpha \in \mathbb{C}$ with $|\alpha| > 1$ for which $[w] := [\mathrm{e}_0 + \alpha \mathrm{e}_1]$ is in the $\mathcal{C}_1$-orbit of $[v]$. Indeed, without loss of generality, $[v] = [\mathrm{e}_0 + \beta \mathrm{e}_1]$ for some complex number $\beta \in \mathbb{C}^\times \setminus \{0, \pm 1, \pm i\}$, because $[v]$ is not a stabilizer state. If $|\beta| > 1$ we are done, and if $|\beta| < 1$ then let $[w] = X[v]$. If $|\beta| = 1$ then either $H[v]$ or $XH[v]$ must have the desired form. When $[v] = [T]$, we can take $[w] = XH[v] = [\mathrm{e}_0 + \frac{i}{\sqrt{2}-1}\mathrm{e}_1]$.

Since $|\alpha| > 1$, there exists $k \in \mathbb{N}$ for which $|\alpha|^k \geq 2$. (When $[v] = [T]$, we can take $k = 1$.) Now observe that the complex numbers $1, \alpha^k, \alpha^{2k}, \ldots, \alpha^{\lfloor m/k \rfloor k}$ all appear as coordinates of $v^{\otimes m}$. By Theorem 22, it follows that

$$\chi([v^{\otimes m}]) \geq \frac{\lfloor m/k \rfloor + 1}{4\log_2(\lfloor m/k \rfloor + 1)}.$$

This completes the proof. □

5.1.3 Lower bounds on approximate stabilizer rank

In this subsection, we use Theorem 21 to prove that, for any non-stabilizer qubit state $[v]$, there exists a constant $\delta > 0$ for which $\chi_\delta([v^{\otimes m}]) \geq \sqrt{m}/(2\log_2 m)$.

Theorem 25. *For any non-stabilizer qubit state $[v] \in \mathbb{P}^1$, there exists a constant $\delta > 0$ such that, for every integer $m \geq 2$,*

$$\chi_\delta([v^{\otimes m}]) \geq \frac{\sqrt{m}}{2\log_2 m}.$$

Proof. As in the proof of Corollary 24, since $[v]$ is not a stabilizer state, there exists $\alpha \in \mathbb{C}$ with $|\alpha| > 1$ for which $[\mathrm{e}_0 + \alpha \mathrm{e}_1] \in \mathcal{C}_1([v])$. Let $\beta = \frac{1}{\sqrt{1+|\alpha|^2}}$, $\gamma = \frac{\alpha}{\sqrt{1+|\alpha|^2}}$, and $w = \beta \mathrm{e}_0 + \gamma \mathrm{e}_1$, so that $w \in \mathcal{S}(\mathbb{C}^2)$ and $[w] = [\mathrm{e}_0 + \alpha \mathrm{e}_1]$. Since the (approximate) stabilizer rank is unchanged under $\mathcal{C}_m$, it suffices to lower bound the approximate stabilizer rank of $[w^{\otimes m}]$.

Let $k \in \mathbb{N}$ be the smallest integer for which $|\alpha|^k > 2$, and let $\lambda = \frac{2}{|\alpha|^k}$. If $[v] = [T]$, then we can take $\alpha = \frac{1}{\sqrt{2}-1}$ and $k = 1$. Let $u \in \mathcal{S}(\mathcal{U}_m)$ be a unit tensor, and let $S \subseteq [m]$ be the set of integers $p \in [m]$ for which

$$|\gamma|^2 m - k\lceil\sqrt{m}\rceil \leq p \leq |\gamma|^2 m + k\lceil\sqrt{m}\rceil, \tag{5.2}$$

and for all $x \in \mathbb{F}_2^m$ of Hamming weight $|x| = p$,

$$|w_x^{\otimes m} - u_x| \geq \left(\frac{1-\lambda}{1+\lambda}\right) |\beta|^{m-p}|\gamma|^p.$$

By DeMoivre-Laplace, there exists a constant $c > 0$ (which may depend on $|\alpha|$, but does not depend on m) for which

$$\left(\frac{1-\lambda}{1+\lambda}\right)^2 \binom{m}{p} |\beta^{m-p}\gamma^p|^2 \geq c/\sqrt{m}$$

for all $p \in S$ (see [56, Section VII, Theorem 1] or [101, Claim 4.6]). It follows that

$$\left\|w^{\otimes m} - u\right\|^2 \geq \sum_{p \in S} \left(\frac{1-\lambda}{1+\lambda}\right)^2 \binom{m}{p} |\beta^{m-p}\gamma^p|^2 \geq |S| \frac{c}{\sqrt{m}}.$$

Let $\delta = \sqrt{ck}$, and suppose that $\left\|w^{\otimes m} - u\right\| < \delta$. Then $|S| \leq k\sqrt{m}$. Let $P \subseteq [m]$ be the set of integers $p \in [m]$ that satisfy (5.2) and are not contained in S. Observe that $|P| \geq k\lceil\sqrt{m}\rceil$. By our definition of P, for each $p \in P$ there exists $x_p \in \mathbb{F}_2^m$ with $|x_p| = p$ and

$$\left|w_{x_p}^{\otimes m} - u_{x_p}\right| \leq \left(\frac{1-\lambda}{1+\lambda}\right) |\beta|^{m-p}|\gamma|^m. \tag{5.3}$$

For any $p, q \in P$ with $p < q$, it holds that

$$\begin{aligned}
\frac{|u_{x_q}|}{|u_{x_p}|} &= \frac{\left|u_{x_q} - w_{x_q}^{\otimes m} + w_{x_q}\right|}{\left|u_{x_p} - w_{x_p}^{\otimes m} + w_{x_p}^{\otimes m}\right|} \\
&\geq \frac{\left|w_{x_q}^{\otimes m}\right| - \left|u_{x_q} - w_{x_q}^{\otimes m}\right|}{\left|w_{x_p}^{\otimes m}\right| + \left|u_{x_p} - w_{x_p}^{\otimes m}\right|} \\
&\geq \lambda \frac{\left|\beta^{m-q}\gamma^q\right|}{\left|\beta^{m-p}\gamma^p\right|} \\
&= \lambda|\alpha|^{q-p},
\end{aligned}$$

where the first line is trivial, the second is the triangle inequality, the third follows from (5.3), and the fourth is obvious. In particular, if $q - p \geq k$, then $\frac{|u_{x_q}|}{|u_{x_p}|} \geq 2$. Since $|P| \geq k\lceil\sqrt{m}\rceil$, there exists a subset $Q \subseteq P$ of size $|Q| \geq \sqrt{m}$ for which $p - q \geq k$ for all $p, q \in Q$ with $p < q$. By Theorem 22,

$$\chi([u]) \geq \frac{\sqrt{m}}{2\log_2(m)}.$$

This completes the proof. □

We note that, using [56, Section VII, Theorem 1] and a similar proof as above, slight improvements to the above bound can be obtained if δ is allowed to decay in m.

5.2 States with multiplicative stabilizer rank under the tensor product

It is a standard fact that the stabilizer rank is *sub-multiplicative* under the tensor product, i.e. $\chi([v \otimes v]) \leq \chi([v])^2$ for any quantum state $[v]$ [102, Section 2.1.3]. In [102, Section 4.4] it was remarked that there are no known examples of quantum states $[v]$ of stabilizer rank greater than one for which equality holds. In this section, we explicitly construct two-qubit states $[v]$ for which $\chi([v]) = 2$ and $\chi([v \otimes v]) = 4$. Note that this is the smallest non-trivial example of multiplicative stabilizer rank that one can hope for, since for any single-qubit state $[v] = [\mathrm{e}_0 + \alpha\mathrm{e}_1] \in \mathbb{P}^1$, it holds that $[v^{\otimes 2}] = [\mathrm{e}_{00} + \alpha(\mathrm{e}_{01} + \mathrm{e}_{10}) + \alpha^2\mathrm{e}_{11}]$, which has stabilizer rank at most 3.

Theorem 26. *Let*

$$v_\alpha = \mathrm{e}_{00} + \alpha(\mathrm{e}_{01} + \mathrm{e}_{10}) \in (\mathbb{C}^2)^{\otimes 2}$$

when $\alpha \in \mathbb{C}^\times$ is a non-zero complex number. Then $\chi([v_\alpha]) = 2$, and for all but finitely many α it holds that $\chi([v_\alpha^{\otimes 2}]) = 4$. In particular, $\chi([v_\alpha^{\otimes 2}]) = 4$ for any α that is transcendental over $\mathbb{Q}$.

Proof. The fact that $\chi([v_\alpha]) = 2$ is obvious, so it suffices to prove that $\chi([v_\alpha^{\otimes 2}]) = 4$. Note that, since the imaginary unit i is algebraic over $\mathbb{Q}$, α is transcendental over $\mathbb{Q}$ if and only if it is transcendental over $\mathbb{Q}(i)$. Since there are only finitely many stabilizer states, it suffices to prove that for any set of three stabilizer states $\{[\sigma_1], [\sigma_2], [\sigma_3]\} \subseteq \mathrm{Stab}_2$, there are at most finitely many α (and no α which are transcendental over $\mathbb{Q}$) for which

$$v_\alpha^{\otimes 2} \in \mathrm{span}\{\sigma_1, \sigma_2, \sigma_3\}. \tag{5.4}$$

For each $j \in [3]$, we may assume that

$$\sigma_j = \sum_{x \in A_j} i^{l_j(x)} \cdot (-1)^{q_j(x)} \cdot \mathrm{e}_x$$

for some affine linear subspace $A_j \subseteq \mathbb{F}_2^2$, linear functional l_j, and quadratic form q_j, as in (2.8). Let $S = (\sigma_1, \sigma_2, \sigma_3) \in \mathbb{Q}(i)^{\{0,1\}^2 \times 3}$. For every α that satisfies (5.4), there exist complex numbers $\beta_\alpha, \gamma_\alpha, \lambda_\alpha \in \mathbb{C}$ for which $v_\alpha^{\otimes 2} = \beta_\alpha \sigma_1 + \gamma_\alpha \sigma_2 + \lambda_\alpha \sigma_3$, and a 3×3 submatrix T of S for which

$$\begin{bmatrix} T & | - \mathbb{1}_3 \end{bmatrix} \begin{bmatrix} \beta_\alpha \\ \gamma_\alpha \\ \lambda_\alpha \\ 1 \\ \alpha \\ \alpha^2 \end{bmatrix} = 0. \tag{5.5}$$

Since there are only finitely many 3×3 submatrices of S, it suffices to prove that for any choice of 3×3 submatrix T, only finitely many α (and no transcendental α) can satisfy (5.5).

If T is singular, then there exists a non-trivial $\mathbb{Q}(i)$-linear combination of $\{1, \alpha, \alpha^2\}$ that equals zero, which at most finitely many α will satisfy (and in particular, no transcendental α will satisfy).

If T is nonsingular, then by applying T^{-1} to both sides we can find polynomials $f, g, h \in \mathbb{Q}(i)[x]$ for which $\beta_\alpha = f(\alpha), \gamma_\alpha = g(\alpha)$, and $\lambda_\alpha = h(\alpha)$ for every choice of α that satisfies (5.5).

Let $R = \{00, 01, 10\}^{\times 2} \subseteq \mathbb{F}_2^4$ be the subset of bitstrings corresponding to non-zero coordinates of $v_\alpha^{\otimes 2}$. If there exists $x \in A_i \setminus R$ for some $i \in [3]$, then there exists a non-trivial $\mathbb{Q}(i)$-linear combination of $\{f, g, h\}$ which equals zero, which at most finitely many α (and no transcendental α) will satisfy. So we may assume $A_i \subseteq R$ for all $i \in [3]$.

Let $R_0 = \{0000\}$, $R_1 = \{0001, 0010, 0100, 1000\}$, and $R_2 = \{0101, 0110, 1001, 1010\}$, so $R = R_0 \cup R_1 \cup R_2$. If there exists $x \in A_1 \cap R_2$ and $y \in R_2 \setminus A_1$, then there exist linear combinations

$$\begin{aligned}\star f(\alpha) + \cdot\, g(\alpha) + \cdot h(\alpha) &= \alpha^2\\ \cdot\, g(\alpha) + \cdot h(\alpha) &= \alpha^2,\end{aligned}$$

where each $\cdot \in \mathbb{C}$ denotes an arbitrary complex number we do not bother to name, and each $\star \in \mathbb{C}^\times$ denotes an arbitrary non-zero complex number we don't bother to name. Subtracting these equations yields a non-trivial $\mathbb{Q}(i)$-linear combination of $\{f, g, h\}$ that equals zero, which at most finitely many α (and no transcendental α) will satisfy. By symmetry, for all $i \in [3]$, either $R_2 \subseteq A_i$ or $R_2 \cap A_i = \emptyset$. If $R_2 \subseteq A_i$, then $R_2 = A_i$ because this is the only affine linear subspace contained in R that contains R_2. Without loss of generality, we may assume $A_1 = R_2$. If $A_2 = R_2$, then $A_3 = R_0 \cup R_1$, which is not an affine linear subspace, a contradiction. Otherwise, $A_2 \cup A_3 = R_0 \cup R_1$, which is also a contradiction, as there are no affine linear subspaces of size greater than 2 contained in $R_0 \cup R_1$. This completes the proof. □

5.3 Generic stabilizer rank

Fact 27. For any positive integer m, all but finitely many qubit states $[v] \in \mathbb{P}^1$ maximize $\chi([v^{\otimes m}])$. Similarly, all but finitely many qubit states $[v] \in \mathbb{P}^1$ maximize $\chi^{\mathbb{R}}([v^{\otimes m}])$.

This fact, the proof of which we defer until after introducing this section, motivates us to define the m-th *generic stabilizer rank* as

$$\chi_m = \max_{[v]\in\mathbb{P}^1} \chi([v^{\otimes m}]),$$

and the m-th *generic real stabilizer rank* as

$$\chi_m^{\mathbb{R}} = \max_{[v]\in\mathbb{P}^1} \chi^{\mathbb{R}}([v^{\otimes m}]).$$

In Proposition 30 we prove that χ_m and $\chi_m^{\mathbb{R}}$ differ by at most a constant factor of two. This motivates the study of $\chi_m^{\mathbb{R}}$, since it has the same scaling as χ_m, but may be easier to work with. Note that χ_m upper bounds $\chi([T^{\otimes m}])$.

In this section, we prove bounds on χ_m and $\chi_m^{\mathbb{R}}$, and observe some useful reductions for studying these quantities. In Proposition 28 we use Fact 27 to obtain a modest improvement of the best-known upper bound on χ_m. In Proposition 29 we prove a useful reduction for studying χ_m (respectively, $\chi_m^{\mathbb{R}}$), namely, that there must exist a single set of χ_m stabilizer states (resp. $\chi_m^{\mathbb{R}}$ real stabilizer states) that span the symmetric subspace. As an immediate consequence, there exists a single set of χ_m stabilizer states (resp. $\chi_m^{\mathbb{R}}$ real stabilizer states) such that, for every qubit state $[v]$, the state $[v^{\otimes m}]$ lies in the span of the set. In Proposition 31 we introduce a technique for upper bounding χ_m when upper bounds on $\chi([v^{\otimes m}])$ are known for sufficiently many (linear in m) qubit states $[v]$, and obtain similar results for $\chi_m^{\mathbb{R}}$. In Proposition 32 we refine Fact 27 to obtain quantitative upper bounds on the (finite) number of states of sub-generic stabilizer rank and sub-generic real stabilizer rank.

Before proceeding, we prove Fact 27. We refer the reader to [65] for the basic algebraic-geometric definitions and arguments used in this proof, and elsewhere in this section. In this proof and in others in this section, we only prove the statement for χ_m, as the proof for $\chi_m^{\mathbb{R}}$ is essentially identical.

Proof of Fact 27. We prove only the statement for χ_m. The set $\nu_m(\mathbb{P}^1)$ forms a 1-dimensional irreducible projective variety, so its intersection with any other projective variety is either empty, zero-dimensional, or equal to $\nu_m(\mathbb{P}^1)$. Hence, for any positive integer r, $\Sigma_r(\mathrm{Stab}_m) \cap \nu_m(\mathbb{P}^1)$ is either a finite set of points, or $\Sigma_r(\mathrm{Stab}_m) \supseteq \nu_m(\mathbb{P}^1)$. This completes the proof. □

Fact 27 immediately implies the following slight improvement of the upper bound $\chi_m = \mathcal{O}((m+1)2^{m/2})$ obtained in [103, Theorem 3]:

Proposition 28. *$\chi_m = \mathcal{O}(2^{m/2})$.*

Proof. This follows directly from our Fact 27 and [103, Theorem 2], which states that equitorial states (an infinite family of states) have stabilizer rank $\mathcal{O}(2^{m/2})$. □

We next prove that the affine cone over a single set of χ_m stabilizer states (respectively, $\chi_m^{\mathbb{R}}$ real stabilizer states) spans $S^m(\mathbb{C}^2)$.

Proposition 29. *For any positive integer m, there exists a single set of stabilizer states $\{[\sigma_1], \dots, [\sigma_{\chi_m}]\} \subseteq \mathrm{Stab}_m$ for which*

$$S^m(\mathbb{C}^2) \subseteq \mathrm{span}\{\sigma_1, \dots, \sigma_{\chi_m}\}.$$

Similarly, there exists a single set of real stabilizer states $\{[\sigma_1], \dots, [\sigma_{\chi_m^{\mathbb{R}}}]\} \subseteq \mathrm{Stab}_m^{\mathbb{R}}$ for which

$$S^m(\mathbb{C}^2) \subseteq \mathrm{span}\{\sigma_1, \dots, \sigma_{\chi_m^{\mathbb{R}}}\}.$$

Since $\dim(S^m(\mathbb{C}^2)) = m+1$, it follows from Proposition 29 that $\chi_m \geq m+1$. We note that a similar proof as below can be used to show that a generic product state of the form $[v_1 \otimes \cdots \otimes v_m]$, for qubit states $v_i \in \mathbb{P}^1$, has stabilizer rank 2^m, where in this context we define *generic* in the same algebraic-geometric sense as in [91]. In Corollary 23 we have constructed an explicit sequence of product states of stabilizer rank at least $\frac{2^m}{4m}$, which is near maximal.

Proof of Proposition 29. We prove only the statement for χ_m. Let $r = \chi_m$, so $\Sigma_r(\mathrm{Stab}_m) \supseteq \nu_m(\mathbb{P}^1)$. Since $\nu_m(\mathbb{P}^1)$ is irreducible, and $\Sigma_r(\mathrm{Stab}_m)$ is reducible into a finite union of projective $(r-1)$-dimensional linear subspaces, one of these subspaces must contain $\nu_m(\mathbb{P}^1)$. To complete the proof, recall that the affine cone over $\mathrm{span}(\nu_m(\mathbb{P}^1))$ is $S^m(\mathbb{C}^2)$. □

The next proposition shows that $\chi_m^{\mathbb{R}}$ and χ_m differ by at most a constant factor of two, which motivates the study of $\chi_m^{\mathbb{R}}$.

Proposition 30. *For any positive integer m, it holds that $\chi_m \leq \chi_m^{\mathbb{R}} \leq 2\chi_m$.*

Proof. The first inequality is obvious, so it suffices to prove $\chi_m^{\mathbb{R}} \leq 2\chi_m$. Since the tensors $\sum_{|x|=k} \mathrm{e}_x$ for $k \in [m]$ form a basis for $S^m(\mathbb{C}^2)$, then by Proposition 29, χ_m is the minimum number for which there exists a set of stabilizer states $\{[\sigma_1], \ldots, [\sigma_{\chi_m}]\} \subseteq \mathrm{Stab}_m$ such that

$$\sum_{\substack{x \in \mathbb{F}_2^m \\ |x|=k}} \mathrm{e}_x \in \mathrm{span}\{\sigma_1, \ldots, \sigma_{\chi_m}\} \tag{5.6}$$

for all $k \in \{0, 1, \ldots, m\}$. Without loss of generality, for each $j \in [\chi_m]$ it holds that

$$\sigma_j = \sum_{x \in A_j} i^{l_j(x)} \cdot (-1)^{q_j(x)} \cdot \mathrm{e}_x$$

for some affine linear subspace $A_j \subseteq \mathbb{F}_2^m$, linear functional l_j, and quadratic form q_j (see (2.8)). For each $j \in [\chi_m]$, we define ρ_j to be the tensor

$$\rho_j = \sum_{x \in A_j} (-1)^{q_j(x) + l_j(x)} \cdot \mathrm{e}_x,$$

and observe that for any real numbers $a, b \in \mathbb{R}$, it holds that

$$\mathrm{Re}((a+ib)\sigma_j) = (a-b)\mathrm{Re}(\sigma_j) + b\rho_j, \tag{5.7}$$

where $\mathrm{Re}(\cdot)$ is the real part with respect to the computational basis. For each $k \in \{0, 1, \dots, m\}$, there exist real numbers $a_1, \dots, a_{\chi_m}, b_1, \dots, b_{\chi_m} \in \mathbb{R}$ for which

$$\begin{aligned}\sum_{|x|=k} \mathrm{e}_x &= \sum_{j\in[\chi_m]} (a_j + ib_j)\sigma_i \\ &= \mathrm{Re}\left(\sum_{j\in[\chi_m]} (a_j + ib_j)\sigma_j\right) \\ &= \sum_{j\in[\chi_m]} \big((a_j - b_j)\mathrm{Re}(\sigma_j) + b\rho_j\big),\end{aligned}$$

where the first line follows from (5.6), the second follows from the fact that the tensor is already real, and the third follows from (5.7). Hence,

$$S^m(\mathbb{C}^2) \in \mathrm{span}_{\mathbb{C}}\{\mathrm{Re}(\sigma_1), \dots, \mathrm{Re}(\sigma_{\chi_m}), \rho_1, \dots, \rho_{\chi_m}\}.$$

Since $[\mathrm{Re}(\sigma_j)]$ and $[\rho_j]$ are clearly real stabilizer states for all $j \in [\chi_m]$, it follows that $\chi_m^{\mathbb{R}} \leq 2\chi_m$. □

Using similar techniques as in the proof of Proposition 29, it is straightforward to show that $\chi([v]) \leq \chi^{\mathbb{R}}([v]) \leq 2\chi([v])$ for any real quantum state $[v] \in \mathbb{P}^{2^m-1}$.

The next proposition tells us that upper bounds on the stabilizer ranks (respectively, real stabilizer ranks) of any set of $m+1$ states of the form $[v^{\otimes m}]$ implies an upper bound on χ_m (resp. $\chi_m^{\mathbb{R}}$).

Proposition 31. *Let r be a positive integer. If there exists a set of $m+1$ distinct qubit states $\{[v_1], \dots, [v_{m+1}]\} \subseteq \mathbb{P}^1$ with $\chi([v_i^{\otimes m}]) \leq r$ for all $i \in [m+1]$, then $\chi_m \leq r(m+1)$. Similarly, if there exists a set of $m+1$ distinct qubit states $\{[v_1], \dots, [v_{m+1}]\} \subseteq \mathbb{P}^1$ with $\chi^{\mathbb{R}}([v_i^{\otimes m}]) \leq r$ for all $i \in [m+1]$, then $\chi_m^{\mathbb{R}} \leq r(m+1)$.*

Proof. We prove only the statement for χ_m. Since $\chi([v_i^{\otimes m}]) \leq r$ for all $i \in [m+1]$, then by taking the span of the union of all stabilizer states appearing in the decompositions of each $[v_i^{\otimes m}]$, we have $\mathrm{span}\{[v_1^{\otimes m}], \dots, [v_{m+1}^{\otimes m}]\} \subseteq \Sigma_{r(m+1)}(\mathrm{Stab}_m)$. The proposition follows from the fact that the affine cone over any set of $m+1$ distinct states of this form spans $S^m(\mathbb{C}^2)$, so $\hat{\Sigma}_{r(m+1)}(\mathrm{Stab}_m) \supseteq S^2(\mathbb{C}^2)$ (see [63, Proposition 3.1] or [94, Corollary 18]). □

We close this section by refining Fact 27 to give a quantitative upper bound on the number of states of sub-generic rank.

Proposition 32. *There are at most*

$$m\binom{|\mathrm{Stab}_m|}{\chi_m - 1}$$

states $[v] \in \mathbb{P}^1$ *for which* $\chi([v^{\otimes m}]) < \chi_m$. *Similarly, there are at most*

$$m\binom{\left|\mathrm{Stab}_m^{\mathbb{R}}\right|}{\chi_m^{\mathbb{R}} - 1}$$

states $[v] \in \mathbb{P}^1$ *for which* $\chi^{\mathbb{R}}([v^{\otimes m}]) < \chi_m^{\mathbb{R}}$.

We have computed the quantities $|\mathrm{Stab}_m|$ and $\left|\mathrm{Stab}_m^{\mathbb{R}}\right|$ at the end of Section 2.2.

Proof. We prove only the statement for χ_m. Note that any set of $m+1$ distinct tensors in $\hat{\nu}_1(\mathbb{P}^1)$ (i.e. tensors of the form $v^{\otimes m}$) span $S^m(\mathbb{C}^2)$ (see [63, Proposition 3.1] or [94, Corollary 18]), so any set of $\chi_m - 1$ linearly independent stabilizer states can contain at most m distinct elements of $\hat{\nu}_1(\mathbb{P}^1)$ in their span. Since there are at most

$$\binom{|\mathrm{Stab}_m|}{\chi_m - 1}$$

distinct sets of $\chi_m - 1$ linearly independent stabilizer states, the bound follows. □

Chapter Acknowledgments I thank Gerry Myerson for pointing me to the work of Moulton [97]. I thank Kieran Mastel and William Slofstra for helpful discussions.

Chapter 6

On decomposable correlation matrices

The contents of this chapter are based on the journal article [90], which was written by myself.

Recall that a *correlation matrix* is a positive semidefinite matrix with ones on the diagonal, and that a *d-decomposable* correlation matrix is one that can be written as a Schur product of correlation matrices of rank at most d (to our knowledge, the reference [90] is the first to define and study d-decomposable correlation matrices).

In this chapter, we obtain several results on decomposable correlation matrices. We prove that for all $d \geq 2$, every $(d+1) \times (d+1)$ correlation matrix is d-decomposable. We construct $(2d+1) \times (2d+1)$ correlation matrices that are not d-decomposable, and prove as a corollary that $\mathrm{Cor}_d(\mathbb{C}^n)$ is not convex when $d \geq 1$ and $n \geq 2d+1$. One question this leaves open is whether every 4×4 correlation matrix is 2-decomposable, which we make partial progress towards resolving. Finally, we observe an application of non-d-decomposability in an entanglement detection scenario. See Chapter 1 for an introduction and relevant motivations in quantum information theory for the d-decomposability question.

6.1 Preliminaries on correlation matrices

In this section, we review some elementary facts about correlation matrices. It is straightforward to verify that a matrix $P \in M_n^{\mathbb{C}}$ is contained in $\mathrm{Cor}(\mathbb{C}^n)$ if and only if $P = A^*A$

for some $s \times n$ matrix $A \in M^{\mathbb{C}}_{s\times n}$ (and positive integer s), the columns of which form unit vectors. We say P is *generated by* some set of unit vectors $\{v_a : a \in [n]\} \subset \mathcal{S}(\mathbb{C}^s)$ if these vectors can be chosen as the columns of A. Note that $P(a,b) = \langle v_a, v_b \rangle$, so P is the matrix of inner products (i.e. the *Gram matrix*) of any generating set of unit vectors. Recall that two sets of unit vectors $\{v_a : a \in [n]\} \subset \mathcal{S}(\mathbb{C}^s)$ and $\{u_a : a \in [n]\} \subset \mathcal{S}(\mathbb{C}^s)$ with $s \leq t$ generate the same correlation matrix if and only if there exists an isometry $U \in \mathrm{U}(\mathbb{C}^s, \mathbb{C}^t)$ such that $Uv_a = u_a$ for all $a \in [n]$. This property follows from the standard result that two operators $A \in \mathrm{L}(\mathbb{C}^n, \mathbb{C}^s)$ and $B \in \mathrm{L}(\mathbb{C}^n, \mathbb{C}^t)$ satisfy $A^*A = B^*B$ if and only if $B = UA$ for some isometry $U \in \mathrm{U}(\mathbb{C}^s, \mathbb{C}^t)$. Note that by linearity, the linear dependence of every generating set is the same. The following proposition reveals a straightforward yet important connection between the Schur product of correlation matrices and the tensor product of the unit vectors in their generating sets, which we state without proof.

Proposition 33. *A correlation matrix $P \in \mathrm{Cor}(\mathbb{C}^n)$ is r-decomposable if and only if there exists a positive integer m, complex vector spaces $\mathcal{V}_1, \ldots, \mathcal{V}_m$ with $\dim \mathcal{V}_i \leq r$ for all $i \in [m]$, and a set of unit product vectors $\{x_a : a \in [n]\} \subset \mathcal{S}\mathrm{Prod}(\mathcal{X}_1 : \cdots : \mathcal{X}_m)$ that generate P.*

6.2 Positive results for d-decomposability

In this section, we prove that for all $d \geq 2$, every $(d+1) \times (d+1)$ correlation matrix is d-decomposable.

Theorem 34. *For any integers $r \geq 2$ and $n \leq d+1$, $\mathrm{Cor}_d(\mathbb{C}^n) = \mathrm{Cor}(\mathbb{C}^n)$. More generally, let $\mathcal{V}$ be a complex Euclidean space and $P \in \mathrm{Cor}(\mathcal{V})$ be a correlation matrix. If* $\mathrm{rank}(P) \geq 3$ *and P is generated by a set of unit vectors that contains a vector that is not in the span of the rest, then P is* $(\mathrm{rank}(P) - 1)$*-decomposable.*

Note that a correlation matrix $P \in \mathrm{Cor}(\mathbb{C}^n)$ with $\mathrm{rank}(P) \leq d$ is trivially contained in $\mathrm{Cor}_d(\mathbb{C}^n)$, so the only non-trivial part of the first sentence is that if an $n \times n$ correlation matrix P has full rank (i.e. it is generated by a linearly independent set of unit vectors), then it is $(n-1)$-decomposable. The second sentence generalizes this statement to say that if P is generated by a set of unit vectors that contains a vector that is not in the span of the rest, then it is $(\mathrm{rank}(P) - 1)$-decomposable.

Proof. We prove the second (more general) statement. Let $\{v_a : a \in [n]\}$ be a set of unit vectors that generate P such that

$$v_c \notin \mathrm{span}\{v_a : a \in [n] \setminus \{c\}\}$$

for some index $c \in [n]$.

If v_c is orthogonal to every other vector, then the construction is easy: the set of vectors with each v_a replaced by $v_a \otimes \mathrm{e}_1$ for $a \neq c$, and v_c replaced by $v_c' \otimes e_2$ for any unit vector $v_c' \in \operatorname{span}\{v_a : a \in [n] \setminus \{c\}\}$ generates P. This is a $(\operatorname{rank}(P) - 1)$-decomposition of P, since

$$\dim \operatorname{span}\{v_a : a \in [n] \setminus \{c\}\} = \operatorname{rank}(P) - 1$$

and

$$\dim \operatorname{span}\{\mathrm{e}_1, \mathrm{e}_2\} = 2 \leq \operatorname{rank}(P) - 1.$$

If v_c is not orthogonal to every other vector, then define

$$\Pi := \operatorname{Proj}\left(\operatorname{span}\{v_a : a \in [n] \setminus \{c\}\}\right),$$

and define two correlation matrices R and Q as

$$R(a, b) = \frac{\langle v_a, \Pi v_b \rangle}{\|\Pi v_a\| \|\Pi v_b\|}$$

and

$$Q(a, b) = \begin{cases} \|\Pi v_c\|, & a \neq b \text{ and } c \in \{a, b\} \\ 1, & \text{otherwise.} \end{cases}$$

It is straightforward to verify that $P = R \odot Q$. Indeed, for $c \notin \{a, b\}$,

$$(R \odot Q)(a, b) = \frac{\langle v_a, \Pi v_b \rangle}{\|\Pi v_a\| \|\Pi v_b\|} = \langle v_a, v_b \rangle.$$

Otherwise,

$$(R \odot Q)(a, c) = \frac{\langle v_a, \Pi v_c \rangle}{\|\Pi v_a\| \|\Pi v_c\|} \|\Pi v_c\| = \langle v_a, v_c \rangle,$$

and similarly, $(R \odot Q)(c, a) = \langle v_c, v_a \rangle$. The correlation matrix R has $\operatorname{rank}(R) = \operatorname{rank}(P) - 1$, and is generated by the unit vectors $\Pi v_a / \|\Pi v_a\|$. The correlation matrix Q has $\operatorname{rank}(Q) = 2$, and is generated by the set of unit vectors $\{u_a : a \in [n]\} \subset \mathcal{S}(\mathbb{C}^2)$, where

$$u_a = \begin{cases} \mathrm{e}_1, & a \neq c \\ \|\Pi v_c\| \mathrm{e}_1 + \sqrt{1 - \|\Pi v_c\|^2} \mathrm{e}_2, & a = c. \end{cases}$$

This completes the proof. □

6.3 Negative results for d-decomposability

In this section, we construct $(2d+1) \times (2d+1)$ correlation matrices that are not d-decomposable, and prove as a corollary that $\mathrm{Cor}_d(\mathbb{C}^n)$ is not convex when $d \geq 1$ and $n \geq 2d+1$. We require the following lemma, which we reference without proof. We note that this lemma holds more generally over an arbitrary field.

Lemma 35 ([128, 74], Corollary 10 in [89])**.** *Let $m \geq 1$ be an integer, let $\mathcal{V}_1, \ldots, \mathcal{V}_m$ be complex Euclidean spaces, and let $x_1, x_2 \in \mathrm{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ be product vectors. Then the following statements are equivalent:*

1. *For all scalars $\alpha_1, \alpha_2 \in \mathbb{C}$, it holds that $\alpha_1 x_1 + \alpha_2 x_2 \in \mathrm{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m) \cup \{0\}$.*
2. *For some non-zero scalars $\alpha_1, \alpha_2 \in \mathbb{C} \setminus \{0\}$, it holds that $\alpha_1 x_1 + \alpha_2 x_2 \in \mathrm{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m) \cup \{0\}$.*
3. *There exists at most a single index $j \in [m]$ for which $\dim \mathrm{span}\{x_{1,j}, x_{2,j}\} = 2$.*

Theorem 36. *For all integers $d \geq 1$ and $n \geq 2d+1$, $\mathrm{Cor}_d(\mathbb{C}^n) \subsetneq \mathrm{Cor}(\mathbb{C}^n)$.*

Proof. For $d = 1$, the statement follows easily from the fact that the Schur product of rank-one correlation matrices is again rank one (see the proof of Lemma 42). Assume $d \geq 2$. We find a correlation matrix $P \in \mathrm{Cor}\left(\mathbb{C}^{2d+1}\right)$ that is not contained in $\mathrm{Cor}_d\left(\mathbb{C}^{2d+1}\right)$. This will prove the claim, as it implies that any correlation matrix in $\mathrm{Cor}(\mathbb{C}^n)$ with principal submatrix P is not d-decomposable.

Let $v_1, \ldots, v_{d+1}$ be any linearly independent collection of unit vectors for which

$$|\langle v_a, v_{a+2} \rangle| > |\langle v_a, v_{a+1} \rangle| \cdot |\langle v_{a+1}, v_{a+2} \rangle|. \tag{6.1}$$

For example, one could choose any $p \in (0,1)$ and let $\langle v_a, v_b \rangle = p$ for all $a \neq b \in [d+1]$. Let $\alpha_1, \ldots, \alpha_d, \beta_2, \ldots, \beta_{d+1} \in \mathbb{C} \setminus \{0\}$ be any collection of non-zero scalars subject to the constraint that for all $a \in [d]$ it holds that $\|\alpha_a v_a + \beta_{a+1} v_{a+1}\| = 1$, and let P be the correlation matrix generated by

$$\{v_1, \ldots, v_{d+1}, \alpha_1 v_1 + \beta_2 v_2, \alpha_2 v_2 + \beta_3 v_3, \ldots, \alpha_d v_d + \beta_{d+1} v_{d+1}\}.$$

Note that $\mathrm{rank}(P) = d+1$. For notational convenience, we extend the definition of v_a to denote the a-th vector in this set for each $a \in [2d+1]$.

We proceed by contradiction. The existence of a d-decomposition of P is equivalent to the existence of a positive integer m, complex Euclidean spaces $\mathcal{V}_1, \ldots, \mathcal{V}_m \cong \mathbb{C}^d$, and unit product vectors $\{x_a : a \in [2d+1]\} \subset \mathcal{S}\mathrm{Prod}\,(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ such that $\langle x_a, x_b\rangle = \langle v_a, v_b\rangle$ for all $a, b \in [2d+1]$. By Lemma 35, this implies that for each $a \in [d]$,

$$\dim \operatorname{span}\{x_{a,i}, x_{a+1,i}\} = 2 \tag{6.2}$$

for at most a single index $i \in [m]$. Furthermore, such an index indeed exists for every $a \in [d]$, since for all $a \in [d]$,

$$\dim \operatorname{span}\{x_a, x_{a+1}\} = \dim \operatorname{span}\{v_a, v_{a+1}\} = 2.$$

For each $a \in [d]$, fix $i_a \in [m]$ to denote the unique index that satisfies (6.2). Since $|\langle x_{a,i}, x_{a+1,i}\rangle| = 1$ for all $i \neq i_a$, it must hold that $|\langle x_{a,i_a}, x_{a+1,i_a}\rangle| = |\langle v_a, v_{a+1}\rangle|$ for all $a \in [d]$. Note that

$$\dim \operatorname{span}\{x_{a,i} : a \in [2d+1]\} \leq d$$

for all $i \in [m]$, and

$$\dim \operatorname{span}\{x_1, \ldots, x_{d+1}\} = \dim \operatorname{span}\{v_1, \ldots, v_{d+1}\} = d+1,$$

so there must exist an index $a \in [d-2]$ such that $i_a \neq i_{a+1}$. Fix a to denote one such index. Note that

$$\begin{aligned}
|\langle x_{a,i_a}, x_{a+1,i_a}\rangle| &= |\langle v_a, v_{a+1}\rangle| \\
|\langle x_{a+1,i_{a+1}}, x_{a+2,i_{a+1}}\rangle| &= |\langle v_{a+1}, v_{a+2}\rangle| \\
|\langle x_{a,i_{a+1}}, x_{a+1,i_{a+1}}\rangle| &= 1 \\
|\langle x_{a+1,i_a}, x_{a+2,i_a}\rangle| &= 1,
\end{aligned}$$

from which it follows that

$$\begin{aligned}
|\langle x_{a,i_a}, x_{a+2,i_a}\rangle| &= |\langle v_a, v_{a+1}\rangle| \\
|\langle x_{a,i_{a+1}}, x_{a+2,i_{a+1}}\rangle| &= |\langle v_{a+1}, v_{a+2}\rangle|,
\end{aligned}$$

but this implies

$$\begin{aligned}
|\langle v_a, v_{a+2}\rangle| &= |\langle x_a, x_{a+2}\rangle| \\
&= \prod_{i=1}^{m} |\langle x_{a,i}, x_{a+2,i}\rangle| \\
&\leq |\langle x_{a,i_a}, x_{a+2,i_a}\rangle| \cdot |\langle x_{a,i_{a+1}}, x_{a+2,i_{a+1}}\rangle| \\
&= |\langle v_a, v_{a+1}\rangle| \cdot |\langle v_{a+1}, v_{a+2}\rangle|,
\end{aligned}$$

a contradiction to (6.1). This completes the proof. □

Corollary 37. *For all integers $d \geq 1$ and $n \geq 2d+1$, $\mathrm{Cor}_d\left(\mathbb{C}^n\right)$ is not convex.*

Proof. We first prove that $\mathrm{Cor}_d\left(\mathbb{C}^{2d+1}\right)$ is not convex. Let $P \in \mathrm{Cor}\left(\mathbb{C}^{2d+1}\right) \setminus \mathrm{Cor}_d\left(\mathbb{C}^{2d+1}\right)$ be any correlation matrix constructed in Theorem 36. Since $\mathrm{Cor}\left(\mathbb{C}^{2d+1}\right)$ is contained in a real affine space of dimension $2d(2d+1)$, then by Carathéodory's theorem [107],

$$P = \sum_{i=1}^{s} p(i) R_i$$

for some positive integer $s \leq 2d(2d+1)+1$, probability vector p, and extreme point correlation matrices R_i. By Corollary 2 in [85], $\mathrm{rank}(R_i) \leq \lfloor \sqrt{2d+1} \rfloor \leq r$ for all $i \in [s]$. It follows that $\mathrm{Cor}_d\left(\mathbb{C}^{2d+1}\right)$ is not convex, since each R_i is d-decomposable and P is not.

For the general statement, let $n \geq 2d+1$ be any integer. For each $i \in [s]$, let $R_i' \in \mathrm{Cor}\left(\mathbb{C}^n\right)$ be any correlation matrix with $\mathrm{rank}(R_i') = \mathrm{rank}(R_i) \leq d$ that contains R_i as the upper-left principal submatrix. Then

$$P' := \sum_{i=1}^{s} p(i) R_i' \in \mathrm{Cor}\left(\mathbb{C}^n\right)$$

contains P as the upper-left principal submatrix, so P' is not d-decomposable. As before, it follows that $\mathrm{Cor}_d\left(\mathbb{C}^n\right)$ is not convex, since each R_i' is d-decomposable and P' is not. □

6.4 Is the containment $\mathrm{Cor}_2\left(\mathbb{C}^4\right) \subseteq \mathrm{Cor}\left(\mathbb{C}^4\right)$ strict?

Theorem 34 implies $\mathrm{Cor}_2\left(\mathbb{C}^3\right) = \mathrm{Cor}\left(\mathbb{C}^3\right)$, while Theorem 36 implies $\mathrm{Cor}_2\left(\mathbb{C}^5\right) \subsetneq \mathrm{Cor}\left(\mathbb{C}^5\right)$. This leaves open the question of whether the containment $\mathrm{Cor}_2\left(\mathbb{C}^4\right) \subseteq \mathrm{Cor}\left(\mathbb{C}^4\right)$ is strict. For a correlation matrix $P \in \mathrm{Cor}\left(\mathbb{C}^4\right)$, it might seem possible that a 2-decomposition (1.1) exists only for large values of m, which could make our problem intractable. Theorem 38 allows us to restrict our attention to $m = 2$.

Theorem 38. *The following statements are equivalent:*

1. $\mathrm{Cor}_2\left(\mathbb{C}^4\right) \subsetneq \mathrm{Cor}\left(\mathbb{C}^4\right)$.

2. *There exists a correlation matrix $P \in \mathrm{Cor}\left(\mathbb{C}^4\right)$ such that $\mathrm{rank}(P) = 3$, and a generating set of vectors for P forms a circuit.*

Theorem 38 shows that it suffices to consider rank-three correlation matrices for which the generating set of vectors forms a circuit. In Proposition 39, we construct 2-decompositions of an infinite family of such correlation matrices, thus narrowing our question even further. We speculate that perhaps our construction can inspire a more general construction of all such correlation matrices.

Proposition 39. *Let $P \in \mathrm{Cor}\left(\mathbb{C}^4\right)$ be any correlation matrix generated by a set of unit vectors $\{v_a : a \in [4]\}$ such that there exists a real number $-1/2 < p < 1$ for which $\langle v_a, v_b \rangle = p$ for all $a \neq b \in [3]$, and there exist non-zero scalars $\alpha_1, \alpha_2 \in \mathbb{C} \setminus \{0\}$ for which $v_4 = \alpha_1(v_1 + v_3) + \alpha_2 v_2$. Then $P \in \mathrm{Cor}_2\left(\mathbb{C}^4\right)$.*

In the remainder of this section, we prove Theorem 38 and Proposition 39. For Theorem 38, (1 ⇒ 2) will follow from Lemma 40, and (2 ⇒ 1) will follow from Lemma 42. We now prove these lemmas.

Lemma 40. *For all integers $n \geq 3$ and $2 \leq d \leq n-1$, if $\mathrm{Cor}_d\left(\mathbb{C}^n\right) \subsetneq \mathrm{Cor}\left(\mathbb{C}^n\right)$, then there exists a correlation matrix $P \in \mathrm{Cor}\left(\mathbb{C}^n\right) \setminus \mathrm{Cor}_d\left(\mathbb{C}^n\right)$ such that every vector in a generating set of P is contained in the span of the rest.*

Proof. By assumption, there exists $P \in \mathrm{Cor}\left(\mathbb{C}^n\right)$ that is not d-decomposable. If there exists a vector in a generating set of P that is not contained in the span of the rest, then by the proof of Theorem 34 there exists a decomposition $P = Q \odot R$ where $\mathrm{rank}(Q) = 2$ and $\mathrm{rank}(R) = \mathrm{rank}(P) - 1$. If there exists a vector in a generating set of R that is not contained in the span of the rest, then this process can be repeated until we have a decomposition

$$P = Q_1 \odot \cdots \odot Q_m \odot R'$$

for which each Q_i has rank 2 and every vector in a generating set of R' is contained in the span of the rest. Furthermore, R' is not d-decomposable, for otherwise P would be d-decomposable. □

To prove Lemma 42, we require Lemma 41, which is proven in Chapter 7 as a consequence of Theorem 59. We note that this lemma holds more generally over an arbitrary field.

Lemma 41 (Corollary 60)**.** *Let n and m be positive integers, let $\mathcal{V}_1, \ldots, \mathcal{V}_m$ be complex vector spaces, and let $\{x_a : a \in [n]\} \subset \mathrm{Prod}\left(\mathcal{V}_1 : \cdots : \mathcal{V}_m\right)$ be a multiset of product tensors. If $\{x_a : a \in [n]\}$ forms a circuit, then $\dim \mathrm{span}\{x_{a,j} : a \in [n]\} > 1$ for at most $n-2$ indices $j \in [m]$.*

Lemma 42. *Let $n \geq 3$ and $d \in [n-1] \setminus \{1\}$ be integers, and let $P \in \mathrm{Cor}\,(\mathbb{C}^n)$ be any correlation matrix of rank $n-1$ generated by a set of unit vectors $\{v_a : a \in [n]\}$ that forms a circuit. If $P \in \mathrm{Cor}_d\,(\mathbb{C}^n)$, then P is d-decomposable as the Schur product of $n-2$ correlation matrices of rank $\leq d$.*

Proof. By assumption, there exists a positive integer $m \geq 2$, complex vector spaces $\mathcal{V}_1, \ldots, \mathcal{V}_m$, and unit product vectors $\{u_a : a \in [n]\} \subset \mathcal{S}\mathrm{Prod}\,(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ that generate P and form a circuit. By Lemma 41, this implies $\dim \mathrm{span}\{u_{a,i} : a \in [n]\} > 1$ for at most $n-2$ indices $i \in [m]$. For each $i \in [m]$, let R_i be the correlation matrix generated by $\{u_{a,i} : a \in [n]\}$, so that

$$P = R_1 \odot \cdots \odot R_m. \tag{6.3}$$

Then $\mathrm{rank}\, R_i > 1$ for at most $n-2$ indices $i \in [m]$.

We conclude by showing that for any correlation matrix R and rank-one correlation matrix R', $R \odot R'$ is a correlation matrix with $\mathrm{rank}(R \odot R') = \mathrm{rank}(R)$. This will complete the proof, since all the rank-one correlation matrices in the d-decomposition (6.3) can be absorbed into the $\leq n-2$ correlation matrices of rank > 1 to construct the desired decomposition.

It follows from Schur's product theorem that $R \odot R'$ is a correlation matrix [70]. Since R' is positive semidefinite and rank-one, then $R' = xx^*$ for some vector x. Furthermore, since R' has ones on the diagonal, each element of x has unit modulus. It follows that

$$R \odot R' = R \odot xx^* = \mathrm{Diag}\,(x)\, R\, \mathrm{Diag}\,(x)^*\,,$$

where $\mathrm{Diag}\,(x)$ is the diagonal unitary matrix with $\mathrm{Diag}\,(x)\,(a,a) = x(a)$. Since $\mathrm{Diag}\,(x)$ has full rank, then $\mathrm{rank}(\mathrm{Diag}\,(x)\, R\, \mathrm{Diag}\,(x)^*) = \mathrm{rank}(R)$, which completes the proof. □

Theorem 38 follows easily from Lemma 40 and Lemma 42. Now we prove Proposition 39.

Proof of Proposition 39. We have

$$P = \begin{bmatrix} 1 & p & p & \alpha_1 + (\alpha_1 + \alpha_2)p \\ p & 1 & p & \alpha_2 + 2\alpha_1 p \\ p & p & 1 & \alpha_1 + (\alpha_1 + \alpha_2)p \\ \overline{\alpha_1} + (\overline{\alpha_1} + \overline{\alpha_2})p & \overline{\alpha_2} + 2\overline{\alpha_1}p & \overline{\alpha_1} + (\overline{\alpha_1} + \overline{\alpha_2})p & 1 \end{bmatrix}.$$

We construct P as

$$P = Q_1 \odot Q_2,$$

where

$$Q_1 = \begin{bmatrix} 1 & \sqrt{\frac{1+p}{2}} & p & \sqrt{\frac{1+p}{2}} \\ \sqrt{\frac{1+p}{2}} & 1 & \sqrt{\frac{1+p}{2}} & 1 \\ p & \sqrt{\frac{1+p}{2}} & 1 & \sqrt{\frac{1+p}{2}} \\ \sqrt{\frac{1+p}{2}} & 1 & \sqrt{\frac{1+p}{2}} & 1 \end{bmatrix},$$
$$Q_2 = \begin{bmatrix} 1 & p\sqrt{\frac{2}{1+p}} & 1 & \sqrt{\frac{2}{1+p}}(\alpha_1+(\alpha_1+\alpha_2)p) \\ p\sqrt{\frac{2}{1+p}} & 1 & p\sqrt{\frac{2}{1+p}} & \alpha_2+2\alpha_1 p \\ 1 & p\sqrt{\frac{2}{1+p}} & 1 & \sqrt{\frac{2}{1+p}}(\alpha_1+(\alpha_1+\alpha_2)p) \\ \sqrt{\frac{2}{1+p}}(\overline{\alpha_1}+(\overline{\alpha_1}+\overline{\alpha_2})p) & \overline{\alpha_2}+2\overline{\alpha_1}p & \sqrt{\frac{2}{1+p}}(\overline{\alpha_1}+(\overline{\alpha_1}+\overline{\alpha_2})p) & 1 \end{bmatrix}.$$

The equality is clear; it only remains to show that Q_1 and Q_2 are positive semidefinite and rank two.

First, it is easily verified that Q_1 is the correlation matrix generated by the unit vectors

$$\begin{aligned} q_{1,1} &= \mathrm{e}_1 \\ q_{1,2} &= \sqrt{\frac{1+p}{2}}\mathrm{e}_1 + \sqrt{\frac{1-p}{2}}\mathrm{e}_2 \\ q_{1,3} &= p\mathrm{e}_1 + \sqrt{1-p^2}\mathrm{e}_2 \\ q_{1,4} &= q_{1,2}, \end{aligned}$$

which implies Q_1 is positive semidefinite. Furthermore, $\operatorname{rank}(Q_1) \leq 2$, since these vectors span at most a two-dimensional space.

Second, we verify that Q_2 is the correlation matrix of the unit vectors

$$\begin{aligned} q_{2,1} &= p\sqrt{\frac{2}{1+p}}\mathrm{e}_1 + \frac{\overline{\alpha_1}}{|\alpha_1|}\sqrt{\frac{1+p-2p^2}{1+p}}\mathrm{e}_2 \\ q_{2,2} &= \mathrm{e}_1 \\ q_{2,3} &= q_{2,1} \\ q_{2,4} &= (\alpha_2 + 2\alpha_1 x)\mathrm{e}_1 + |\alpha_1|\sqrt{2(1+p-2p^2)}\mathrm{e}_2, \end{aligned}$$

which will complete the proof, since it implies $\operatorname{rank}(Q_2) \leq 2$ as above. The vectors $q_{2,1}, q_{2,2}, q_{2,3}$ are easily seen to be normalized. For $q_{2,4}$, recall the normalization condition on v_4

$$\langle v_4, v_4\rangle = 2|\alpha_1|^2(p+1) + |\alpha_2|^2 + (\overline{\alpha_1}\alpha_2 + \alpha_1\overline{\alpha_2})2p = 1,$$

which implies

$$1 - |\alpha_2 + 2\alpha_1 p|^2 = 1 - \big(|\alpha_2|^2 + 4|\alpha_1|^2 p^2 + (\overline{\alpha_1}\alpha_2 + \alpha_1\overline{\alpha_2})2p\big) = 2|\alpha_1|^2(1 + p - 2p^2).$$

It follows that $q_{2,4}$ is normalized. Now we show that the inner products between $q_{2,1}, \ldots, q_{2,4}$ reproduce Q_2. All are easily seen except $\langle q_{2,1}, q_{2,4}\rangle$, which we now verify:

$$\begin{aligned}\langle q_{2,1}, q_{2,4}\rangle &= p\sqrt{\frac{2}{1+p}}(\alpha_2 + 2\alpha_1 p) + \frac{\alpha_1}{|\alpha_1|}\sqrt{\frac{1+p-2p^2}{1+p}}|\alpha_1|\sqrt{2(1+p-2p^2)}\\ &= \sqrt{\frac{2}{1+p}}\,\big(p(\alpha_2 + 2\alpha_1 p) + \alpha_1(1+p-2p^2)\big)\\ &= \sqrt{\frac{2}{1+p}}(\alpha_1 + (\alpha_1 + \alpha_2)p).\end{aligned}$$

This completes the proof. □

6.5 Application of non-d-decomposability in an entanglement detection scenario

In this section, we apply our results to an entanglement detection scenario. Say we are given many copies of unknown pure states $v_1 v_1^*, \ldots, v_n v_n^*$, with $v_1, \ldots, v_n \in \mathcal{S}(\mathcal{V})$ for an unknown complex vector space $\mathcal{V}$. Suppose further that we are allowed to perform any of the measurements

$$\{v_1 v_1^*, \mathbb{1} - v_1 v_1^*\}, \ldots, \{v_n v_n^*, \mathbb{1} - v_n v_n^*\}$$

on any of the states $v_1 v_1^*, \ldots, v_n v_n^*$, and we wish to detect entanglement in the following sense. For some positive integer d, we wish to detect that for any complex Euclidean space $\mathcal{V}$, any set of unit vectors $v_1, \ldots, v_n \in \mathcal{S}(\mathcal{V})$ that are consistent with the measurement outcomes observed in the above scenario, and any decomposition $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ of $\mathcal{V}$ into spaces of dimension $\dim(\mathcal{V}_i) \leq d$, at least one of the vectors $v_1, \ldots, v_n$ must be entangled (i.e. not in $\operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$).

In the above scenario, the only meaningful information that can be gained from the measurement outcomes is precisely the Gram matrix of $\{v_1 v_1^*, \ldots, v_n v_n^*\}$ (the matrix of inner products $\langle v_a v_a^*, v_b v_b^* \rangle$ for $a, b \in [n]$). Note that a correlation matrix R is the Gram matrix of rank-one projectors if and only if $R = P \odot \overline{P}$ for some correlation matrix P. The above scenario is therefore equivalent to being given some correlation matrix R that is the Gram matrix of rank-one projectors, and wishing to detect that for any correlation matrix P, if $R = P \odot \overline{P}$, then P is not d-decomposable. In Proposition 43 we find examples of such entanglement detection.

Proposition 43. *For any integer $d \geq 1$ and real number $0 < p < 1$, there exists a correlation matrix arising from a set of $2d+1$ unit vectors*

$$\{v_1, \ldots, v_{d+1}, v_{(1,2)}, v_{(2,3)}, \ldots, v_{(d,d+1)}\}$$

such that for all $a \neq b \in [d+1]$,

$$|\langle v_a, v_b \rangle|^2 = p^2,$$

and for all $a \in [d]$,

$$|\langle v_a, v_{(a,a+1)} \rangle|^2 = |\langle v_{a+1}, v_{(a,a+1)} \rangle|^2 = \frac{1+p}{2}.$$

Furthermore, any such correlation matrix with $0 < p < \frac{1}{d}$ is not d-decomposable.

Proof. We first prove the existence of such a correlation matrix. The correlation matrix generated by the set of unit vectors

$$\left\{ v_1, \ldots, v_{d+1}, \frac{1}{\sqrt{2(1+p)}}(v_1 + v_2), \ldots, \frac{1}{\sqrt{2(1+p)}}(v_d + v_{d+1}) \right\},$$

with $\langle v_a, v_b \rangle = p$ for all $a \neq b \in [d+1]$, satisfies the desired conditions. Indeed,

$$\begin{aligned} \left\langle v_a, \frac{1}{\sqrt{2(1+p)}}(v_a + v_{a+1}) \right\rangle &= \frac{1}{\sqrt{2(1+p)}}(1+p) \\ &= \sqrt{\frac{1+p}{2}}, \end{aligned}$$

and similarly,

$$\left\langle v_{a+1}, \frac{1}{\sqrt{2(1+p)}}(v_a + v_{a+1}) \right\rangle = \sqrt{\frac{1+p}{2}}.$$

Now we prove that any such correlation matrix with $0 < p < \frac{1}{d}$ is not d-decomposable. For $d = 1$, the statement follows easily from the fact that the Schur product of any two rank-one correlation matrices is again rank one (see the proof of Lemma 42), and that for all $0 < p < 1$, any correlation matrix satisfying the conditions of the proposition has rank ≥ 2. Assume $d \geq 2$. It is clear that

$$|\langle v_a, v_{a+2} \rangle| > |\langle v_a, v_{a+1} \rangle| \cdot |\langle v_{a+1}, v_{a+2} \rangle|$$

for all $a \in [d-2]$. Thus, by the proof of Theorem 36 it suffices to show that the vectors $\{v_1, \ldots, v_{d+1}\}$ are linearly independent, and that for all $a \in [d]$ it holds that $v_{(a,a+1)} = \alpha_a v_a + \beta_{a+1} v_{a+1}$ for some non-zero scalars $\alpha_a, \beta_{a+1} \in \mathbb{C}^{\times}$.

First, by Gershgorin's circle theorem [70], the condition that $|\langle v_a, v_b \rangle|^2 = p^2$ for all $a \neq b \in [d+1]$, along with $0 < p < \frac{1}{d}$, implies that the vectors $\{v_1, \ldots, v_{d+1}\}$ are linearly independent. Second, for each $a \in [d]$ the principal submatrix of P generated by the vectors $\{v_a, v_{a+1}, v_{(a,a+1)}\}$ is of the form

$$P^{(a,a+1)} = \begin{pmatrix} 1 & e^{i\phi_1} p & e^{i\phi_2}\sqrt{\frac{1+p}{2}} \\ e^{-i\phi_1} p & 1 & e^{i\phi_3}\sqrt{\frac{1+p}{2}} \\ e^{-i\phi_2}\sqrt{\frac{1+p}{2}} & e^{-i\phi_3}\sqrt{\frac{1+p}{2}} & 1 \end{pmatrix}$$

for some $\phi_1, \phi_2, \phi_3 \in [0, 2\pi)$. Note that

$$\mathrm{Det}(P^{(a,a+1)}) = p(1+p)(-1 + \cos(\phi_1 - \phi_2 + \phi_3)) \leq 0,$$

and since $P^{(a,a+1)}$ is positive semidefinite,

$$\mathrm{Det}(P^{(a,a+1)}) = 0.$$

This implies that $P^{(a,a+1)}$ has rank one or two. We can deduce $\mathrm{rank}(P^{(a,a+1)}) \neq 1$ because v_a and v_{a+1} are linearly independent. Thus, $\mathrm{rank}(P^{(a,a+1)}) = 2$, which implies $v_{(a,a+1)} = \alpha_a v_a + \beta_{a+1} v_{a+1}$ for some scalars $\alpha_a, \beta_{a+1} \in \mathbb{C}$, both of which must be non-zero because no entry in $P^{(a,a+1)}$ has unit magnitude. □

Chapter Acknowledgments I thank Norbert Lütkenhaus for first proposing this research topic, and for many helpful discussions. I thank John Watrous for helpful discussions and comments on drafts of this manuscript. I thank Chi-Kwong Li, Ashutosh Marwah, and Daniel Puzzuoli for helpful discussions.

Chapter 7

A generalization of Kruskal’s theorem on tensor decomposition

The contents of this chapter are based on the preprint [94], which was written in collaboration with Fedor Petrov.

Kruskal’s theorem states that a sum of product tensors constitutes a unique tensor rank decomposition if the so-called *k-ranks* of the product tensors are large. We prove a “splitting theorem” for sets of product tensors, in which the k-rank condition of Kruskal’s theorem is weakened to the standard notion of rank, and the conclusion of uniqueness is relaxed to the statement that the set of product tensors splits (i.e. is disconnected as a matroid). Our splitting theorem implies a generalization of Kruskal’s theorem. While several extensions of Kruskal’s theorem are already present in the literature, all of these use Kruskal’s original permutation lemma, and hence still cannot certify uniqueness when the k-ranks are below a certain threshold. Our generalization uses a completely new proof technique, contains many of these extensions, and can certify uniqueness below this threshold. We obtain several other useful results on tensor decompositions as consequences of our splitting theorem. We prove sharp lower bounds on tensor rank and Waring rank, which extend Sylvester’s matrix rank inequality to tensors. We also prove novel uniqueness results for non-rank tensor decompositions.

We write $S \cup T$ to denote the union of two sets S and T. If S and T happen to be disjoint, we often write $S \sqcup T$ instead to remind the reader of this fact. For a positive integer t, we say that a collection of subsets $S_1, \dots, S_t \subseteq T$ *partitions* T if $S_p \cap S_q = \{\}$ for all $p \neq q \in [t]$, and $S_1 \sqcup \cdots \sqcup S_t = T$.

For a multiset of non-zero vectors $E = \{v_1, \dots, v_n\} \subseteq \mathcal{V}$, a *connected component* of E is an inclusion-maximal connected subset of E. Any multiset of non-zero vectors E can be (uniquely, up to reordering) partitioned into disjoint connected components $T_1 \sqcup \cdots \sqcup T_t = E$ [99, Proposition 4.1.2]. Observe that

$$\operatorname{span}(E) = \bigoplus_{i \in [t]} \operatorname{span}(T_i),$$

and note that $S \subseteq E$ separates E if and only if

$$\dim \operatorname{span}\{v_1, \dots, v_n\} = \dim \operatorname{span}\{v_a : a \in S\} + \dim \operatorname{span}\{v_a : a \in S^c\}$$

if and only if

$$\operatorname{span}\{v_a : a \in S\} \cap \operatorname{span}\{v_a : a \in S^c\} = \{0\}$$

(see [99, Proposition 4.2.1]).

This chapter is organized as follows: In Section 7.1 we prove our splitting theorem (Theorem 4), introduced in Chapter 1. In Section 7.2 we observe that our splitting theorem implies the generalization of Kruskal's theorem stated in Theorem 2. In Section 7.3 we prove that the inequality appearing in our splitting theorem cannot be weakened. In Section 7.4 we prove a family of statements that interpolate between our generalization of Kruskal's theorem and a natural offshoot of our splitting theorem (briefly mentioned in Section 1.3.1). In Section 7.5 we use our splitting theorem to prove a new lower bound on Waring rank. In Section 7.6 we prove a novel uniqueness result for non-Waring rank decompositions. In Section 7.7 we compare our generalization of Kruskal's theorem to other uniqueness criteria.

7.1 Proving the splitting theorem

In this section, we prove Theorem 4. We first observe the following basic fact.

Proposition 44. *Let $n \geq 2$ be an integer, let $\mathcal{V} = \mathcal{V}_1 \otimes \mathcal{V}_2$ be a vector space over a field $\mathbb{F}$, and let*

$$E = \{x_a \otimes y_a : a \in [n]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \mathcal{V}_2)$$

be a multiset of product tensors. If E is connected, then $\{x_a : a \in [n]\}$ and $\{y_a : a \in [n]\}$ are both connected.

Proof. Suppose toward contradiction that E is connected and $\{x_a : a \in [n]\}$ splits, i.e.

$$\dim \operatorname{span}\{x_a : a \in [n]\} = \operatorname{span}\{x_a : a \in S\} \oplus \operatorname{span}\{x_a : a \in S^c\} \tag{7.1}$$

for some non-empty proper subset $S \subsetneq [n]$. Since E is connected, there exists a non-zero vector

$$v \in \operatorname{span}\{x_a \otimes y_a : a \in S\} \cap \operatorname{span}\{x_a \otimes y_a : a \in S^c\}.$$

Let $f \in \mathcal{V}_2^*$ be any linear functional such that $(\mathbb{1} \otimes f)v \neq 0$. Then $(\mathbb{1} \otimes f)v$ is a non-zero element of

$$\operatorname{span}\{x_a : a \in S\} \cap \operatorname{span}\{x_a : a \in S^c\},$$

contradicting (7.1). The result is obviously symmetric under permutation of $\mathcal{V}_1$ and $\mathcal{V}_2$. □

It is not difficult to see that Theorem 4 follows directly from the $m = 2$ case of Theorem 4, Proposition 44, and an inductive argument (we omit this proof). We therefore need only prove the $m = 2$ case of Theorem 4, which we now explicitly state for clarity.

Theorem 45 ($m = 2$ case of Theorem 4)**.** *Let $n \geq 2$ be an integer, let $\mathcal{V} = \mathcal{V}_1 \otimes \mathcal{V}_2$ be a vector space over a field $\mathbb{F}$, and let*

$$E = \{x_a \otimes y_a : a \in [n]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \mathcal{V}_2)$$

be a multiset of product tensors. Let

$$d_1 = \dim \operatorname{span}\{x_a : a \in [n]\}$$

and

$$d_2 = \dim \operatorname{span}\{y_a : a \in [n]\}.$$

If E is connected, then $\dim \operatorname{span}(E) \geq d_1 + d_2 - 1$.

To prove Theorem 45, we require a matroid-theoretic construction called the *ear decomposition* of a connected matroid (see, e.g. [44]). For completeness, we review the construction here. We refer the reader to [99] for the basic matroid-theoretic arguments used in this proof.

Lemma 46 (Ear decomposition)**.** *Let $n \geq 2$ be an integer, let $\mathcal{V}$ be a vector space over a field $\mathbb{F}$, and let $E = \{v_1, \dots, v_n\} \subseteq \mathcal{V} \setminus \{0\}$ be a multiset of non-zero vectors. If E is connected, then there exists a collection of circuits $C_1, \dots, C_t \subseteq E$ such that*

$$E = C_1 \cup C_2 \cup \cdots \cup C_t,$$

and for each $p \in [t]$, the multisets C_p and $E_p := C_1 \cup \cdots \cup C_p$ satisfy the following two properties:

1. $C_p \cap E_{p-1} \neq \{\}$

2. $\dim \operatorname{span}(E_p) - \dim \operatorname{span}(E_{p-1}) = |E_p \setminus E_{p-1}| - 1$

Proof. Let $C_1 \subseteq E$ be an arbitrary circuit, which must exist because E is non-empty and connected, and assume by induction that $C_1, \dots, C_p$ have already been constructed to satisfy properties 1 and 2. Let $B \subseteq E_p$ be a basis for $\operatorname{span}(E_p)$, and choose vectors $u_1, u_2, \dots \in E \setminus E_p$ sequentially such that at each step q, $\{u_1, \dots, u_q\}$ is linearly independent. Terminate when

$$\dim \operatorname{span}\{B \cup \{u_1, \dots, u_q\}\} = |B| + q - 1.$$

Note that this process must terminate, otherwise E would split. Fixing q to be the terminating step of this process, note that if u_q is removed from $B \cup \{u_1, \dots, u_q\}$, then the resulting multiset is linearly independent, so $B \cup \{u_1, \dots, u_q\}$ contains a unique circuit containing u_q. Call this circuit C_{p+1}, and observe that properties 1 and 2 hold for $E_{p+1} := C_1 \cup \cdots \cup C_{p+1}$. The lemma follows by repeating this process until the circuits cover E. □

Now we prove Theorem 45.

Proof of Theorem 45. For a subset $S \subseteq [n]$, let

$$\begin{aligned} d^S &= \dim \operatorname{span}\{x_a \otimes y_a : a \in S\}, \\ d_1^S &= \dim \operatorname{span}\{x_a : a \in S\}, \\ d_2^S &= \dim \operatorname{span}\{y_a : a \in S\}. \end{aligned}$$

In a slight change of notation from Lemma 46, let $C_1, \dots, C_t \subseteq [n]$ be the index sets corresponding to an ear decomposition of E, and let $E_p = C_1 \cup \cdots \cup C_p \subseteq [n]$ for each $p \in [t]$. The theorem follows from the following two claims

Claim 47.
$$d^{E_1} \geq d_1^{E_1} + d_2^{E_1} - 1.$$

Claim 48. For each $p \in \{2, \ldots, t\}$,

$$|E_p \setminus E_{p-1}| - 1 \geq d_1^{E_p} - d_1^{E_{p-1}} + d_2^{E_p} - d_2^{E_{p-1}}.$$

Before proving these claims, let us first use them to complete the proof. Note that

$$\begin{aligned} d^{E_2} &= d^{E_1} + |E_2 \setminus E_1| - 1 \\ &\geq d_1^{E_1} + d_2^{E_1} - 1 + |E_2 \setminus E_1| - 1 \\ &\geq d_1^{E_2} + d_2^{E_2} - 1. \end{aligned}$$

The first line is a property of the ear decomposition, the second line follows from Claim 47, and the third line follows from Claim 48. So Claim 47 holds with E_1 replaced with E_2. Repeating this process inductively gives $d^{[n]} \geq d_1^{[n]} + d_2^{[n]} - 1$, which is what we wanted to prove. This completes the proof, modulo proving the claims.

Proof of Claim 47. By permuting $[n]$, we may assume that $C_1 = [q]$ for some $q \in [n]$, and that $\{x_a : a \in [d_1^{[q]}]\}$ is a basis for $\operatorname{span}\{x_a : a \in [q]\}$. Let $s = d_1^{[q]}$.

Suppose that there exists $b \in [s]$ such that $y_b \notin \operatorname{span}\{y_a : a \in [q] \setminus [s]\}$. Let $f \in \mathcal{V}_1^*$, $g \in \mathcal{V}_2^*$ be linear functionals such that $f(x_b) = g(y_b) = 1$, $f(x_a) = 0$ for all $a \in [s] \setminus \{b\}$, and $g(y_a) = 0$ for all $a \in [q] \setminus [s]$. So

$$(f \otimes g)(x_a \otimes y_a) = \begin{cases} 1, & a = b \\ 0, & a \neq b \end{cases}.$$

It follows that $x_b \otimes y_b \notin \operatorname{span}\{x_a \otimes y_a : a \in [q] \setminus \{b\}\}$, contradicting the fact that C_1 indexes a circuit. So $\{y_a : a \in [s]\} \subseteq \operatorname{span}\{y_a : a \in [q] \setminus [s]\}$, which implies

$$\begin{aligned} d_2^{C_1} &\leq q - s \\ &= d^{C_1} + 1 - d_1^{C_1}, \end{aligned}$$

completing the proof. △

Now we prove Claim 48.

Proof of Claim 48. Let $B \subseteq E_{p-1}$ be such that $\{x_a : a \in B\}$ is a basis for $\operatorname{span}\{x_a : a \in E_{p-1}\}$. By permuting $[n]$, we may assume that $E_p \setminus E_{p-1} = [q]$ for some $q \in [n]$, and that

$B \cup \{x_a : a \in [s]\}$ is a basis for $\operatorname{span}\{x_a : a \in E_p\}$, where $s = d^{E_p} - d^{E_{p-1}}$. If there exists $b \in [s]$ for which $y_b \notin \operatorname{span}\{y_a : a \in [q] \setminus [s]\}$, then, as in the proof of Claim 47,

$$x_b \otimes y_b \notin \operatorname{span}\{x_a \otimes y_a : a \in E_p \setminus \{b\}\}.$$

But this contradicts connectedness of E, a contradiction. It follows that $d_2^{E_p \setminus E_{p-1}} \leq q - s$, so

$$\begin{aligned} d_2^{E_p} - d_2^{E_{p-1}} &\leq d_2^{C_p} - d_2^{C_p \cap E_{p-1}} \\ &\leq d_2^{C_p \setminus (C_p \cap E_{p-1})} - 1 \\ &= d_2^{E_p \setminus E_{p-1}} - 1 \\ &\leq q - s - 1 \\ &= |E_p \setminus E_{p-1}| - (d_1^{E_p} - d_1^{E_{p-1}}) - 1. \end{aligned}$$

The first line is easy to verify (in matroid-theoretic terms, this is submodularity of the rank function). The second line follows from the fact that $\{y_a : a \in C_p\}$ is connected. The third line is obvious, the fourth line we proved above, and the fifth line follows from our definitions. This completes the proof. △

The proofs of Claims 47 and 48 complete the proof of the theorem. □

7.2 Using our splitting theorem to generalize Kruskal's theorem

In this section we use our splitting theorem (Theorem 4) to prove our generalization of Kruskal's theorem (Theorem 2). We then introduce a reshaped version of Theorem 2, which has many more degrees of freedom than the standard reshaping of Kruskal's theorem.

To prove Theorem 2, we first observe the following useful corollary to our splitting theorem.

Corollary 49. *Let $n \geq 2$ and $m \geq 2$ be integers, let $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ be a vector space over a field $\mathbb{F}$, let*

$$E = \{x_a : a \in [n]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$$

be a multiset of product tensors, and for each $j \in [m]$, let

$$d_j = \dim \operatorname{span}\{x_{a,j} : a \in [n]\}.$$

If $n \leq \sum_{j=1}^{m}(d_j - 1) + 1$, then E splits.

Proof. If E is linearly independent, then it obviously splits. Otherwise,

$$\dim \operatorname{span}(E) \leq n - 1,$$

and the result follows immediately from our splitting theorem. □

Now we use this corollary to prove our generalization of Kruskal's theorem.

Proof of Theorem 2. Let $x_a = x_{a,1} \otimes \cdots \otimes x_{a,m}$ for each $a \in [n]$, and suppose that $\sum_{a \in [n]} x_a = \sum_{a \in [r]} y_a$ for some non-negative integer $r \leq n$ and multiset of product tensors $\{y_a : a \in [r]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$. For notational convenience, for each $a \in [r]$ let $x_{n+a} = -y_a$, so that $\sum_{a \in [n+r]} x_a = 0$. Let $T_1 \sqcup \cdots \sqcup T_t = [n+r]$ be the index sets of the connected components of $\{x_a : a \in [n+r]\}$. Since $\sum_{a \in [n+r]} x_a = 0$, it follows that $\sum_{a \in T_p} x_a = 0$ for all $p \in [t]$, so $|T_p| \geq 2$ for all $p \in [t]$.

For each $p \in [t]$, if

$$\big|T_p \cap [n]\big| \geq \big|T_p \cap [n+r] \setminus [n]\big|, \tag{7.2}$$

then it must hold that

$$\big|T_p \cap [n]\big| = \big|T_p \cap [n+r] \setminus [n]\big| = 1, \tag{7.3}$$

otherwise $\{x_a : a \in T_p\}$ would split by Corollary 49, a contradiction. Since $r \leq n$ and the inequality (7.2) can never be strict, it follows that $r = n$ and (7.3) holds for all $p \in [t]$. This completes the proof. □

For $m \geq 4$, both Kruskal's theorem and our Theorem 2 can be "reshaped" by regarding multiple subsystems as a single subsystem, to give potentially stronger uniqueness criteria. It is worth noting that the reshaped version of Theorem 2 has quite a different flavour from the reshaped version of Kruskal's theorem; in particular, there are many more degrees of freedom to choose from. We omit the proof of the following reshaped version of Theorem 2, because it is similar to the proof of Theorem 2.

Theorem 50 (Reshaped generalization of Kruskal's theorem)**.** *Let $n \geq 2$ and $m \geq 3$ be integers, let $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ be a vector space over a field $\mathbb{F}$, and let*

$$\{x_a : a \in [n]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$$

be a multiset of product tensors. For each $S \subseteq [n]$ and $J \subseteq [m]$, let

$$d_J^S = \dim \operatorname{span}\Big\{ \bigotimes_{j \in J} x_{a,j} : a \in S \Big\}.$$

If for every subset $S \subseteq [n]$ with $2 \leq |S| \leq n$ there exists a partition $J_1 \sqcup \cdots \sqcup J_t = [m]$ (which may depend on S) such that $2|S| \leq \sum_{i \in [t]} (d_{J_i}^S - 1) + 1$, then $\sum_{a \in [n]} x_a$ constitutes a unique tensor rank decomposition.

It is instructive to compare Theorem 50 to the standard reshaping of Kruskal's theorem:

Theorem 51 (Reshaped Kruskal's theorem)**.** *Let $n \geq 2$ and $m \geq 3$ be integers, let $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ be a vector space over a field $\mathbb{F}$, and let*

$$\{x_a : a \in [n]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$$

be a multiset of product tensors. For each $J \subseteq [m]$, let

$$k_J = \text{k-rank}(\bigotimes_{j \in J} x_{a,j} : a \in [n]).$$

If there exists a partition of $[m]$ into three disjoint subsets $J \sqcup K \sqcup L = [m]$ such that $2n \leq k_J + k_K + k_L - 2$, then $\sum_{a \in [n]} x_a$ constitutes a unique tensor rank decomposition.

Theorem 51 clearly follows from our Theorem 50. In Theorem 51, one could of course consider more general partitions of $[m]$ into more than three subsets, but since the k-rank satisfies $k_{J \cup K} \geq \min\{n, k_J + k_K - 1\}$ for any disjoint subsets $J, K \subseteq [m]$ (See Lemma 1 in [110]), it suffices to consider tripartitions $J \sqcup K \sqcup L = [m]$. In contrast, it is not clear that one can restrict to tripartitions in Theorem 50. There is another major difference between these two theorems: In Theorem 51, one chooses a single partition of $[m]$, whereas in Theorem 50, one is free to choose a different partition of $[m]$ for every S.

We remark that many other statements in this chapter (for example, the splitting theorem itself) can be reshaped similarly to Theorem 50. We do not explicitly state these reshapings.

7.3 The inequality appearing in our splitting theorem cannot be weakened

In this section, we find a connected multiset of product tensors $E = \{x_a : a \in [n]\}$ that satisfies $\dim \operatorname{span}(E) = \sum_{j=1}^m (d_j - 1) + 1$. In fact, we prove that this multiset of product tensors forms a circuit, which is stronger than being connected. This proves that the bound in Corollary 60, and the inequality $\dim \operatorname{span}(E) \leq \sum_{j=1}^m (d_j - 1)$ appearing in Theorem 4, cannot be weakened. The example we use is Derksen's [49], which he used to prove that the inequality appearing in Kruskal's theorem cannot be weakened.

Fact 52. For any field $\mathbb{F}$ with $\operatorname{Char}(\mathbb{F}) = 0$, and positive integers $d_1, \ldots, d_m$ with $n - 1 = \sum_{j=1}^m (d_j - 1) + 1$, there exist vector spaces $\mathcal{V}_1, \ldots, \mathcal{V}_m$ over $\mathbb{F}$ and a multiset of product tensors $\{x_a : a \in [n]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ that forms a circuit, and satisfies

$$\dim \operatorname{span}\{x_{a,j} : a \in [n]\} \geq d_j$$

for all $a \in [n]$.

We note that if $d_1 = \cdots = d_m$, then the multiset of product tensors $\{x_a : a \in [n]\}$ can be taken to be symmetric in the sense introduced in Section 2 (this is obvious from Derksen's construction [49]). As a result, our splitting theorem is also sharp for symmetric product tensors. We use this fact in Sections 7.5 and 7.6 to prove optimality of our results on symmetric decompositions. We remark that the assumption $\operatorname{Char}(\mathbb{F}) = 0$ can be weakened, see [49].

Proof of Fact 52. By Theorem 2 of [49], there exist vector spaces $\mathcal{V}_1, \ldots, \mathcal{V}_m$ over $\mathbb{F}$, a positive integer $\tilde{n} \leq n$, and product tensors $\{x_a : a \in [\tilde{n}]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ with k-ranks $d_j = \text{k-rank}(x_{1,j}, \ldots, x_{\tilde{n},j})$ such that $\sum_{a \in [\tilde{n}]} x_a = 0$. If $\tilde{n} < n$, then $\tilde{n} \leq \sum_{j=1}^m (d_j - 1) + 1$, which implies $\{x_a : a \in [\tilde{n}]\}$ is linearly independent by Corollary 57 (or Proposition 3.1 in [63]). But this contradicts $\sum_{a \in [\tilde{n}]} x_a = 0$, so $\tilde{n} = n$. The equality $n = \sum_{j=1}^m (d_j - 1) + 2$ implies that $d_j \leq n - 1$ for all $j \in [m]$. It follows that for any subset $S \subseteq [n]$ of size $|S| = n - 1$, it holds that $\text{k-rank}(x_{a,j} : a \in S) \geq d_j$. Since $n - 1 = \sum_{j=1}^m (d_j - 1) + 1$, then by Corollary 57, $\{x_a : a \in S\}$ is linearly independent. It follows that $\{x_a : a \in [n]\}$ is a circuit. □

7.4 Interpolating between our generalization of Kruskal's theorem and an offshoot of our splitting theorem

For the entirety of this section, we fix non-negative integers $n \geq 2$ and $m \geq 2$, a vector space $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ over a field $\mathbb{F}$, and a multiset of product tensors $\{x_a : a \in [n]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$. For each subset $S \subseteq [n]$ and index $j \in [m]$, we define

$$d_j^S = \dim \operatorname{span}\{x_{a,j} : a \in S\},$$

and use the shorthand $d_j = d_j^{[n]}$ for all $j \in [m]$.

As a consequence of our splitting theorem, if $n \leq \sum_{j=1}^m (d_j - 1) + 1$, then $\{x_a : a \in [n]\}$ splits (Corollary 49). Our generalization of Kruskal's theorem states that if $2|S| \leq \sum_{j=1}^m (d_j^S - 1) + 1$ for every subset $S \subseteq [n]$ with $2 \leq |S| \leq n$, then $\sum_{a \in [n]} x_a$ constitutes a unique tensor rank decomposition. It is natural to ask what happens when other, similar inequalities hold. In particular, suppose that

$$|S| + \mathcal{R}(|S|) \leq \sum_{j=1}^{m} (d_j^S - 1) + 1 \tag{7.4}$$

for all $S \subseteq [n]$ with $s + 1 \leq |S| \leq n$, for some $s \in [n-1]$ and function $\mathcal{R} : [n] \setminus [s] \to \mathbb{Z}$. What can be said about the tensors $v \in \operatorname{span}\{x_a : a \in [n]\}$?

In this section, we use our splitting theorem to answer this question for choices of s and $\mathcal{R}$ that produce useful results on tensor decompositions. In Section 7.4.1 we prove uniqueness results for low-rank tensors in $\operatorname{span}\{x_a : a \in [n]\}$. These results can be viewed as an interpolation between the two extreme choices of parameters in Corollary 49 (where $s = n - 1$ and $\mathcal{R}(n) = n$) and our generalization of Kruskal's theorem (where $s = 1$ and $\mathcal{R} = \mathbb{1}$). We use this interpolation to extend several recent results in [63, 10, 9]. In Section 7.4.2 we prove uniqueness results for non-rank decompositions of $\sum_{a \in [n]} x_a$ (i.e., decompositions into a non-minimal number of product tensors), which appear to be the first known results of this kind.

We will make use of the following terminology.

Definition 53. For positive integers n and r, multisets of product tensors

$$\{x_a : a \in [n]\}, \{y_a : a \in [r]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m),$$

and non-zero scalars

$$\{\alpha_a : a \in [n]\}, \{\beta_a : a \in [r]\} \subseteq \mathbb{F}^\times,$$

for which

$$\sum_{a\in[n]} \alpha_a x_a = \sum_{a\in[r]} \beta_a y_a,$$

we say that the (ordered) pair of decompositions $(\sum_{a\in[n]} \alpha_a x_a, \sum_{a\in[r]} \beta_a y_a)$ has an (s,l)-*subpartition* for some positive integers s and l if there exist pairwise disjoint subsets $Q_1,\dots,Q_l \subseteq [n]$ and pairwise disjoint subsets $R_1,\dots,R_l \subseteq [r]$ for which

$$\max\{1, |R_p|\} \leq |Q_p| \leq s$$

and $\sum_{a\in Q_p} \alpha_a x_a = \sum_{a\in R_p} \beta_a y_a$ for all $p \in [l]$. We say that the pair $(\sum_{a\in[n]} \alpha_a x_a, \sum_{a\in[r]} \beta_a y_a)$ has an (s,l)-*partition* if the sets $Q_1,\dots,Q_l \subseteq [n]$ and $R_1,\dots,R_l \subseteq [r]$ can be chosen to partition $[n]$ and $[r]$, respectively.

We say that the pair $(\sum_{a\in[n]} \alpha_a x_a, \sum_{a\in[r]} \beta_a y_a)$ is *reducible* if there exist subsets $Q \subseteq [n]$ and $R \subseteq [r]$ for which $|Q| > |R|$ and $\sum_{a\in Q} \alpha_a x_a = \sum_{a\in R} \beta_a y_a$. We say that the pair is *irreducible* if it is not reducible.

(Technically, the linear combinations appearing in the pair $(\sum_{a\in[n]} \alpha_a x_a, \sum_{a\in[r]} \beta_a y_a)$ should be regarded formally, so that they contain the data of the decompositions, and the linear combinations appearing elsewhere should be regarded as standard linear combinations in $\mathcal{V}$.)

For brevity, we will often abuse notation and say that $\sum_{a\in[n]} \alpha_a x_a = \sum_{a\in[r]} \beta_a y_a$ has an (s,l)-subpartition (or is reducible) to mean that $(\sum_{a\in[n]} \alpha_a x_a, \sum_{a\in[r]} \beta_a y_a)$ has an (s,l)-subpartition (or is reducible). Note that the properties of (s,l)-subpartitions and reducibility are not symmetric with respect to permutation of the first and second decompositions. Typically, the first decomposition $\sum_{a\in[n]} \alpha_a x_a$ will be known, and the second decomposition $\sum_{a\in[r]} \beta_a y_a$ will be some unknown decomposition that we want to control.

An immediate consequence of Corollary 49 is that if $\sum_{a\in[n]} x_a = \sum_{a\in[r]} y_a$ for some $r \leq n$, and the inequality (7.4) holds for $s = n-1$ and $\mathcal{R}(n) = r$, then this pair of decompositions has an $(n-1,1)$-subpartition (see Corollary 59 for a slight extension of this statement). By comparison, our generalization of Kruskal's theorem states that if $r \leq n$, and (7.4) holds for $s = 1$ and $\mathcal{R} = \mathbb{1}$, then $r = n$ and this pair of decompositions has a $(1,n)$-subpartition. In Section 7.4.1 we prove statements on the existence of (s,l)-subpartitions for $r \leq n$, which interpolate between these two statements by trading stronger assumptions for stronger notions of uniqueness. In Section 7.4.2 we prove a similar family of statements for $r \geq n+1$, obtaining novel uniqueness results for non-rank decompositions.

We conclude the introduction to this section by making a few notes about our definitions of (s,l)-subpartitions and reducibility. It may seem a bit strange at first that the inequality $|R_p| \leq |Q_p|$ appears in our definition of an (s,l)-subpartition. We have chosen to include this inequality because we typically want to reduce the number of product tensors that appear a decomposition. Our definition of reducibility captures a similar idea: If $n \leq r$ and $(\sum_{a\in[n]} \alpha_a x_a, \sum_{a\in[r]} \beta_a y_a)$ is reducible, then these decompositions can easily be combined to produce a decomposition into fewer than n product tensors. (When $r \leq n$, reducibility of $(\sum_{a\in[r]} \beta_a y_a, \sum_{a\in[n]} \alpha_a x_a)$ captures a similar idea.) Assuming irreducibility will allow us to avoid certain pathological cases. Note that if $\sum_{a\in[n]} \alpha_a x_a$ is a tensor rank decomposition, then $(\sum_{a\in[n]} \alpha_a x_a, \sum_{a\in[r]} \beta_a y_a)$ is automatically irreducible.

Note that when $(\sum_{a\in[n]} \alpha_a x_a, \sum_{a\in[r]} \beta_a y_a)$ is irreducible, the existence of an (s,l)-subpartition is equivalent to the existence of pairwise disjoint subsets $Q_1, \dots, Q_l \subseteq [n]$ and pairwise disjoint subsets $R_1, \dots, R_l \subseteq [r]$ for which

$$1 \leq |R_p| = |Q_p| \leq s$$

and $\sum_{a\in Q_p} \alpha_a x_a = \sum_{a\in R_p} \beta_a y_a$ for all $p \in [l]$. When $s = 1$, these statements are equivalent even without the irreducibility assumption.

7.4.1 Low-rank tensors in the span of a set of product tensors

In this subsection, we prove statements about low-rank tensors in $\operatorname{span}\{x_a : a \in [n]\}$. Most of our results in this section are consequences of Theorem 54, which is a somewhat complicated statement on the existence of (s,l)-partitions. For $s = 1$, and any $r \in \{0, 1, \dots, n\}$ we obtain a condition on $\{x_a : a \in [n]\}$ for which the only rank $\leq r$ tensors in $\operatorname{span}\{x_a : a \in [n]\}$ are those that can be written (uniquely) as a linear combination of $\leq r$ elements of $\{x_a : a \in [n]\}$. For $s = 1, r = 0$ we obtain a sufficient condition for linear independence of $\{x_a : a \in [n]\}$. For $s = 1, r = 1$ we obtain a sufficient condition for the only product tensors in $\operatorname{span}\{x_a : a \in [n]\}$ to be scalar multiples of $x_1, \dots, x_n$. These generalize Proposition 3.1 and Theorem 3.2 in [63], respectively. The case $s = 1, r = n$ reproduces our generalization of Kruskal's theorem. For $s = n - 1$, we strengthen recent results in [10, 9] on circuits of product tensors.

Most of the statements in this subsection are consequences of the following theorem, which is complicated to state, but easy to prove with our splitting theorem.

Theorem 54. *Let $s \in [n-1]$, and $r \in \{0, 1, \dots, n\}$ be integers. Suppose that for every* *subset $S \subseteq [n]$ with $s+1 \leq |S| \leq n$, it holds that*

$$\min\{2|S|, |S|+r\} \leq \sum_{j=1}^{m}(d_j^S - 1) + 1. \tag{7.5}$$

Then for any $v \in \operatorname{span}\{x_a : a \in [n]\}$ with $\operatorname{rank}(v) \leq r$, and any decomposition $v = \sum_{a \in [\tilde{r}]} y_a$ of v into $\tilde{r} \leq r$ product tensors $\{y_a : a \in [\tilde{r}]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$, the following holds: For any subset $S \subseteq [n]$ for which $|S| \geq s+1$, and non-zero scalars $\{\alpha_a : a \in S\} \subseteq \mathbb{F}^\times$ for which it holds that

$$\sum_{a \in S} \alpha_a x_a = \sum_{a \in [\tilde{r}]} y_a$$

and $(\sum_{a \in [\tilde{r}]} y_a, \sum_{a \in S} \alpha_a x_a)$ is irreducible, the pair of decompositions $(\sum_{a \in [n]} \alpha_a x_a, \sum_{a \in [\tilde{r}]} y_a)$ has an (s,l)-partition, for $l = \lceil |S|/s \rceil$.

Proof. For each $a \in [\tilde{r}]$, let $x_{n+a} = -y_a$, and let $E = S \cup ([n+\tilde{r}] \setminus [n]) \subseteq [n+\tilde{r}]$. Let $T_1 \sqcup \cdots \sqcup T_t = E$ be a partition of E into index sets corresponding to the connected components of $\{x_a : a \in E\}$. Since $(\sum_{a \in [\tilde{r}]} y_a, \sum_{a \in S} \alpha_a x_a)$ is irreducible, it must hold that

$$\bigl|T_p \cap S\bigr| \geq \bigl|T_p \cap (E \setminus S)\bigr|$$

for all $p \in [t]$, and hence

$$|T_p| \leq \min\bigl\{2\bigl|T_p \cap S\bigr|, \bigl|T_p \cap S\bigr| + r\bigr\}.$$

If $|T_p \cap S| \geq s+1$, then $\{x_a : a \in T_p\}$ splits by (7.5) and Corollary 49, a contradiction. So it must hold that $|T_p \cap S| \leq s$ for all $p \in [t]$. It follows that $t \geq \lceil |S|/s \rceil$ by the pigeonhole principle, and one can take $Q_p = T_p \cap S$ and

$$R_p = \{a \in [\tilde{r}] : n+a \in T_p \cap (E \setminus S)\}$$

for all $p \in [t]$ to conclude. □

$s = 1$ case of Theorem 54

The $s = 1$ case of Theorem 54 gives a sufficient condition for which the only tensor rank $\leq r$ elements of $\operatorname{span}\{x_a : a \in [n]\}$ are those which can be written (uniquely) as a linear combination of $\leq r$ elements of $\operatorname{span}\{x_a : a \in [n]\}$. In this subsection, we state this case explicitly, and observe several consequences of this case. In particular, we observe a lower bound on tensor rank and a sufficient condition for a set of product tensors to be linearly independent.

Corollary 55 ($s = 1$ case of Theorem 54)**.** *Let $r \in \{0, 1, \dots, n\}$ be an integer. Suppose that for every subset $S \subseteq [n]$ such that $2 \leq |S| \leq n$, it holds that*

$$|S| + \min\{|S|, r\} \leq \sum_{j=1}^{m} (d_j^S - 1) + 1. \tag{7.6}$$

Then any non-zero linear combination of more than r elements of $\{x_a : a \in [n]\}$ has tensor rank greater than r, and every tensor $v \in \operatorname{span}\{x_a : a \in [n]\}$ of tensor rank at most r has a unique tensor rank decomposition into a linear combination of elements of $\{x_a : a \in [n]\}$.

Note that a sufficient condition for the inequality (7.6) to hold is that

$$n + r \leq \sum_{j=1}^{m} (k_j - 1) + 1,$$

where $k_j = \text{k-rank}(x_{1,j}, \dots, x_{n,j})$ for all $j \in [m]$. This recovers Proposition 3.1 and Theorem 3.2 in [63] in the $r = 0$ and $r = 1$ cases, respectively, and interpolates between Kruskal's theorem and these results. For clarity, we will explicitly state the $r = 0$ and $r = 1$ cases of Corollary 55 at the end of this subsection.

Proof of Corollary 55. Let $S \subseteq [n]$ be a subset, let $\{\alpha_a : a \in S\} \subseteq \mathbb{F}^\times$ be a multiset of non-zero scalars, let $\tilde{r} = \operatorname{rank}[\sum_{a \in S} \alpha_a x_a]$, and let $\{y_a : a \in [\tilde{r}]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \dots : \mathcal{V}_m)$ be such that $\sum_{a \in S} \alpha_a x_a = \sum_{a \in [\tilde{r}]} y_a$. If $\tilde{r} \leq r$, then by the $s = 1$ case of Theorem 54, this pair of decompositions has a $(1, |S|)$-partition. It follows that $|S| = \tilde{r}$. Hence, every linear combination of more than r elements of $\{x_a : a \in [n]\}$ has tensor rank greater than r.

Let $v \in \operatorname{span}\{x_a : a \in [n]\}$ have tensor rank $\tilde{r} \leq r$. Then $v = \sum_{a \in Q} \alpha_a x_a$ for some set $Q \subseteq [n]$ of size $|Q| = \tilde{r}$ and non-zero scalars $\{\alpha_a : a \in Q\}$. It follows from (7.6) and Theorem 2 that this is the unique tensor rank decomposition of v. □

Corollary 55 immediately implies the following lower bound on $\operatorname{rank}[\sum_{a \in [n]} x_a]$.

Corollary 56. *If for every subset $S \subseteq [n]$ for which $2 \leq |S| \leq n$, it holds that*

$$|S| + \min\{|S|, r\} \leq \sum_{j=1}^{m} (d_j^S - 1) + 1, \tag{7.7}$$

then $\operatorname{rank}[\sum_{a \in [n]} x_a] \geq r + 1$.

In particular, Corollary 56 implies that

$$\operatorname{rank}\left[\sum_{a\in[n]} x_a\right] \geq \min\left\{n, \sum_{j=1}^{m}(k_j - 1) + 2 - n\right\}. \tag{7.8}$$

In Section 7.5 we prove that when the Kruskal ranks are sufficiently balanced, two of the k-ranks k_i, k_j appearing in the bound (7.8) can be replaced with standard ranks d_i, d_j (Theorem 67). Our Theorem 67 is independent of the bound in Corollary 56 (see Example 68).

We close this subsection by stating the $r = 0$ and $r = 1$ cases of Corollary 55, which generalize Proposition 3.1 and Theorem 3.2 in [63], respectively. We remark that the $m = 2$ subcase of Corollary 57 was proven by Pierpaola Santarsiero in unpublished work, using a different proof technique.

Corollary 57 ($s = 1$, $r = 0$ case of Theorem 54)**.** *If for every subset $S \subseteq [n]$ for which $2 \leq |S| \leq n$, it holds that*

$$|S| \leq \sum_{j=1}^{m}(d_j^S - 1) + 1,$$

then $\{x_a : a \in [n]\}$ is linearly independent.

Corollary 58 ($s = 1$, $r = 1$ case of Theorem 54)**.** *If for every subset $S \subseteq [n]$ for which $2 \leq |S| \leq n$, it holds that*

$$|S| \leq \sum_{j=1}^{m}(d_j^S - 1),$$

then

$$\operatorname{span}\{x_a : a \in [n]\} \cap \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m) = \mathbb{C}^\times x_1 \sqcup \cdots \sqcup \mathbb{C}^\times x_n.$$

$s = n - 1$ case of Theorem 54

In this subsection we state a slight adaptation of the $s = n - 1$ case of Theorem 54, which gives sufficient conditions for a pair of decompositions to have an $(n - 1, 1)$-subpartition. After stating this case, we observe that the subcase $r = 1$ improves recent results in [10, 9] concerning circuits of product tensors. We then remark on applications of this special case in quantum information theory.

Corollary 59 ($s = n-1$ case of Theorem 54)**.** *Let $r \in \{0, 1, \dots, n\}$ be an integer. If $n + r \leq \sum_{j=1}^m (d_j - 1) + 1$, then for any non-negative integer $\tilde{r} \leq r$ and multiset of product tensors $\{y_a : a \in [\tilde{r}]\}$ for which $\sum_{a \in [n]} x_a = \sum_{a \in [\tilde{r}]} y_a$, the pair of decompositions $(\sum_{a\in[n]} x_a, \sum_{a\in[\tilde{r}]} y_a)$ has an $(n-1, 1)$-subpartition.*

Moreover, if $n + r \leq \sum_{j=1}^m (d_j - 1) + 1$, $\tilde{r} = \operatorname{rank}[\sum_{a\in[n]} x_a]$, and $1 \leq \tilde{r} \leq \min\{r, n-1\}$, then there exists a subset $S \subseteq [n]$ with $\tilde{r} \leq |S| \leq n-1$ for which

$$\operatorname{rank}\left[\sum_{a\in S} x_a\right] < \tilde{r}.$$

Proof. The statement of the first paragraph is slightly different from the $s = n-1$ case of Theorem 54, and it follows easily from Corollary 49. To prove the statement of the second paragraph, let $\{z_a : a \in [\tilde{r}]\} \in \operatorname{Prod}(\mathcal{V}_1 : \dots : \mathcal{V}_m)$ be any multiset of product tensors for which $\sum_{a\in[n]} x_a = \sum_{a\in[\tilde{r}]} z_a$, and let $Q \subseteq [n]$, $R \subseteq [\tilde{r}]$ be subsets for which

$$\max\{|R|, 1\} \leq |Q| \leq n - 1$$

and $\sum_{a\in Q} x_a = \sum_{a\in R} z_a$. If $|R| < |Q|$ and $|Q| \geq \tilde{r}$, then we can take $S = Q$. If $|R| < |Q|$ and $|Q| \leq \tilde{r} - 1$, then we can take $S \subseteq [n]$ to be any subset for which $S \supseteq Q$ and $|S| = \tilde{r}$. It remains to consider the case $|R| = |Q|$. In this case, it must hold that $\big|[\tilde{r}] \setminus R\big| < \big|[n] \setminus Q\big|$, so we can find S using the same arguments as in the case $|R| < |Q|$. □

A special case of the $r = 1$ case of Corollary 59 gives an upper bound of $n - 2$ on the number of subsystems $j \in [m]$ for which a circuit of product tensors can have $d_j > 1$. This bound improves those obtained in [9, Theorem 1.1] and [10, Lemma 4.5], and is sharp (see Section 7.3).

Corollary 60. *If $\{x_a : a \in [n]\}$ forms a circuit, then $d_j > 1$ for at most $n - 2$ indices $j \in [m]$.*

Proof. This follows immediately from Corollary 49, since circuits are connected. Alternatively, this follows from the second paragraph in the statement of Corollary 59, since for any circuit it holds that $\sum_{a\in S} x_a \neq 0$ for all $S \subseteq [n]$ with $1 \leq |S| \leq n-1$. □

As an immediate consequence of Corollary 60, a sum of two product tensors is again a product tensor if and only if $d_j > 1$ for at most a single subsystem index $j \in [m]$ (see Corollary 15 in [88]). This statement is well-known. In particular, it was used in [128, 74] to characterize the invertible linear operators that preserve the set of product tensors. In [90, 88] the first author used this statement to study decomposable correlation matrices, and observed that it directly provides an elementary proof of a recent result in quantum information theory [24] (see Corollary 16 in [88]).

7.4.2 Uniqueness results for non-rank decompositions

In this subsection we prove uniqueness results for decompositions of $\sum_{a\in[n]} x_a$ into $r \geq n+1$ product tensors. Namely, we provide conditions on $\{x_a : a \in [n]\}$ for which whenever $\sum_{a\in[n]} x_a = \sum_{a\in[r]} y_a$ for some multiset of product tensors $\{y_a : a \in [r]\}$, this pair of decompositions has an (s,l)-subpartition. In particular, for $s=1$ we obtain sufficient conditions for the existence of subsets $Q \subseteq [n]$, $R \subseteq [r]$ of size $|Q| = |R| = l$ for which $\{x_a : a \in Q\} = \{y_a : a \in R\}$. We refer the reader also to Section 7.6, in which we prove uniqueness results on non-Waring rank decompositions of symmetric tensors, and identify applications of our non-rank uniqueness results.

In Theorem 61 we give sufficient conditions for which whenever $(\sum_{a\in[n]} x_a, \sum_{a\in[r]} y_a)$ is irreducible, it has an (s,l)-subpartition. We then observe that for $s=1$ we can drop the irreducibility assumption and obtain the result described in the previous paragraph. We then prove a modified version of Theorem 61, which drops the irreducibility assumption for arbitrary $s \in [n-1]$. At the end of this subsection, we review these statements in the $s=n-1$ case.

Theorem 61. *Let $n \geq 2$, $q \in [n-1]$, $s \in [q]$, and r be positive integers for which*

$$n+1 \leq r \leq n + \left\lceil \frac{n-q}{s} \right\rceil, \tag{7.9}$$

and let $l = \lfloor q/s \rfloor$. If for every subset $S \subseteq [n]$ for which $s+1 \leq |S| \leq n$, it holds that

$$2|S| + \max\left\{0, (r-n) - \left\lceil \frac{n-q+s}{|S|} \right\rceil + 1\right\} \leq \sum_{j=1}^{m} (d_j^S - 1) + 1, \tag{7.10}$$

then for any multiset of product tensors $\{y_a : a \in [r]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ for which $\sum_{a\in[n]} x_a = \sum_{a\in[r]} y_a$ and $(\sum_{a\in[n]} x_a, \sum_{a\in[r]} y_a)$ is irreducible, this pair of decompositions has an (s,l)-subpartition.

One may be concerned about whether the complicated collection of inequalities (7.10) can ever be satisfied. The answer is yes, simply because the righthand side can depend on m, whereas the lefthand side does not. So for m large enough, one can always find $\{x_a : a \in [n]\}$ that satisfies these inequalities. In fact, they can even be satisfied non-trivially for $m = 3$, as we observe in Example 65.

Proof of Theorem 61. For each $a \in [r]$, let $x_{n+a} = -y_a$, and let $T_1 \sqcup \cdots \sqcup T_t = [n+r]$ be the index sets of the decomposition of $\{x_a : a \in [n+r]\}$ into connected components. Note that for each $p \in [t]$, it must hold that

$$\big|T_p \cap [n+r] \setminus [n]\big| \geq \big|T_p \cap [n]\big|,$$

otherwise we would contradict irreducibility. For each $p \in [t]$, if

$$\big|T_p \cap [n+r] \setminus [n]\big| = \big|T_p \cap [n]\big|,$$

then $\big|T_p \cap [n]\big| \leq s$, otherwise $\{x_a : a \in T_p\}$ would split by (7.10) and Corollary 49. Assume without loss of generality that

$$\begin{aligned} \big|T_1 \cap [n]\big| - \big|T_1 \cap [n+r] \setminus [n]\big| &\geq \big|T_2 \cap [n]\big| - \big|T_2 \cap [n+r] \setminus [n]\big| \\ &\vdots \\ &\geq \big|T_t \cap [n]\big| - \big|T_t \cap [n+r] \setminus [n]\big|. \end{aligned}$$

If

$$\big|T_1 \cap [n]\big| = \big|T_1 \cap [n+r] \setminus [n]\big|,$$

then let $\tilde{l} \in [t]$ be the largest integer for which

$$\big|T_{\tilde{l}} \cap [n]\big| = \big|T_{\tilde{l}} \cap [n+r] \setminus [n]\big|. \tag{7.11}$$

Otherwise, let $\tilde{l} = 0$. Then for all $p \in [t] \setminus [\tilde{l}]$ it holds that

$$|T_p \cap [n]| < |T_p \cap [n+r] \setminus [n]| \tag{7.12}$$

(recall that we define $[0] = \{\}$). To complete the proof, we will show that $\tilde{l} \geq l$, for then we can take $Q_p = T_p \cap [n]$ and $R_p = T_p \cap [n+r] \setminus [n]$ for all $p \in [l]$ to conclude.

Suppose toward contradiction that $\tilde{l} < l$. We require the following two claims:

Claim 62. It holds that $\tilde{l} < t$, $\left\lceil \frac{n - s\tilde{l}}{t - \tilde{l}} \right\rceil \geq s + 1$, and there exists $p \in [t] \setminus [\tilde{l}]$ for which

$$\big|T_p \cap [n]\big| \geq \left\lceil \frac{n - s\tilde{l}}{t - \tilde{l}} \right\rceil. \tag{7.13}$$

Claim 63. For all $p \in [t] \setminus [\tilde{l}]$, it holds that

$$\big|T_p \cap [n+r] \setminus [n]\big| \le \big|T_p \cap [n]\big| + r - n + \tilde{l} - t + 1 \tag{7.14}$$

Before proving these claims, we first use them to complete the proof of the theorem. Let $p \in [t] \setminus [\tilde{l}]$ be as in Claim 62. Then,

$$\begin{aligned}
|T_p| &= \big|T_p \cap [n]\big| + \big|T_p \cap [n+r] \setminus [n]\big| \\
&\le 2\big|T_p \cap [n]\big| + r - n + \tilde{l} - t + 1 \\
&\le 2\big|T_p \cap [n]\big| + r - n - \left\lceil \frac{n - s\tilde{l}}{\big|T_p \cap [n]\big|} \right\rceil + 1 \\
&\le 2\big|T_p \cap [n]\big| + r - n - \left\lceil \frac{n - q + s}{\big|T_p \cap [n]\big|} \right\rceil + 1 \\
&\le \sum_{j=1}^{m} (d_j^{T_p \cap [n]} - 1) + 1,
\end{aligned}$$

where the first line is obvious, the second follows from Claim 63, the third follows from Claim 62, the fourth follows from $\tilde{l} < l$, and the fifth follows from (7.10) and the fact that $\big|T_p \cap [n]\big| \ge s + 1$. So $\{x_a : a \in T_p\}$ splits, a contradiction. This completes the proof, modulo proving the claims.

Proof of Claim 62. To prove the claim, we first observe that $n > st$. Indeed, if $n \le st$ then

$$\begin{aligned}
r &= \sum_{p=1}^{t} \big|T_p \cap [n+r] \setminus [n]\big| \\
&\ge n + t - \tilde{l} \\
&\ge n + \frac{n-q}{s} + 1,
\end{aligned}$$

where the first line is obvious, the second follows from (7.11) and (7.12), and the third follows from $n \le st$ and $\tilde{l} < l$. This contradicts (7.9), so it must hold that $n > st$.

Note that $\tilde{l} < t$, for otherwise we would have $n \le st$ by the fact that $\big|T_p \cap [n]\big| \le s$ for all $p \in [\tilde{l}]$. To verify that $\left\lceil \frac{n - s\tilde{l}}{t - \tilde{l}} \right\rceil \ge s + 1$, it suffices to prove $\frac{n - s\tilde{l}}{t - \tilde{l}} > s$, which follows from $n > st$. To verify (C.2), since $\big|T_p \cap [n]\big| \le s$ for all $p \in [\tilde{l}]$, by the pigeonhole principle there exists $p \in [t] \setminus [\tilde{l}]$ for which

$$\big|T_p \cap [n]\big| \ge \left\lceil \frac{n - s\tilde{l}}{t - \tilde{l}} \right\rceil.$$

This proves the claim. △

Proof of Claim 63. Suppose toward contradiction that the inequality (7.14) does not hold for some $\tilde{p} \in [t] \setminus [\tilde{l}]$. Then

$$\begin{aligned} r &= \sum_{p=1}^{t} \big|T_p \cap [n+r] \setminus [n]\big| \\ &\geq \sum_{p \neq \tilde{p}} \big|T_p \cap [n+r] \setminus [n]\big| + \big|T_{\tilde{p}} \cap [n]\big| + (r-n) + \tilde{l} - t + 2 \\ &\geq r+1, \end{aligned}$$

where the first two lines are obvious, and the last line follows from (7.11) and (7.12), a contradiction. △

The proofs of Claims 62 and 63 complete the proof of the theorem. □

$s = 1$ case of Theorem 61

In the $s = 1$ case of Theorem 61, we can drop the assumption that the pair of decompositions is irreducible. This is because the other assumptions already imply that $\sum_{a \in [n]} x_a$ constitutes a (unique) tensor rank decomposition by Theorem 2, so $\sum_{a \in [n]} x_a = \sum_{a \in [r]} y_a$ will automatically be irreducible (see the discussion at the beginning of Section 7.4).

Corollary 64 ($s = 1$ case of Theorem 61)**.** *Let $q \in [n-1]$ and r be positive integers for which $n+1 \leq r \leq 2n-q$. If for every subset $S \subseteq [n]$ with $2 \leq |S| \leq n$ it holds that*

$$2|S| + \max\left\{0, (r-n) - \left\lceil \frac{n-q+1}{|S|} \right\rceil + 1\right\} \leq \sum_{j=1}^{m} (d_j^S - 1) + 1, \tag{7.15}$$

then for any multiset of product tensors $\{y_a : a \in [r]\} \subseteq \mathrm{Prod}\,(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ for which $\sum_{a \in [n]} x_a = \sum_{a \in [r]} y_a$, there exist subsets $Q \subseteq [n]$ and $R \subseteq [r]$ of size $|Q| = |R| = q$ for which $\{x_a : a \in Q\} = \{y_a : a \in R\}$ (in other words, this pair of decompositions has a $(1, q)$-subpartition).

It is worth noting that although the assumptions of Corollary 64 require $\sum_{a \in [n]} x_a$ to constitute a unique tensor rank decomposition, this result can also be applied to arbitrary decompositions $\sum_{a\in[n]} x_a$, provided that $\sum_{a\in S} x_a$ constitutes a unique tensor rank decomposition for some subset $S \subseteq [n]$ with $2 \leq |S| \leq n$, as one can simply apply Corollary 64 to the pair of decompositions $(\sum_{a\in S} x_a, \sum_{a\in[r]} y_a - \sum_{a\in[n]\setminus S} x_a)$. It is not difficult to produce explicit examples in which Corollary 64 can be applied in this way (for instance, by modifying Example 65).

As an example, we now use Corollary 64 to prove uniqueness of non-rank decompositions of the *identity tensor* $\sum_{a\in[n]} e_a^{\otimes 3}$.

Example 65. Let $n \geq 2$, $q \in [n-1]$, and r be positive integers for which $n+1 \leq r \leq 2n-q$ and

$$q \leq n+1-\frac{1}{4}\left((r-n+2)^2+1\right).$$

If

$$\sum_{a\in[n]} e_a^{\otimes 3} = \sum_{a\in[r]} y_a$$

for some multiset of product tensors $\{y_a : a \in [r]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \mathcal{V}_2 : \mathcal{V}_3)$, then there exist subsets $Q \subseteq [n]$ and $R \subseteq [n+r]$ of sizes $|Q| = |R| = q$ such that $\{x_a : a \in Q\} = \{y_a : a \in R\}$. For example, if $r = n+1$ then we can take $q = n-2$ for any $n \geq 3$.

To verify Example 65, it suffices to show that the inequality (7.15) holds for all $S \subseteq [n]$ with $2 \leq |S| \leq n$. This reduces to proving that

$$|S|(r-n+2-|S|) - (n-q+1) < 0,$$

which occurs whenever the polynomial in $|S|$ on the lefthand side has no real roots, i.e. whenever

$$(r-n+2)^2 \leq 4(n-q+1)-1.$$

Modifying Theorem 61 to apply to reducible pairs of decompositions

A drawback to Theorem 61 is that it only applies to irreducible pairs of decompositions. We now present a modification of this result, which can certify the existence of an (s,l)-subpartition even for reducible decompositions, at the cost of stricter assumptions. We defer this proof to the appendix, as it is very similar to that of Theorem 61.

Theorem 66. *Let $q \in [n-1]$, $s \in [q]$, and r be positive integers for which*

$$n+1 \leq r \leq \left\lceil \left(\frac{s+1}{s}\right)(n-q+s)\right\rceil - 1,$$

and let $l = \lfloor q/s \rfloor$. If for every subset $S \subseteq [n]$ for which $s+1 \leq |S| \leq n$, it holds that

$$2|S| + \max\left\{0, (r-n+q-s) - \left\lceil \frac{n-q+s}{|S|} \right\rceil + 1\right\} \leq \sum_{j=1}^{m}(d_j^S - 1) + 1,$$

then for any multiset of product tensors $\{y_a : a \in [r]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ for which $\sum_{a\in[n]} x_a = \sum_{a \in [r]} y_a$, this pair of decompositions has an (s,l)-subpartition.

$s = n-1$ case of Theorem 66

When $s = n-1$, then it necessarily holds that $r = n+1$ and $q = n-1$, and Theorem 66 simply says that if $2n+1 \leq \sum_{j=1}^m (d_j - 1) + 1$, then $\sum_{a\in[n]} x_a = \sum_{a\in[n+1]} y_a$ has an $(n-1,1)$-subpartition. Theorem 61 yields a weaker statement.

7.5 A lower bound on tensor rank

In Section 7.4.1 we saw that for a multiset of product tensors $\{x_a : a \in [n]\}$ with k-ranks $k_j = \text{k-rank}(x_{a,j} : a \in [n])$, it holds that

$$\operatorname{rank}\left[\sum_{a\in[n]} x_a\right] \geq \min\left\{n, \sum_{j=1}^m (k_j - 1) + 2 - n\right\}. \tag{7.16}$$

In this section, we prove that when the k-ranks are sufficiently balanced, two of the k-ranks k_i, k_j appearing in this bound can be replaced with standard ranks d_i, d_j, which improves this bound when the k-ranks and ranks are not equal, and specializes to Sylvester's matrix rank inequality when $m = 2$. We prove that this improved bound is independent of a different lower bound on tensor rank that we observed in Corollary 56. We furthermore observe that this improved bound is sharp in a wide parameter regime. As a consequence, we obtain a lower bound on Waring rank, which we also prove is sharp.

Theorem 67 (Tensor rank lower bound)**.** *Let $n \geq 2$ and $m \geq 2$ be integers, let $\mathcal{V} = \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_m$ be a vector space over a field $\mathbb{F}$, and let*

$$E = \{x_a : a \in [n]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$$

be a multiset of product tensors. For each index $j \in [m]$, let $k_j = \text{k-rank}(x_{a,j} : a \in [n])$ and $d_j = \dim \operatorname{span}\{x_{a,j} : a \in [n]\}$. Define

$$\mu = \max_{\substack{i,j \in [m] \\ i \neq j}} \{d_i - k_i + d_j - k_j\}. \tag{7.17}$$

If for every index $i \in [m]$ it holds that

$$k_i \leq \sum_{\substack{j \in [m] \\ j \neq i}} (k_j - 1) + 1, \tag{7.18}$$

then

$$\operatorname{rank}\left[\sum_{a \in [n]} x_a\right] \geq \min\left\{n, \mu + \sum_{j=1}^{m} (k_j - 1) + 2 - n\right\}. \tag{7.19}$$

Intuitively, the condition (7.18) ensures that the k-ranks are sufficiently balanced. This inequality is satisfied, for example, when the product tensors are symmetric. While we are unaware whether the precise inequality (7.18) is necessary for the lower bound (7.19) to hold, the following example illustrates that some inequality of this form must hold:

Example 68. The set of product tensors

$$E = \{e_1^{\otimes 3}, e_2^{\otimes 3}, e_3^{\otimes 3}, e_4^{\otimes 3}, e_5 \otimes (e_1 + e_2)^{\otimes 2}, e_6 \otimes (e_1 - e_2)^{\otimes 2}\}$$

does not satisfy (7.19). Indeed,

$$\begin{aligned} \operatorname{rank}[\Sigma(E)] &= 5 \\ &< q + k_1 + k_2 + k_3 - 1 - n \\ &= d_2 + d_3 - 1 \\ &= 7. \end{aligned}$$

This example illustrates that in order for the bound (7.19) to hold, the k-ranks must be sufficiently "balanced" in order to avoid cases such as this. In particular, some inequality resembling (7.18) is necessary. We remark that this example can be extended to further parameter regimes using Derksen's example [49], and similar arguments as in Sections 7.5.1 and 7.6.1.

Note that when $m = 2$, Theorem 67 states that

$$\operatorname{rank}\left[\sum_{a\in[n]} x_a\right] \geq d_1 + d_2 - n,$$

provided that $k_1 = k_2$. This is Sylvester's matrix rank inequality (although Sylvester's result holds also when $k_1 \neq k_2$) [70].

The following example demonstrates that our two lower bounds on tensor rank in Theorem 67 and Corollary 56 are independent.

Example 69. By Theorem 67, the sum of the set of product tensors

$$\{e_1^{\otimes 3}, e_2^{\otimes 3}, (e_1+e_2)^{\otimes 2} \otimes e_3, e_3^{\otimes 2} \otimes (e_1+e_2+e_3)\}$$

has tensor rank 4. Note that this bound cannot be achieved with the flattening rank lower bound, nor with Corollary 56, as the first three vectors do not satisfy (7.7). Many more such examples can be obtained using the construction in Section 7.5.1.

Conversely, the sum of the set of product tensors

$$\{e_1^{\otimes 3}, e_2^{\otimes 3}, e_3^{\otimes 3}, e_4^{\otimes 3}, (e_2+e_3)\otimes(e_2+e_4)\otimes(e_1+e_4)\}$$

has tensor rank 5 by Corollary 56, while Theorem 67 only certifies that this sum has tensor rank at least 4.

Now we prove Theorem 67.

Proof of Theorem 67. Let $r = \operatorname{rank}[\sum_{a\in[n]} x_a]$, and let $\{y_a : a \in [r]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \cdots : \mathcal{V}_m)$ be a multiset of product tensors for which $\sum_{a\in[n]} x_a = \sum_{a\in[r]} y_a$ is a tensor rank decomposition. We need to prove that r satisfies the inequality (7.19). For each $a \in [r]$, let $x_{n+a} = -y_a$, and let $T_1 \sqcup \cdots \sqcup T_t = [n+r]$ be the index sets of the connected components of $\{x_a : a \in [n+r]\}$. For each subset $S \subseteq [n]$ and index $j \in [m]$, let

$$d_j^S = \dim \operatorname{span}\{x_{a,j} : a \in S\}.$$

We first consider the case $t = 1$, i.e. $\{x_a : a \in [n+r]\}$ is connected. By the splitting theorem, it holds that

$$\begin{aligned} n + r &\geq \sum_{j=1}^m (d_j - 1) + 2 \\ &\geq \mu + \sum_{j=1}^m (k_j - 1) + 2, \end{aligned}$$

completing the proof in this case.

We proceed by induction on t. Suppose the theorem holds whenever the number of connected components is less than t. Assume without loss of generality that

$$\begin{aligned}\big|T_1 \cap [n]\big| - \big|T_1 \cap [n+r] \setminus [n]\big| &\geq \big|T_2 \cap [n]\big| - \big|T_2 \cap [n+r] \setminus [n]\big| \\ &\vdots \\ &\geq \big|T_t \cap [n]\big| - \big|T_t \cap [n+r] \setminus [n]\big| \\ &\geq 0,\end{aligned}$$

where the last line follows from the fact that $\sum_{a\in[r]} y_a$ is a tensor rank decomposition. If

$$\big|T_1 \cap [n]\big| = \big|T_1 \cap [n+r] \setminus [n]\big|,$$

then $r = n$ and we are done. Otherwise,

$$\big|T_{[t-1]} \cap [n]\big| > \big|T_{[t-1]} \cap [n+r] \setminus [n]\big|,$$

where $T_{[t-1]} = T_1 \sqcup \cdots \sqcup T_{t-1}$.

Observe that $k_j < \big|T_{[t-1]} \cap [n]\big|$ for all $j \in [m]$. Indeed, since

$$\operatorname{rank}\Bigg[\sum_{a \in T_{[t-1]} \cap [n]} x_a\Bigg] < \big|T_{[t-1]} \cap [n]\big|,$$

it must hold that

$$2\big|T_{[t-1]} \cap [n]\big| - 1 \geq \sum_{j=1}^{m} \Big(\min\big\{\big|T_{[t-1]} \cap [n]\big|, k_j\big\} - 1\Big) + 2,$$

by (7.16). If $k_i \geq \big|T_{[t-1]} \cap [n]\big|$ for some $i \in [m]$, then this inequality implies that $k_j < \big|T_{[t-1]} \cap [n]\big|$ for all $j \neq i$, and hence

$$k_i \geq \big|T_{[t-1]} \cap [n]\big| \geq \sum_{\substack{j\in[m] \\ j\neq i}} (k_j - 1) + 2,$$

contradicting (7.18). So $k_j < \big|T_{[t-1]} \cap [n]\big|$ for all $j \in [m]$.

Since $k_j < \left|T_{[t-1]} \cap [n]\right|$ for all $j \in [m]$, the k-ranks of $\{x_a : a \in T_{[t-1]} \cap [n]\}$ satisfy (7.18), so by the induction hypothesis,

$$\left|T_{[t-1]}\right| \geq \mu^{T_{[t-1]} \cap [n]} + \sum_{j=1}^{m} (k_j - 1) + 2, \tag{7.20}$$

where

$$\mu^{T_{[t-1]} \cap [n]} = \max_{\substack{i,j \in [m] \\ i \neq j}} \left\{ d_i^{T_{[t-1]} \cap [n]} - k_i + d_j^{T_{[t-1]} \cap [n]} - k_j \right\}.$$

To complete the proof, we will show that

$$\left|T_{[t-1]}\right| + |T_t| \geq \mu + \sum_{j=1}^{m} (k_j - 1) + 2.$$

Let $i, i' \in [m]$ be such that $\mu = d_i - k_i + d_{i'} - k_{i'}$. Then

$$\begin{aligned}
\left|T_{[t-1]}\right| + |T_t| &\geq d_i^{T_{[t-1]} \cap [n]} - k_i + d_{i'}^{T_{[t-1]} \cap [n]} - k_{i'} + \sum_{j=1}^{m} (k_j - 1) + \sum_{j=1}^{m} (d_j^{T_t \cap [n]} - 1) + 4 \\
&\geq d_i^{T_{[t-1]} \cap [n]} - k_i + d_{i'}^{T_{[t-1]} \cap [n]} - k_{i'} + \sum_{j=1}^{m} (k_j - 1) + d_i^{T_t \cap [n]} + d_{i'}^{T_t \cap [n]} + 2 \\
&\geq d_i - k_i + d_{i'} - k_{i'} + \sum_{j=1}^{m} (k_j - 1) + 2 \\
&= \mu + \sum_{j=1}^{m} (k_j - 1) + 2,
\end{aligned}$$

where the first line follows from (7.20) and the fact that $\{x_a : a \in T_t\}$ is connected, the second is obvious, the third is easy to verify (in matroid-theoretic terms, this is submodularity of the rank function), and the fourth is by definition. This completes the proof. □

As an immediate corollary to Theorem 67, we obtain the following lower bound on the Waring rank of a symmetric tensor, in terms of a known symmetric decomposition.

Corollary 70 (Waring rank lower bound)**.** *Let $n \geq 2$, and $m \geq 2$ be integers, let $\mathcal{W}$ be a vector space over a field $\mathbb{F}$ with $\mathrm{Char}(\mathbb{F}) = 0$ or $\mathrm{Char}(\mathbb{F}) > m$, and let $\{v_a : a \in [n]\} \subseteq \mathcal{W} \setminus \{0\}$ be a multiset of non-zero vectors. Let*

$$k = \text{k-rank}(v_a : a \in [n])$$

and

$$d = \dim \operatorname{span}\{v_a : a \in [n]\}.$$

Then for any multiset of non-zero scalars

$$\{\alpha_a : a \in [n]\} \subseteq \mathbb{F}^\times,$$

it holds that

$$\text{WaringRank}\left[\sum_{a \in [n]} \alpha_a v_a^{\otimes m}\right] \geq \min\{n, 2d + (m-2)(k-1) - n\}. \tag{7.21}$$

7.5.1 Our tensor rank lower bound is sharp

In this subsection, we observe that, in a wide parameter regime, the inequalities (7.19) and (7.21) appearing in Theorem 67 and Corollary 70 cannot be improved.

Let $\mathbb{F}$ be a field with $\text{Char}(\mathbb{F}) = 0$, let $n \geq 2$, $m \geq 2$,

$$2 \leq d_1, \ldots, d_m \leq n,$$

and

$$k_1 \leq d_1, \ldots, k_m \leq d_m$$

be positive integers, and let

$$\lambda = \sum_{j=1}^{m} (k_j - 1) + 2.$$

Suppose that the following conditions hold:

1. $\mu = 2(d_i - k_i)$ for some index $i \in [m]$, where μ is defined as in (7.17).
2. $\max\{k_j : j \in [m]\} + d_i - k_i + 1 \leq n \leq d_i - k_i + \lambda$
3. The inequality (7.18) is satisfied.

Then there exists a multiset of product tensors E corresponding to these choices of parameters that satisfies (7.19) with equality. Indeed, the bound $\text{rank}[\Sigma(E)] \geq n$ is trivial to attain with equality, and the bound

$$\text{rank}[\Sigma(E)] \geq 2(d_i - k_i) + \lambda - n \tag{7.22}$$

can be attained with equality as follows. Let

$$\{x_a : a \in [\lambda]\} \subseteq \text{Prod}\left(\mathbb{F}^{d_1} : \cdots : \mathbb{F}^{d_m}\right)$$

be a multiset of product tensors that forms a circuit and satisfies

$$\dim \text{span}\{x_{a,j} : a \in [\lambda]\} = \text{k-rank}(x_{a,j} : a \in [\lambda]) = k_j \tag{7.23}$$

for all $j \in [m]$. An example of such a circuit is presented in [49], and reviewed in Section 7.3. Now, let

$$\{x_a : a \in [\lambda + d_i - k_i] \setminus [\lambda]\} \subseteq \text{Prod}\left(\mathbb{F}^{d_1} : \cdots : \mathbb{F}^{d_m}\right)$$

be any multiset of product tensors for which

$$\dim \text{span}\{x_{a,j} : a \in [\lambda + d_i - k_i]\} = d_j \tag{7.24}$$

and

$$\text{k-rank}(x_{a,j} : a \in [\lambda + d_i - k_i]) = k_j$$

for all $j \in [m]$, which is guaranteed to exist since $\mathbb{F}$ is infinite. Let

$$E = \{x_a : a \in [n - d_i + k_i]\} \sqcup \{x_a : a \in [\lambda + d_i - k_i] \setminus [\lambda]\}$$

and

$$F = \{x_a : a \in [\lambda] \setminus [n - d_i + k_i]\} \sqcup \{x_a : a \in [\lambda + d_i - k_i] \setminus [\lambda]\}.$$

Recall that $n \leq d_i - k_i + \lambda$ by assumption, so the set $[\lambda] \setminus [n - d_i + k_i]$ that appears in the definition of F is well-defined. Since $n - d_i + k_i \geq k_j + 1$ for all $j \in [m]$, E has k-ranks $k_1, \ldots, k_m$, as desired. It is also clear that E has ranks $d_1, \ldots, d_m$, by (7.23) and (7.24). Since $\{x_a : a \in [\lambda]\}$ forms a circuit, some non-zero linear combination of E is equal to a non-zero linear combination of F. Since $|F|$ is equal to the right hand side of (7.22), this completes the proof.

Out of the three conditions required for our construction, $\mu = 2(d_i - k_i)$ seems the most restrictive. Unfortunately, our methods appear to require this condition. A nearly identical construction shows that the inequality (7.21) appearing in Corollary 70 cannot be improved (and our restrictive condition on μ is automatically satisfied in this case). The only difference in the construction is to choose the product tensors $\{x_a : a \in [\lambda + d_i - k_i]\}$ to be symmetric in this case, which can always be done (in particular, the product tensors appearing in Derksen's example can be taken to be symmetric).

7.6 A uniqueness result for non-Waring rank decompositions

In this section, we prove a sufficient condition on a symmetric decomposition

$$v = \sum_{a \in [n]} \alpha_a v_a^{\otimes m}$$

under which any distinct decomposition $v = \sum_{a \in [r]} \beta_a u_a^{\otimes m}$ must have r lower bounded by some quantity, which we call $r_{\min}$ for now. When $r_{\min} \leq n$, this yields a lower bound on $\operatorname{WaringRank}(v)$ that is contained in Corollary 70. When $r_{\min} = n+1$, this yields a uniqueness criterion for symmetric decompositions that is contained in Theorem 2, but improves Kruskal's theorem in a wide parameter regime. The main result in this section is the case $r_{\min} > n+1$, where we obtain an even stronger statement than uniqueness: Every symmetric decomposition of v into less than $r_{\min}$ terms must be equal to $\sum_{a \in [n]} \alpha_a v_a^{\otimes m}$ (in the language introduced in Section 2, $\sum_{a \in [n]} \alpha_a v_a^{\otimes m}$ is the *unique symmetric decomposition of v into less than $r_{\min}$ terms*). In Section 7.6.1 we prove that our bound $r_{\min}$ cannot be improved. In Section 7.6.2 we identify potential applications of our non-rank uniqueness results.

Our results in this section were inspired by, and generalize, Theorem 6.8 and Remark 6.14 in [38]. Our results in this section should be compared with those of Section 7.4.2 on uniqueness of non-rank decompositions of tensors that are not necessarily symmetric.

Theorem 71. *Let $n \geq 2$ and $m \geq 2$ be integers, let $\mathcal{W}$ be a vector space over a field $\mathbb{F}$ with $\operatorname{Char}(\mathbb{F}) = 0$ or $\operatorname{Char}(\mathbb{F}) > m$, let $E = \{v_a : a \in [n]\} \subseteq \mathcal{W} \setminus \{0\}$ be a multiset of non-zero vectors with $\text{k-rank}(v_a : a \in [n]) \geq 2$, and let*

$$d = \dim \operatorname{span}\{v_a : a \in [n]\}.$$

Then for any non-negative integer $r \geq 0$, multiset of non-zero vectors $F = \{u_a : a \in [r]\} \subseteq \mathcal{W} \setminus \{0\}$ with k-rank$(u_a : a \in [r]) \geq \min\{2, r\}$*, and multisets of non-zero scalars*

$$\{\alpha_a : a \in [n]\}, \{\beta_a : a \in [r]\} \subseteq \mathbb{F}^\times$$

for which

$$\{\alpha_a v_a^{\otimes m} : a \in [n]\} \neq \{\beta_a u_a^{\otimes m} : a \in [r]\} \tag{7.25}$$

and

$$\sum_{a \in [n]} \alpha_a v_a^{\otimes m} = \sum_{a \in [r]} \beta_a u_a^{\otimes m}, \tag{7.26}$$

it holds that

$$n + r \geq m + 2d - 2. \tag{7.27}$$

In the language of the introduction to this section, $r_{\min} = m+2d-2-n$. For comparison, the result we have referred to in [38] asserts that, under the condition $n \leq m$, it holds that $n + r \geq m + d$, which is weaker than our bound (7.27).

Proof of Theorem 71. By subtracting terms from both sides of (7.26), and combining parallel product tensors into single terms (or to zero), it is clear that it suffices to prove the statement when E is linearly independent (so $d = n$).

Note that $r \geq n$ by Kruskal's theorem. For each $a \in [r]$, let $v_{n+a} = u_a$, and let $T_1 \sqcup \cdots \sqcup T_t = [n + r]$ be the index sets of the connected components of $\{v_a^{\otimes m} : a \in [n + r]\}$. Assume without loss of generality that $\left|T_1 \cap [n]\right| \geq \cdots \geq \left|T_t \cap [n]\right|$, and let $\tilde{t} \in [t]$ be the largest integer for which $\left|T_{\tilde{t}} \cap [n]\right| \geq 1$. By (7.25), there must exist $\tilde{p} \in [\tilde{t}]$ for which $|T_{\tilde{p}}| \geq 3$. Note that

$$\dim \operatorname{span}\{v_a : a \in T_{\tilde{p}}\} \geq \max\left\{2, \left|T_{\tilde{p}} \cap [n]\right|\right\}.$$

Since $\{v_a^{\otimes m} : a \in T_{\tilde{p}}\}$ is connected, it follows from our splitting theorem that

$$|T_{\tilde{p}}| \geq m(\max\left\{2, \left|T_{\tilde{p}} \cap [n]\right|\right\} - 1) + 2. \tag{7.28}$$

Now,

$$
\begin{aligned}
n + r &\geq \sum_{p \in [\tilde{t}]} |T_p| \\
&\geq \sum_{p \neq \tilde{p}} \left[m \left(\left| T_p \cap [n] \right| - 1 \right) + 2 \right] + m \left(\max \left\{ 2, \left| T_{\tilde{p}} \cap [n] \right| \right\} - 1 \right) + 2 \\
&= m \left(n - |T_{\tilde{p}} \cap [n]| \right) - (m-2) \left(\tilde{t} - 1 \right) + m \left(\max \left\{ 2, \left| T_{\tilde{p}} \cap [n] \right| \right\} - 1 \right) + 2 \\
&\geq m \left(n - |T_{\tilde{p}} \cap [n]| \right) - (m-2) \left(n - \left| T_{\tilde{p}} \cap [n] \right| \right) + m \left(\max \left\{ 2, \left| T_{\tilde{p}} \cap [n] \right| \right\} - 1 \right) + 2 \\
&= 2n - 2 \left| T_{\tilde{p}} \cap [n] \right| + m \left(\max \left\{ 2, \left| T_{\tilde{p}} \cap [n] \right| \right\} - 1 \right) + 2 \\
&\geq 2n + m - 2.
\end{aligned}
$$

The first line is obvious, the second follows from (7.28) and the fact that every multiset $\{v_a^{\otimes m} : a \in T_p\}$ is connected, the third is algebra, the fourth uses the fact that $\left|T_p \cap [n]\right| \geq 1$ for all $p \in [\tilde{t}]$, and the rest is algebra. This completes the proof. □

Theorem 71 immediately implies the following uniqueness result for non-Waring rank decompositions.

Corollary 72 (Uniqueness result for non-Waring rank decompositions)**.** *Let $n \geq 2$ and $m \geq 2$ be integers, let $\mathcal{W}$ be a vector space over a field $\mathbb{F}$ with $\mathrm{Char}(\mathbb{F}) = 0$ or $\mathrm{Char}(\mathbb{F}) > m$, let $\{v_a : a \in [n]\} \subseteq \mathcal{W} \setminus \{0\}$ be a multiset of non-zero vectors with $\text{k-rank}(v_a : a \in [n]) \geq 2$, let $\{\alpha_a : a \in [n]\} \subseteq \mathbb{F}^\times$ be a multiset of non-zero scalars, and let $d = \dim \operatorname{span}\{v_a : a \in [n]\}$. If*

$$2n + 1 \leq m + 2d - 2,$$

then $\sum_{a \in [n]} \alpha_a v_a^{\otimes m}$ constitutes a unique Waring rank decomposition. More generally, if

$$n + r + 1 \leq m + 2d - 2,$$

for some $r \geq n$, then $\sum_{a \in [n]} \alpha_a v_a^{\otimes m}$ is the unique symmetric decomposition of this tensor into at most r terms.

Note that the $r = n$ case of Corollary 72 improves Kruskal's theorem for symmetric decompositions as soon as $2d > m(k-2) + 4$, where $k = \text{k-rank}(v_a : a \in [n])$. This case of Corollary 72 is in fact contained in our generalization of Kruskal's theorem (Theorem 2), since for every subset $S \subseteq [n]$ with $2 \leq |S| \leq n$, it holds that

$$\begin{aligned} 2|S| &= 2n - 2\big|[n] \setminus S\big| \\ &\leq m + 2d - 2\big|[n] \setminus S\big| - 3 \\ &\leq m + 2d^S - 3 \\ &\leq m(d^S - 1) + 1, \end{aligned}$$

where $d^S = \dim \operatorname{span}\{v_a : a \in S\}$. This demonstrates that our generalization of Kruskal's theorem is stronger than Kruskal's theorem, even for symmetric tensor decompositions.

Our main result in this section is the $r > n$ case of Corollary 72, which yields uniqueness results for non-Waring rank decompositions of $\sum_{a\in[n]} \alpha_a v_a^{\otimes n}$. The following example illustrates this case in practice.

Example 73. It follows from Corollary 72 that for any positive integers $m \geq 3$ and $n \geq 2$, $\sum_{a\in[n]} e_a^{\otimes m}$ is the unique symmetric decomposition of this tensor into at most $m + n - 3$ terms.

It is natural to ask if Corollary 72 can be improved under further restrictions on $\text{k-rank}(v_a : a \in [n])$. At the end of Section 7.6.1 we prove that this cannot be done, at least in a particular parameter regime.

7.6.1 The inequality appearing in our uniqueness result is sharp

In this subsection we prove that the inequality (7.27) that appears in Theorem 71 cannot be improved, by constructing explicit multisets of symmetric product tensors that satisfy this bound with equality.

Let $\mathbb{F}$ be a field with $\operatorname{Char}(\mathbb{F}) = 0$. We will prove that for any choice of positive integers $m \geq 2$, $d \geq 2$, $r \geq d - 2$, and $n \geq d$ for which $n + r = m + 2d - 2$, there exist multisets of non-zero vectors E and F that satisfy the assumptions of Theorem 71. Note that the inequality $r \geq d - 2$ automatically holds when $r \geq n$, so this assumption does not restrict the parameter regime in which the inequality appearing in our uniqueness result (Corollary 72) is sharp as a consequence.

We first consider the case $d = 2$. Let $\{v_a^{\otimes m} : a \in [m+2]\} \subseteq \operatorname{Prod}(\mathbb{F}^2 : \cdots : \mathbb{F}^2)$ be a circuit of symmetric product tensors for which

$$\text{k-rank}(v_a : a \in [m+2]) = 2.$$

An example of such a circuit is given in [49], and reviewed in Section 7.3. So there exist non-zero scalars $\{\alpha_a : a \in [m+2]\} \subseteq \mathbb{F}^\times$ for which $\sum_{a\in[m+2]} \alpha_a v_a^{\otimes m} = 0$, and we can take the multisets $E = \{v_a : a \in [n]\}$ and $F = \{v_a : a \in [m+2] \setminus [n]\}$ to conclude.

For $d \geq 3$, let $\{v_a^{\otimes m} : a \in [m+2]\} \subseteq \operatorname{Prod}\left(\mathbb{F}^d : \cdots : \mathbb{F}^d\right)$ be the same multiset of symmetric product tensors as above, embedded in a larger space. Let

$$\{v_a : a \in [d+m] \setminus [m+2]\} \subseteq \mathbb{F}^d \setminus \{0\}$$

be any multiset of non-zero vectors for which

$$\dim \operatorname{span}\{v_a : a \in [d+m]\} = d$$

and

$$\text{k-rank}\{v_a : a \in [d+m]\} \geq 2,$$

which is guaranteed to exist since $\mathbb{F}$ is infinite. Since $r \geq d-2$, we can take the multisets

$$E = \{v_a : a \in [n-d+2]\} \sqcup \{v_a : a \in [d+m] \setminus [m+2]\}$$

and

$$F = \{v_a : a \in [m+2] \setminus [n-d+2]\} \sqcup \{v_a : a \in [d+m] \setminus [m+2]\}$$

to conclude.

Somewhat surprisingly, the inequality (7.27) is very nearly sharp even when the k-rank condition is tightened to $\text{k-rank}(v_a : a \in [n]) \geq k$ for some $k \geq 3$, under certain parameter constraints. More specifically, for any $k \in \{3, 4, \ldots, d-1\}$, it is almost sharp under the choice $n = d+1$ and $r = m+d-1$. Let

$$E = \{v_a : a \in [d+m] \setminus [m]\} \sqcup \left\{\sum_{a\in[k]} v_a\right\},$$

and

$$F = \{v_a : a \in [m]\} \sqcup \{v_a : a \in [d+m] \setminus [m+2]\} \sqcup \left\{\sum_{a\in[k]} v_a\right\}.$$

Here, $|E| + |F| = 2d + m$, exceeding our lower bound by 2. When $k = d$, take the same multisets E and F, with $\sum_{a\in[d]} v_a$ removed, to observe that our bound is sharp under the choice $n = d$ and $r = m+d-2$. Note that the k-rank is brought down to k because of a single vector in the multiset. This is a concrete demonstration of the fact that the k-rank is a very crude measure of genericity. We emphasize that this construction relies on the particular choice of parameters $n = d+1$, and $r = m+d-1$. It is possible that the inequality (7.27) could be significantly stengthened for other choices of n and r. Indeed, we have exhibited such an improvement for $r \leq n$ in Corollary 70.

7.6.2 Applications of non-rank uniqueness results

In this subsection, we identify potential applications of our results on uniqueness of non-rank decompositions. For concreteness, we focus on the symmetric case and our non-Waring rank uniqueness result in Corollary 72, however similar comments can be applied to our analogous results in Section 7.4.2 in the non-symmetric case.

We say a symmetric tensor v is *identifiable* if it has a unique Waring rank decomposition. For the purposes of this discussion, we will say that v is r*-identifiable* for some $r \geq \text{rank}(v)$ if the Waring rank decomposition of v is the unique symmetric decomposition of v into at most r terms (see Section 2). Corollary 72 provides a sufficient condition for a symmetric tensor v to be r-identifiable for $r > \text{rank}(v)$, and Example 73 demonstrates the existence of symmetric tensors satisfying this condition. We can thus define a hierarchy of identifiable symmetric tensors (of some fixed rank), where those that are r-identifiable for larger r can be thought of as "more identifiable." We suggest that studying this hierarchy could be a useful tool for studying symmetric tensor decompositions. For example, although most symmetric tensors of sub-generic rank are identifiable, it is notoriously difficult to find the rank decomposition of such tensors [81, 16, 41]. Perhaps one can leverage the additional structure of r-identifiable symmetric tensors to find efficient decompositions.

In applications, one often has a symmetric decomposition of a tensor, and wants to control the possible symmetric decompositions with fewer terms. Uniqueness results for non-rank decompositions can be turned around to apply in this setting: Suppose we know that if a symmetric decomposition into n terms satisfies some condition, call it C, then it is the unique symmetric decomposition into at most r terms, for some $r > n$. Then if one starts with a symmetric decomposition of a symmetric tensor v into r terms, she knows that there are no symmetric decompositions of v into $n < r$ terms that satisfies condition C. In this way, one can use a non-rank uniqueness result to control the possible decompositions of v into fewer than r symmetric product tensors. Applying this reasoning to our Corollary 72 simply yields a special case of Theorem 71. However, applying analogous reasoning to Corollary 64 in the non-symmetric case seems to produce new results.

7.7 Comparing our generalization of Kruskal's theorem to the uniqueness criteria of Domanov, De Lathauwer, and Sørensen

In this section we compare our generalization of Kruskal's theorem to uniqueness criteria obtained by Domanov, De Lathauwer, and Sørensen (DLS) in the case of three subsystems [51, 52, 53, 116, 115], which are the only previously known extensions of Kruskal's theorem that we are aware of. A drawback to the uniqueness criteria of DLS is that, similarly to Kruskal's theorem, they require the k-ranks to be above a certain threshold. In Section 7.7.1 we make this statement precise, and show by example that our generalization of Kruskal's theorem can certify uniqueness below this threshold. Moreover, in Section 7.7.2 we observe that our generalization of Kruskal's theorem contains many of the uniqueness criteria of DLS. The uniqueness criteria of DLS are spread across five papers, and can be difficult to keep track of. For clarity and future reference, in Theorem 75 we combine all of these criteria into a single statement. In Section 7.7.3 we use insight gained from this synthesization and our Theorem 2 as evidence to support a conjectural uniqueness criterion that would contain and unify every uniqueness criteria of DLS into a single, elegant statement.

For the remainder of this section, we fix a vector space $\mathcal{V} = \mathcal{V}_1 \otimes \mathcal{V}_2 \otimes \mathcal{V}_3$ over a field $\mathbb{F}$, and a multiset of product tensors

$$\{x_a : a \in [n]\} \subseteq \text{Prod}\,(\mathcal{V}_1 : \mathcal{V}_2 : \mathcal{V}_3)$$

with k-ranks $k_j = \text{k-rank}(x_{a,j} : a \in [n])$ for each $j \in [3]$. For each subset $S \subseteq [n]$ with $2 \leq |S| \leq n$ and index $j \in [3]$, we let

$$d_j^S = \dim \text{span}\{x_{a,j} : a \in [n]\}.$$

We also let $d_j = d_j^{[n]}$ for all $j \in [3]$.

7.7.1 Uniqueness below the k-rank threshold of DLS

All of the uniqueness criteria of DLS require the k-ranks to be above a certain threshold. In this subsection, we show by example that our generalization of Kruskal's theorem can certify uniqueness below this threshold.

Making this threshold precise, the uniqueness criteria of DLS cannot be applied whenever

$$\begin{aligned} &\min\{k_2,k_3\} \leq n-d_1+1,\\ \text{and } &\min\{k_1,k_3\} \leq n-d_2+1,\\ \text{and } &\min\{k_1,k_2\} \leq n-d_3+1. \end{aligned} \tag{7.29}$$

For example, if $k_2 = k_3 = 2$, then the uniqueness criteria of DLS can only certify uniqueness if $d_1 = n$. The following example shows that our generalization of Kruskal's theorem (Theorem 2) can certify uniqueness even if (7.29) holds.

Example 74. Consider the multiset of product tensors

$$\{\alpha_1 e_1^{\otimes 3}, \alpha_2 e_2^{\otimes 3}, \alpha_3 e_3^{\otimes 3}, \alpha_4 e_4^{\otimes 3}, \alpha_5 (e_2+e_3)\otimes(e_2+e_4)\otimes(e_1+e_4)\} \quad \text{for} \quad \alpha_1,\dots,\alpha_5 \in \mathbb{F}^\times.$$

In this example, $k_1 = k_2 = k_3 = 2$, $d_1 = d_2 = d_3 = 4$, and $n - d_j + 2 = 3$ for all $j \in [3]$, so (7.29) holds. Nevertheless, for arbitrary $\alpha_1,\dots,\alpha_5 \in \mathbb{F}^\times$, our generalization of Kruskal's theorem certifies that the sum of these product tensors constitutes a unique tensor rank decomposition. We note that uniqueness for $\alpha_2 = \dots = \alpha_5 = 1$ was proven in [52, Example 5.2], using a proof specific to this case, in order to demonstrate that their uniqueness criteria are not also necessary for uniqueness.

Example 74 shows that Theorem 2 is strictly stronger than Kruskal's theorem, and is independent of the uniqueness criteria of DLS. It is natural to ask if Theorem 2 is stronger than Kruskal's theorem even for symmetric tensor decompositions. We have observed in Section 7.6 that this is indeed the case.

7.7.2 Extending several uniqueness criteria of DLS

In this subsection, we observe that several of the uniqueness criteria of DLS are contained in our generalization of Kruskal's theorem, and prove a further, independent uniqueness criterion. The uniqueness criteria of DLS are numerous, and can be difficult to keep track of. To more easily analyze these criteria, in Theorem 75 we combine them all into a single statement.

Conditions U, H, C, and S

Here we introduce several different conditions on multisets of product tensors, which will make the uniqueness criteria of DLS easier to state, and also make them easier to relate to our generalization of Kruskal's theorem. We first recall Conditions U, H, and C from [51, 52]. For notational convenience, we have changed these definitions slightly from [51, 52]. For example, our Condition U is their Condition U_{n-d_1+2}, with the added condition that $k_1 \geq 2$. After reviewing Conditions U, H, and C, we introduce Condition S, which captures the conditions of our generalization of Kruskal's theorem in the case $m = 3$. Unlike Conditions U, H, and C, our Condition S does not appear in [51, 52], nor anywhere else that we are aware of.

For a vector $\alpha \in \mathbb{F}^n$, we let $\omega(\alpha)$ denote the number of non-zero entries in α.

Condition U. *It holds that $k_1 \geq 2$, and for all $\alpha \in \mathbb{F}^n$,*

$$\operatorname{rank}\Big[\sum_{a\in[n]} \alpha_a x_{a,2} \otimes x_{a,3}\Big] \geq \min\{\omega(\alpha), n - d_1 + 2\}. \tag{7.30}$$

Condition H. *It holds that $k_1 \geq 2$, and*

$$d_2^S + d_3^S - |S| \geq \min\{|S|, n - d_1 + 2\}$$

for all $S \subseteq [n]$ with $2 \leq |S| \leq n$.

Condition C takes a bit more work to describe. We use coordinates for this condition, in order to avoid having to introduce further multilinear algebra notation. For positive integers q, r, and t, and matrices

$$Y = (y_1, \dots, y_t) \in \mathrm{L}(\mathbb{F}^t, \mathbb{F}^q)$$
$$Z = (z_1, \dots, z_t) \in \mathrm{L}(\mathbb{F}^t, \mathbb{F}^r),$$

let

$$Y \odot Z = (y_1 \otimes z_1, \dots, y_t \otimes z_t) \in \mathrm{L}(\mathbb{F}^t, \mathbb{F}^{qr})$$

denote the *Khatri-Rao product* of Y and Z. Suppose $\mathcal{V}_j = \mathbb{F}^{d_j}$ for each $j \in [3]$, and consider the matrices

$$X_j = (x_{1,j}, \dots, x_{n,j}) \in \mathrm{L}(\mathbb{F}^n, \mathbb{F}^{d_j})$$

for $j \in [3]$. For a positive integer $s \leq d_j$, let $\mathcal{C}_s(X_j)$ be the $\binom{d_j}{s} \times \binom{n}{s}$ matrix of $s \times s$ minors of X_j, with rows and columns arranged according to the lexicographic order on the size s subsets of $[d_j]$ and $[n]$, respectively. Define the matrix

$$C_s = \mathcal{C}_s(X_2) \odot \mathcal{C}_s(X_3) \in \mathrm{L}(\mathbb{F}^{\binom{n}{s}}, \mathbb{F}^q),$$

where $q = \binom{d_2}{s}\binom{d_3}{s}$. Now we can state Condition C.

Condition C. *It holds that* $k_1 \geq 2$, $\min\{d_2, d_3\} \geq n - d_1 + 2$, *and*

$$\operatorname{rank}(C_{n-d_1+2}) = \tbinom{n}{n-d_1+2}\,.$$

To more easily compare our generalization of Kruskal's theorem to the uniqueness criteria of DLS, we give a name (Condition S) to the condition of our Theorem 2 in the case $m = 3$.

Condition S. *It holds that*

$$2|S| \leq d_1^S + d_2^S + d_3^S - 2$$

for all $S \subseteq [n]$ *with* $2 \leq |S| \leq n$.

These conditions are related to each other as follows:

$$\begin{tikzcd}
\text{Condition H} \arrow[r, Rightarrow] \arrow[dr, Rightarrow] & \text{Condition S} \\
& \text{Condition U} \\
\text{Condition C} \arrow[ur, Rightarrow] &
\end{tikzcd} \tag{7.31}$$

All of the implications in (7.31) except (Condition H $\Rightarrow$ Condition S) were proven in [51]. To see that Condition H $\Rightarrow$ Condition S, note that for any subset $S \subseteq [n]$ with $2 \leq |S| \leq n$, the condition $k_1 \geq 2$ implies

$$d_1^S \geq \max\{2, d_1 - (n - |S|)\},$$

so by Condition H,

$$\begin{aligned}
d_1^S + d_2^S + d_3^S &\geq \max\{2, d_1 - (n - |S|)\} + |S| + \min\{|S|, n - d_1 + 2\} \\
&\geq 2|S| + 2,
\end{aligned}$$

and Condition S holds. It is easy to find examples that certify Condition C $\not\Rightarrow$ Condition S. By Example 74, Condition S $\not\Rightarrow$ Condition U. In [51] it is asked whether Condition H $\Rightarrow$ Condition C. Condition U is theoretically computable, as it can be phrased as an ideal membership problem, however we are unaware of an efficient implementation. By comparison, Conditions C, H, and S are easy to check.

In the case of three subsystems, our Theorem 2 states that Condition S implies uniqueness. Since Condition H $\Rightarrow$ Condition S, then a corollary to Theorem 2 is that Condition H implies uniqueness. Similarly, Theorem 75 below states that Condition U + extra assumptions implies uniqueness. By (7.31), this implies that Condition H + the same extra assumptions implies uniqueness, and similarly, Condition C + the same extra assumptions implies uniqueness. Since we have proven that Condition H alone implies uniqueness, it is natural to ask whether Conditions C or U alone imply uniqueness. We reiterate this line of reasoning in Section 7.7.3, and pose this question formally.

Synthesizing the uniqueness criteria of DLS

The following theorem contains every uniqueness criterion of DLS for which we are aware of an efficient implementation. This theorem is stated in terms of Condition U to maintain generality, however only the implied statements in which Condition U is replaced by Conditions H or C have an efficient implementation. Note that our Theorem 2 generalizes the Condition H version of this theorem, to the statement that Condition S alone implies uniqueness (so in particular, Condition H alone implies uniqueness).

Theorem 75. *Suppose that Condition U holds, and any one of the following conditions holds:*

1. $$k_1 + \min\{k_2, k_3 - 1\} \geq n + 1.$$

2. *It holds that $k_2 \geq 2$ and for all $\alpha \in \mathbb{F}^n$,*
$$\mathrm{rank}\Big[\sum_{a \in [n]} \alpha_a x_{a,1} \otimes x_{a,3}\Big] \geq \min\{\omega(\alpha), n - d_2 + 2\}.$$
(Note that this is just Condition U with the first subsystem replaced by the second).

3. *There exists a subset $S \subseteq [n]$ with $0 \leq |S| \leq d_1$ such that the following three conditions hold:*

   (a) $$d_1^S = |S|.$$

   (b) $$d_2^{[n]\setminus S} = n - |S|.$$

(c) For any linear map $\Pi \in \mathrm{L}(\mathcal{V}_1)$ with $\ker(\Pi) = \operatorname{span}\{x_{a,1} : a \in S\}$, scalars $\alpha_1, \ldots, \alpha_n \in \mathbb{F}$, and index $b \in [n] \setminus S$ such that

$$\sum_{a \in [n] \setminus S} \alpha_a \Pi x_{a,1} \otimes x_{a,3} = \Pi x_{b,1} \otimes z$$

for some $z \in \mathcal{V}_{\sigma(3)}$, it holds that $\omega(\alpha) \leq 1$.

4. *There exists a permutation $\tau \in S_n$ for which the matrix*

$$X_1^\tau = (x_{\tau(1),1}, \ldots, x_{\tau(n),n})$$

has reduced row echelon form

$$Y = \left[\begin{array}{ccc|c} 1 & & & \\ & \ddots & & Z \\ & & 1 & \end{array}\right],$$

where $Z \in \mathrm{L}(\mathbb{F}^{n-d_1}, \mathbb{F}^{d_1})$ and the blank entries are zero. Furthermore, for each $a \in [d_1 - 1]$, the columns of the submatrix of Y with row index $\{a, a+1, \ldots, d_1\}$ and column index $\{a, a+1, \ldots, n\}$ have k-rank at least two.

5. $$k_1 = d_1.$$

6. *For all $\alpha \in \mathbb{F}^n$,*

$$\operatorname{rank}\Big[\sum_{a \in [n]} \alpha_a x_{a,2} \otimes x_{a,3}\Big] \geq \min\{\omega(\alpha), n - k_1 + 2\}.$$

(Note that this is a stronger statement than Condition U, as it replaces the quantity $n - d_1 + 2$ with the possibly larger quantity $n - k_1 + 2$.)

Then $\sum_{a \in [n]} x_a$ constitutes a unique tensor rank decomposition.

For each $i \in [5]$, we will refer to Theorem 75.i as the statement that Condition U and the i-th condition appearing in Theorem 75 imply uniqueness. Theorems 75.1 and 75.2 are Corollary 1.23 and Proposition 1.26 in [52, 53]. The Condition C version of Theorem 75.3 is stated in Theorem 2.2 in [115], although the proof is contained in [51, 52, 116]. Condition 3b in Theorem 75 can be formulated as checking the rank of a certain matrix (see [115]). Theorem 75.4 is a new result that we will prove (see Proposition 76 for a coordinate-free statement). The Condition C version of Theorems 75.5 and 75.6 are Theorems 1.6 and 1.7 in [53]. It is easy to see that our Theorem 75.4 contains Theorem 75.5, which in turn contains Theorem 75.6, by the arguments used in [53].

Most of these statements have previously only been formulated for $\mathbb{F} = \mathbb{R}$ or $\mathbb{F} = \mathbb{C}$, however in all of these cases the proof can be adapted to hold over an arbitrary field. The first step in proving all of these statements is to show that Condition U implies uniqueness in the first subsystem. This is Proposition 4.3 in [51], and it is proven using Kruskal's permutation lemma [79] (the proof of the permutation lemma in [81] holds word-for-word over an arbitrary field). In fact, uniqueness in the first subsystem holds even with the assumption $k_1 \geq 2$ removed from Condition U [51].

A less-restrictive condition than Condition U, which we would call Condition W, also appears in [51, 52], and is the same as Condition U except that it only requires (7.30) to hold when $\alpha = (f(x_{1,1}), \dots, f(x_{n,1}))$ for some linear functional $f \in \mathcal{V}_1^*$. We note that Theorem 75 also holds with Condition U replaced by Condition W. Although the Condition W version of Theorem 75 is slightly stronger than the Condition U version, we are not aware of an efficient algorithm to check either Condition U or Condition W, and the existence of such an algorithm seems unlikely.

We conclude this subsection by proving Theorem 75.4. For this we require the following proposition, which restates Condition 4 in a coordinate-free manner.

Proposition 76. *Condition 4 in Theorem 75 holds if and only if there exists a permutation $\tau \in S_n$ such that for each $a \in [d_1 - 1]$ there is a linear operator $\Pi_a \in \mathrm{L}(\mathcal{V}_1)$ for which*

$$\Pi_a(x_{\tau(b),1}) = 0$$

for all $b \in [a-1]$, and

$$\text{k-rank}(\Pi_a x_{\tau(a),1}, \dots, \Pi_a x_{\tau(n),1}) \geq 2. \tag{7.32}$$

Proof. Assume without loss of generality that $\mathcal{V}_1 = \mathbb{F}^{d_1}$. To see that the first statement implies the second, for each $a \in [d_1 - 1]$ let $\Pi_a = D_a P$, where $P \in \mathrm{L}(\mathbb{F}^{d_1})$ is the invertible matrix for which $P X_1^\tau = Y$, and $D_a \in \mathrm{L}(\mathbb{F}^{d_1})$ is the diagonal matrix with the first $a-1$ entries zero and the remaining entries 1. It is easy to verify that (7.32) holds.

Conversely, suppose that the reduced row echelon form of X_1^τ, given by $P X_1^\tau$ for some invertible matrix $P \in \mathrm{L}(\mathbb{F}^{d_1})$, does not have the specified form. Then there exists $a \in [d_1 - 1]$ for which the columns of $D_a P X_1^\tau$ have k-rank at most one. Any matrix $\Pi_a \in \mathrm{L}(\mathbb{F}^{d_1})$ for which $\Pi_a(x_{\tau(b),1}) = 0$ for all $b \in [a-1]$ satisfies

$$\Pi_a = \Pi_a P^{-1} D_a P.$$

Since the k-rank is non-increasing under matrix multiplication from the left, (7.32) does not hold. □

With Proposition 76 in hand, we can now prove Theorem 75.4.

Proof of Theorem 75.4. The question of whether or not the decomposition $\sum_{a\in[n]} x_a$ constitutes a unique tensor rank decomposition is invariant under permutations $\tau \in S_n$ of the tensors, so it suffices to prove the statement under the assumption that the permutation τ appearing in Condition 4 is trivial. We prove the statement by induction on d_1. If $d_1 = 2$, then Condition U implies $k_2 = k_3 = n$, so uniqueness follows from Kruskal's theorem. For $d_1 > 2$, suppose $\sum_{a\in[n]} x_a = \sum_{a\in[r]} y_a$ for some non-negative integer $r \leq n$ and multiset of product tensors

$$\{y_a : a \in [r]\} \subseteq \operatorname{Prod}(\mathcal{V}_1 : \mathcal{V}_2 : \mathcal{V}_3)\,.$$

By Proposition 4.3 in [51] (or rather, the extension of this result to an arbitrary field), $r = n$, and there exists a permutation $\sigma \in S_n$ and nonegative integers $\alpha_1, \dots, \alpha_n \in \mathbb{F}^\times$ such that $\alpha_a x_{a,1} = y_{\sigma(a),1}$ for all $a \in [n]$. Let $\Pi_1 \in \mathrm{L}(\mathcal{V}_1)$ be any operator for which $\ker(\Pi_1) = \operatorname{span}\{x_{a,1}\}$ and (7.32) holds (recall that τ is trivial). Then

$$\sum_{a\in[n]\setminus\{1\}} (\Pi_1 x_{a,1}) \otimes x_{a,2} \otimes x_{a,3} = \sum_{a\in[n]\setminus\{1\}} (\alpha_a \Pi_1 x_{a,1}) \otimes y_{\sigma(a),2} \otimes y_{\sigma(a),3}.$$

Now, $\dim \operatorname{span}\{\Pi_1 x_{a,1} : a \in [n] \setminus \{1\}\} = d_1 - 1$, and Condition U again holds for the multiset of product tensors

$$\{(\Pi_1 x_{a,1}) \otimes x_{a,2} \otimes x_{a,3} : a \in [n] \setminus \{1\}\}.$$

Furthermore, these product tensors again satisfy Condition 4 of Theorem 75, so by the induction hypothesis

$$(\Pi_1 x_{a,1}) \otimes x_{a,2} \otimes x_{a,3} = (\alpha_a \Pi_1 x_{a,1}) \otimes y_{\sigma(a),2} \otimes y_{\sigma(a),3} \quad \text{for all} \quad a \in [n] \setminus \{1\}.$$

It follows that $x_a = y_{\sigma(a)}$ for all $a \in [n] \setminus \{1\}$, so $x_1 = y_{\sigma(1)}$. This completes the proof. □

7.7.3 Conjectural generalization of all uniqueness criteria of DLS

In the case of three subsystems, our generalization of Kruskal's theorem states that Condition S implies uniqueness. Since Condition H ⇒ Condition S, then a corollary to Theorem 2 is that Condition H implies uniqueness. Similarly, Theorem 75 above states that Condition U + extra assumptions implies uniqueness, which implies that Condition H + the same extra assumptions implies uniqueness. Since we have proven that Condition H alone implies uniqueness, it is natural to ask whether Condition U alone implies uniqueness. We now state this question formally. A positive answer to Question 77 would generalize and unify all of the uniqueness criteria of DLS (synthesized in Theorem 75) into a single, elegant statement.

Question 77. *Does Condition U imply that $\sum_{a\in[n]} x_a$ constitutes a unique tensor rank decomposition?*

Chapter Acknowledgments I thank Edoardo Ballico, Luca Chiantini, Matthias Christandl, Harm Derksen, Dragomir Djoković, Ignat Domanov, Timothy Duff, Joshua A. Grochow, Nathaniel Johnston, Joseph M. Landsberg, Lieven De Lathauwer, Chi-Kwong Li, Daniel Puzzuoli, Pierpaola Santarsiero, William Slofstra, Hans De Sterck, and John Watrous for helpful discussions and comments on drafts of this manuscript. I thanks Luca Chiantini and Pierpaola Santarsiero for helpful feedback on Sections 7.5 and 7.6. In previous work [89], I conjectured a (slightly weaker) "non-minimal version" of Theorem 4. I thank Harm Derksen for suggesting the splitting version that appears here. I thank Dragomir Djoković for first suggesting a connection to Kruskal's theorem, and for suggesting that these results might hold over an arbitrary field. I thank Joshua A. Grochow for first asking about uniqueness results for non-rank decompositions, which inspired our results in Sections 7.4.2 and 7.6.

Chapter 8

Conclusion

In this thesis, we used algebraic-geometric and combinatorial techniques to study tensor decompositions, with applications in quantum information theory, algebraic complexity theory, and algebraic statistics.

- In Chapter 3 we used algebraic-geometric techniques to determine "how useful" a given quantum state is for local state discrimination, in terms of the dimension of a certain algebraic variety. It would be interesting to determine the utility of quantum states for other fundamental tasks in the multi-party setting.

- In Chapter 4 we generalized previous notions of entangled subspaces to higher, multiparty entanglement. We determined the maximum dimension of such r-entangled subspaces, derived explicit constructions of such subspaces, and observed a close connection to entanglement witnesses. It would be interesting to study a robust variant of r-entangled subspaces: A subspace for which every state in the subspace is at least ϵ far from any state of border rank $\leq r$. Such study could have implications for the bound entanglement problem, and for constructing better entanglement witnesses [75].

- In Chapter 5 we introduced number-theoretic and algebraic-geometric techniques for studying the stabilizer rank, and obtained simpler proofs of the best-known lower bounds on stabilizer rank and approximate stabilizer rank, up to a log factor. We also identified a potential avenue to prove better bounds, by finding states with low T-count and long exponentially increasing sequences. It would be nice to know if this is a dead end, boulevard, or something in-between.

- In Chapter 6 we studied decomposable correlation matrices, (or, as they are now called in the literature, *(non)-absolutely entangled sets*. It would be nice to get a better picture of which correlation matrices are decomposable. The simplest question we have left open is whether every 4×4 correlation matrix is 2-decomposable.

- In Chapter 7 we proved what we call the *splitting theorem:* a matroidal statement about sets of product tensors. We used the splitting theorem to generalize Kruskal's theorem on uniqueness of tensor rank decompositions. A decomposition of a tensor v into product tensors is completely described by the set of product tensors that appear in the decomposition. Since matroid theory is dedicated to studying (abstractions of) sets of vectors, a natural approach to studying tensor decompositions is to characterize the matroidal structure of sets of product tensors. Despite this, and quite surprisingly, this angle remains largely unexplored. We think of the splitting theorem as a first observation in this direction. It would be interesting to see if any further matroidal structure for sets of product tensors can be obtained.

From entanglement theory to statistical machine learning, scientists from different disciplines are encountering tensor decompositions and studying them from different angles. It is hoped that this interdisciplinary thesis fosters increased communication and sharing of knowledge between communities.

References

[1] Most quantum states are too entangled to be useful as computational resources. *Physical Review Letters*, 102:190501, 2009.

[2] Hirotachi Abo, Giorgio Ottaviani, and Chris Peterson. Induction for secant varieties of Segre varieties. *Transactions of the American Mathematical Society*, 361(2):767–792, 2008.

[3] James Alexander and André Hirschowitz. Polynomial interpolation in several variables. *Journal of Algebraic Geometry*, 4(2):201–222, 1995.

[4] Elizabeth S. Allman, Catherine Matias, and John A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *The Annals of Statistics*, 37(6A):3099–3132, 2009.

[5] Jinpeng An. Rigid geometric structures, isometric actions, and algebraic quotients. *Geometriae Dedicata*, 157(1):153–185, 2012.

[6] Elena Angelini, Luca Chiantini, and Nick Vannieuwenhoven. Identifiability beyond Kruskal's bound for symmetric tensors of degree 4. *Atti della Accademia Nazionale dei Lincei, Classe di Scienze Fisiche, Matematiche e Naturali, Rendiconti Lincei Matematica E Applicazioni*, 29, 2018.

[7] Juan Miguel Arrazola and Norbert Lütkenhaus. Quantum fingerprinting with coherent states and a constant mean number of photons. *Physical Review A*, 89:062305, 2014.

[8] Remigiusz Augusiak, Jordi Tura, and Maciej Lewenstein. A note on the optimality of decomposable entanglement witnesses and completely entangled subspaces. *Journal of Physics A: Mathematical and Theoretical*, 44(21):212001, 2011.

[9] Edoardo Ballico. Linearly dependent and concise subsets of a Segre variety depending on k factors. *Bulletin of the Korean Mathematical Society*, 58(1):253–267, 2021.

[10] Edoardo Ballico, Alessandra Bernardi, Luca Chiantini, and Elena Guardo. Bounds on the tensor rank. *Annali di Matematica Pura ed Applicata (1923 -)*, 197(6):1771–1785, 2018.

[11] Somshubhro Bandyopadhyay, Saronath Halder, and Michael Nathanson. Entanglement as a resource for local state discrimination in multipartite systems. *Physical Review A*, 94:022311, 2016.

[12] M Barbieri, F De Martini, G Di Nepi, Paolo Mataloni, GM D'Ariano, and C Macchiavello. Detection of entanglement with polarized photons: Experimental realization of an entanglement witness. *Physical Review Letters*, 91(22):227901, 2003.

[13] Karin Baur, Jan Draisma, and Willem A. de Graaf. Secant dimensions of minimal orbits: Computations and conjectures. *Experimental Mathematics*, 16(2):239–250, 2007.

[14] Alessandra Bernardi, Enrico Carlini, Maria Catalisano, Alessandro Gimigliano, and Alessandro Oneto. The hitchhiker guide to: Secant varieties and tensor decomposition. *Mathematics*, 6(12):314, 2018.

[15] Alessandra Bernardi and Davide Vanzo. A new class of non-identifiable skew-symmetric tensors. *Annali di Matematica Pura ed Applicata (1923-)*, 197(5):1499–1510, 2018.

[16] Aditya Bhaskara, Moses Charikar, Ankur Moitra, and Aravindan Vijayaraghavan. Open problem: Tensor decompositions: Algorithms up to the uniqueness threshold? In *Conference on Learning Theory*, pages 1280–1282. Proceedings of Machine Learning Research, 2014.

[17] Aditya Bhaskara, Moses Charikar, and Aravindan Vijayaraghavan. Uniqueness of tensor decompositions with applications to polynomial identifiability. *Journal of Machine Learning Research*, (35):742–778, 2014.

[18] B. V. Rajarama Bhat. A completely entangled subspace of maximal dimension. *International Journal of Quantum Information*, 04(02):325–330, 2006.

[19] Ada Boralevi. A note on secants of grassmannians. *Rendiconti dell'Istituto di Matematica dell'Universita di Trieste*, 45:67–72, 2013.

[20] Armand Borel. *Linear Algebraic Groups*. Graduate Texts in Mathematics. Springer New York, 2012.

[21] Max Born and Emil Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013.

[22] Sougato Bose, Anupam Mazumdar, Gavin W Morley, Hendrik Ulbricht, Marko Toroš, Mauro Paternostro, Andrew A Geraci, Peter F Barker, MS Kim, and Gerard Milburn. Spin entanglement witness for quantum gravity. *Physical Review Letters*, 119(24):240401, 2017.

[23] Mohamed Bourennane, Manfred Eibl, Christian Kurtsiefer, Sascha Gaertner, Harald Weinfurter, Otfried Gühne, Philipp Hyllus, Dagmar Bruß, Maciej Lewenstein, and Anna Sanpera. Experimental detection of multipartite entanglement using witness operators. *Physical Review Letters*, 92(8):087902, 2004.

[24] Michel Boyer, Rotem Liss, and Tal Mor. Geometry of entanglement in the Bloch sphere. *Physical Review A*, 95:032308, 2017.

[25] Kamil Brádler, Patrick Hayden, Dave Touchette, and Mark Wilde. Trade-off capacities of the quantum Hadamard channels. *Physical Review A*, 81:062312, 2010.

[26] Sergey Bravyi, Dan Browne, Padraic Calpin, Earl Campbell, David Gosset, and Mark Howard. Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum*, 3:181, 2019.

[27] Sergey Bravyi and David Gosset. Improved classical simulation of quantum circuits dominated by Clifford gates. *Physical Review Letters*, 116:250501, 2016.

[28] Sergey Bravyi, Graeme Smith, and John A. Smolin. Trading classical and quantum computational resources. *Physical Review X*, 6:021043, 2016.

[29] Michel Brion. Introduction to actions of algebraic groups. *Les cours du CIRM*, 1(1):1–22, 2010.

[30] Harry Buhrman, Richard Cleve, John Watrous, and Ronald De Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, 2001.

[31] Harry Buhrman, Richard Cleve, and Avi Wigderson. Quantum vs. classical communication and computation. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, pages 63–68, New York, NY, USA, 1998. ACM.

[32] Peter Bürgisser, Michael Clausen, and Mohammad A Shokrollahi. *Algebraic complexity theory*, volume 315. Springer Science & Business Media, 2013.

[33] Yu Cai, Baichu Yu, Pooja Jayachandran, Nicolas Brunner, Valerio Scarani, and Jean-Daniel Bancal. Entanglement for any definition of two subsystems. *Physical Review A*, 103:052432, 2021.

[34] Anthony Chefles, Richard Jozsa, and Andreas Winter. On the existence of physical transformations between sets of quantum states. *International Journal of Quantum Information*, 2(01):11–21, 2004.

[35] Jianxin Chen, Hillary Dawkins, Zhengfeng Ji, Nathaniel Johnston, David Kribs, Frederic Shultz, and Bei Zeng. Uniqueness of quantum states compatible with given measurement results. *Physical Review A*, 88:012109, 2013.

[36] Lin Chen, Eric Chitambar, Runyao Duan, Zhengfeng Ji, and Andreas Winter. Tensor rank and stochastic entanglement catalysis for multipartite pure states. *Physical Review Letters*, 105:200501, 2010.

[37] Lin Chen, Yu Yang, and Wai-Shing Tang. Schmidt number of bipartite and multipartite states under local projections. *Quantum Information Processing*, 16:75, 2017.

[38] Luca Chiantini. *Hilbert Functions and Tensor Analysis*, pages 125–151. Springer International Publishing, Cham, 2019.

[39] Luca Chiantini and Ciro Ciliberto. On the concept of k-secant order of a variety. *Journal of the London Mathematical Society*, 73(2):436–454, 04 2006.

[40] Luca Chiantini, Giorgio Ottaviani, and Nick Vannieuwenhoven. Effective criteria for specific identifiability of tensors and forms. *SIAM Journal on Matrix Analysis and Applications*, 38(2):656–681, 2017.

[41] Luca Chiantini, Giorgio Ottaviani, and Nick Vannieuwenhoven. On generic identifiability of symmetric tensors of subgeneric rank. *Transactions of the American Mathematical Society*, 369(6):4021–4042, 2017.

[42] Andrzej Cichocki, Danilo Mandic, Lieven De Lathauwer, Guoxu Zhou, Qibin Zhao, Cesar Caiafa, and Huy Anh Phan. Tensor decompositions for signal processing applications: From two-way to multiway component analysis. *IEEE signal processing magazine*, 32(2):145–163, 2015.

[43] Richard Cleve, William Slofstra, Falk Unger, and Sarvagya Upadhyay. Perfect parallel repetition theorem for quantum XOR proof systems. *Computational Complexity*, 17(2):282–299, 2008.

[44] Collette R. Coullard and Lisa Hellerstein. Independence and port oracles for matroids, with an application to computational learning theory. *Combinatorica*, 16(2):189–208, 1996.

[45] Toby Cubitt, Ashley Montanaro, and Andreas Winter. On the dimension of subspaces with bounded schmidt rank. *Journal of Mathematical Physics*, 49(2):022107, 2008.

[46] Julio I. de Vincente, Cornelia Spee, and Barbara Kraus. Maximally entangled set of multipartite quantum states. *Physical Review Letters*, 111:110502, 2013.

[47] Jeroen Dehaene and Bart De Moor. Clifford group, stabilizer states, and linear and quadratic operations over GF(2). *Physical Review A*, 68:042318, 2003.

[48] Maciej Demianowicz and Remigiusz Augusiak. From unextendible product bases to genuinely entangled subspaces. *Physical Review A*, 98:012313, 2018.

[49] Harm Derksen. Kruskal's uniqueness inequality is sharp. *Linear Algebra and its Applications*, 438(2):708 – 712, 2013.

[50] Igor Devetak and Peter Shor. The capacity of a quantum channel for simultaneous transmission of classical and quantum information. *Communications in Mathematical Physics*, 256(2):287–303, 2005.

[51] Ignat Domanov and Lieven De Lathauwer. On the uniqueness of the canonical polyadic decomposition of third-order tensors—Part I: Basic results and uniqueness of one factor matrix. *SIAM Journal on Matrix Analysis and Applications*, 34(3):855–875, 2013.

[52] Ignat Domanov and Lieven De Lathauwer. On the uniqueness of the canonical polyadic decomposition of third-order tensors—Part II: Uniqueness of the overall decomposition. *SIAM Journal on Matrix Analysis and Applications*, 34(3):876–903, 2013.

[53] Ignat Domanov and Lieven De Lathauwer. Canonical polyadic decomposition of third-order tensors: Reduction to generalized eigenvalue decomposition. *SIAM Journal on Matrix Analysis and Applications*, 35(2):636–660, 2014.

[54] Wolfgang Dür, Guifre Vidal, and J. Ignacio Cirac. Three qubits can be entangled in two inequivalent ways. *Physical Review A*, 62:062314, 2000.

[55] Shaun M. Fallat. Bidiagonal factorizations of totally nonnegative matrices. *The American Mathematical Monthly*, 108(8):697–712, 2001.

[56] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. John Wiley and Sons, third edition, 1991.

[57] Jay Goldman and Gian-Carlo Rota. On the foundations of combinatorial theory IV: Finite vector spaces and Eulerian generating functions. Technical report, Harvard University, 1970.

[58] Daniel Gottesman. The Heisenberg representation of quantum computers. In *22nd International Colloquium on Group Theoretical Methods in Physics*, pages 32–43, 1998.

[59] Gilad Gour and Nolan R. Wallach. Entanglement of subspaces and error-correcting codes. *Physical Review A*, 76:042309, 2007.

[60] Janusz Grabowski, Marek Kuś, and Giuseppe Marmo. Segre maps and entanglement for multipartite systems of indistinguishable particles. *Journal of Physics A: Mathematical and Theoretical*, 45(10):105301, 2012.

[61] Daniel R. Grayson and Michael E. Stillman. Macaulay2, a software system for research in algebraic geometry. Available at http://www.math.uiuc.edu/Macaulay2/.

[62] David Gross. Hudson's theorem for finite-dimensional quantum systems. *Journal of Mathematical Physics*, 47(12):122107, 2006.

[63] Kil-Chan Ha and Seung-Hyeok Kye. Multi-partite separable states with unique decompositions and construction of three qubit entanglement with positive partial transpose. *Journal of Physics A: Mathematical and Theoretical*, 48(4):045303, 2015.

[64] Uffe Haagerup and Magdalena Musat. Factorization and dilation problems for completely positive maps on von Neumann algebras. *Communications in Mathematical Physics*, 303(2):555–594, 2011.

[65] Joe Harris. *Algebraic Geometry: A First Course*. Graduate Texts in Mathematics. Springer New York, 2013.

[66] Robin Hartshorne. *Algebraic Geometry.* Graduate Texts in Mathematics. Springer New York, 2013.

[67] Patrick Hayden, Debbie W. Leung, and Andreas Winter. Aspects of generic entanglement. *Communications in Mathematical Physics*, 265(1):95–117, 2006.

[68] Sven M. Hein, Jens Eisert, and Hans J. Briegel. Multiparty entanglement in graph states. *Physical Review A*, 69:062311, Jun 2004.

[69] Teiko Heinosaari, Luca Mazzarella, and Michael M. Wolf. Quantum tomography under prior information. *Communications in Mathematical Physics*, 318(2):355–374, 2013.

[70] Roger Horn and Charles Johnson. *Matrix Analysis.* Cambridge University Press, 2013.

[71] Zejun Huang, Chi-Kwong Li, Edward Poon, and Nung-Sing Sze. Physical transformations between quantum states. *Journal of mathematical physics*, 53(10):102209, 2012.

[72] James E. Humphreys. *Linear Algebraic Groups.* Graduate Texts in Mathematics. Springer New York, 2012.

[73] Anthony Iarrobino and Vassil Kanev. *Power sums, Gorenstein algebras, and determinantal loci.* Springer Science & Business Media, 1999.

[74] Nathaniel Johnston. Characterizing operations preserving separability measures via linear preserver problems. *Linear and Multilinear Algebra*, 59(10):1171–1187, 2011.

[75] Nathaniel Johnston. *Norms and Cones in the Theory of Quantum Entanglement.* PhD thesis, University of Guelph, 2012.

[76] Nathaniel Johnston. Non-positive-partial-transpose subspaces can be as large as any entangled subspace. *Physical Review A*, 87(6), 2013.

[77] Nathaniel Johnston, Benjamin Lovitz, and Daniel Puzzuoli. The non-m-positive dimension of a positive linear map. *Quantum*, 3:172, 2019.

[78] Wilhelmus Petrus Krijnen. *The analysis of three-way arrays by constrained PARAFAC methods.* DSWO Press, Leiden University, 1993.

[79] Joseph 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.

[80] Joseph Landsberg. Tensors: Geometry and applications. volume 128 of *Graduate Studies in Mathematics*. AMS, 2011.

[81] Joseph Landsberg. *Tensors: Geometry and Applications*. Graduate studies in mathematics. American Mathematical Society, 2012.

[82] Joseph Landsberg and Laurent Manivel. On the ideals of secant varieties of segre varieties. *Foundations of Computational Mathematics*, 4(4):397–422, 2004.

[83] Joseph Landsberg and Mateusz Michałek. Towards finding hay in a haystack: explicit tensors of border rank greater than 2.02m in $\mathbb{C}^m \otimes \mathbb{C}^m \otimes \mathbb{C}^m$. *arXiv preprint*, cs.CC/1912.11927, 2019.

[84] Lieven De Lathauwer. A short introduction to tensor-based methods for factor analysis and blind source separation. *ISPA 2011 - 7th International Symposium on Image and Signal Processing and Analysis*, 2011.

[85] Chi-Kwong Li and Bit-Shun Tam. A note on extreme correlation matrices. *SIAM Journal on Matrix Analysis and Applications*, 15(3):903–908, 1994.

[86] Mao-Sheng Li and Man-Hong Yung. Absolutely entangled set of pure states. *arXiv preprint*, quant-ph/2011.04903, 2020.

[87] Xiangqian Liu and Nicholas Sidiropoulos. Cramér-Rao lower bounds for low-rank decomposition of multidimensional arrays. *IEEE Transactions on Signal Processing*, 49(9):2074–2086, 2001.

[88] Benjamin Lovitz. Toward a generalization of Kruskal's theorem on tensor decomposition. *arXiv preprint*, math.CO/1812.00264, 2018.

[89] Benjamin Lovitz. Toward an analog of Kruskal's theorem on tensor decomposition. *arXiv preprint*, math.CO/1812.00264v1, 2018.

[90] Benjamin Lovitz. On decomposable correlation matrices. *Linear and Multilinear Algebra*, 69(11):2115–2129, 2021.

[91] Benjamin Lovitz and Nathaniel Johnston. Entangled subspaces and generic local state discrimination with pre-shared entanglement. *arXiv preprint*, quant-ph/2010.02876, 2020.

[92] Benjamin Lovitz and Nathaniel Johnston. Entangled-subspaces-code. `https://github.com/benjamin-lovitz/Entangled-subspaces-code`, 2020.

[93] Benjamin Lovitz and Norbert Lütkenhaus. Families of quantum fingerprinting protocols. *Physical Review A*, 97:032340, 2018.

[94] Benjamin Lovitz and Fedor Petrov. A generalization of Kruskal's theorem on tensor decomposition. *arXiv preprint*, math.CO/2103.15633, 2021.

[95] Benjamin Lovitz and Vincent Steffan. New techniques for bounding stabilizer rank. *arXiv preprint*, quant-ph/2110.07781, 2021.

[96] Ashutosh S. Marwah and Norbert Lütkenhaus. Characterization of gram matrices of multimode coherent states. *Physical Review A*, 99:012346, 2019.

[97] David Petrie Moulton. Representing powers of numbers as subset sums of small sets. *Journal of Number Theory*, 89(2):193–211, 2001.

[98] Michael Nielsen. Conditions for a class of entanglement transformations. *Physical Review Letters*, 83:436–439, 1999.

[99] James G. Oxley. *Matroid theory*, volume 3. Oxford University Press, USA, 2006.

[100] Kalyanapuram Rangachari Parthasarathy. On the maximal dimension of a completely entangled subspace for finite level quantum systems. *Proceedings Mathematical Sciences*, 114(4):365–374, 2004.

[101] Shir Peleg, Amir Shpilka, and Ben Lee Volk. Lower bounds on stabilizer rank. *arXiv preprint*, quant-ph/2106.03214, 2021.

[102] Hammam Qassim. *Classical simulations of quantum systems using stabilizer decompositions*. PhD thesis, University of Waterloo, 2021.

[103] Hammam Qassim, Hakop Pashayan, and David Gosset. Improved upper bounds on the stabilizer rank of magic states. *Quantum*, 5:606, 2021.

[104] Stephan Rabanser, Oleksandr Shchur, and Stephan Günnemann. Introduction to tensor decompositions and their applications in machine learning. *arXiv preprint*, stat.ML/1711.10781, 2017.

[105] Aleksandr A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical notes of the Academy of Sciences of the USSR*, 41:333–338, 1987.

[106] John A. Rhodes. A concise proof of Kruskal's theorem on tensor decomposition. *Linear Algebra and its Applications*, 432(7):1818 – 1824, 2010.

[107] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton Landmarks in Mathematics and Physics. Princeton University Press, 2015.

[108] Anna Sanpera, Dagmar Bruß, and Maciej Lewenstein. Schmidt-number witnesses and bound entanglement. *Physical Review A*, 63:050301, 2001.

[109] Andrew J. Scott. Multipartite entanglement, quantum-error-correcting codes, and entangling power of quantum evolutions. *Physical Review A*, 69:052330, 2004.

[110] Nicholas Sidiropoulos and Rasmus Bro. On the uniqueness of multilinear decomposition of n-way arrays. *Journal of Chemometrics: A Journal of the Chemometrics Society*, 14(3):229–239, 2000.

[111] Nicholas Sidiropoulos, Lieven De Lathauwer, Xiao Fu, Kejun Huang, Evangelos E. Papalexakis, and Christos Faloutsos. Tensor decomposition for signal processing and machine learning. *IEEE Transactions on Signal Processing*, 65(13):3551–3582, 2017.

[112] Łukasz Skowronek, Erling Størmer, and Karol Życzkowski. Cones of positive maps and their duality relations. *Journal of Mathematical Physics*, 50:062106, 2009.

[113] Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, pages 77–82, 1987.

[114] Roman Smolensky. On representations by low-degree polynomials. In *Proceedings of the 1993 IEEE 34th Annual Foundations of Computer Science*, pages 130–138, 1993.

[115] Mikael Sørensen and Lieven De De Lathauwer. Coupled canonical polyadic decompositions and (coupled) decompositions in multilinear rank-(L_r,n,L_r,n,1) terms—Part I: Uniqueness. *SIAM Journal on Matrix Analysis and Applications*, 36(2):496–522, 2015.

[116] Mikael Sørensen and Lieven De Lathauwer. New uniqueness conditions for the canonical polyadic decomposition of third-order tensors. *SIAM Journal on Matrix Analysis and Applications*, 36(4):1381–1403, 2015.

[117] Volker Strassen. Gaussian elimination is not optimal. *Numerische mathematik*, 13(4):354–356, 1969.

[118] Volker Strassen. Vermeidung von divisionen. *Journal für die reine und angewandte Mathematik*, 264:184–202, 1973.

[119] Volker Strassen. Rank and optimal computation of generic tensors. *Linear Algebra and its Applications*, 52-53:645 – 685, 1983.

[120] Volker Strassen. Relative bilinear complexity and matrix multiplication. *Journal für die reine und angewandte Mathematik*, 375/376:406–443, 1987.

[121] Barbara M. Terhal and Paweł Horodecki. Schmidt number for density matrices. *Physical Review A*, 61:040301(R), 2000.

[122] Dave Touchette, Benjamin Lovitz, and Norbert Lütkenhaus. Practical quantum appointment scheduling. *Physical Review A*, 97:042320, 2018.

[123] Boris Tsirelson. Quantum analogues of the Bell inequalities. The case of two spatially separated domains. *Journal of Soviet Mathematics*, 36(4):557–570, 1987.

[124] Ledyard R. Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3):279–311, 1966.

[125] Maarten Van den Nest. Classical simulation of quantum computation, the Gottesman-Knill theorem, and slightly beyond. *Quantum Information and Computation*, 10(3):258–271, 2010.

[126] Nolan R. Wallach. An unentangled Gleason's theorem. *Contemporary Mathematics*, 305:291–298, 2002.

[127] John Watrous. *The Theory of Quantum Information.* Cambridge University Press, 2018.

[128] Roy Westwick. Transformations on tensor spaces. *Pacific Journal of Mathematics*, 23(3):613–620, 1967.

[129] Andreas Winter and Dong Yang. Potential capacities of quantum channels. *IEEE Transactions on Information Theory*, 62(3):1415–1424, 2016.

[130] Baichu Yu, Pooja Jayachandran, Adam Burchardt, Yu Cai, Nicolas Brunner, and Valerio Scarani. Absolutely entangled sets of pure states for bipartitions and multipartitions. *Physical Review A*, 104:032414, 2021.

[131] Nengkun Yu, Runyao Duan, and Quanhua Xu. Bounds on the distance between a unital quantum channel and the convex hull of unitary channels. *IEEE Transactions on Information Theory*, 63(2):1299–1310, 2017.

APPENDICES

Appendix A

Appendix to Chapter 3

A.1 Proofs of facts from Section 3.1

In this appendix we prove Facts 5, 6, 8, and 9, which were presented in Section 3.1.

Fact 5. *A generic n-tuple of states spans a generic projective $(n-1)$-plane, and vice versa. In more details, let n be a positive integer, let $\mathcal{V}$ be a $\mathbb{C}$-vector space, and let*

$$\tilde{\pi} : \mathbb{P}(\mathcal{V})^{\times n} \dashrightarrow \mathrm{Gr}(n-1, \mathbb{P}\mathcal{V})$$

be the rational map defined by $\tilde{\pi}([v_1], \ldots, [v_n]) = [v_1 \wedge \cdots \wedge v_n]$. Then a subset $U \subseteq \mathrm{Gr}(n-1, \mathbb{P}\mathcal{V})$ is open-dense if and only if $\tilde{\pi}^{-1}(U) \subseteq \mathbb{P}(\mathcal{V})^{\times n}$ is open-dense.

Proof. Consider the rational map

$$\pi : \mathcal{V}^{\times n} \dashrightarrow \mathrm{Gr}(n-1, \mathbb{P}\mathcal{V}), \tag{A.1}$$

defined on the open subset $U \subseteq \mathcal{V}^{\times n}$ of linearly independent n-tuples, and given by $\pi(v_1, \ldots, v_n) = [v_1 \wedge \cdots \wedge v_n]$. We require the following claim, which is standard. For completeness, we prove this claim at the end of the proof.

Claim 78. The map π defines a quotient (in the sense of Section II.6.3 in [20]) of U by the algebraic group $\mathrm{GL}(\mathbb{C}^n)$ under the action $A \cdot (v_1, \ldots, v_n) = (v_1, \ldots, v_n)A^{-1}$ for all $A \in \mathrm{GL}(\mathbb{C}^n)$ and $(v_1, \ldots, v_n) \in U$, where the product on the right is matrix multiplication.

It follows that the Zariski topology on $\mathrm{Gr}(n-1,\mathbb{P}\mathcal{V})$ is precisely the quotient topology of U by the equivalence relation induced by the group action. In other words, a subset $V \subseteq \mathrm{Gr}(n-1,\mathbb{P}\mathcal{V})$ is open if and only if $\pi^{-1}(V)$ is open. Passing to $\mathbb{P}(\mathcal{V})^{\times n}$, it follows that a subset $V \subseteq \mathrm{Gr}(n-1,\mathbb{P}\mathcal{V})$ is open if and only if $\tilde{\pi}^{-1}(V)$ is open. This completes the proof, modulo proving the claim.

Proof of Claim 78. To prove that π is a quotient map, note that both U and $\mathrm{Gr}(n-1,\mathbb{P}\mathcal{V})$ are irreducible and smooth, so it suffices to verify that π is a surjective, open map by [20, Lemma II.6.2]. Surjectivity is obvious. To prove that π is open, it suffices to prove that it is flat by [66, Exercise III.9.1], which in turn follows from the fact that U and $\mathrm{Gr}(n-1,\mathbb{P}\mathcal{V})$ are smooth, and the fibers of π are equidimensional [66, Exercise III.10.9]. △

□

Fact 6. *The bijection $\mathrm{Gr}(n-1,\mathbb{P}^d) \cong \mathrm{Gr}(d-n,\mathbb{P}^d)$, which sends a subspace to its orthogonal complement with respect to some non-degenerate bilinear form $\langle\cdot,\cdot\rangle$, defines an isomorphism of projective varieties.*

Proof. Recall the map π defined in (A.1). After change of basis, we may assume $\langle u,v\rangle = u^\intercal v$ for all $u,v \in \mathbb{C}^{d+1}$. For any element $[v]=[v_1\wedge\cdots\wedge v_n] \in \mathrm{Gr}(n-1,\mathbb{P}^d)$, there exists an element of $\pi^{-1}([v])$ of the form $\left(\frac{\mathbb{1}_n}{A}\right)$ for some $A \in \mathrm{L}(\mathbb{C}^n,\mathbb{C}^{d+1-n})$. It is straightforward to verify that $\left(\frac{A^\intercal}{-\mathbb{1}_{d+1-n}}\right) \in \pi^{-1}([v]^\perp)$, where $[v]^\perp$ denotes the orthogonal complement of the projective $(n-1)$-plane $[v]$. One can also verify that, up to sign, the $n\times n$ minor of $\left(\frac{\mathbb{1}_n}{A}\right)$ corresponding to a column index set $S \subseteq [d+1]$ of size $|S|=n$ is precisely the $(d+1-n)\times(d+1-n)$ minor of the matrix $\left(\frac{A^\intercal}{-\mathbb{1}_{d+1-n}}\right)$ corresponding to the column index set $[d+1]\setminus S$. The result follows from the fact that these minors are exactly the coordinates of $[v]\in\mathbb{P}(\bigwedge^n\mathbb{C}^d)$ and $[v]^\perp\in\mathbb{P}(\bigwedge^{d+1-n}\mathbb{C}^d)$ in the Plücker embedding, respectively [65]. □

Fact 8. *Let $\mathcal{V}=\bigotimes_{j=1}^m\mathbb{C}^{d_j}$ and $\mathcal{W}=\bigotimes_{j=1}^m\mathbb{C}^{c_j}$ be vector spaces, and let $[w]\in\mathbb{P}\mathcal{W}$ and $[v]\in\mathbb{P}\mathcal{W}$ be states. Then the sets $\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])$ and $\mathcal{O}_{[v]}$ are both irreducible and constructible in the Zariski topology.*

Proof. Note that $\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])$ is the image of the irreducible quasiprojective variety

$$Z=\{[A_1\otimes\cdots\otimes A_m] : A_i\in\mathrm{L}(\mathbb{C}^{c_i},\mathbb{C}^{d_i}) \quad \text{for all} \quad i\in[m] \quad \text{and} \quad (A_1\otimes\cdots\otimes A_m)w\neq 0\}$$

under the morphism $w : Z \to \mathbb{P}\mathcal{V}$ that sends $(A_1 \otimes \cdots \otimes A_m)$ to $[(A_1 \otimes \cdots \otimes A_m)w]$. Since Z is constructible, $\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])$ is constructible by Chevalley's theorem [65, Theorem 3.16]. Since Z is irreducible, $\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])$ is irreducible in the subspace topology. Nearly identical arguments show that $\mathcal{O}_{[v]}$ is also constructible and irreducible. □

Fact 9. *Let $\mathcal{V} = \bigotimes_{j=1}^m \mathbb{C}^{d_j}$ and $\mathcal{W} = \bigotimes_{j=1}^m \mathbb{C}^{c_j}$ be vector spaces with $c_j \leq d_j$ for all $j \in [m]$. Then $\dim(\overline{\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])})$ is maximized for a generic state $[w] \in \mathbb{P}\mathcal{W}$.*

Proof. Let $\iota : \mathbb{P}\mathcal{W} \to \mathbb{P}\mathcal{V}$ be the canonical inclusion map, which acts on product tensors $[x_1 \otimes \cdots \otimes x_m]$ by appending $d_j - c_j$ zeroes to x_j for each $j \in [m]$. Then for any state $[w] \in \mathbb{P}\mathcal{W}$, it holds that

$$\begin{aligned}\overline{\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])} &= \overline{\mathrm{Im}_{\mathbb{P}\mathcal{V}}(\iota([w]))}\\ &= \overline{\mathcal{O}_{\iota([w])}} \subseteq \mathbb{P}\mathcal{V}.\end{aligned}$$

It is well known that for any non-negative integer k, the set

$$S_k := \{[v] \in \mathbb{P}\mathcal{V} : \dim(\overline{\mathcal{O}_{[v]}}) \leq k\} \subseteq \mathbb{P}\mathcal{V}$$

is Zariski closed [29, Lemma 1.14]. It follows that $S_k \cap \iota(\mathbb{P}\mathcal{W}) \subseteq \mathbb{P}\mathcal{V}$ is Zariski closed, so

$$\iota^{-1}(S_k \cap \iota(\mathbb{P}\mathcal{W})) = \{[w] \in \mathbb{P}\mathcal{W} : \dim(\overline{\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])}) \leq k\} \subseteq \mathbb{P}\mathcal{W}$$

is Zariski closed. It follows that the set of $[w] \in \mathbb{P}\mathcal{W}$ that maximize $\dim(\overline{\mathrm{Im}_{\mathbb{P}\mathcal{V}}([w])})$ is open-dense in $\mathbb{P}\mathcal{W}$. This completes the proof. □

Appendix B

Appendix to Chapter 4

B.1 The maximum dimensions of entangled subspaces

In this appendix, we write down the maximum dimensions of entangled subspaces, by invoking Theorem 7 and the known dimensions of secant varieties reviewed in Section 2.1. Corollary 79 extends results in [100] and [45].

Corollary 79. *Let $\mathcal{V} = \bigotimes_{j=1}^{m} \mathbb{C}^{d_j}$, and let*

$$Y = \mathrm{Seg}(\mathbb{P}^{d_1-1} \times \cdots \times \mathbb{P}^{d_m-1}) \subseteq \mathbb{P}(\mathcal{V})$$

denote the Segre variety of product states. The maximum dimension of an r-entangled subspace of $\mathbb{P}(\mathcal{V})$ is

$$d_1 \cdots d_m - \dim(\Sigma_r(Y)) - 2,$$

and a generic projective linear subspace of this dimension is r-entangled.

As a result, there always exists an r-entangled subspace of dimension

$$d_1 \cdots d_m - r \sum_{j=1}^{m} (d_j - 1) - r - 1,$$

whenever this quantity is non-negative. Furthermore, this is often the maximum dimension of an r-entangled subspace, with a conjecturally complete set of exceptions [2, 14]. If $r = 1$ then this is the maximum dimension, which gives Wallach's result [126].

In the bipartite case $\mathcal{V} = \mathbb{C}^{d_1} \otimes \mathbb{C}^{d_2}$, combining Corollary 79 with Equation (2.7) gives that the maximum dimension of an r-entangled subspace is

$$(d_1 - r)(d_2 - r) - 1,$$

whenever $r \leq \min\{d_1, d_2\}$ (this is [45, Theorem 11]). Under the isomorphism $\mathbb{C}^{d_1} \otimes \mathbb{C}^{d_2} \cong \mathrm{L}((\mathbb{C}^{d_1})^*, \mathbb{C}^{d_2})$, this is the maximum dimension of a projective linear subspace of $d_2 \times d_1$ matrices of rank greater than r.

If $r = 2$ and $m \geq 3$, then $\Sigma_r(Y)$ has the expected dimension, so the largest dimension of a 2-entangled subspace in this case is

$$d_1 \cdots d_m - 2\sum_{j=1}^{m}(d_j - 1) - 3. \tag{B.1}$$

We explicitly construct maximal 2-entangled subspaces in Section 4.1.

Corollary 80. *Let m and d be positive integers, and let*

$$\nu_m(\mathbb{P}^{d-1}) \subseteq \mathbb{P}(S^m\mathbb{C}^d)$$

denote the Veronese variety of unentangled states in the symmetric space. The maximum dimension of a symmetric r-entangled subspace of $\mathbb{P}(S^m\mathbb{C}^d)$ is

$$\binom{d-1+m}{m} - \dim(\Sigma_r(\nu_m(\mathbb{P}^{d-1}))) - 2$$

and a generic projective linear subspace of this dimension is symmetric r-entangled.

As a result, there always exists a symmetric r-entangled subspace of dimension

$$\binom{d-1+m}{m} - rd - 1, \tag{B.2}$$

whenever this quantity is non-negative. This is the maximum dimension of a symmetric r-entangled subspace in many cases, with a known set of exceptions [3, 14]. If $r = 1$, then (B.2) is the maximum dimension. If $m = 2$, then the maximum dimension is given by

$$\binom{d-r+1}{2} - 1. \tag{B.3}$$

We explicitly construct maximal symmetric r-entangled subspaces in these two cases in Section 4.1.

Corollary 81. *Let m and d be positive integers with $m \geq d$, and let*

$$\mathrm{Gr}(m-1, \mathbb{P}^{d-1}) \subseteq \mathbb{P}(\textstyle\bigwedge^m \mathbb{C}^d)$$

denote the Grassmannian variety of unentangled states in the antisymmetric space. The maximum dimension of an antisymmetric r-entangled subspace of $\mathbb{P}(\bigwedge^m(\mathbb{C}^d))$ is

$$\binom{d-1+m}{m} - \dim(\Sigma_r(\mathrm{Gr}(m-1, \mathbb{P}^{d-1}))) - 2,$$

and a generic projective linear subspace of this dimension is antisymmetric r-entangled.

As a result, there always exists an antisymmetric r-entangled subspace of dimension

$$\binom{d-1+m}{m} - 1 - rm(d-m) - r, \tag{B.4}$$

whenever this quantity is non-negative. This is often the maximum dimension of an antisymmetric r-entangled subspace, with a conjecturally complete set of exceptions [13, 14]. If $r = 1$, then (B.4) is the maximum dimension. If $m = 2$, then the maximum dimension is given by

$$\binom{d-2r}{2} - 1. \tag{B.5}$$

We explicitly construct maximal antisymmetric r-entangled subspaces in these two cases in Section 4.1.

B.2 A 2-Entangled Qutrit-Qutrit-Qubit Subspace

In this appendix, we prove that the span of the set B from Equation (4.2) is 2-entangled. In order to show this, we prove that each member of that span has a flattening with rank 3. Indeed, since $\mathbb{C}^3 \otimes \mathbb{C}^3 \otimes \mathbb{C}^2$ is naturally isomorphic to the space of 3×6 matrices, we can think of this subspace as consisting of block matrices of the form

$$M = \left[\begin{array}{ccc|ccc} \lambda & \alpha_1 & \beta_1 & \gamma & 0 & \alpha_4 \\ \alpha_2 & \lambda & 0 & 0 & \gamma & \beta_4 \\ \beta_2 & 0 & \lambda & \alpha_3 & \beta_3 & \gamma \end{array}\right],$$

where $\alpha_j = \delta_j + \epsilon_j$ and $\beta_j = \theta_j + \kappa_j$ for all $1 \leq j \leq 4$ (and $\{\delta, \epsilon, \theta, \kappa\}$ is as in Section 4.1.5). The partial transpose of M is

$$M^\Gamma = \left[\begin{array}{ccc|ccc} \lambda & \alpha_2 & \beta_2 & \gamma & 0 & \alpha_3 \\ \alpha_1 & \lambda & 0 & 0 & \gamma & \beta_3 \\ \beta_1 & 0 & \lambda & \alpha_4 & \beta_4 & \gamma \end{array}\right],$$

which is another flattening of this same state, so our goal is to show that $\operatorname{rank}(M) \geq 3$ or $\operatorname{rank}(M^\Gamma) \geq 3$.

Importantly, because of how we chose $\{\delta, \epsilon, \theta, \kappa\}$, we know if that if $\alpha_i \neq 0$ for any $i \in [4]$, then $\alpha_i \neq 0$ for at least three $i \in [4]$, and similarly for the β_j's. Indeed, we saw the desirability of this property in Section 4.1, where we repeatedly used Lemma 15.

To show that $\operatorname{rank}(M) \geq 3$ or $\operatorname{rank}(M^\Gamma) \geq 3$, we now split into several cases depending on which of λ, γ, α_i and β_j equal 0.

Case 1(a): $\lambda = 0$, $\alpha_1, \alpha_2, \alpha_3 \neq 0$.

The submatrix of M corresponding to its 1st, 2nd, and 4th columns, up to permutation similarity, has the form

$$\begin{bmatrix} \alpha_2 & 0 & 0 \\ \beta_2 & \alpha_3 & 0 \\ 0 & \gamma & \alpha_1 \end{bmatrix}$$

which clearly has rank 3 since it is triangular with non-zero diagonal entries.

The above case contains the flavor of most of the cases that we will consider, so from now on we just list which columns of M or M^Γ give rise to a submatrix that is (up to permutation similarity) triangular with non-zero diagonal entries, and thus has rank 3. For example, for Case 1(a) we would just now just say "$M(1,2,4)$".

Case 1: $\lambda = 0$.

(a): $\alpha_1, \alpha_2, \alpha_3 \neq 0$. $M(1,2,4)$.

(b): $\alpha_1, \alpha_2, \alpha_4 \neq 0$. $M^\Gamma(1,2,4)$.

(c): $\alpha_1, \alpha_3, \alpha_4 \neq 0$.

(i): $\gamma = 0$. $M^\Gamma(1,4,6)$.

(ii): $\gamma \neq 0$. $M(2,4,5)$.

(d): $\alpha_2, \alpha_3, \alpha_4 \neq 0$.

(i): $\gamma = 0$. $M(1,4,6)$.

(ii): $\gamma \neq 0$. $M^\Gamma(2,4,5)$.

(e): $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 0$.

(i): $\gamma = 0$. This case is identical to Case 1(a–d) via symmetry (just rotate M by 180 degrees, which does not change its rank).

(ii): $\gamma \neq 0$, $\beta_1 \neq 0$. $M^\Gamma(1,4,5)$.

(iii): $\gamma \neq 0$, $\beta_2, \beta_3, \beta_4 \neq 0$. $M(1,4,6)$.

Case 2: $\lambda \neq 0$, $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 0$.

(a): $\beta_3 = \beta_4 = 0$ (and thus $\beta_1 = \beta_2 = 0$ too). $M(1,2,3)$.

(b): $\beta_3 \neq 0$.

(i): $\gamma = 0$. $M(1,2,5)$.

(ii): $\gamma \neq 0$. $M(2,4,6)$.

(c): $\beta_4 \neq 0$.

(i): $\gamma = 0$. $M^\Gamma(1,2,5)$.

(ii): $\gamma \neq 0$. $M(2,4,6)$.

Case 3: $\lambda \neq 0$ and $\alpha_i \neq 0$ for at least three $i \in [4]$.

(a): $\beta_3 = \beta_4 = 0$.

(i): $\alpha_3 \neq 0$. $M^\Gamma(2,3,6)$.

(ii): $\alpha_4 \neq 0$. $M(2,3,6)$.

(b): $\beta_3 \neq 0$, $\beta_4 = 0$.

(i): $\alpha_4 = 0$, $\gamma = 0$. $M(1,3,4)$.

(ii): $\alpha_4 = 0$, $\gamma \neq 0$. $M^\Gamma(4,5,6)$.

(iii): $\alpha_4 \neq 0$, $\gamma = 0$. $M(2,3,6)$.

(iv): $\alpha_4 \neq 0$, $\gamma \neq 0$, $\alpha_2 \neq 0$. $M^\Gamma(2,3,5)$.

(v): $\alpha_4 \neq 0$, $\gamma \neq 0$, $\alpha_2 = 0$. This case is much more difficult, so we leave it until after the remaining cases are dealt with.

(c): $\beta_3 = 0$, $\beta_4 \neq 0$. This case is identical to Case 3(b) by taking the partial transpose (i.e., replace M with M^Γ and vice-versa).

(d): $\beta_3, \beta_4 \neq 0$.

(i): $\alpha_4 = 0$, $\gamma = 0$. $M(1, 4, 6)$.

(ii): $\alpha_4 = 0$, $\gamma \neq 0$. $M^\Gamma(3, 4, 5)$.

(iii): $\alpha_4 \neq 0$, $\beta_2 = 0$, $\gamma \neq 0$. $M^\Gamma(3, 4, 5)$.

(iv): $\alpha_4 \neq 0$, $\beta_2 = 0$, $\gamma = 0$, $\alpha_3 = 0$. $M^\Gamma(1, 3, 6)$.

(v): $\alpha_4 \neq 0$, $\beta_2 = 0$, $\gamma = 0$, $\alpha_3 \neq 0$. $M(2, 3, 4)$.

(vi): $\alpha_4 \neq 0$, $\beta_2 \neq 0$, $\gamma = 0$. $M^\Gamma(2, 3, 4)$.

(vii): $\alpha_4 \neq 0$, $\beta_2 \neq 0$, $\gamma \neq 0$. This is another difficult case that we deal with separately.

The only two remaining cases from above are 3(b)(v) and 3(d)(vii). These cases require a more intricate argument to demonstrate that at least one of M or M^Γ has rank 3, which we now provide.

Case 3(b)(v)

For this case, $\gamma, \lambda \neq 0$, $\alpha_2 = \beta_4 = 0$, $\alpha_1, \alpha_3, \alpha_4 \neq 0$, and $\beta_1, \beta_2, \beta_3 \neq 0$. Let $\alpha = (\alpha_1, \alpha_2, \alpha_3, \alpha_4)$ and $\beta = (\beta_1, \beta_2, \beta_3, \beta_4)$. Because of how we chose $\{\delta, \epsilon, \theta, \kappa\}$, we have $\alpha = (\alpha_1, \alpha_1 + \alpha_4, 2\alpha_1 + \alpha_4, \alpha_4)$ and $\beta = (\beta_1, \beta_1 + 2\beta_4, \beta_1 + \beta_4, \beta_4)$. Furthermore, using the facts that $0 = \alpha_2 = \alpha_1 + \alpha_4$ and $0 = \beta_4$ shows that $\alpha = (\alpha_1, 0, \alpha_1, -\alpha_1)$ and $\beta = (\beta_1, \beta_1, \beta_1, 0)$. The matrix M^Γ thus has the form

$$M^\Gamma = \left[\begin{array}{ccc|ccc} \lambda & 0 & \beta_1 & \gamma & 0 & \alpha_1 \\ \alpha_1 & \lambda & 0 & 0 & \gamma & \beta_1 \\ \beta_1 & 0 & \lambda & -\alpha_1 & 0 & \gamma \end{array}\right].$$

We can then see that $\mathrm{rank}(M^\Gamma) = 3$ as follows. If it had rank ≤ 2 then it would be the case that $\mathrm{Det}(M^\Gamma(1, 2, 3)) = \lambda(\lambda^2 - \beta_1^2) = 0$, which implies $\beta_1 = \pm\lambda$. We would similarly have $\mathrm{Det}(M^\Gamma(4, 5, 6)) = \gamma(\gamma^2 + \alpha_1^2) = 0$, which implies $\alpha_1 = \pm i\gamma$. Finally, we would also have $\mathrm{Det}(M^\Gamma(3, 4, 6)) = -\beta_1(\gamma\lambda + \alpha_1\beta_1) = 0$, which (since $\alpha_1 = \pm i\gamma$ and $\beta_1 = \pm\lambda$) implies $\gamma\lambda \pm i\gamma\lambda = 0$, which is impossible (recall that all of these variables are non-zero). It follows that at least one of these determinants is non-zero, so $\mathrm{rank}(M^\Gamma) = 3$, which completes this case.

Case 3(d)(vii)

For this case, $\gamma, \lambda \neq 0$, $\alpha_4 \neq 0$, and $\beta_2, \beta_3, \beta_4 \neq 0$. We now split into four subcases:

- $\alpha_1 = \beta_3 = 0$. This case is identical to Case 3(b)(v) from Section B.2 by taking the partial transpose, swapping the roles of γ and λ, and swapping the roles of α and β.

- $\alpha_1 = 0$, $\beta_3 \neq 0$. $M(2, 4, 5)$.

- $\alpha_1 \neq 0$, $\beta_3 = 0$. $M(2, 3, 5)$.

- $\alpha_1, \beta_3 \neq 0$. This is the difficult subcase. Similarly to Case 3(b)(v), we have $\alpha = (\alpha_1, \alpha_1 + \alpha_4, 2\alpha_1 + \alpha_4, \alpha_4)$ and $\beta = (\beta_1, \beta_1 + 2\beta_4, \beta_1 + \beta_4, \beta_4)$, so M and M^Γ have the form

  $$M = \left[\begin{array}{ccc|ccc} \lambda & \alpha_1 & \beta_1 & \gamma & 0 & \alpha_4 \\ \alpha_1 + \alpha_4 & \lambda & 0 & 0 & \gamma & \beta_4 \\ \beta_1 + 2\beta_4 & 0 & \lambda & 2\alpha_1 + \alpha_4 & \beta_1 + \beta_4 & \gamma \end{array}\right] \quad \text{and}$$
  $$M^\Gamma = \left[\begin{array}{ccc|ccc} \lambda & \alpha_1 + \alpha_4 & \beta_1 + 2\beta_4 & \gamma & 0 & 2\alpha_1 + \alpha_4 \\ \alpha_1 & \lambda & 0 & 0 & \gamma & \beta_1 + \beta_4 \\ \beta_1 & 0 & \lambda & \alpha_4 & \beta_4 & \gamma \end{array}\right].$$

  Now suppose that $\operatorname{rank}(M) = \operatorname{rank}(M^\Gamma) = 0$. Then we would have

  $$\operatorname{Det}(M(2, 3, 6)) = \lambda(\lambda\alpha_4 - \alpha_1\beta_4 - \gamma\beta_1) = 0,$$

  which implies $\lambda\alpha_4 - \gamma\beta_1 = \alpha_1\beta_4$. We would also have

  $$\operatorname{Det}(M^\Gamma(1, 4, 5)) = \gamma(\gamma\beta_1 - \lambda\alpha_4 - \alpha_1\beta_4) = 0,$$

  which implies $\lambda\alpha_4 - \gamma\beta_1 = -\alpha_1\beta_4$. Combining these two expressions for $\lambda\alpha_4 - \gamma\beta_1$ shows that $\alpha_1\beta_4 = 0$, which contradicts the fact that we are assuming that $\alpha_1, \beta_4 \neq 0$ in this case. It follows that at least one of these determinants is non-zero, which completes this final subcase and the proof.

Numerics suggest that almost all choices of $\{\delta, \epsilon, \theta, \kappa\}$ lead to this subspace being 2-entangled. Indeed, the primary property of $\{\delta, \epsilon, \theta, \kappa\}$ that we made use of was that a non-zero linear combination of δ and ϵ must never have more than one 0 entry, and similarly for θ and κ (this property is generic). The only other place where the particular entries of these vectors was used was in Case 3(b)(v) (and the symmetric first subcase of Case 3(d)(vii)), where non-invertibility of a 3×3 submatrix of M or M^Γ was more delicate.

To illustrate why Case 3(b)(v) is more delicate, notice that if we had instead chosen $\delta = (1, 1, 1, 2)$, $\epsilon = (0, -1, 1, 1)$, $\theta = (2, 1, 1, 0)$, and $\kappa = (1, 1, 0, 1)$, then it would still be the case that any non-zero linear combination of δ and ϵ would never have more than one 0 entry, and similarly for θ and κ (so all of the other cases still work fine). However, the proof would fall apart in Case 3(b)(v), since we would get the pair of matrices

$$M = \left[\begin{array}{ccc|ccc} 1 & -1 & 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 & 1 & 1 \end{array}\right] \quad \text{and} \quad M^\Gamma = \left[\begin{array}{ccc|ccc} 1 & 0 & 1 & 1 & 0 & 1 \\ -1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 1 & 0 & 1 \end{array}\right],$$

both of which have rank 2.

Appendix C

Appendix to Chapter 7

C.1 Proof of Theorem 66

In this appendix we prove Theorem 66. The proof is very similar to that of Theorem 61.

Proof of Theorem 66. For each $a \in [r]$, let $x_{n+a} = -y_a$, and let $T_1 \sqcup \cdots \sqcup T_t = [n+r]$ be the index sets of the decomposition of $\{x_a : a \in [n+r]\}$ into connected components. Note that for each $p \in [t]$, if

$$\big|T_p \cap [n+r] \setminus [n]\big| \leq \big|T_p \cap [n]\big|,$$

then $\big|T_p \cap [n]\big| \leq s$, otherwise $\{x_a : a \in T_p\}$ would split. Assume without loss of generality that

$$\begin{aligned}\big|T_1 \cap [n]\big| - \big|T_1 \cap [n+r] \setminus [n]\big| &\geq \big|T_2 \cap [n]\big| - \big|T_2 \cap [n+r] \setminus [n]\big| \\ &\vdots \\ &\geq \big|T_t \cap [n]\big| - \big|T_t \cap [n+r] \setminus [n]\big|,\end{aligned}$$

If

$$\big|T_1 \cap [n]\big| \geq \big|T_1 \cap [n+r] \setminus [n]\big|,$$

then let $\tilde{l} \in [t]$ be the largest integer for which

$$\big|T_{\tilde{l}} \cap [n]\big| \geq \big|T_{\tilde{l}} \cap [n+r] \setminus [n]\big|. \tag{C.1}$$

Otherwise, let $\tilde{l} = 0$. Then for all $p \in [t] \setminus [\tilde{l}]$ it holds that

$$\big|T_p \cap [n]\big| < \big|T_p \cap [n+r] \setminus [n]\big|.$$

To complete the proof, we will show that $\tilde{l} \geq l$, for then we can take $Q_p = T_p \cap [n]$ and $R_p = T_p \cap [n+r] \setminus [n]$ for all $p \in [l]$ to conclude.

Suppose toward contradiction that $\tilde{l} < l$. We will require the following two claims:

Claim 82. It holds that $\tilde{l} < t$, $\lceil \frac{n - s\tilde{l}}{t - \tilde{l}} \rceil \geq s + 1$, and there exists $p \in [t] \setminus [\tilde{l}]$ for which

$$\big|T_p \cap [n]\big| \geq \left\lceil \frac{n - s\tilde{l}}{t - \tilde{l}} \right\rceil. \tag{C.2}$$

Claim 83. For all $p \in [t] \setminus [\tilde{l}]$, it holds that

$$\big|T_p \cap [n+r] \setminus [n]\big| \leq \big|T_p \cap [n]\big| + (r - n) + (s+1)\tilde{l} - t + 1. \tag{C.3}$$

Before proving these claims, we first use them to complete the proof of the theorem. Let $p \in [t] \setminus [\tilde{l}]$ be as in Claim 82. Then,

$$\begin{aligned}
|T_p| &= \big|T_p \cap [n]\big| + \big|T_p \cap [n+r] \setminus [n]\big| \\
&\leq 2\big|T_p \cap [n]\big| + r - n + (s+1)\tilde{l} - t + 1 \\
&\leq 2\big|T_p \cap [n]\big| + r - n + s\tilde{l} - \left\lceil \frac{n - s\tilde{l}}{\big|T_p \cap [n]\big|} \right\rceil + 1 \\
&\leq 2\big|T_p \cap [n]\big| + (r - n + q - s) - \left\lceil \frac{n - q + s}{\big|T_p \cap [n]\big|} \right\rceil + 1 \\
&\leq \sum_{j=1}^{m} (d_j^{T_p \cap [n]} - 1) + 1,
\end{aligned}$$

where the first line is obvious, the second follows from Claim 83, the third follows from Claim 82, the fourth follows from $\tilde{l} < l$, and the fifth follows from the assumptions of the theorem and the fact that $|T_p \cap [n]| \geq s + 1$. So $\{x_a : a \in T_p\}$ splits, a contradiction. This completes the proof, modulo proving the claims.

Proof of Claim 62. To prove the claim, we first observe that $n > st$. Indeed, if $n \leq st$, then

$$\begin{aligned}
r &\geq \sum_{p=\tilde{l}+1}^{t} \left|T_p \cap [n+r] \setminus [n]\right| \\
&\geq \sum_{p=\tilde{l}+1}^{t} \left(\left|T_p \cap [n]\right| + 1\right) \\
&= n - \left|(T_1 \sqcup \cdots \sqcup T_{\tilde{l}}) \cap [n]\right| + t - \tilde{l} \\
&\geq n + t - (s+1)\tilde{l} \\
&\geq n + \left\lceil \frac{n}{s} - (s+1)(q/s - 1) \right\rceil \\
&= \left\lceil \left(\frac{s+1}{s}\right)(n - q + s) \right\rceil,
\end{aligned}$$

where the first line is obvious, the second follows from (7.12), the third is obvious, the fourth follows from $|T_p \cap [n]| \leq s$ for all $p \in [\tilde{l}]$, the fifth follows from $n \leq st$ and $\tilde{l} < l$, and the sixth is algebra. This contradicts the assumptions of the theorem, so it must hold that $n > st$.

Note that $\tilde{l} < t$, for otherwise we would have $n \leq st$ by the fact that $\left|T_p \cap [n]\right| \leq s$ for all $p \in [\tilde{l}]$. To verify that $\left\lceil \frac{n - s\tilde{l}}{t - \tilde{l}} \right\rceil \geq s + 1$, it suffices to prove $\frac{n - s\tilde{l}}{t - \tilde{l}} > s$, which follows from $n > st$. To verify (C.2), since $|T_p \cap [n]| \leq s$ for all $p \in [\tilde{l}]$, by the pigeonhole principle there exists $p \in [t] \setminus [\tilde{l}]$ for which

$$\left|T_p \cap [n]\right| \geq \left\lceil \frac{n - s\tilde{l}}{t - \tilde{l}} \right\rceil.$$

This proves the claim. △

Proof of Claim 83. Suppose toward contradiction that the inequality (C.3) does not hold for some $\tilde{p} \in [t] \setminus [\tilde{l}]$. Then

$$\begin{aligned} r &\geq \sum_{p=\tilde{l}+1}^{t} \big|T_p \cap [n+r] \setminus [n]\big| \\ &\geq \sum_{p \neq \tilde{p}} \big(\big|T_p \cap [n]\big| + 1\big) + \big|T_{\tilde{p}} \cap [n]\big| + (r-n) + (s+1)\tilde{l} - t + 2 \\ &= \sum_{p=\tilde{l}+1}^{t} \big|T_p \cap [n]\big| + (r-n) + s\tilde{l} + 1 \\ &\geq r + 1, \end{aligned}$$

where the first three lines are obvious, and the fourth follows from (C.1), a contradiction. $\triangle$

The proofs of Claims 82 and 83 complete the proof of the theorem. $\square$