

Overcomplete Tensor Decomposition via Koszul–Young Flattenings

Pravesh K. Kothari^{*1}, Ankur Moitra^{†2,3}, and Alexander S. Wein^{‡4}

¹Computer Science Department, Princeton

²Department of Mathematics, MIT

³Computer Science and Artificial Intelligence Lab, MIT

⁴Department of Mathematics, UC Davis

Abstract

Motivated by connections between algebraic complexity lower bounds and tensor decompositions, we investigate Koszul–Young flattenings, which are the main ingredient in recent lower bounds for matrix multiplication. Based on this tool we give a new algorithm for decomposing an $n_1 \times n_2 \times n_3$ tensor as the sum of a minimal number of rank-1 terms, and certifying uniqueness of this decomposition. For $n_1 \leq n_2 \leq n_3$ with $n_1 \rightarrow \infty$ and $n_3/n_2 = O(1)$, our algorithm is guaranteed to succeed when the tensor rank is bounded by $r \leq (1 - \epsilon)(n_2 + n_3)$ for an arbitrary $\epsilon > 0$, provided the tensor components are generically chosen. For any fixed ϵ , the runtime is polynomial in n_3 . When $n_2 = n_3 = n$, our condition on the rank gives a factor-of-2 improvement over the classical simultaneous diagonalization algorithm, which requires $r \leq n$, and also improves on the recent algorithm of Koiran (2024) which requires $r \leq 4n/3$. It also improves on the PhD thesis of Persu (2018) which solves rank detection for $r \leq 3n/2$.

We complement our upper bounds by showing limitations, in particular that no flattening of the style we consider can surpass rank $n_2 + n_3$. Furthermore, for $n \times n \times n$ tensors, we show that an even more general class of degree- d polynomial flattenings cannot surpass rank Cn for a constant $C = C(d)$. This suggests that for tensor decompositions, the case of generic components may be fundamentally harder than that of random components, where efficient decomposition is possible even in highly overcomplete settings.

^{*}Email: kothari@cs.princeton.edu. Partially supported by NSF CAREER Award #2047933, Alfred P. Sloan Fellowship and a Google Research Scholar Award.

[†]Email: moitra@mit.edu. Partially supported by a Microsoft Trustworthy AI Grant, an ONR grant and a David and Lucile Packard Fellowship.

[‡]Email: aswein@ucdavis.edu. Partially supported by an Alfred P. Sloan Research Fellowship and NSF CAREER Award CCF-2338091.

Contents

1	Introduction	3
2	Main Results	7
2.1	Preliminaries	7
2.1.1	Tensor Decomposition and Rank	7
2.1.2	Genericity and Symbolic Rank	7
2.2	Upper Bounds	8
2.2.1	Trivial Flattenings	8
2.2.2	Koszul–Young Flattenings	9
2.2.3	Rank Detection	10
2.2.4	Decomposition Algorithm	11
2.2.5	Uniqueness Theorem and Algorithmic Guarantees	13
2.3	Lower Bounds	15
2.3.1	Flattenings of Koszul–Young Type	15
2.3.2	Linear Flattenings	16
2.3.3	Low-Degree Flattenings	16
2.4	Open Problems	18
3	Proofs for Upper Bounds	19
3.1	Proof of Proposition 2.3	19
3.2	Proof of Theorem 2.4	19
3.3	Proof of Theorem 2.7	21
3.4	Proof of Lemma 3.2	24
3.5	Proof of Lemma 3.3	24
3.6	Proof of Theorem 2.8	26
3.7	Proof of Lemma 3.6	28
4	Proofs for Lower Bounds	30
4.1	Matrix Decomposition Lemma	30
4.2	Proof of Theorem 2.9	31
4.3	Proof of Theorem 2.10	31
4.4	Proof of Theorem 2.12	32
A	Trivial Flattenings	34
B	Commuting Extensions	35
	Acknowledgments	38
	References	38

1 Introduction

A *tensor* is a multi-way array $T \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_k}$ with entries denoted T_{i_1, i_2, \dots, i_k} for $i_j \in [n_j] := \{1, \dots, n_j\}$, where k is called the *order* of the tensor. For vectors $a^{(1)} \in \mathbb{R}^{n_1}$, $a^{(2)} \in \mathbb{R}^{n_2}, \dots, a^{(k)} \in \mathbb{R}^{n_k}$, the associated *rank-1* tensor is denoted $T = a^{(1)} \otimes a^{(2)} \otimes \dots \otimes a^{(k)}$ and has entries $T_{i_1, i_2, \dots, i_k} = a_{i_1}^{(1)} a_{i_2}^{(2)} \dots a_{i_k}^{(k)}$. Also, $a^{\otimes k}$ is shorthand for $a \otimes a \otimes \dots \otimes a$ (k times). While various notions of tensor rank exist, we consider the *CP (canonical polyadic)* decomposition, that is, we define the rank of a tensor T to be the minimum r such that T can be expressed as the sum of r rank-1 tensors. In the matrix case $k = 2$, this coincides with the usual notion of matrix rank.

Tensor decomposition is the algorithmic task of breaking down a given rank- r tensor into its r constituent rank-1 summands. For matrices, this is ill-posed due to the “rotation problem” (see [Moi18]). In contrast, minimum-rank tensor decompositions for $k \geq 3$ tend to be unique, provided the rank is small enough. This uniqueness is key to the usefulness of tensor methods in various applications throughout statistics, data science, and machine learning. For instance, minimum rank decompositions of the third/fourth *moment* tensor can be used to learn the parameters of expressive statistical models, particularly ones with latent variables. We refer the reader to [KB09, Moi18, RSG17, Vij20] for more background and applications of tensor decompositions.

In this work, we study tensor decompositions for an $n_1 \times \dots \times n_k$ tensor T of rank r . Any such tensor can be written as

$$T = \sum_{\ell=1}^r a^{(1,\ell)} \otimes \dots \otimes a^{(k,\ell)}. \quad (1)$$

The vectors $a^{(i,\ell)} \in \mathbb{R}^{n_i}$ are called the *components* of the tensor T . In this work, we focus on the well-studied setting of “large” dimensional tensors of a fixed order; that is, we imagine a sequence of problem sizes where k remains fixed but some or all of the dimensions n_1, \dots, n_k tend to infinity.

Three natural and interrelated questions of interest arise when studying decompositions for T .

- **Uniqueness Theorem:** A uniqueness theorem specifies an explicit list of conditions (ideally, efficiently checkable) on the collection of components $\{a^{(i,\ell)} : i \in [k], \ell \in [r]\}$ that guarantees the uniqueness of the decomposition (1) up to natural ambiguities in specifying the components (see Definition 2.2). E.g., note that we cannot hope to resolve the norm of each component $a^{(i,\ell)}$, since $a \otimes b \otimes c = (2a) \otimes (b/2) \otimes c$.
- **Rank Detection:** Given a tensor T of the form (1) generated according to some model, the goal is to output r . Our lower bounds (i.e., impossibility results) will apply to this goal, which is, in general, easier than computing a minimum rank decomposition for T .
- **Decomposition:** Given a tensor T of the form (1) generated according to some model, the goal is to recover the individual rank-1 terms $a^{(1,\ell)} \otimes \dots \otimes a^{(k,\ell)}$ up to re-ordering (i.e., the r terms can be listed in any order) and rescalings (again, see Definition 2.2).

Tensors pose a number of computational challenges because basic operations such as computing the rank or finding the best rank-1 approximation are NP-hard for tensors of order 3 and up [Hås90, HL13, Swe18]. As a result, we do not hope for an efficient (i.e., polynomial-time) tensor decomposition algorithm that works on all possible inputs, and it will be important to impose some non-degeneracy assumptions on the input tensor. It is typical to assume the components of T satisfy:

- **Generic components:** Here, the components are assumed to be chosen “generically” in the sense of algebraic geometry (see Section 2.1.2). This amounts to asking the algorithm to succeed for “almost all” choices of components (all but a measure-zero set).

- **Smoothed analysis model:** Here, the components are obtained by starting with an arbitrary collection of components and then perturbing them by a small amount of random noise. Additional noise may also be added to the observed tensor. While results in the generic setting tend to assume exact access to the input tensor and exact computation of linear-algebraic primitives, results in the smoothed setting tend to provide a more detailed analysis of the algorithm’s numerical stability and robustness to noise.
- **Random components:** Here, the components are assumed to be drawn independently from a particular prior, such as i.i.d. Gaussian or uniform on the unit sphere. As compared to generic components, this induces rather strong assumptions on the components (e.g., near-orthogonality).

Additionally, some prior work assumes the rank-1 terms in the decomposition are *symmetric* tensors $(a^{(\ell)})^{\otimes k}$, in place of the non-symmetric setup we have presented above. We will be interested in the *overcomplete* setting, which in the “square” case $n_1 = n_2 = \dots = n_k = n$ means $r > n$. While much is known about overcomplete third-order tensor decomposition in the random model, there are wide gaps in our understanding in the smoothed and generic settings. This will be our main focus in this paper.

Prior work. First we give an overview of prior works on tensor decomposition, focusing mainly on tensors of order $k = 3, 4$ and, for ease of exposition, restricting our attention to square tensors, i.e., those of format $n \times n \times \dots \times n$ with n large. A first question is that of identifiability: when are minimum-rank tensor decompositions unique? A simple argument shows that an order- k tensor of format $n \times \dots \times n$ cannot have rank larger than n^{k-1} (for instance, an $n \times n \times n$ tensor is comprised of n matrix “slices” and each is $n \times n$ with rank $\leq n$). This means if T is generated according to (1) with $r > n^{k-1}$, the planted decomposition will not be the minimum-rank one, so we do not hope to find it. On the other hand, if $r \leq cn^{k-1}$ for a constant $c = c(k) > 0$ and the components are generically chosen, then the planted decomposition will be the unique rank- r decomposition (for $k \geq 3$) [BCO14]; so in this regime we might hope to solve decomposition (and related problems) in principle, but this does not guarantee an efficient algorithm.

Known polynomial-time algorithms require r to be significantly below the identifiability threshold. For order-3 tensors, a classical method called *simultaneous diagonalization* or *Jennrich’s algorithm* can decompose generic tensors of rank $r \leq n$ (the “undercomplete” case) [Har70, LRA93] (see [Moi18]). Furthermore, the analysis comes with a uniqueness theorem: if $T = \sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)}$ where $\{a^{(\ell)}\}$ and $\{b^{(\ell)}\}$ are each linearly independent and $\{c^{(\ell)}\}$ are pairwise linearly independent, then the decomposition is unique and the algorithm will find it. Another well-known uniqueness theorem due to Kruskal [Kru77] has the advantage of tolerating higher rank $r \leq 3n/2$ but the disadvantage that the conditions on the components are not efficiently checkable and no algorithm is provided to find the decomposition; see also [LP23] and references within for various extensions of Kruskal’s theorem.

For order-4 tensors, the *FOOBI* algorithm and variants [Car91, DL06, DLCC07, MSS16, HSS19, JLV23] can decompose generic tensors of rank $r \leq cn^2$ for a constant $c > 0$. See [JLV23] for a recent account, which generalizes the setting and corrects an error in the original analysis of [DL06]; this result also certifies uniqueness of the decomposition, providing an efficiently-checkable uniqueness theorem. Noise-robustness and numerical precision of (variants of) the above algorithms have been studied (for instance, in the smoothed analysis model) for order-3 [BCMV14b, GVX14, KS23] and order-4 [MSS16, HSS19] tensors.

If the components are taken to be random rather than generic, striking improvements over the above are possible for order-3 tensors, with poly-time algorithms known to reach $r \approx n^{3/2}$

(up to logarithmic factors in n) [GM15, MSS16, HSS16, KP20, DdL⁺22]. It appears likely that this threshold — and more generally, $n^{k/2}$ for order- k tensors — is the fundamental limit for efficient algorithms in the random model, as evidenced by lower bounds against low-degree polynomials [Wei23]. Note that for order-4 tensors, the algorithms for the generic model match the lower bounds for the (easier) random model, meaning we (likely) cannot expect further improvements in the generic setting, and the generic setting is no harder than the random setting. In contrast, known algorithms for order-3 tensors display a gap between the generic ($r \approx n$) and random ($r \approx n^{3/2}$) settings.

Generic overcomplete order-3 tensors. The above discussion suggests the possibility of finding improved algorithms for *generic* $n \times n \times n$ tensors. For some time, the question of finding algorithms (or an efficient uniqueness theorem) in the “overcomplete” regime $r > n$ remained a notable open problem (highlighted at COLT 2014 [BCMV14a] and also mentioned in [MSS16, Moi18]), and recently some progress has been made. The work of [CR22] gave an algorithm for decomposing slightly overcomplete tensors, namely $r = n + \ell$ where the runtime is polynomial in n but exponential in ℓ . More recently, Koiran [Koi24a] gave both a decomposition algorithm and an efficient uniqueness theorem for $r \leq 4n/3$; this algorithm only uses 4 “slices” of the tensor, and works more generally for tensors of format $n \times n \times p$ for $p \geq 4$.

The current work aims to push these approaches to the limit, improving the condition on r as much as possible and also investigating whether there are fundamental limitations for efficient algorithms. Even a modest improvement on the value of r has potential consequences for downstream applications in learning. For example, consider the *heterogeneous multi-reference alignment* [PWB⁺19] problem. Here the goal is to learn a collection of K “signals” (n -dimensional vectors), given a mixture of noisy copies of these signals that have been circularly shifted by unknown offsets. When applying the tensor decomposition approach of [PWB⁺19] using the third moment tensor, the rank of the $n \times n \times n$ tensor that needs to be decomposed is Kn . This would be interesting even for $K = 2$. Since no such decomposition algorithms are known (except for $K = 1$), [PWB⁺19] instead uses the order-5 moment tensor, which is more expensive in terms of runtime and sample complexity.

Even more tantalizing is the prospect of potentially finding significantly better algorithms for generic tensors, say, matching the threshold $r \approx n^{3/2}$ that appears for random tensors. On one hand, this appears plausible because order-4 tensors have no gap between the generic and random settings, and it would perhaps be unexpected to have such a qualitative difference between odd-order and even-order tensors. For instance, order- k tensor decomposition may appear superficially similar to refutation of k -XOR equations — which can be viewed as a tensor-valued problem — where we know the odd- and even-order cases both have the threshold $n^{k/2}$ for both random and smoothed settings, even though the odd- k case is much more technically challenging (see [GKM22]). The same thresholds also arise for the problem of completing low-rank tensors, with similar complexities in the odd- k case [BM16, KP20, LM20]. On the other hand, tensor decomposition appears related to the problem of proving lower bounds on the rank of explicit tensors, which is connected to arithmetic complexity theory (where “explicit” has a precise meaning) [AFT11, Raz13, LO15, Lan15]. Known barriers suggest that it may be hard to go beyond $r \approx n$ for order-3 tensors [EGOW18, GMOW19].

Our contributions: upper and lower bounds. Our starting point is the work of Garg, Kayal and Saha [GKS20] who showed a surprising connection between algebraic complexity theory and learning. In particular they showed how a powerful lower bound method, called the method of *shifted partial derivatives* [CKW11], could be leveraged to give algorithms for decomposing power

sums of polynomials, which is a close relative of tensor decomposition. Specifically we will be interested in *Koszul–Young flattenings* [LO13] (see also [Lan11]) which are the main ingredient in recent lower bounds on the border rank of the matrix multiplication tensor [LO15]. Their algorithmic applications were studied in the PhD thesis of Persu [Per18] which gave an efficient algorithm for rank detection (but not decomposition) when $r \leq 3n/2$.

Our algorithmic results (“upper bounds”) apply to order-3 tensors of format $n_1 \times n_2 \times n_3$, where (without loss of generality) $n_1 \leq n_2 \leq n_3$. To simplify the discussion, suppose further that we are in an asymptotic regime where $n_1 \rightarrow \infty$ and $n_3/n_2 = O(1)$. In the setting of generic components, we give algorithms for both the rank detection and decomposition tasks, as well as an efficient uniqueness theorem. These results tolerate rank as large as $r \leq (1 - \epsilon)(n_2 + n_3)$ for an arbitrary $\epsilon > 0$. For any fixed ϵ , the runtime is polynomial in n_3 . For comparison, the previous state-of-the-art results for rank detection [Per18] and decomposition [Koi24a] assume $n_2 = n_3 = n$ and require $r \leq 3n/2$ and $r \leq 4n/3$ respectively. Also, our efficient uniqueness theorem is the first uniqueness theorem (efficient or not) to surpass Kruskal’s classical bound $r \leq 3n/2$ [Kru77].

In addition to the Koszul–Young flattenings, a key ingredient in our approach will be the algorithm of [JLV23] (see also the precursors [Car91, DL06, DLCC07, BVD⁺18]) for finding rank-1 matrices in a linear subspace. In contrast, the decomposition algorithm of [Koi24a] takes a different approach, based on connections with the problem of computing a *commuting extension*. We can exploit this connection in the other direction: Appendix B shows how our decomposition algorithm implies an algorithm for commuting extensions that is (in some sense) nearly optimal.

The blueprint for our algorithm is to solve rank detection by “flattening” the tensor to a particular matrix whose rank can be related to the original tensor rank. We also investigate limitations of this approach (“lower bounds”), building on [EGOW18]. We give a various results for successively broader classes of flattenings. First we identify a subclass of linear flattenings that provably cannot surpass rank $n_2 + n_3$, making our method essentially optimal within this class. We also define a much more general class of degree- d polynomial flattenings and prove that for $n \times n \times n$ tensors, these cannot surpass rank Cn for a constant $C = C(d)$. This extends [EGOW18], which handles the case $d = 1$. We note that this style of lower bound is distinct from the low-degree polynomial framework used by [Wei23] for the case of random components, as that approach only applies to fully-Bayesian random models.

Our lower bounds formalize the sense in which reaching $r \gg n$ would require a rather different approach, and might be fundamentally out of reach for efficient algorithms. Of course, there are other approaches that are not ruled out by our lower bounds. For instance, there is the commuting extension approach of [Koi24a], but curiously this appears to have a natural breaking point at $r = 2n$ (see Section 2 of [Koi24b]), matching our algorithm. We note that there is an explicit tensor with a rank lower bound of roughly $3n$ [AFT11], but this does not appear to suggest a decomposition algorithm.

Model of computation. Our algorithm only uses standard linear algebra primitives such as solving linear equations, computing eigenvalues, and intersecting linear subspaces (all applied to polynomial-sized inputs). The eigenvalue computation is used within simultaneous diagonalization, which is a subroutine of [JLV23]. Following [Koi24a, Section 1.4] (and numerous prior works in the tensor decomposition literature), we treat these as atomic operations and assume they can be implemented exactly. Our algorithm can be implemented to give guarantees in the computational model that allows rational arithmetic with polynomial bits of precision.

We leave for future work the question of obtaining a detailed analysis of the stability and robustness of our algorithm, e.g., in the style of [BCMV14b, GVX14]. Certainly our description

and analysis of the algorithm in exact terms is a necessary step toward this goal. Another seemingly necessary step would be to obtain a robust analysis of the [JLV23] algorithm (which we use as a subroutine), but currently this only exists for the “non-planted” variant of the problem [BESV24].

2 Main Results

2.1 Preliminaries

2.1.1 Tensor Decomposition and Rank

Definition 2.1 (Tensor Rank). *For vectors $a^{(1)} \in \mathbb{R}^{n_1}, a^{(2)} \in \mathbb{R}^{n_2}, \dots, a^{(k)} \in \mathbb{R}^{n_k}$, the rank-1 tensor $a^{(1)} \otimes a^{(2)} \otimes \dots \otimes a^{(k)}$ has order k and format $n_1 \times \dots \times n_k$, with entries $(a^{(1)} \otimes \dots \otimes a^{(k)})_{i_1, \dots, i_k} = a_{i_1}^{(1)} a_{i_2}^{(2)} \dots a_{i_k}^{(k)}$. In general, the rank of an $n_1 \times \dots \times n_k$ tensor T is the minimum r for which T can be expressed as the sum of r rank-1 tensors: $T = \sum_{\ell=1}^r a^{(1,\ell)} \otimes \dots \otimes a^{(k,\ell)}$ where $a^{(i,\ell)} \in \mathbb{R}^{n_i}$.*

Definition 2.2 (Unique Decomposition). *For an $n_1 \times \dots \times n_k$ tensor T and a positive integer r , we say that a decomposition $T = \sum_{\ell=1}^r a^{(1,\ell)} \otimes \dots \otimes a^{(k,\ell)}$ into r rank-1 terms is unique if for any decomposition $T = \sum_{\ell=1}^r \tilde{a}^{(1,\ell)} \otimes \dots \otimes \tilde{a}^{(k,\ell)}$ there exists a permutation $\pi : [r] \rightarrow [r]$ such that $\tilde{a}^{(1,\ell)} \otimes \dots \otimes \tilde{a}^{(k,\ell)} = a^{(1,\pi(\ell))} \otimes \dots \otimes a^{(k,\pi(\ell))}$ for all $\ell \in [r]$. When we speak of recovering such a decomposition, we mean the task of finding the r rank-1 tensors (which can be listed in any order), given T as input.*

Note that if T has a unique decomposition into r rank-1 terms then none of these terms can be zero, and the rank of T must be exactly r .

2.1.2 Genericity and Symbolic Rank

Throughout, we use the standard meaning of the term “generic” from algebraic geometry. Namely, for symbolic variables $x = (x_1, \dots, x_p)$, we say that a property holds true “generically” (or holds for “generically chosen x ”) if there exists a nonidentically-zero polynomial $P(x_1, \dots, x_p) \in \mathbb{R}[x]$ such that the property holds true for all $x \in \mathbb{R}^n$ such that $P(x) \neq 0$. If a property holds true generically, it holds for “almost all” x in the sense that the set of x for which the property fails is measure zero. If a finite collection of properties each holds true generically, then it is also generically true that all the properties hold simultaneously; this can be seen by multiplying together all the associated polynomials $P(x)$.

For us, a common use case for genericity will be in reference to a tensor decomposition $T = \sum_{\ell=1}^r a^{(1,\ell)} \otimes \dots \otimes a^{(k,\ell)}$ where $a^{(i,\ell)} \in \mathbb{R}^{n_i}$. When we say a property holds true for “generically chosen components,” we mean the notion of generic from above, where the variables x are all the entries of all the vectors $a^{(i,\ell)}$ (so $r \sum_{i=1}^k n_i$ variables in total).

Suppose $M = M(x)$ is a matrix whose entries are polynomials in the variables $x = (x_1, \dots, x_p)$, that is, $M \in \mathbb{R}[x]^{m \times n}$. Every minor of M (determinant of a square submatrix) is a polynomial in the variables x . The *symbolic rank* of M is the largest integer r for which M has an $r \times r$ minor that is not identically zero. Equivalently, this is the rank of M over the fraction field $\mathbb{R}(x)$. Note that if M has symbolic rank r then $M(x)$ has rank r for generically chosen x ; this follows from the characterization of rank as the size of the largest nonzero minor. This allows us to speak of the “generic rank” of a matrix (which equals the symbolic rank). For specific values of x , the rank of $M(x)$ may be lower than the generic rank of M , but never higher (since all minors larger than the generic rank are identically zero).

2.2 Upper Bounds

2.2.1 Trivial Flattenings

As a benchmark, we first discuss the “trivial” flattenings of tensors to matrices. Suppose T is a tensor of format $n_1 \times \cdots \times n_k$. Choose a subset $S \subseteq [k]$ and let $\bar{S} := [k] \setminus S$ denote its complement. We can flatten T to a matrix $M^{\text{triv}} = M^{\text{triv}}(T; S)$ of format $(\prod_{i \in S} n_i) \times (\prod_{i \in \bar{S}} n_i)$ in a straightforward way:

$$M_{(i_j : j \in S), (i_j : j \in \bar{S})}^{\text{triv}} := T_{i_1, \dots, i_k} \quad \text{where } i_j \in [n_j]. \quad (2)$$

A rank-1 tensor $a^{(1)} \otimes \cdots \otimes a^{(k)}$ flattens to a rank-1 matrix

$$M^{\text{triv}}(a^{(1)} \otimes \cdots \otimes a^{(k)}; S) = \left(\bigotimes_{i \in S} a^{(i)} \right) \left(\bigotimes_{i \in \bar{S}} a^{(i)} \right)^\top.$$

A rank- r tensor $T = \sum_{\ell=1}^r a^{(1,\ell)} \otimes \cdots \otimes a^{(k,\ell)}$ flattens to the matrix

$$\sum_{\ell=1}^r M^{\text{triv}}(a^{(1,\ell)} \otimes \cdots \otimes a^{(k,\ell)}; S),$$

which has rank at most r . For rank detection, it would be useful if the rank were exactly r , i.e., if the rank of the flattening were *additive* across the rank-1 terms of the tensor, so that we could read off the rank of T from the rank of M^{triv} . The rank of $M^{\text{triv}}(T; S)$ certainly cannot exceed the minimum of its two dimensions, namely

$$R_S := \min \left\{ \prod_{i \in S} n_i, \prod_{i \in \bar{S}} n_i \right\}.$$

Conceivably, additivity could hold as long as $r \leq R_S$. Indeed, this is the case as long as the components $a^{(i,\ell)}$ are generically chosen in the sense of Section 2.1.2; we include the proof in Appendix A for completeness. This gives a solution to the rank detection problem for tensors of rank up to R_S (with generic components). Taking the best choice of S , the trivial flattening can solve rank detection up to rank

$$n_* := \max_{S \subseteq [k]} \min \left\{ \prod_{i \in S} n_i, \prod_{i \in \bar{S}} n_i \right\}.$$

For order- k tensors of format $n \times \cdots \times n$, this gives rank detection up to rank $n^{\lfloor k/2 \rfloor}$.

Section 6 of [JLV23] uses the trivial flattening to give an algorithm for recovering the minimum-rank decomposition (and certifying its uniqueness), under slightly stronger assumptions. The idea is to search the column span of the flattening for vectors that have rank-1 structure. For instance, Corollary 27 of [JLV23] can decompose an $n_1 \times n_2 \times n_3$ tensor with $n_1 \leq n_2 \leq n_3$ and generically chosen components, of rank r up to

$$r \leq \min \left\{ \frac{1}{4}(n_1 - 1)(n_2 - 1), n_3 \right\}. \quad (3)$$

For comparison, n_* takes the value $\min\{n_1 n_2, n_3\}$ in this setting.

2.2.2 Koszul–Young Flattenings

Our new algorithm is inspired by an idea first introduced by [LO13] as “Young flattenings” [LO13, Lan11, Lan15] and later referred to as “Koszul flattenings” [OO13, LM17, HOOS19] or “Koszul–Young flattenings” [Far16, Gual17]. Essentially, Koszul–Young flattenings provide a non-trivial linear map from tensors to matrices. These ideas have already appeared in tensor decomposition algorithms [OO13], but for $n \times n \times n$ tensors, these results only handle the undercomplete case $r \leq n$. More recently, [Per18] used Koszul–Young flattenings to solve rank detection for overcomplete tensors with generic components, namely $n \times n \times n$ tensors of rank up to $3n/2$. Furthermore, [Per18] proposed a hierarchy of poly-time approaches for rank detection that conjecturally approaches rank $2n$, as well as an approach for decomposition that conjecturally tolerates rank $(1 + \epsilon)n$ for a small constant $\epsilon > 0$. We will resolve and extend these conjectures: for sufficiently large $n \times n \times n$ tensors with generic components, we solve both rank detection and decomposition in polynomial time, up to rank $(2 - \epsilon)n$ for an arbitrary constant $\epsilon > 0$. We will also handle non-square tensors.

While the Koszul–Young flattenings are typically described in abstract algebraic terms, we will give instead a very explicit combinatorial description of a matrix based on these ideas. Suppose T is $n_1 \times n_2 \times n_3$ and fix integers p, q with $0 \leq p \leq p + 1 \leq q \leq n_1$. Define a matrix $M = M(T; p, q)$ with rows indexed by

$$\{(S, j) : S \subseteq [q], |S| = p, j \in [n_2]\},$$

columns indexed by

$$\{(U, k) : U \subseteq [q], |U| = p + 1, k \in [n_3]\},$$

and entries

$$M_{Sj, Uk} := \sum_{i=1}^q \mathbb{1}_{U=S \sqcup \{i\}} \cdot \sigma(U, i) \cdot T_{ijk}. \quad (4)$$

Here, $\mathbb{1}_A$ denotes the $\{0, 1\}$ -valued indicator for an event A , and $\sigma(U, i) \in \{\pm 1\}$ is the parity of i 's position in U , that is,

$$\sigma(U, i) := (-1)^{|\{j \in U : j < i\}|}. \quad (5)$$

While we have written (4) as a sum over $i \in [q]$, at most one term in this sum can be nonzero. Note that only q slices of T (along the first mode) are used.

The dimensions of M are $\binom{q}{p} n_2 \times \binom{q}{p+1} n_3$. We will leave p, q as free variables for now, but eventually it will be advantageous to take q to be a (large) constant so that M has polynomial size, and $p \approx q \frac{n_3}{n_2 + n_3}$ which makes M as square as possible. One special case of interest is $n_2 = n_3$, in which case we can take q odd with $q = 2p + 1$ so that M is exactly square, since $\binom{2p+1}{p} = \binom{2p+1}{p+1}$; the prior work [Per18] considers only this square case and mainly focuses on $p = 1$. We remark that the extreme cases $p = 0$ and $p = q - 1$ essentially reduce to “trivial” flattenings from the previous section (but with additional signs in the case $p = q - 1$).

The flattening (4) bears resemblance to a signed variant of the *Kikuchi matrices* that were introduced by [WEM19] (with a variation independently discovered by [Has20]) and have subsequently appeared in results on constraint satisfaction [GKM22], error-correcting codes [AGKM23, KM24a, KM24b], and extremal combinatorics [GKM22, HKM23]. In particular, the related matrix A defined in (6) below resembles a Kikuchi matrix in the sense that its rows and columns are indexed by subsets and it only contains a nonzero entry when the two subsets have their symmetric difference of a particular size (in this case, 1).

2.2.3 Rank Detection

If T is a rank-1 tensor $T = a \otimes b \otimes c$ then the flattening $M = M(T; p, q)$ from (4) is the Kronecker product $M = A \otimes (bc^\top)$ where $A = A(a; p, q)$ is the $\binom{q}{p} \times \binom{q}{p+1}$ matrix

$$A_{SU} = \sum_{i=1}^q \mathbb{1}_{U=S \sqcup \{i\}} \cdot \sigma(U, i) \cdot a_i. \quad (6)$$

Therefore, as long as neither b nor c is the zero vector, M has the same rank as A . We will see that the matrix A is rank-deficient, with generic rank $\binom{q-1}{p}$. If, for instance, $q = 2p + 1$ then A is square since $\binom{2p+1}{p} = \binom{2p+1}{p+1}$, but crucially, its generic rank $\binom{2p}{p}$ is roughly half its dimension (for large p). The following is proved in Section 3.1.

Proposition 2.3. *For a rank-1 tensor $T = a \otimes b \otimes c$ with a, b, c generically chosen, the matrix $M(T; p, q)$ defined in (4) has rank exactly $\binom{q-1}{p}$.*

The transformation $T \mapsto M(T; p, q)$ is linear, so a rank- r tensor $T = \sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)}$ flattens to the matrix

$$M(T; p, q) = \sum_{\ell=1}^r M(a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)}; p, q).$$

If the rank were additive across these r terms then $M(T; p, q)$ would need to have rank $r \binom{q-1}{p}$, in light of Proposition 2.3. Conceivably, additivity could hold as long as $r \binom{q-1}{p}$ does not exceed the minimum dimension of $M(T; p, q)$, which is $\min\{\binom{q}{p}n_2, \binom{q}{p+1}n_3\}$. Dividing through by $\binom{q-1}{p}$, this condition becomes $r \leq \min\{\frac{q}{q-p}n_2, \frac{q}{p+1}n_3\}$. Choosing q large and $p \approx q \frac{n_3}{n_2+n_3}$ to balance the two terms in $\min\{\dots\}$, this heuristic suggests that rank detection might be possible for r as large as, roughly, $n_2 + n_3$. We show that indeed this pans out.

Theorem 2.4 (Rank detection). *Let $1 \leq q \leq n_1$ and*

$$p = \left\lfloor q \cdot \frac{n_3}{n_2 + n_3} \right\rfloor.$$

If $T = \sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)}$ is $n_1 \times n_2 \times n_3$ with generically chosen components, and

$$r \leq (n_2 + n_3) \left(1 - \frac{1 + \alpha}{q}\right) - q \quad \text{where } \alpha := \max\left\{\frac{n_2}{n_3}, \frac{n_3}{n_2}\right\},$$

then the matrix $M(T; p, q)$ defined in (4) has rank exactly $r \binom{q-1}{p}$.

The proof can be found in Section 3.2. Of course, one has the option to permute the 3 modes of the tensor before applying Theorem 2.4, and it is advantageous to make n_1 the smallest mode. Suppose $n_1 \leq n_2 \leq n_3$ and consider the regime $n_1 \rightarrow \infty$ with α of constant order. For any constant $\epsilon > 0$ we can handle r as large as $(1 - \epsilon)(n_2 + n_3)$ by taking q to be a sufficiently large constant (depending on α, ϵ). For comparison, the trivial flattening of Section 2.2.1 works up to rank n_3 in this regime. Our advantage over the trivial flattening diminishes as the ratio $\alpha = n_3/n_2$ increases. As long as q is a constant, the size of M and the complexity of computing its entries are polynomial in the largest dimension n_3 .

Remark 2.5 (Certifying lower bounds on tensor rank). *For an arbitrary tensor T , we always have $\text{rank}(M(T; p, q)) \leq \binom{q-1}{p} \cdot \text{rank}(T)$. This follows from the fact that the rank of A from (6) never exceeds its generic rank $\binom{q-1}{p}$; see Lemma 3.1. As a result, M provides a certificate that $\text{rank}(T) \geq \text{rank}(M(T; p, q)) / \binom{q-1}{p}$ for any given tensor T .*

2.2.4 Decomposition Algorithm

We now outline our strategy for how to use the flattening (4) to extract the components of the tensor. The aim of this section is to describe the algorithm and the intuition behind it, while Section 2.2.5 below will contain the main theorems that guarantee this algorithm's success when the tensor components are generically chosen.

Consider an $n_1 \times n_2 \times n_3$ tensor $T = \sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)}$. Build the matrix $M = M(T; p, q)$ defined in (4) and compute its column span. Recall that each rank-1 term $a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)}$ flattens to $A(a^{(\ell)}; p, q) \otimes (b^{(\ell)} c^{(\ell)\top})$, where $A(a; p, q)$ is defined in (6). Generically, we have the additivity property of Theorem 2.4, implying that the column span of M is equal to the span of all columns of these individual flattenings, that is,

$$\text{colspan}(M) = \text{span}\{z^{(U)}(a^{(\ell)}; p, q) \otimes b^{(\ell)} : U \subseteq [q], |U| = p + 1, \ell \in [r]\} \quad (7)$$

where $z^{(U)}(a; p, q) \in \mathbb{R}^{\binom{q}{p}}$ denotes column U of $A(a; p, q)$.

Fix one particular column of A , namely $V := [p + 1]$. The column $z^{(V)}$ has a particular sparsity pattern, with nonzeros only allowed in the $p + 1$ rows S where $S \subseteq V$. Let $Z_{p,q} \subseteq \mathbb{R}^{\binom{q}{p}}$ denote the $(p + 1)$ -dimensional subspace comprised of all vectors with this sparsity pattern, that is,

$$Z_{p,q} := \text{span}\{e^{(S)} : S \subseteq [q], |S| = p, S \subseteq [p + 1]\} \subseteq \mathbb{R}^{\binom{q}{p}},$$

where $e^{(S)}$ denotes the standard unit basis vector. We will compute the intersection of the two subspaces $\text{colspan}(M)$ and $Z_{p,q} \otimes \mathbb{R}^{n_2}$. By design, the r vectors $z^{(V)}(a^{(\ell)}; p, q) \otimes b^{(\ell)}$ for $\ell \in [r]$ lie in the intersection, and we will show that, generically, the intersection does not contain any additional “spurious” elements, that is,

$$\text{colspan}(M) \cap (Z_{p,q} \otimes \mathbb{R}^{n_2}) = \text{span}\{z^{(V)}(a^{(\ell)}; p, q) \otimes b^{(\ell)} : \ell \in [r]\}. \quad (8)$$

It will be convenient to change basis so as to remove the zero entries of $z^{(V)}$. To this end, we will define a specific isomorphism $\varphi_{p,q} : Z_{p,q} \rightarrow \mathbb{R}^{p+1}$, namely the linear map defined on unit basis vectors as follows:

$$\varphi_{p,q}(e^{([p+1] \setminus \{i\})}) = (-1)^{i-1} e^{(i)} \quad \text{for } i \in [p + 1]. \quad (9)$$

This is designed so that $\varphi_{p,q}(z^{(V)}(a; p, q)) = (a_1, \dots, a_{p+1})^\top$, since entry $V \setminus \{i\}$ of $z^{(V)}(a; p, q)$ is equal to $\sigma(V, i) a_i = (-1)^{i-1} a_i$. Defining $d^{(\ell)}$ as the first $p + 1$ entries of $a^{(\ell)}$,

$$d^{(\ell)} := (a_1^{(\ell)}, \dots, a_{p+1}^{(\ell)})^\top,$$

this allows us to write (8) as

$$(\varphi_{p,q} \otimes I)(\text{colspan}(M) \cap (Z_{p,q} \otimes \mathbb{R}^{n_2})) = \text{span}\{d^{(\ell)} \otimes b^{(\ell)} : \ell \in [r]\}, \quad (10)$$

where I is the identity map $\mathbb{R}^{n_2} \rightarrow \mathbb{R}^{n_2}$.

To recap, we can compute the left-hand side of (10) from the flattening M , which gives us access to the right-hand side. This means we know the span of r different rank-1 tensors $d^{(\ell)} \otimes b^{(\ell)}$ (which can also be viewed as rank-1 matrices of dimension $(p + 1) \times n_2$) and our next objective will be to extract the individual rank-1 elements $d^{(\ell)} \otimes b^{(\ell)}$. Since our rank-1 tensors are generic, the work of [JLV23] provides an algorithm for precisely this task, under the condition $r \leq p(n_2 - 1)/4$; see Corollary 3 of [JLV23], which considers a generalization of this task (in their notation, we only need the case $s = R$). This allows us to recover vectors $\hat{d}^{(\ell)}$ and $\hat{b}^{(\ell)}$ for $\ell \in [r]$, such that $\hat{d}^{(\ell)} = \alpha_\ell \cdot d^{(\pi(\ell))}$ and $\hat{b}^{(\ell)} = \beta_\ell \cdot b^{(\pi(\ell))}$ where $\alpha_\ell, \beta_\ell \in \mathbb{R}$ are nonzero scalars and π is a permutation $[r] \rightarrow [r]$.

At this point, we repeat the entire algorithm so far with the “ b ” and “ c ” modes switched. We use the same value for q , but use $q - p - 1$ in place of p so that the new flattening remains as square as possible. This allows us to obtain vectors $\hat{c}^{(\ell)}$ and $\hat{f}^{(\ell)}$ which serve as estimates for $c^{(\ell)}$ and

$$f^{(\ell)} := (a_1^{(\ell)}, \dots, a_{q-p}^{(\ell)})^\top,$$

respectively (again, up to re-ordering and scalar multiple). By comparing the $\hat{d}^{(\ell)}$ and $\hat{f}^{(\ell)}$ vectors we can identify the correct pairing between $\hat{b}^{(\ell)}$ and $\hat{c}^{(\ell)}$, and re-index the $\hat{c}^{(\ell)}$ so that $\hat{b}^{(1)}$ is paired with $\hat{c}^{(1)}$, etc. Now we know the decomposition takes the form

$$T = \sum_{\ell=1}^r w^{(\ell)} \otimes \hat{b}^{(\ell)} \otimes \hat{c}^{(\ell)}$$

for unknown vectors $w^{(\ell)}$, which can be recovered by solving a linear system of equations; we will show that, generically, this linear system has a unique solution, namely the desired decomposition.

The full algorithm is presented as Algorithm 1 below. The second phase of the algorithm involves a flattening $M' = M'(T; p, q)$ that is identical to M but with the “ b ” and “ c ” modes switched, and with $q - p - 1$ in place of p . Formally, $M'(T; p, q)$ is indexed by

$$\{(S, k) : S \subseteq [q], |S| = q - p - 1, k \in [n_3]\} \times \{(U, j) : U \subseteq [q], |U| = q - p, j \in [n_2]\}$$

with entries

$$M'_{S k, U j} := \sum_{i=1}^q \mathbb{1}_{U=S \sqcup \{i\}} \cdot \sigma(U, i) \cdot T_{ijk}. \quad (11)$$

Algorithm 1 (Overcomplete third-order tensor decomposition).

- *Input:* A tensor $T \in \mathbb{R}^{n_1 \times n_2 \times n_3}$.
 - *Input:* Parameters $p, q \in \mathbb{Z}$ with $q \leq n_1$ and $\bar{p} := \min\{p + 1, q - p\} \geq 2$.
 - *Input:* A parameter r (or learn r from $\text{rank}(M)$ using Theorem 2.4).
1. Build the matrix $M = M(T; p, q)$ as in (4).
 2. Construct the subspace $B := (\varphi_{p,q} \otimes I)(\text{colspan}(M) \cap (Z_{p,q} \otimes \mathbb{R}^{n_2})) \subseteq \mathbb{R}^{p+1} \otimes \mathbb{R}^{n_2}$.
 3. Run the algorithm of [JLV23, Corollary 3] to find rank-1 tensors in B , call them $\hat{d}^{(1)} \otimes \hat{b}^{(1)}, \dots, \hat{d}^{(r)} \otimes \hat{b}^{(r)}$. If the number of rank-1 tensors found is not r , output “fail.”
 4. Build the matrix $M' = M'(T; p, q)$ as in (11).
 5. Construct the subspace $C := (\varphi_{q-p-1,q} \otimes I)(\text{colspan}(M') \cap (Z_{q-p-1,q} \otimes \mathbb{R}^{n_3})) \subseteq \mathbb{R}^{q-p} \otimes \mathbb{R}^{n_3}$.
 6. Run the algorithm of [JLV23, Corollary 3] to find rank-1 tensors in C , call them $\hat{f}^{(1)} \otimes \hat{c}^{(1)}, \dots, \hat{f}^{(r)} \otimes \hat{c}^{(r)}$. If the number of rank-1 tensors found is not r , output “fail.”
 7. Define a permutation $\tau : [r] \rightarrow [r]$ so that the first \bar{p} entries of $\hat{f}^{(\tau(\ell))}$ form a scalar multiple of the first \bar{p} entries of $\hat{d}^{(\ell)}$ for each $\ell \in [r]$. If this is not possible, output “fail.”

8. Solve the following linear system of equations in variables $w^{(\ell)} \in \mathbb{R}^{n_1}$ for $\ell \in [r]$:

$$\sum_{\ell=1}^r w^{(\ell)} \otimes \hat{b}^{(\ell)} \otimes \hat{c}^{(\tau(\ell))} = T, \quad (12)$$

and output the resulting rank- r decomposition of T . If there is no solution, output “fail.”

We have assumed $\bar{p} \geq 2$, since otherwise Step 7 will not have a unique solution. If $\bar{p} = 1$, this means either $p = 0$ or $p = q - 1$, which are the extreme cases where $M(T; p, q)$ essentially reduces to a trivial flattening. In these cases, one should use a simpler decomposition algorithm of [JLV23] based on the trivial flattening, as we mentioned in Section 2.2.1; this algorithm works under the condition (3).

2.2.5 Uniqueness Theorem and Algorithmic Guarantees

We will see that the discussion above gives not only a decomposition algorithm that works for generic rank- r tensors but also an explicit list of efficiently-checkable conditions on the components that guarantees uniqueness of the decomposition. We first present these conditions (Theorem 2.7) and then show that generic components satisfy these conditions when r is small enough (Theorem 2.8).

Some of the conditions below may appear rather opaque, so we first briefly motivate what role they play. It is natural for the matrix M to appear (condition (vi)), given its use as a rank detection device (Theorem 2.4) and its role in the decomposition algorithm, but we will also define two additional matrices N and P . The condition (viii) on N is used to ensure that (8) holds (no spurious elements in the subspace intersection), while the condition (x) on P is used to ensure that the [JLV23] algorithm succeeds. The matrices M', N', P' are identical to M, N, P but with the “ b ” and “ c ” modes of the tensor switched and with p replaced by $q - p - 1$.

Definition 2.6. Consider an $n_1 \times n_2 \times n_3$ tensor T with decomposition

$$T = \sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)} \quad (13)$$

and parameters $p, q \in \mathbb{Z}$ with $q \leq n_1$ and $\bar{p} := \min\{p + 1, q - p\} \geq 2$.

- The matrix N has rows indexed by

$$\{(S, j) : S \subseteq [q], |S| = p, S \not\subseteq [p + 1], j \in [n_2]\}$$

and columns indexed by

$$\{(U, \ell) : U \subseteq [q], |U| = p + 1, 1 \in U, U \neq [p + 1], \ell \in [r]\},$$

with entries

$$N_{Sj, U\ell} = b_j^{(\ell)} \cdot \sum_{i=1}^q \mathbb{1}_{U=S \sqcup \{i\}} \cdot \sigma(U, i) \cdot a_i^{(\ell)}, \quad (14)$$

where σ is defined in (5).

- The matrix N' has rows indexed by

$$\{(S, k) : S \subseteq [q], |S| = q - p - 1, S \not\subseteq [q - p], k \in [n_3]\}$$

and columns indexed by

$$\{(U, \ell) : U \subseteq [q], |U| = q - p, 1 \in U, U \neq [q - p], \ell \in [r]\},$$

with entries

$$N'_{Sk, U\ell} = c_k^{(\ell)} \cdot \sum_{i=1}^q \mathbb{1}_{U=S \sqcup \{i\}} \cdot \sigma(U, i) \cdot a_i^{(\ell)}.$$

- The matrix P is indexed by

$$\{(i_1, i_2, j_1, j_2) : 1 \leq i_1 < i_2 \leq p + 1, 1 \leq j_1 < j_2 \leq n_2\} \times \{(\ell_1, \ell_2) : 1 \leq \ell_1 < \ell_2 \leq r\}$$

with entries

$$P_{i_1 i_2 j_1 j_2, \ell_1 \ell_2} = a_{i_1}^{(\ell_1)} b_{j_1}^{(\ell_1)} a_{i_2}^{(\ell_2)} b_{j_2}^{(\ell_2)} + a_{i_1}^{(\ell_2)} b_{j_1}^{(\ell_2)} a_{i_2}^{(\ell_1)} b_{j_2}^{(\ell_1)} - a_{i_1}^{(\ell_1)} b_{j_2}^{(\ell_1)} a_{i_2}^{(\ell_2)} b_{j_1}^{(\ell_2)} - a_{i_1}^{(\ell_2)} b_{j_2}^{(\ell_2)} a_{i_2}^{(\ell_1)} b_{j_1}^{(\ell_1)}. \quad (15)$$

- The matrix P' is indexed by

$$\{(i_1, i_2, j_1, j_2) : 1 \leq i_1 < i_2 \leq q - p, 1 \leq j_1 < j_2 \leq n_3\} \times \{(\ell_1, \ell_2) : 1 \leq \ell_1 < \ell_2 \leq r\}$$

with entries

$$P'_{i_1 i_2 j_1 j_2, \ell_1 \ell_2} = a_{i_1}^{(\ell_1)} c_{j_1}^{(\ell_1)} a_{i_2}^{(\ell_2)} c_{j_2}^{(\ell_2)} + a_{i_1}^{(\ell_2)} c_{j_1}^{(\ell_2)} a_{i_2}^{(\ell_1)} c_{j_2}^{(\ell_1)} - a_{i_1}^{(\ell_1)} c_{j_2}^{(\ell_1)} a_{i_2}^{(\ell_2)} c_{j_1}^{(\ell_2)} - a_{i_1}^{(\ell_2)} c_{j_2}^{(\ell_2)} a_{i_2}^{(\ell_1)} c_{j_1}^{(\ell_1)}.$$

Theorem 2.7 (Uniqueness theorem). *Consider an $n_1 \times n_2 \times n_3$ tensor T with decomposition*

$$T = \sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)} \quad (16)$$

and parameters $p, q \in \mathbb{Z}$ with $q \leq n_1$ and $\bar{p} := \min\{p + 1, q - p\} \geq 2$. Suppose the components satisfy the following conditions:

- (i) $a_1^{(\ell)} \neq 0$ for all $\ell \in [r]$.
- (ii) The r vectors $(a_1^{(\ell)}, \dots, a_{\bar{p}}^{(\ell)})^\top$ for $\ell \in [r]$ are pairwise linearly independent, i.e., no one of these vectors is a scalar multiple of another.
- (iii) The r vectors $\{d^{(\ell)} \otimes b^{(\ell)} : \ell \in [r]\}$ are linearly independent, where $d^{(\ell)} := (a_1^{(\ell)}, \dots, a_{p+1}^{(\ell)})^\top$.
- (iv) The r vectors $\{f^{(\ell)} \otimes c^{(\ell)} : \ell \in [r]\}$ are linearly independent, where $f^{(\ell)} := (a_1^{(\ell)}, \dots, a_{q-p}^{(\ell)})^\top$.
- (v) The r vectors $\{b^{(\ell)} \otimes c^{(\ell)} : \ell \in [r]\}$ are linearly independent.
- (vi) The matrix $M = M(T; p, q)$ from (4) has rank exactly $r \binom{q-1}{p}$.
- (vii) The matrix $M' = M'(T; p, q)$ from (11) has rank exactly $r \binom{q-1}{p}$.
- (viii) The matrix N from Definition 2.6 has full column rank.
- (ix) The matrix N' from Definition 2.6 has full column rank.
- (x) The matrix P from Definition 2.6 has full column rank.

(xi) The matrix P' from Definition 2.6 has full column rank.

Then (16) is the unique rank- r decomposition of T (in the sense of Definition 2.2), and this decomposition is recovered by Algorithm 1 (with input T, p, q, r).

The proof of Theorem 2.7 can be found in Section 3.3. It may be possible to simplify the above conditions or to find conditions that are “minimal” in some sense, but we have not attempted this here. If the above conditions fail to hold for a particular decomposition, one can potentially remedy this by permuting the modes of the tensor or by applying an invertible change of basis on one or more of the modes. Next we show that the above conditions hold for “almost all” tensors of low enough rank.

Theorem 2.8 (Success for generic tensors). *Consider the setting of Theorem 2.7 with parameters $q \leq n_1$ and*

$$p = \left\lfloor q \cdot \frac{n_3}{n_2 + n_3} \right\rfloor.$$

Suppose

$$q \geq (4 + 5\alpha) \left(1 + \frac{1}{\alpha}\right) \quad \text{where } \alpha := \max \left\{ \frac{n_2}{n_3}, \frac{n_3}{n_2} \right\}.$$

If the components $a^{(\ell)}, b^{(\ell)}, c^{(\ell)}$ are generically chosen with

$$r \leq (n_2 + n_3) \left(1 - \frac{3 + \alpha}{q}\right) - \frac{q^3}{4} \tag{17}$$

then $\bar{p} \geq 2$ and conditions (i)–(xi) are satisfied.

The proof can be found in Section 3.6. This is qualitatively the same behavior as our rank detection result. Suppose $n_1 \leq n_2 \leq n_3$ and consider the regime $n_1 \rightarrow \infty$ with α of constant order. For any constant $\epsilon > 0$ we can handle r as large as $(1 - \epsilon)(n_2 + n_3)$ by taking q to be a sufficiently large constant (depending on α, ϵ). As long as q is a constant, the runtime of Algorithm 1 is polynomial in the largest dimension n_3 .

2.3 Lower Bounds

We have given an algorithm for tensor decomposition based on flattening the tensor to a matrix in a non-trivial manner: $M = M(T)$. We now explore the inherent limitations of this style of approach. A seemingly necessary prerequisite to finding a better decomposition algorithm is to find a better algorithm for the easier task of rank detection. We therefore focus on rank detection, and specifically we will show that certain types of flattenings cannot have “additivity of rank” when the tensor rank is too large. This additivity property was key to our rank detection algorithm (Theorem 2.4) and was also crucial for our decomposition algorithm because it ensures that the column span of $M(\sum_{\ell} a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)})$ retains all vectors from the individual column spans of the terms $M(a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)})$. We will give three different lower bounds that apply to increasingly general families of flattenings.

2.3.1 Flattenings of Koszul–Young Type

The flattenings used in our algorithm have the property that two modes are flattened in the trivial manner while the third is flattened according to a non-trivial linear map. Formally, we have a linear

map $M : \mathbb{R}^{n_1 \times n_2 \times n_3} \rightarrow \mathbb{R}^{N_1 \times N_2}$ (or in other words, the entries of M are degree-1 homogeneous polynomials in the variables T_{ijk}) with the property

$$M(a \otimes b \otimes c) = A(a) \otimes (bc^\top) \quad (18)$$

for a matrix $A(a)$ whose entries are degree-1 homogeneous polynomials in the entries of a . On the right-hand side above, the symbol \otimes denotes the Kronecker product of matrices. We will first see that any flattening of this form cannot surpass rank $n_2 + n_3$, making our algorithm optimal up to a factor $1 - \epsilon$ among such flattenings (in the regime $n_1 \leq n_2 \leq n_3$ with $n_1 \rightarrow \infty$ and $n_3/n_2 = O(1)$).

Theorem 2.9. *Suppose $M : \mathbb{R}^{n_1 \times n_2 \times n_3} \rightarrow \mathbb{R}^{N_1 \times N_2}$ is a linear map with the property (18). Suppose further that for generically chosen components $a^{(\ell)}, b^{(\ell)}, c^{(\ell)}$ and some $r \geq 1$, we have rank- r additivity, that is,*

$$\text{rank} \left(M \left(\sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)} \right) \right) = r \cdot \text{rank} \left(M \left(a^{(1)} \otimes b^{(1)} \otimes c^{(1)} \right) \right). \quad (19)$$

Then

$$r \leq n_2 + n_3.$$

The proof can be found in Section 4.2. Recall from Section 2.1.2 that symbolic matrices have a well-defined notion of generic rank, so the expressions on the left- and right-hand sides of (19) must each take one specific value when the components are chosen generically. (We cannot have a situation where some positive-measure set of components leads to rank r_1 while some other positive-measure set of components leads to a different rank r_2 .) The right-hand side of (19) could equivalently be written as

$$\sum_{\ell=1}^r \text{rank} \left(M \left(a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)} \right) \right),$$

since for generic components, the r terms in this sum are all equal.

2.3.2 Linear Flattenings

We now consider a more general class of flattenings, namely all linear flattenings (without the requirement (18)). Our bound on the rank will be weaker than above.

Theorem 2.10. *Suppose $M : \mathbb{R}^{n_1 \times n_2 \times n_3} \rightarrow \mathbb{R}^{N_1 \times N_2}$ is a linear map. Suppose further that for generically chosen components and some $r \geq 1$, we have rank- r additivity as defined in (19). Then*

$$r \leq 2(n_1 + n_2 + n_3 + 1).$$

The proof can be found in Section 4.3. This result is similar to Theorem 4.4 of [EGOW18], where the bound $r \leq 8n$ is proven for “square” tensors, i.e., tensors of format $n \times n \times n$. For square tensors, our result gives an improvement from $8n$ to $6n + 2$. Our argument is similar to that of [EGOW18] and we describe the difference in Remark 4.2.

2.3.3 Low-Degree Flattenings

In the introduction, we argued why it might seem plausible to expect decomposition algorithms that reach rank $r \approx n^{3/2}$ for $n \times n \times n$ tensors with generic components. In the previous section we saw that linear flattenings cannot surpass rank $O(n)$ (and this was known already by the work

of [EGOW18]). One can imagine a much broader class of flattenings where the entries of $M(T)$ are higher-degree polynomials in the variables T_{ijk} , and this is stated as an open problem in Section 6 of [EGOW18]. We will define such a class of strategies and show that even these will not surpass rank $O(n)$, where $O(\cdot)$ hides a constant depending on the polynomial degree.

For motivation, a common strategy for handling a third-order $n \times n \times n$ tensor is to apply a particular degree-2 polynomial map to produce a fourth-order $n \times n \times n \times n$ tensor, and then flatten this to an $n^2 \times n^2$ matrix. This trick is used, for instance, in work on refuting random constraint satisfaction problems [CGL07, AOW15], and within the sum-of-squares proof system it can be viewed as an application of the Cauchy–Schwarz inequality.

Inspired by this, here is a seemingly plausible approach for rank detection of third-order tensors. First, build some matrix $M(T)$ whose entries are degree-2 homogeneous polynomials in the variables T_{ijk} . Now we would like to ask for some version of “rank additivity” to hold, but since $T \mapsto M(T)$ is no longer linear, $M(\sum_{\ell} a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)})$ no longer breaks down as the sum of the individual terms $M(a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)})$. Instead, write $T = \sum_{\ell=1}^r T^{(\ell)}$ where $T^{(\ell)} = a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)}$, and view the entries of $M(T)$ as (homogeneous degree-2) polynomials in the variables $T_{ijk}^{(\ell)}$. Decompose $M = M(T)$ as

$$M = \sum_{1 \leq \ell \leq r} \hat{M}(T^{(\ell)}) + \sum_{1 \leq \ell_1 < \ell_2 \leq r} \tilde{M}(T^{(\ell_1)}, T^{(\ell_2)}) =: M_{\text{repeat}} + M_{\text{distinct}},$$

where \hat{M} has entries that are homogeneous of degree 2 in its input tensor while \tilde{M} has entries that are homogeneous with degree 1 in each of its two input tensors (so degree 2 in total). We have now written M as the sum of r “repeat” terms and $\binom{r}{2}$ “distinct” terms, and we might hope for the rank to be additive across these $\Theta(r^2)$ terms. It is tantalizing to imagine a scenario where M has dimensions $N \times N$ with $N = \Theta(n^3)$, and each of the individual terms $\hat{M}(T^{(\ell)})$ or $\tilde{M}(T^{(\ell_1)}, T^{(\ell_2)})$ has generic rank $\Theta(1)$, as this would potentially allow additive rank (and therefore successful rank detection) up to $r = \Theta(n^{3/2})$. Unfortunately we will see that this cannot pan out, and to prove a lower bound we will focus on only the “distinct” terms (which make up most of the terms) and show that even these alone cannot have additive rank.

We now define our setting more formally, and we will consider a generalization of the above discussion to any degree $d \geq 1$ and any tensor format $n_1 \times \dots \times n_k$. We think of d as a constant, so that the entries of M can be computed in polynomial time.

Definition 2.11 (Decomposition of M). *For $T \in \mathbb{R}^{n_1 \times \dots \times n_k}$, suppose $M = M(T)$ is an $N_1 \times N_2$ matrix whose entries are homogeneous degree- d polynomials in the variables T_{i_1, \dots, i_k} . Set $T = \sum_{\ell=1}^r T^{(\ell)}$ and, viewing the entries of $M(T)$ as (homogeneous degree- d) polynomials in the variables $T_{i_1, \dots, i_k}^{(\ell)}$, consider the unique decomposition of $M = M(T)$ as*

$$M = M_{\text{repeat}} + \sum_{1 \leq \ell_1 < \dots < \ell_d \leq r} \tilde{M}(T^{(\ell_1)}, \dots, T^{(\ell_d)}),$$

where \tilde{M} has entries that are homogeneous with degree 1 in each of its d input tensors (so degree d in total), and M_{repeat} consists of those monomials in M that involve fewer than d of the tensors $T^{(1)}, \dots, T^{(r)}$.

Theorem 2.12. *Suppose M is a matrix whose entries are homogeneous degree- d polynomials in the entries of an $n_1 \times \dots \times n_k$ tensor T . Set $T = \sum_{\ell=1}^r T^{(\ell)}$, and decompose $M(T)$ as in Definition 2.11. Suppose that for $T^{(\ell)} = a^{(1, \ell)} \otimes \dots \otimes a^{(k, \ell)}$ with generically chosen components $\{a^{(i, \ell)} : i \in [k], \ell \in [r]\}$*

$[r]$, we have rank- r additivity in the sense that

$$\text{rank} \left(\sum_{1 \leq \ell_1 < \dots < \ell_d \leq r} \tilde{M}(T^{(\ell_1)}, \dots, T^{(\ell_d)}) \right) = \binom{r}{d} \cdot \text{rank} \left(\tilde{M}(T^{(1)}, \dots, T^{(d)}) \right).$$

Then

$$r \leq 4^k d^2 n_*$$

where

$$n_* := \max_{S \subseteq [k]} \min \left\{ \prod_{i \in S} n_i, \prod_{i \in \bar{S}} n_i \right\}.$$

The proof can be found in Section 4.4. Recall from Section 2.2.1 that the trivial flattening achieves rank detection up to rank n_* , so Theorem 2.12 shows that for any fixed k, d , a general class of higher-degree approaches cannot beat the trivial flattening by more than a constant factor $4^k d^2$.

Remark 2.13. We discuss a few sanity checks on the definition of rank- r additivity used in Theorem 2.12. First note that for $d = 1$, this definition reduces to the one from (19), since $M_{\text{repeat}} = 0$. Next, let us give an example of a flattening that satisfies rank- r additivity for arbitrary d . Start with any linear (homogeneous degree-1) flattening $T \mapsto L(T)$ and consider the flattening $M(T) = L(T)^{\otimes d}$ (the d -fold Kronecker product), which is homogeneous of degree d . Suppose that (for generic components), L has rank- r additivity in the sense of (19). We claim that M has rank- r additivity in the sense of Theorem 2.12 (again, for generic components). We sketch the proof in the special case $k = 3, d = 2$, but the proof can be extended to all k and d . First, a basic property of the Kronecker product is $\text{rank}(A \otimes B) = \text{rank}(A) \cdot \text{rank}(B)$. When $T = \sum_{\ell=1}^r T^{(\ell)} = \sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)}$, we can decompose $M(T) = \sum_{\ell_1, \ell_2 \in [r]} L(T^{(\ell_1)}) \otimes L(T^{(\ell_2)})$. Using the additivity for L , the r^2 terms in this sum must have additive rank. This implies additivity for the $\binom{r}{2}$ “distinct” terms, which in this case are $\tilde{M}(T^{(\ell_1)}, T^{(\ell_2)}) = L(T^{(\ell_1)}) \otimes L(T^{(\ell_2)}) + L(T^{(\ell_2)}) \otimes L(T^{(\ell_1)})$ for $1 \leq \ell_1 < \ell_2 \leq r$.

2.4 Open Problems

We suggest the following questions as interesting directions for future work.

- For an $n \times n \times n$ tensor with generic components, what is the maximum rank for which a linear flattening can have additivity? We have shown the answer is asymptotically cn for a constant $c \in [2, 6]$. This improved the previous range of $[3/2, 8]$ by [Per18] and [EGOW18], but the exact value of c remains unknown.
- Can rank detection improve with the degree of flattening? Our lower bounds do not rule out this possibility, but it is possible the Koszul–Young flattening is already optimal.
- For order-4 tensors of format $n \times n \times n \times n$ (with generic components), the trivial flattening gives rank detection up to rank n^2 , but existing decomposition results stop at cn^2 for a constant $c < 1$ [DLCC07, MSS16, HSS19, JLV23]. Can we improve the decomposition results to rank $(1 - \epsilon)n^2$? This might require improving the constant $1/4$ in [JLV23], which may be sub-optimal (but curiously was not a bottleneck in the order-3 case). We remark that the constant has been improved from $1/4$ to $1 - \epsilon$ for a non-planted variant of the problem [DJL24].
- Can our algorithmic guarantees be extended to *symmetric* tensors of format $n \times n \times n$?

3 Proofs for Upper Bounds

3.1 Proof of Proposition 2.3

The genericity assumptions we will need on the components are that neither b nor c is the zero vector and that a_1, a_2, \dots, a_q are not all zero. Recall $M = A \otimes (bc^\top)$, so $\text{rank}(M) = \text{rank}(A)$. Now Proposition 2.3 follows immediately from Lemma 3.1 below, which shows that A generically has $\text{rank} \binom{q-1}{p}$. (Lemma 3.1 also includes additional facts about the structure of A that will be needed later.)

Lemma 3.1. *Consider the matrix $A = A(a; p, q)$ defined in (6). If a_1, \dots, a_q are all zero then $A = 0$ and so $\text{rank}(A) = 0$. Otherwise, $\text{rank}(A) = \binom{q-1}{p}$. Furthermore, if $a_i \neq 0$ for some $i \in [q]$ then the columns of A indexed by $\{U : i \in U\}$ form a basis for the column span of A , and similarly the rows $\{S : i \notin S\}$ form a basis for the row span.*

Proof. Fix $i \in [q]$ such that $a_i \neq 0$. We first show $\text{rank}(A) \geq \binom{q-1}{p}$ by demonstrating a $\binom{q-1}{p} \times \binom{q-1}{p}$ submatrix with full rank. Specifically, consider the submatrix indexed by $\{S : i \notin S\} \times \{U : i \in U\}$. Note that this submatrix has exactly one nonzero entry per row and exactly one nonzero entry per column, and all the nonzero entries are equal to $\pm a_i$. This shows $\text{rank}(A) \geq \binom{q-1}{p}$.

The full-rank submatrix demonstrated above implies that the columns $\{U : i \in U\}$ are linearly independent, and same for the columns $\{S : i \notin S\}$. It remains to show $\text{rank}(A) \leq \binom{q-1}{p}$. To do this, we will show that every column of A lies in the span of the $\binom{q-1}{p}$ columns $\{U : i \in U\}$. Fix a column W with $i \notin W$, and let $Y = W \sqcup \{i\}$. We claim that the columns indexed by $\{Y \setminus \{j\} : j \in Y\}$ are linearly dependent, where the coefficients are $\sigma(Y, j) a_j$. To see this, consider entry S of this linear combination. If $S \not\subseteq Y$ then entry S is equal to zero, and otherwise $S = Y \setminus \{j, k\}$ for some $j, k \in Y$ so entry S is equal to

$$\sigma(Y, j) a_j A_{S, Y \setminus \{j\}} + \sigma(Y, k) a_k A_{S, Y \setminus \{k\}} = \sigma(Y, j) a_j \sigma(Y \setminus \{j\}, k) a_k + \sigma(Y, k) a_k \sigma(Y \setminus \{k\}, j) a_j,$$

which is zero because $\sigma(Y, j) \sigma(Y \setminus \{j\}, k)$ and $\sigma(Y, k) \sigma(Y \setminus \{k\}, j)$ have opposite signs. Now take the linear dependence from above and solve for column W (which has nonzero coefficient $\sigma(Y, i) a_i$); this allows column W to be expressed as a linear combination of the columns $\{U : i \in U\}$. \square

3.2 Proof of Theorem 2.4

Provided $a_i \neq 0$ for all $i \in [q]$ (which is true for generic a), the matrix $A = A(a; p, q)$ defined in (6) can be factored as $A = \text{diag}(v) \cdot \tilde{A} \cdot \text{diag}(w)$ where $\tilde{A}_{SU} = \sum_{i \in [q]} \mathbb{1}_{U=S \sqcup \{i\}} \cdot \sigma(U, i)$, $v_S = \prod_{i \in S} a_i^{-1}$, and $w_U = \prod_{i \in U} a_i$. Note that $\tilde{A} = A(\mathbb{1})$ where $\mathbb{1} := (1, \dots, 1)^\top$, so by Lemma 3.1 we have $\text{rank}(\tilde{A}) = \binom{q-1}{p}$. We can therefore factor $\tilde{A} = \tilde{Q} \tilde{R}^\top$ where \tilde{Q}, \tilde{R} are $\binom{q}{p} \times \binom{q-1}{p}$ and $\binom{q}{p+1} \times \binom{q-1}{p}$, respectively. Now write $A = QR^\top$ where $Q = \text{diag}(v) \cdot \tilde{Q}$ and $R = \text{diag}(w) \cdot \tilde{R}$. We use the superscript (ℓ) to mark that a matrix is built from the tensor component $a^{(\ell)}$, so for instance, $A^{(\ell)} := A(a^{(\ell)}; p, q)$. Factor M as

$$M = \sum_{\ell=1}^r A^{(\ell)} \otimes (b^{(\ell)} c^{(\ell)\top}) = \begin{bmatrix} Q^{(1)} & & Q^{(r)} \\ \otimes & \cdots & \otimes \\ b^{(1)} & & b^{(r)} \end{bmatrix} \begin{bmatrix} R^{(1)} & & R^{(r)} \\ \otimes & \cdots & \otimes \\ c^{(1)} & & c^{(r)} \end{bmatrix}^\top \quad (20)$$

where, e.g., $P^{(1)} \otimes b^{(1)}$ denotes the $\binom{q}{p} n_2 \times \binom{q-1}{p}$ Kronecker product. This gives a factorization of M as the product of two matrices with inner dimension $r \binom{q-1}{p}$, so it suffices to show that the two

factors on the right-hand side of (20) have full column rank and full row rank, respectively. We will show this for the first factor and then explain the (minor) changes needed to adapt the proof for the second factor.

Let m be the smallest integer for which $r \leq mq$. It suffices to verify the extreme case $r = mq$, i.e., we will show that even when r is increased to add additional columns, all the columns are linearly independent. Focusing on the first factor in (20), the first $r \binom{q-1}{p}$ rows — call this square submatrix Q' — constitute an $m \times m$ grid of square blocks, where each square block has dimension $(q-p) \binom{q}{p} = q \binom{q-1}{p}$ consisting of (scaled) copies of Q arranged in a $(q-p) \times q$ grid. We require $n_2 \geq m(q-p)$ so that there are enough rows available to form the submatrix Q' .

Our goal is to show that $\det(Q')$ is generically nonzero, and it suffices to show this after plugging in values for some of the b variables. Namely, by setting the appropriate b variables to zero, we can keep only the m square blocks on the diagonal of Q' , and thus it suffices to show invertibility for a single square block of the form

$$Q'' = \begin{bmatrix} b_{11}Q^{(1)} & b_{12}Q^{(2)} & \cdots & b_{1,q}Q^{(q)} \\ b_{21}Q^{(1)} & b_{22}Q^{(2)} & & \vdots \\ \vdots & & \ddots & \\ b_{q-p,1}Q^{(1)} & \cdots & & b_{q-p,q}Q^{(q)} \end{bmatrix},$$

where $b_{j\ell} := b_j^{(\ell)}$, and we will similarly write $a_{i\ell} := a_i^{(\ell)}$. Partition Q'' into “cells” of size $1 \times \binom{q-1}{p}$: a cell is indexed by (S, j, ℓ) where $S \subseteq [q]$, $|S| = p$, $j \in [q-p]$, and $\ell \in [q]$, and the content of the (S, j, ℓ) cell is $(b_{j\ell} \prod_{i \in S} a_{i\ell}^{-1}) \tilde{Q}_S$ where \tilde{Q}_S denotes row S of \tilde{Q} .

We aim to show $\det(Q'')$ is nonzero as a polynomial in the variables $b_{j\ell}, a_{i\ell}^{-1}$. In the usual expansion of $\det(Q'')$ as a sum over permutations, each term chooses exactly 1 entry per row and column of Q'' . This choice can be broken down as first choosing a set of cells (S, j, ℓ) with exactly 1 cell per row (S, j) and exactly $\binom{q-1}{p}$ cells per “pillar” ℓ (where “pillar” describes a vertical stack of cells that are all indexed by the same ℓ value), and then choosing one entry per cell in a way that each column appears once. Let H denote a subset of tuples (S, j, ℓ) representing the choice at the first stage. For a given H , define matrices F_1^H, \dots, F_q^H , each of size $\binom{q-1}{p} \times \binom{q-1}{p}$, where F_ℓ^H is formed by collecting the rows \tilde{Q}_S for each $(S, j, \ell) \in H$ (i.e., the cells selected in pillar ℓ). Now

$$\det(Q'') = \sum_H \tau_H \cdot (a, b)^H \cdot \prod_{\ell=1}^q \det(F_\ell^H), \quad (21)$$

where $(a, b)^H$ denotes the monomial $\prod_{(S, j, \ell) \in H} b_{j\ell} \prod_{i \in S} a_{i\ell}^{-1}$, and $\tau_H \in \{\pm 1\}$ are signs whose values will not matter to us. If F_ℓ^H has a repeated row \tilde{Q}_S then its determinant is 0, so we need only consider H for which the cells in each pillar have distinct S values.

To complete the proof, we will identify a particular monomial that has nonzero coefficient in $\det(Q'')$. To this end, define a “good” set of cells G as follows, which (like H) will have 1 cell per row and $\binom{q-1}{p}$ cells per pillar. Within each pillar $\ell \in [q]$, for each S with $\ell \notin S$, include the cell (S, j, ℓ) where j is the first available index (i.e., the minimum $j \in [q-p]$ such that (S, j, ℓ') has not already been declared a good cell for some $\ell' < \ell$). Each S is used $q-p$ times in total (across all pillars), so the j values will not “run out.” We claim that the only term in (21) that produces the monomial $(a, b)^G$ (with nonzero coefficient) is $H = G$: within each pillar $\ell = 1, 2, \dots$, the powers of a (particularly the absence of $a_{\ell\ell}^{-1}$) determine the collection of S values that H must use (namely one copy of each S for which $\ell \notin S$), and the powers of b determine the multiset of j values to use. The choice of “first available j ” used in G constrains a unique pairing between S and j values.

We conclude that the coefficient of $(a, b)^G$ in $\det(Q'')$ is $\pm \prod_{\ell=1}^q \det(F_\ell^G)$ and it remains to show that each $\det(F_\ell^G)$ is nonzero. This follows because F_ℓ^G has rows \tilde{Q}_S for all S with $\ell \notin S$, and the corresponding rows of $\tilde{A} = A(\mathbb{1}) = \tilde{Q}\tilde{R}^\top$ are linearly independent by Lemma 3.1.

To handle the second factor in (20), we also need to show invertibility of the analogous square block

$$R'' = \begin{bmatrix} c_{11}R^{(1)} & c_{12}R^{(2)} & \cdots & c_{1,q}R^{(q)} \\ c_{21}R^{(1)} & c_{22}R^{(2)} & & \vdots \\ \vdots & & \ddots & \\ c_{p+1,1}R^{(1)} & \cdots & & c_{p+1,q}R^{(q)} \end{bmatrix},$$

whose entries are polynomials in the variables $a_{i\ell} := a_i^{(\ell)}$ and $c_{j\ell} := c_j^{(\ell)}$. We again partition R'' into cells of size $1 \times \binom{q-1}{p}$ indexed by (U, j, ℓ) where $U \subseteq [q]$, $|U| = p+1$, $j \in [p+1]$, and $\ell \in [q]$, and the content of the (U, j, ℓ) cell is $(c_{j\ell} \prod_{i \in U} a_{i\ell}) \tilde{R}_U$ where \tilde{R}_U denotes row U of \tilde{R} . The remainder of the argument is essentially identical to above, with one minor change: in the construction of G , each pillar ℓ includes the U values for which $\ell \in U$ (in place of the previous condition $\ell \notin S$).

Finally, the above constructions required an integer m such that $r \leq mq$, $n_2 \geq m(q-p)$, and $n_3 \geq m(p+1)$, or equivalently,

$$\frac{r}{q} \leq m \leq \min \left\{ \frac{n_2}{q-p}, \frac{n_3}{p+1} \right\}.$$

For this, it suffices to have

$$r \leq q \left(\min \left\{ \frac{n_2}{q-p}, \frac{n_3}{p+1} \right\} - 1 \right). \quad (22)$$

Now using the choice of $p = \lfloor q \cdot n_3 / (n_2 + n_3) \rfloor$, we can bound

$$\frac{qn_2}{q-p} \geq \frac{qn_2}{q - \left(q \cdot \frac{n_3}{n_2+n_3} - 1 \right)} = \frac{n_2 + n_3}{1 + \frac{1}{q} \left(1 + \frac{n_3}{n_2} \right)} \geq (n_2 + n_3) \left[1 - \frac{1}{q} \left(1 + \frac{n_3}{n_2} \right) \right] \quad (23)$$

and

$$\frac{qn_3}{p+1} \geq \frac{qn_3}{q \cdot \frac{n_3}{n_2+n_3} + 1} = \frac{n_2 + n_3}{1 + \frac{1}{q} \left(1 + \frac{n_2}{n_3} \right)} \geq (n_2 + n_3) \left[1 - \frac{1}{q} \left(1 + \frac{n_2}{n_3} \right) \right]. \quad (24)$$

Therefore it suffices to have

$$r \leq (n_2 + n_3) \left[1 - \frac{1}{q} \left(1 + \max \left\{ \frac{n_2}{n_3}, \frac{n_3}{n_2} \right\} \right) \right] - q.$$

3.3 Proof of Theorem 2.7

Suppose T has an alternative decomposition $T = \sum_{\ell=1}^r \tilde{a}^{(\ell)} \otimes \tilde{b}^{(\ell)} \otimes \tilde{c}^{(\ell)}$. Recalling the setup in Section 2.2.3,

$$M(T; p, q) = \sum_{\ell=1}^r A(a^{(\ell)}; p, q) \otimes (b^{(\ell)} c^{(\ell)\top}) = \sum_{\ell=1}^r A(\tilde{a}^{(\ell)}; p, q) \otimes (\tilde{b}^{(\ell)} \tilde{c}^{(\ell)\top}).$$

By assumption (vi), this matrix has rank exactly $r \binom{q-1}{p}$. By Lemma 3.1, the matrix $A(a; p, q)$ always has rank at most $\binom{q-1}{p}$, which means each term $A(a^{(\ell)}; p, q) \otimes (b^{(\ell)} c^{(\ell)\top})$ must have rank

exactly $\binom{q-1}{p}$, and in particular, $b^{(\ell)} \neq 0$ and $c^{(\ell)} \neq 0$ (i.e., for each ℓ , neither $b^{(\ell)}$ nor $c^{(\ell)}$ is the zero vector). Furthermore, every column of $A(a^{(\ell)}; p, q) \otimes (b^{(\ell)} c^{(\ell)\top})$ must lie in the column span of M , and so

$$\text{colspan}(M) = \text{span}\{z^{(U)}(a^{(\ell)}; p, q) \otimes b^{(\ell)} : U \subseteq [q], |U| = p + 1, \ell \in [r]\}, \quad (25)$$

where, recall, $z^{(U)}(a; p, q)$ denotes column U of $A(a; p, q)$. Similarly, we have the same conclusion for the alternative decomposition: $\tilde{b}^{(\ell)} \neq 0$ and $\tilde{c}^{(\ell)} \neq 0$ for each $\ell \in [r]$, and

$$\text{colspan}(M) = \text{span}\{z^{(U)}(\tilde{a}^{(\ell)}; p, q) \otimes \tilde{b}^{(\ell)} : U \subseteq [q], |U| = p + 1, \ell \in [r]\}. \quad (26)$$

Now recall that we define $V := [p + 1]$ and denote by $Z_{p,q} \subseteq \mathbb{R}^{\binom{q}{p}}$ the subspace consisting of vectors with the same sparsity pattern as $z^{(V)}$, that is, nonzeros are only allowed in the $p + 1$ entries S where $S \subseteq V$.

Lemma 3.2. *Under the assumptions of Theorem 2.7,*

$$\text{colspan}(M) \cap (Z_{p,q} \otimes \mathbb{R}^{n_2}) = \text{span}\{z^{(V)}(a^{(\ell)}; p, q) \otimes b^{(\ell)} : \ell \in [r]\}.$$

The proof of Lemma 3.2 is deferred to Section 3.4. The proof relies on assumption (viii).

Recall that $\varphi_{p,q}$, defined in (9), is an isomorphism $Z_{p,q} \rightarrow \mathbb{R}^{p+1}$ designed so that $\varphi_{p,q}(z^{(V)}(a; p, q)) = (a_1, \dots, a_{p+1})^\top$. From Lemma 3.2 we now conclude

$$(\varphi_{p,q} \otimes I)(\text{colspan}(M) \cap (Z_{p,q} \otimes \mathbb{R}^{n_2})) = \text{span}\{d^{(\ell)} \otimes b^{(\ell)} : \ell \in [r]\}, \quad (27)$$

where, recall, $d^{(\ell)} := (a_1^{(\ell)}, \dots, a_{p+1}^{(\ell)})^\top$.

Lemma 3.3. *Under the assumptions of Theorem 2.7, the vectors $\{d^{(\ell)} \otimes b^{(\ell)} : \ell \in [r]\}$ are the only rank-1 tensors, up to scalar multiple, in their span, $\text{span}\{d^{(\ell)} \otimes b^{(\ell)} : \ell \in [r]\}$. Furthermore, the algorithm of [JLV23, Corollary 3] recovers (scalar multiples of) these r rank-1 tensors given their span.*

The proof of Lemma 3.3 is based on [JLV23] and can be found in Section 3.5. The proof relies on assumptions (iii) and (x).

Define $\tilde{d}^{(\ell)} := (\tilde{a}_1^{(\ell)}, \dots, \tilde{a}_{p+1}^{(\ell)})^\top$. From (26) we have $z^{(V)}(\tilde{a}^{(\ell)}; p, q) \otimes \tilde{b}^{(\ell)} \in \text{colspan}(M) \cap (Z_{p,q} \otimes \mathbb{R}^{n_2})$ for all $\ell \in [r]$, and so $\tilde{d}^{(\ell)} \otimes \tilde{b}^{(\ell)} \in (\varphi_{p,q} \otimes I)(\text{colspan}(M) \cap (Z_{p,q} \otimes \mathbb{R}^{n_2}))$. But by (27) and Lemma 3.3, the only rank-1 tensors in this subspace are scalar multiples of $d^{(\ell)} \otimes b^{(\ell)}$. We therefore conclude that for every $\ell \in [r]$ there exists $\ell' \in [r]$ such that $\tilde{d}^{(\ell)} \otimes \tilde{b}^{(\ell)}$ is a scalar multiple of $d^{(\ell')} \otimes b^{(\ell')}$.

Next we need to repeat the argument with the “ b ” and “ c ” modes switched and with p replaced by $q - p - 1$, which will use the assumptions (iv), (vii), (ix), (xi). Due to symmetry, the argument is identical to before, and we only state the main claims. First, the flattening M' takes the form

$$M'(T; p, q) = \sum_{\ell=1}^r A(a^{(\ell)}; q - p - 1, q) \otimes (c^{(\ell)} b^{(\ell)\top}) = \sum_{\ell=1}^r A(\tilde{a}^{(\ell)}; q - p - 1, q) \otimes (\tilde{c}^{(\ell)} \tilde{b}^{(\ell)\top}),$$

and by Lemma 3.1, the maximum possible rank of $A(a; q - p - 1, q)$ is $\binom{q-1}{q-p-1} = \binom{q-1}{p}$. Analogous to (25) and (26), we have

$$\text{colspan}(M') = \text{span}\{z^{(U)}(a^{(\ell)}; q - p - 1, q) \otimes c^{(\ell)} : U \subseteq [q], |U| = q - p, \ell \in [r]\}$$

and

$$\text{colspan}(M') = \text{span}\{z^{(U)}(\tilde{a}^{(\ell)}; q-p-1, q) \otimes \tilde{c}^{(\ell)} : U \subseteq [q], |U| = q-p, \ell \in [r]\}.$$

Analogous to Lemma 3.2 and (27), we have, with $V' := [q-p]$,

$$\text{colspan}(M') \cap (Z_{q-p-1, q} \otimes \mathbb{R}^{n_3}) = \text{span}\{z^{(V')}(a^{(\ell)}; q-p-1, q) \otimes c^{(\ell)} : \ell \in [r]\}.$$

and

$$(\varphi_{q-p-1, q} \otimes I)(\text{colspan}(M') \cap (Z_{q-p-1, q} \otimes \mathbb{R}^{n_3})) = \text{span}\{f^{(\ell)} \otimes c^{(\ell)} : \ell \in [r]\},$$

where, recall, $f^{(\ell)} := (a_1^{(\ell)}, \dots, a_{q-p}^{(\ell)})^\top$. Analogous to Lemma 3.3, the vectors $\{f^{(\ell)} \otimes c^{(\ell)} : \ell \in [r]\}$ are the only rank-1 tensors in their span, up to scalar multiple, and the algorithm of [JLV23] recovers them from their span. Finally, defining $\tilde{f}^{(\ell)} := (\tilde{a}_1^{(\ell)}, \dots, \tilde{a}_{q-p}^{(\ell)})^\top$, we have $\tilde{f}^{(\ell)} \otimes \tilde{c}^{(\ell)} \in (\varphi_{q-p-1, q} \otimes I)(\text{colspan}(M') \cap (Z_{q-p-1, q} \otimes \mathbb{R}^{n_3}))$ and so every $\tilde{f}^{(\ell)} \otimes \tilde{c}^{(\ell)}$ is a scalar multiple of some $f^{(\ell'')} \otimes c^{(\ell')}$.

Let $R \subseteq [r]$ be the set of ℓ for which $\tilde{a}_1^{(\ell)} \neq 0$. For each $\ell \in [r]$ we have $\tilde{d}^{(\ell)} \neq 0$ and $\tilde{f}^{(\ell)} \neq 0$. Recall that $b^{(\ell)}, \tilde{b}^{(\ell)}, c^{(\ell)}, \tilde{c}^{(\ell)}, d^{(\ell)}, f^{(\ell)}$ are all nonzero (where $d^{(\ell)} \neq 0$ since $a_1^{(\ell)} \neq 0$ by assumption (i), and similarly for $f^{(\ell)}$). For each $\ell \in R$ we have established that $\tilde{d}^{(\ell)} \otimes \tilde{b}^{(\ell)}$ is a nonzero scalar multiple of some $d^{(\ell')} \otimes b^{(\ell')}$, and $\tilde{f}^{(\ell)} \otimes \tilde{c}^{(\ell)}$ is a nonzero scalar multiple of some $f^{(\ell'')} \otimes c^{(\ell')}$. In particular, $\tilde{d}^{(\ell)}$ is a nonzero scalar multiple of $d^{(\ell')}$, and $\tilde{f}^{(\ell)}$ is a nonzero scalar multiple of $f^{(\ell')}$. Since $\tilde{d}^{(\ell)}$ and $\tilde{f}^{(\ell)}$ agree on the first \bar{p} coordinates (and similarly for $d^{(\ell)}, f^{(\ell)}$), Assumption (ii) implies that $\ell' = \ell''$. This means we can define a function $\pi : R \rightarrow [r]$ such that $\tilde{b}^{(\ell)} \otimes \tilde{c}^{(\ell)} = \eta_\ell \cdot b^{(\pi(\ell))} \otimes c^{(\pi(\ell))}$ for a scalar $\eta_\ell \neq 0$, for each $\ell \in R$.

Let $T_1 := (T_{1jk} : j \in [n_2], k \in [n_3])$ denote the first $n_2 \times n_3$ matrix slice of T , namely

$$T_1 = \sum_{\ell=1}^r a_1^{(\ell)} \cdot b^{(\ell)} \otimes c^{(\ell)}.$$

Similarly for the alternative decomposition,

$$T_1 = \sum_{\ell=1}^r \tilde{a}_1^{(\ell)} \cdot \tilde{b}^{(\ell)} \otimes \tilde{c}^{(\ell)} = \sum_{\ell \in R} \tilde{a}_1^{(\ell)} \cdot \eta_\ell \cdot b^{(\pi(\ell))} \otimes c^{(\pi(\ell))},$$

since $\tilde{a}_1^{(\ell)} = 0$ for $\ell \notin R$. By assumptions (i) and (v), $a_1^{(\ell)} \neq 0$ for all $\ell \in [r]$, and the vectors $\{b^{(\ell)} \otimes c^{(\ell)} : \ell \in [r]\}$ are linearly independent. Comparing the previous expressions for T_1 , this means $R = [r]$, π is a permutation of $[r]$, and $\tilde{a}_1^{(\ell)} = \eta_\ell^{-1} \cdot a_1^{(\pi(\ell))}$. Now for each slice $T_i := (T_{ijk} : j \in [n_2], k \in [n_3])$ with $i \in [n_1]$, we can write

$$T_i = \sum_{\ell=1}^r a_i^{(\ell)} \cdot b^{(\ell)} \otimes c^{(\ell)} = \sum_{\ell=1}^r \tilde{a}_i^{(\ell)} \cdot \eta_\ell \cdot b^{(\pi(\ell))} \otimes c^{(\pi(\ell))},$$

and again using the linear independence of $\{b^{(\ell)} \otimes c^{(\ell)} : \ell \in [r]\}$ we must have $\tilde{a}_i^{(\ell)} = \eta_\ell^{-1} \cdot a_i^{(\pi(\ell))}$. This means $\tilde{a}^{(\ell)} = \eta_\ell^{-1} \cdot a^{(\pi(\ell))}$ and so the alternative decomposition has the same rank-1 terms as the original: $\tilde{a}^{(\ell)} \otimes \tilde{b}^{(\ell)} \otimes \tilde{c}^{(\ell)} = a^{(\pi(\ell))} \otimes b^{(\pi(\ell))} \otimes c^{(\pi(\ell))}$, with π a permutation of $[r]$. This completes the proof of uniqueness.

The above argument also implies success of Algorithm 1. From (27) we know the subspace B constructed by the algorithm is equal to $\text{span}\{d^{(\ell)} \otimes b^{(\ell)} : \ell \in [r]\}$, and similarly, $C = \text{span}\{f^{(\ell)} \otimes c^{(\ell)} : \ell \in [r]\}$. By Lemma 3.3, the rank-1 tensors $\hat{d}^{(\ell)} \otimes \hat{b}^{(\ell)}$ recovered by the [JLV23] algorithm

satisfy $\hat{d}^{(\ell)} = \alpha_\ell \cdot d^{(\pi_1(\ell))}$ and $\hat{b}^{(\ell)} = \beta_\ell \cdot b^{(\pi_1(\ell))}$ for scalars $\alpha_\ell \neq 0$, $\beta_\ell \neq 0$ and a permutation π_1 of $[r]$. Similarly, $\hat{f}^{(\ell)} = \gamma_\ell \cdot f^{(\pi_2(\ell))}$ and $\hat{c}^{(\ell)} = \delta_\ell \cdot c^{(\pi_2(\ell))}$ for scalars $\gamma_\ell \neq 0$, $\delta_\ell \neq 0$ and a permutation π_2 of $[r]$. Using Assumption (ii), the (unique) permutation τ found by Algorithm 1 is $\tau = \pi_2^{-1} \circ \pi_1$. The final step is to solve the linear system of equations (12) for $w^{(\ell)}$, which reduces to

$$\sum_{\ell=1}^r w_i^{(\ell)} \cdot (\beta_\ell \cdot b^{(\pi_1(\ell))}) \otimes (\delta_{\tau(\ell)} \cdot c^{(\pi_1(\ell))}) = T_i \quad \text{for } i \in [n_1].$$

One solution is $w^{(\ell)} = \beta_\ell^{-1} \cdot \delta_{\tau(\ell)}^{-1} \cdot a^{(\pi_1(\ell))}$, which recovers the desired decomposition. This solution is unique because $\{b^{(\ell)} \otimes c^{(\ell)} : \ell \in [r]\}$ are linearly independent by assumption (v).

3.4 Proof of Lemma 3.2

Using (25), the vectors $z^{(V)}(a^{(\ell)}; p, q) \otimes b^{(\ell)}$ for $\ell \in [r]$ lie in $\text{colspan}(M) \cap (Z_{p,q} \otimes \mathbb{R}^{n_2})$, so the inclusion “ \supseteq ” in Lemma 3.2 holds. For the reverse inclusion, assume on the contrary that “ \subseteq ” fails. By assumption (i) we have $a_1^{(\ell)} \neq 0$ for all $\ell \in [r]$, so using (25) and Lemma 3.1, the vectors $\{z^{(U)}(a^{(\ell)}; p, q) \otimes b^{(\ell)} : U \subseteq [q], |U| = p+1, 1 \in U, \ell \in [r]\}$ span $\text{colspan}(M)$; in fact, since the number of these vectors matches $\text{rank}(M)$, they must form a basis for $\text{colspan}(M)$. Define $\mathcal{U}_1 := \{U \subseteq [q] : |U| = p+1, 1 \in U, U \neq V\}$. Now, failure of “ \subseteq ” is equivalent to the existence of coefficients $\{\gamma_{U,\ell} \in \mathbb{R} : U \in \mathcal{U}_1, \ell \in [r]\}$, not all zero, such that

$$\sum_{U \in \mathcal{U}_1, \ell \in [r]} \gamma_{U,\ell} \cdot (z^{(U)}(a^{(\ell)}; p, q) \otimes b^{(\ell)}) \in Z_{p,q} \otimes \mathbb{R}^{n_2}. \quad (28)$$

Membership in $Z_{p,q} \otimes \mathbb{R}^{n_2}$ amounts to having a particular sparsity pattern, namely the entries (S, j) with $S \not\subseteq V = [p+1]$ must be zero. This means we can write (28) in matrix form as $N\gamma = 0$ where N is defined in (14). But N has full column rank by assumption (viii), a contradiction.

3.5 Proof of Lemma 3.3

The following arguments are based on [JLV23]. We will unpack some details of [JLV23] in order to extract an explicit condition under which their algorithm succeeds.

Define the set of rank-1 tensors,

$$\mathcal{X} := \{d \otimes b : d \in \mathbb{R}^{p+1}, b \in \mathbb{R}^{n_2}\},$$

and the subspace

$$\mathcal{Y} := \text{span}\{d^{(\ell)} \otimes b^{(\ell)} : \ell \in [r]\}.$$

For any $z \in \mathcal{X} \cap \mathcal{Y}$, the vector $w := z \otimes z$ must belong to the subspace $\text{span}\{y \otimes y : y \in \mathcal{Y}\}$. Since a basis for \mathcal{Y} is given by $y^{(1)}, \dots, y^{(r)}$ with $y^{(\ell)} := d^{(\ell)} \otimes b^{(\ell)}$ (which are linearly independent by assumption (iii)), a basis for the subspace $\text{span}\{y \otimes y : y \in \mathcal{Y}\}$ is given by $\{y^{(\ell_1)} \otimes y^{(\ell_2)} + y^{(\ell_2)} \otimes y^{(\ell_1)} : 1 \leq \ell_1 \leq \ell_2 \leq r\}$, so we can write

$$w = \sum_{1 \leq \ell_1 \leq \ell_2 \leq r} \delta_{\ell_1, \ell_2} (d^{(\ell_1)} \otimes b^{(\ell_1)} \otimes d^{(\ell_2)} \otimes b^{(\ell_2)} + d^{(\ell_2)} \otimes b^{(\ell_2)} \otimes d^{(\ell_1)} \otimes b^{(\ell_1)}), \quad (29)$$

for some coefficients $\delta_{\ell_1, \ell_2} \in \mathbb{R}$. Also, viewing z as a $(p+1) \times n_2$ matrix, since z has rank 1 we know every 2-by-2 minor must vanish:

$$z_{i_1, j_1} z_{i_2, j_2} - z_{i_1, j_2} z_{i_2, j_1} = 0 \quad \text{for } 1 \leq i_1 < i_2 \leq p+1 \text{ and } 1 \leq j_1 < j_2 \leq n_2.$$

Equivalently,

$$\langle w, E_{i_1, i_2; j_1, j_2} \rangle = 0 \quad \text{for } 1 \leq i_1 < i_2 \leq p+1 \text{ and } 1 \leq j_1 < j_2 \leq n_2, \quad (30)$$

where

$$E_{i_1, i_2; j_1, j_2} = e_{i_1} \otimes e_{j_1} \otimes e_{i_2} \otimes e_{j_2} - e_{i_1} \otimes e_{j_2} \otimes e_{i_2} \otimes e_{j_1}.$$

Now w belongs to the intersection of two subspaces $\mathcal{A} \cap \mathcal{B}$ where \mathcal{A} is the set of vectors expressible in the form (29) and \mathcal{B} is the set of solutions to (30). A key step in the algorithm of [JLV23] is to compute this intersection of subspaces (which in their notation is $S^2(\mathcal{Y}) \cap I_2^\perp$). The hope is to have no ‘‘spurious’’ vectors in this subspace, that is,

$$\mathcal{A} \cap \mathcal{B} = \text{span}\{d^{(\ell)} \otimes b^{(\ell)} \otimes d^{(\ell)} \otimes b^{(\ell)} : \ell \in [r]\}. \quad (31)$$

The containment ‘‘ \supseteq ’’ in (31) always holds, so failure of (31) is equivalent to the existence of coefficients δ_{ℓ_1, ℓ_2} , not all zero, but with $\delta_{\ell, \ell} = 0$, such that the vector (29) satisfies (30):

$$\sum_{1 \leq \ell_1 < \ell_2 \leq r} \delta_{\ell_1, \ell_2} \langle d^{(\ell_1)} \otimes b^{(\ell_1)} \otimes d^{(\ell_2)} \otimes b^{(\ell_2)} + d^{(\ell_2)} \otimes b^{(\ell_2)} \otimes d^{(\ell_1)} \otimes b^{(\ell_1)}, E_{i_1, i_2; j_1, j_2} \rangle = 0$$

for $1 \leq i_1 < i_2 \leq p+1$ and $1 \leq j_1 < j_2 \leq n_2$. This reduces to

$$\sum_{1 \leq \ell_1 < \ell_2 \leq r} \delta_{\ell_1, \ell_2} \left(d_{i_1}^{(\ell_1)} b_{j_1}^{(\ell_1)} d_{i_2}^{(\ell_2)} b_{j_2}^{(\ell_2)} + d_{i_1}^{(\ell_2)} b_{j_1}^{(\ell_2)} d_{i_2}^{(\ell_1)} b_{j_2}^{(\ell_1)} - d_{i_1}^{(\ell_1)} b_{j_2}^{(\ell_1)} d_{i_2}^{(\ell_2)} b_{j_1}^{(\ell_2)} - d_{i_1}^{(\ell_2)} b_{j_2}^{(\ell_2)} d_{i_2}^{(\ell_1)} b_{j_1}^{(\ell_1)} \right) = 0,$$

which can be written in matrix form as $P\delta = 0$ with P defined in (15). Since P has full column rank by assumption (x), there is no nonzero solution to $P\delta = 0$, so (31) holds.

We have now shown that for any $z \in \mathcal{X} \cap \mathcal{Y}$, the corresponding $w := z \otimes z$ lies in the subspace (31). Our goal is to show that z is a scalar multiple of some $d^{(\ell)} \otimes b^{(\ell)}$. Since the vectors $d^{(\ell)} \otimes b^{(\ell)}$ are linearly independent by assumption (iii), Lemma 3.4 below implies that the vectors $d^{(\ell)} \otimes b^{(\ell)} \otimes d^{(\ell)} \otimes b^{(\ell)}$ are the only vectors of the form $y \otimes y$ in their span, up to scalar multiple. This means w is a scalar multiple of some $d^{(\ell)} \otimes b^{(\ell)} \otimes d^{(\ell)} \otimes b^{(\ell)}$, and thus z is a scalar multiple of some $d^{(\ell)} \otimes b^{(\ell)}$, as desired.

Lemma 3.4. *If $z^{(1)}, \dots, z^{(r)} \in \mathbb{R}^N$ are linearly independent vectors then $z^{(1)} \otimes z^{(1)}, \dots, z^{(r)} \otimes z^{(r)}$ are the only rank-1 tensors in their span, $\text{span}\{z^{(\ell)} \otimes z^{(\ell)} : \ell \in [r]\}$, up to scalar multiple, and furthermore there is an algorithm based on simultaneous diagonalization to recover (scalar multiples of) the vectors $z^{(\ell)} \otimes z^{(\ell)}$ given their span.*

Proof. First we give a simple self-contained proof of the uniqueness claim. View each $z^{(\ell)} \otimes z^{(\ell)}$ as a rank-1 matrix, and apply a change of basis (to both rows and columns) so that $z^{(1)}, \dots, z^{(r)}$ become standard basis vectors $e^{(1)}, \dots, e^{(r)}$. Now each $z^{(\ell)} \otimes z^{(\ell)}$ becomes a diagonal matrix with a single nonzero entry at position (ℓ, ℓ) . Any linear combination $\sum_{\ell=1}^r \alpha_\ell \cdot z^{(\ell)} \otimes z^{(\ell)}$ is diagonal and has rank equal to the number of nonzero coefficients α_ℓ (and in particular, the vectors $z^{(\ell)} \otimes z^{(\ell)}$ are linearly independent). Therefore the only rank-1 matrices in this span are scalar multiples of a single $z^{(\ell)} \otimes z^{(\ell)}$.

There is also an algorithmic proof using simultaneous diagonalization, which appears in [JLV23]. Given an arbitrary basis $p^{(1)}, \dots, p^{(r)}$ for $\text{span}\{z^{(\ell)} \otimes z^{(\ell)} : \ell \in [r]\}$, construct the third-order tensor $R := \sum_{\ell=1}^r e^{(\ell)} \otimes p^{(\ell)} \in \mathbb{R}^r \otimes \mathbb{R}^N \otimes \mathbb{R}^N$ where $e^{(\ell)}$ denotes the ℓ th standard basis vector. There are

coefficients $\beta_{\ell m}$ such that $p^{(\ell)} = \sum_{m=1}^r \beta_{\ell m} \cdot z^{(m)} \otimes z^{(m)}$ for each $\ell \in [r]$, where $\beta = (\beta_{\ell m} : \ell, m \in [r])$ is an invertible matrix. This means R admits the decomposition

$$R = \sum_{\ell=1}^r e^{(\ell)} \otimes \left(\sum_{m=1}^r \beta_{\ell m} \cdot z^{(m)} \otimes z^{(m)} \right) = \sum_{m=1}^r \left(\sum_{\ell=1}^r \beta_{\ell m} \cdot e^{(\ell)} \right) \otimes z^{(m)} \otimes z^{(m)}. \quad (32)$$

The vectors $\{z^{(m)} : m \in [r]\}$ are linearly independent by assumption, and the vectors $\{\sum_{\ell=1}^r \beta_{\ell m} \cdot e^{(\ell)} : m \in [r]\}$ are linearly independent because they are the columns of the invertible matrix β . This verifies the conditions for simultaneous diagonalization (see [Moi18, Theorem 3.1.3]) so the decomposition on the right-hand side of (32) is unique and can be recovered by an efficient algorithm. This allows recovery of the vectors $z^{(\ell)} \otimes z^{(\ell)}$ (up to re-ordering and scalar multiple). \square

Finally, the algorithmic claim of Lemma 3.3 also follows from the above, since the [JLV23] algorithm first computes the intersection of subspaces $\mathcal{A} \cap \mathcal{B}$, which is equal to (31), and then applies simultaneous diagonalization as in Lemma 3.4 to recover the vectors $d^{(\ell)} \otimes b^{(\ell)}$ from this subspace.

3.6 Proof of Theorem 2.8

First we verify the claim $\bar{p} \geq 2$. To show $p \geq 1$, it suffices by the choice of p to have $qn_3/(n_2+n_3) \geq 1$. This follows by combining two facts: $q > 1 + \alpha$, which is a consequence of the assumption on q , and $n_3/(n_2 + n_3) \geq 1/(1 + \alpha)$, which is a consequence of the definition of α . Similarly, to show $q - p \geq 2$, it suffices by the choice of p to have $qn_3/(n_2 + n_3) < q - 1$, i.e., $qn_2/(n_2 + n_3) > 1$, which again follows from $q > 1 + \alpha$ and $n_2/(n_2 + n_3) \geq 1/(1 + \alpha)$.

Now conditions (i) and (ii) are clearly met for generic components, since $\bar{p} \geq 2$. The following lemma establishes (iii)–(v), provided $r \leq \min\{(p + 1)n_2, (q - p)n_3, n_2n_3\}$.

Lemma 3.5. *Suppose $x^{(1)} \otimes y^{(1)}, \dots, x^{(r)} \otimes y^{(r)} \in \mathbb{R}^m \otimes \mathbb{R}^n$ with the components $x^{(\ell)}, y^{(\ell)}$ generically chosen. If $r \leq mn$ then $\{x^{(\ell)} \otimes y^{(\ell)} : \ell \in [r]\}$ are linearly independent.*

Proof. Form the $mn \times r$ matrix whose columns are $x^{(\ell)} \otimes y^{(\ell)}$. It suffices to identify a square submatrix using all the columns whose determinant is a nonzero polynomial in the entries of $x^{(\ell)}, y^{(\ell)}$. For this, it suffices to demonstrate a choice of the vectors $x^{(\ell)}, y^{(\ell)}$ for which $\{x^{(\ell)} \otimes y^{(\ell)} : \ell \in [r]\}$ are linearly independent. Fix an injection $[r] \rightarrow [m] \times [n]$. For each $\ell \in [r]$, if $\ell \mapsto (i, j)$ then set $x^{(\ell)} = e^{(i)}$ and $y^{(\ell)} = e^{(j)}$ where $e^{(i)}$ is the i th standard unit basis vector. \square

Condition (vi) on M follows from the proof of Theorem 2.4, namely (22), provided

$$r \leq q \left(\min \left\{ \frac{n_2}{q-p}, \frac{n_3}{p+1} \right\} - 1 \right).$$

This condition is unchanged under the operation that swaps n_2 with n_3 and replaces p with $q - p - 1$. Therefore, condition (vii) on M' also holds under the same condition.

The following lemma establishes condition (viii) on N .

Lemma 3.6. *In the setting of Theorem 2.7 with parameters $q \leq n_1$ and $\bar{p} := \min\{p + 1, q - p\} \geq 2$, if the components $a^{(\ell)}, b^{(\ell)}, c^{(\ell)}$ are generically chosen and*

$$r \leq \frac{qn_2}{q-p} \left(1 - \frac{2}{q} \right) - \frac{q^3}{4},$$

then N has full column rank.

The proof is deferred to Section 3.7. By symmetry, condition (ix) on N' also holds under the analogous condition

$$r \leq \frac{qn_3}{p+1} \left(1 - \frac{2}{q}\right) - \frac{q^3}{4}.$$

For condition (x) on P we appeal to the results of [JLV23]. As explained in Section 3.5, condition (x) is equivalent to (31), which is established by [JLV23, Corollary 3] for generic components under the condition $r \leq p(n_2 - 1)/4$. Similarly, condition (xi) on P' holds provided $r \leq (q - p - 1)(n_3 - 1)/4$.

To summarize, we have the conditions

- (a) $r \leq \min\{(p+1)n_2, (q-p)n_3\}$,
- (b) $r \leq n_2n_3$,
- (c) $r \leq \min\left\{\frac{qn_2}{q-p}, \frac{qn_3}{p+1}\right\} - q$,
- (d) $r \leq \min\left\{\frac{qn_2}{q-p}, \frac{qn_3}{p+1}\right\} \left(1 - \frac{2}{q}\right) - \frac{q^3}{4}$,
- (e) $r \leq \frac{1}{4} \min\{p(n_2 - 1), (q - p - 1)(n_3 - 1)\}$.

Condition (a) is subsumed by (e) because, e.g., $p(n_2 - 1)/4 \leq pn_2 \leq (p+1)n_2$. The assumption $\bar{p} \geq 2$ implies $p \geq 1$ and $q \geq p+2 \geq 3$, so $q^3/4 > q$. This means condition (c) is subsumed by (d). Also, (e) implies $n_2, n_3 \geq 2$ (or else $r \leq 0$), so (d) implies

$$r \leq \min\left\{\frac{qn_2}{q-p}, \frac{qn_3}{p}\right\} \leq \max_{x \in (0,1)} \min\left\{\frac{n_2}{1-x}, \frac{n_3}{x}\right\} = n_2 + n_3 \leq 2 \max\{n_2, n_3\} \leq n_2n_3,$$

so (b) is also subsumed. This leaves only (d) and (e).

Now using the choice of $p = \lfloor q \cdot n_3 / (n_2 + n_3) \rfloor$ and the bounds (23),(24) we have

$$\min\left\{\frac{qn_2}{q-p}, \frac{qn_3}{p+1}\right\} \left(1 - \frac{2}{q}\right) \geq (n_2 + n_3) \left(1 - \frac{1+\alpha}{q}\right) \left(1 - \frac{2}{q}\right) \geq (n_2 + n_3) \left(1 - \frac{3+\alpha}{q}\right),$$

so condition (d) can be replaced with the sufficient condition

$$r \leq (n_2 + n_3) \left(1 - \frac{3+\alpha}{q}\right) - \frac{q^3}{4}. \quad (33)$$

Finally, we will show that the previous condition subsumes (e), given the assumption on q . Using the choice of p ,

$$\frac{1}{4}p(n_2 - 1) \geq \frac{1}{4} \left(q \cdot \frac{n_3}{n_2 + n_3} - 1\right) (n_2 - 1) = \frac{(q-1)n_3 - n_2}{4(n_2 + n_3)} (n_2 - 1) \geq \frac{(q-1-\alpha)n_3}{4(n_2 + n_3)} (n_2 - 1)$$

and similarly,

$$\frac{1}{4}(q-p-1)(n_3-1) \geq \frac{1}{4} \left(q - q \cdot \frac{n_3}{n_2 + n_3} - 1\right) (n_3 - 1) \geq \frac{(q-1-\alpha)n_2}{4(n_2 + n_3)} (n_3 - 1).$$

Therefore, the right-hand side of condition (e) can be replaced by

$$\frac{q-1-\alpha}{4(n_2+n_3)} (n_2-1)(n_3-1) \geq \frac{q-1-\alpha}{4} \left(\frac{n_2n_3}{n_2+n_3} - 1\right) = \frac{q-1-\alpha}{4} \left[\frac{\alpha}{(1+\alpha)^2} (n_2+n_3) - 1\right]$$

where the last step follows by checking the cases $n_3 = \alpha n_2$ and $n_2 = \alpha n_3$

$$\geq \frac{(q-1-\alpha)\alpha}{4(1+\alpha)^2}(n_2+n_3) - \frac{q}{4} \geq n_2+n_3 - \frac{q}{4}$$

using the assumption $q \geq (4+5\alpha)(1+1/\alpha)$, so condition (e) is subsumed by (33).

3.7 Proof of Lemma 3.6

It will be convenient to work not with N but with an equivalent matrix, L . Recall the derivation of N from Section 3.4. There, we used the columns $\{U : 1 \in U\}$ as a basis for the column span of $A(a^{(\ell)})$. By Lemma 3.1, and assuming $a_i^{(\ell)} \neq 0$ due to generic components, we can just as easily use the basis of columns $\{U : i \in U\}$ for any $i \in V := [p+1]$. To define L , we will use a different choice of $i \in [p+1]$ for each ℓ . The purpose of this is to create symmetry among the elements $i \in [p+1]$, so that element 1 no longer plays a distinguished role. (However, there remains a distinction between $i \in [p+1]$ versus $i \in [q] \setminus [p+1]$, since we have fixed $V = [p+1]$.)

Formally, choose the smallest integer \bar{r} for which $r \leq \bar{r}q(p+1)(q-p-1)$. It suffices to verify the extreme case $r = \bar{r}q(p+1)(q-p-1)$ (i.e., we will show that even when r is increased to add additional columns, all the columns are linearly independent). We will index $[r]$ by tuples (i, j, ℓ) with $i \in [p+1]$, $j \in [q] \setminus [p+1]$, and $\ell \in [\bar{r}q]$; this way, each element of $[r]$ has an assigned (i, j) pair, where i plays a role in choosing the basis described above, and the role of j will appear later. Define the matrix L with rows indexed by

$$\{(m, S) : m \in [n_2], S \subseteq [q], |S| = p, S \not\subseteq [p+1]\}$$

and columns indexed by

$$\{(i, j, \ell, U) : i \in [p+1], j \in [q] \setminus [p+1], \ell \in [\bar{r}q], U \subseteq [q], |U| = p+1, i \in U, U \neq [p+1]\}, \quad (34)$$

with entries

$$L_{mS, ij\ell U} = b_m^{(i,j,\ell)} \cdot \sum_{k=1}^q \mathbb{1}_{U=S \sqcup \{k\}} \cdot \sigma(U, k) \cdot a_k^{(i,j,\ell)},$$

where we have indexed $[r]$ by tuples (i, j, ℓ) as described above. Our goal is to show that L has full column rank, which implies the same for N , since both these conditions are equivalent to (8) via the argument in Section 3.4 (along with the discussion above).

It suffices to identify a nonzero minor of L that uses all the columns. To this end, we will describe a pairing that assigns every column (i, j, ℓ, U) to some row (m, S) in such a way that each row is used 0 or 1 times. Imagine the columns grouped into $\bar{r}q$ ‘‘epochs,’’ one for each ℓ value, and each epoch is further subdivided into $(p+1)(q-p-1)$ ‘‘pillars,’’ one for each (i, j) pair. The ordering of the epochs is unimportant, but let’s fix one arbitrary ordering, and similarly for the ordering of pillars within an epoch and the ordering of columns within a pillar. Call the first $\bar{r}(p+1)$ epochs ‘‘type 1’’ and the remaining $\bar{r}(q-p-1)$ epochs ‘‘type 2.’’ We first describe which S value gets assigned to a given column (i, j, ℓ, U) :

- For a type 1 epoch (i.e., $1 \leq \ell \leq \bar{r}(p+1)$), remove i from U to produce S : let $S = U \setminus \{i\}$. Note that $i \in U$ is guaranteed by (34).
- For a type 2 epoch (i.e., $\bar{r}(p+1) < \ell \leq \bar{r}q$), remove j if possible, and otherwise remove i . That is, if $j \in U$ and $|U \cap [p+1]| < p$ then let $S = U \setminus \{j\}$, and otherwise let $S = U \setminus \{i\}$. (Recall that $S \subseteq [p+1]$ is forbidden, so j cannot be removed if $|U \cap [p+1]| = p$.)

Now that the S values are assigned, the m values are assigned based on the “first available” rule. That is, column (i, j, ℓ, U) gets assigned to row (m, S) where S is specified above and m is the minimum m' for which (m', S) has not already been assigned to some previous column. Consider the submatrix L' of L formed by all columns of L and only those rows that were assigned to some column by the above procedure. Our goal is to show that $\det(L')$ is a nonzero polynomial in the variables $b_m^{(i,j,\ell)}$ and $a_k^{(i,j,\ell)}$.

Expand the determinant $\det(L')$ as a sum over permutations. Each term in the sum gives a single monomial (multiplied by a coefficient ± 1 or 0). Consider the term in the sum corresponding to the pairing of rows and columns described above. We claim that this term produces a nonzero monomial that is unique in the sense that no other term produces a (nonzero) scalar multiple of this monomial. Once we establish this claim, we are done. To prove the claim we need to argue that, given our monomial of interest, it is possible to deduce the unique pairing of rows and columns that produced it. Within a fixed pillar (i, j, ℓ) , the powers of $a_k^{(i,j,\ell)}$ in the monomial reveal which multiset of k values are removed from the U 's in this pillar to form the associated S 's. By design, there are only 1 or 2 distinct k values, namely $\{i\}$ or $\{i, j\}$ (for a type 1 or type 2 epoch, respectively). In either case, the specific pairing between U 's and S 's can be deduced (using the fact that a type 2 epoch removes j whenever possible). The S 's appearing within a pillar are all distinct (when j is removed, the resulting S contains i , but when i is removed, it does not). Now the powers of $b_m^{(i,j,\ell)}$ in the monomial reveal which multiset of m values are used in each pillar. Recalling the “first available” rule, this allows the m value for each column to be deduced. This completes the proof that $\det(L')$ is a nonzero polynomial.

Finally we need to determine what value of n_2 is required for the above construction. The minimum allowable value for n_2 is the largest m value used in the construction. For each S , we need to count how many columns (i, j, ℓ, U) get assigned to S .

We first consider the special case where $|S \cap [p+1]|$ takes its maximum possible value, $p-1$. This case needs separate consideration due to the special rule for type 2 epochs: if $|U \cap [p+1]| = p$ then j cannot be removed. Fix S with $|S \cap [p+1]| = p-1$. In a type 1 epoch ℓ , the number of tuples (i, j, U) assigned to S is $2(q-p-1)$ because there are 2 choices for $i \notin S$, then $q-p-1$ choices for j , and then $U = S \sqcup \{i\}$. In a type 2 epoch, we have two different contributions to count. First consider U with $|U \cap [p+1]| = p$, and recall that these are never allowed to remove j . This gives $2(q-p-1)$ tuples (i, j, U) , since there are again 2 choices for $i \notin S$, then $q-p-1$ choices for j , and $U = S \sqcup \{i\}$. We also need to consider U with $|U \cap [p+1]| = p-1$ (i.e., cases where j is removed). This gives $(p-1)(q-p-2)$ tuples (i, j, U) , since there are $p-1$ choices for $i \in S$, then $q-p-2$ choices for $j \notin S$, and $U = S \sqcup \{j\}$ (we need $i \in S$ so that $i \in U$, as required by (34)). In total, the number of times that S is used is

$$\begin{aligned} & \bar{r}(p+1) \cdot 2(q-p-1) + \bar{r}(q-p-1) \cdot [2(q-p-1) + (p-1)(q-p-2)] \\ &= \bar{r}(q-p-1)[2q + (p-1)(q-p-2)] \\ &= \bar{r}(p+1)(q-p-1) \left(q-p + \frac{2}{p+1} \right). \end{aligned}$$

Now consider the remaining case: fix S with $h := |S \cap [p+1]| < p-1$. In a type 1 epoch ℓ , the number of tuples (i, j, U) assigned to S is $(p+1-h)(q-p-1)$ because there are $p+1-h$ choices for $i \notin S$, then $q-p-1$ choices for j , and $U = S \sqcup \{i\}$. In a type 2 epoch, we again have two contributions to count. First consider U with $|U \cap [p+1]| = h+1$ (i.e., i is removed). This gives $(p+1-h)(q-2p+h-1)$ tuples (i, j, U) , since there are $p+1-h$ choices for $i \notin S$, then $q-2p+h-1$ choices for $j \notin S$, and $U = S \sqcup \{i\}$. Now consider U with $|U \cap [p+1]| = h$ (i.e., j

is removed). This gives $h(q - 2p + h - 1)$ tuples (i, j, U) , since there are h choices for $i \in S$, then $q - 2p + h - 1$ choices for $j \notin S$, and $U = S \sqcup \{j\}$. In total, the number of times that S is used is

$$\begin{aligned} & \bar{r}(p+1) \cdot (p+1-h)(q-p-1) + \bar{r}(q-p-1) \cdot [(p+1-h)(q-2p+h-1) + h(q-2p+h-1)] \\ &= \bar{r}(q-p-1)[(p+1)(p+1-h) + (p+1)(q-2p+h-1)] \\ &= \bar{r}(p+1)(q-p-1)(q-p), \end{aligned}$$

which is dominated by the first case.

To recap, we need an integer \bar{r} such that $r \leq \bar{r}q(p+1)(q-p-1)$ and $n_2 \geq \bar{r}(p+1)(q-p-1)(q-p+2/(p+1))$, or equivalently,

$$\frac{r}{q(p+1)(q-p-1)} \leq \bar{r} \leq \frac{n_2}{(p+1)(q-p-1)(q-p+2/(p+1))}.$$

It suffices to have

$$\frac{r}{q(p+1)(q-p-1)} \leq \frac{n_2}{(p+1)(q-p-1)(q-p+2/(p+1))} - 1,$$

that is,

$$r \leq \frac{qn_2}{q-p+2/(p+1)} - q(p+1)(q-p-1).$$

We have the bounds $(p+1)(q-p-1) \leq \max_{x \in \mathbb{R}} x(q-x) = q^2/4$, and

$$\frac{qn_2}{q-p+2/(p+1)} \geq \frac{qn_2}{q-p} \left(1 - \frac{2}{(p+1)(q-p)}\right) \geq \frac{qn_2}{q-p} \left(1 - \frac{2}{q}\right).$$

Thus it suffices to have

$$r \leq \frac{qn_2}{q-p} \left(1 - \frac{2}{q}\right) - \frac{q^3}{4}$$

as claimed.

4 Proofs for Lower Bounds

4.1 Matrix Decomposition Lemma

The following result of [EGOW18] is a core component of our lower bounds. (In [EGOW18], only square matrices are considered, but the same proof applies for rectangular matrices.) Recall the notion of symbolic rank from Section 2.1.2.

Lemma 4.1 ([EGOW18], Lemma 3.2). *Let $M(x) \in \mathbb{R}[x]^{m_1 \times m_2}$ be an $m_1 \times m_2$ matrix whose entries are homogeneous degree- d polynomials in the variables x_1, \dots, x_p . If $M(x)$ has symbolic rank r then there are vectors $f_1(x), \dots, f_r(x) \in \mathbb{R}[x]^{m_1}$ and $g_1(x), \dots, g_r(x) \in \mathbb{R}[x]^{m_2}$ such that*

$$M(x) = \sum_{t=1}^r H_d[f_t(x)g_t(x)^\top],$$

where $H_d(\cdot)$ denotes the degree- d homogeneous part (applied entry-wise).

4.2 Proof of Theorem 2.9

Let s be the symbolic rank of $A(a)$, which is also the generic rank of $M(a^{(1)} \otimes b^{(1)} \otimes c^{(1)})$. The entries of $A(a)$ are degree-1 homogeneous polynomials in a , so by Lemma 4.1 there are vectors of polynomials $f_1(a), \dots, f_s$ and $g_1(a), \dots, g_s(a)$ such that

$$A(a) = \sum_{t=1}^s H_1[f_t(a)g_t(a)^\top].$$

Note that

$$H_1[f_t(a)g_t(a)^\top] = H_0[f_t]H_1[g_t(a)]^\top + H_1[f_t(a)]H_0[g_t]^\top,$$

where we write, e.g., $H_0[f_t]$ instead of $H_0[f_t(a)]$ to emphasize that this is a constant (degree-0) vector. Now we can write

$$M(a \otimes b \otimes c) = A(a) \otimes (bc^\top) = \sum_{t=1}^s \{(H_0[f_t] \otimes b)(H_1[g_t(a)] \otimes c)^\top + (H_1[f_t(a)] \otimes b)(H_0[g_t] \otimes c)^\top\}$$

and so, since $T \mapsto M(T)$ is linear,

$$M\left(\sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)}\right) = M_0 + M_1$$

where

$$M_0 := \sum_{t=1}^s \sum_{\ell=1}^r (H_0[f_t] \otimes b^{(\ell)})(H_1[g_t(a^{(\ell)})] \otimes c^{(\ell)})^\top$$

and

$$M_1 := \sum_{t=1}^s \sum_{\ell=1}^r (H_1[f_t(a^{(\ell)})] \otimes b^{(\ell)})(H_0[g_t] \otimes c^{(\ell)})^\top.$$

Note that

$$\text{colspan}(M_0) \subseteq \text{span}\{H_0[f_t] \otimes \mathbb{R}^{n_2} : t \in [s]\}$$

and

$$\text{rowspan}(M_1) \subseteq \text{span}\{H_0[g_t] \otimes \mathbb{R}^{n_3} : t \in [s]\},$$

so for any setting of the variables $a^{(\ell)}, b^{(\ell)}, c^{(\ell)}$ we have

$$\text{rank}\left(M\left(\sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)}\right)\right) \leq \text{rank}(M_0) + \text{rank}(M_1) \leq sn_2 + sn_3.$$

Additivity would imply that this rank must also equal rs , from which we conclude $r \leq n_2 + n_3$.

4.3 Proof of Theorem 2.10

Let s be the generic rank of $M(a^{(1)} \otimes b^{(1)} \otimes c^{(1)})$, which is also the symbolic rank of the matrix $M(a \otimes b \otimes c)$, viewed as a matrix of polynomials in the variables $x = (a_1, \dots, a_{n_1}, b_1, \dots, b_{n_2}, c_1, \dots, c_{n_3})$. The entries of $M(a \otimes b \otimes c)$ are degree-3 homogeneous polynomials, so by Lemma 4.1 there are vectors of polynomials $f_1(x), \dots, f_s(x)$ and $g_1(x), \dots, g_s(x)$ such that

$$M(a \otimes b \otimes c) = \sum_{t=1}^s H_3[f_t(x)g_t(x)^\top].$$

The variables x are naturally partitioned into $X_1 := \{a_i : i \in [n_1]\}$, $X_2 := \{b_i : i \in [n_2]\}$, and $X_3 := \{c_i : i \in [n_3]\}$. For $S \subseteq [3]$, let $H_S(\cdot)$ denote the terms that are degree 1 in X_i for each $i \in S$ and degree 0 in X_i for each $i \notin S$. Note that all terms in $M(a \otimes b \otimes c)$ are degree 1 in each of X_1, X_2, X_3 , so we can rewrite the above as

$$M(a \otimes b \otimes c) = \sum_{t=1}^s H_{\{1,2,3\}}[f_t(x)g_t(x)^\top].$$

We can write

$$H_{\{1,2,3\}}[f_t(x)g_t(x)^\top] = \sum_{S \subseteq [3]} H_S[f_t(x)]H_{[3] \setminus S}[g_t(x)^\top],$$

and so

$$M \left(\sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)} \right) = \sum_{S \subseteq [3]} M_S$$

where

$$M_S := \sum_{t=1}^s \sum_{\ell=1}^r H_S[f_t(x^{(\ell)})]H_{[3] \setminus S}[g_t(x^{(\ell)})^\top].$$

Now, for any setting of the variables $x^{(\ell)}$, we bound the rank of each M_S . First, $\text{rank}(M_\emptyset) \leq s$ because $\text{colspan}(M_\emptyset) \subseteq \text{span}\{H_\emptyset[f_t] : t \in [s]\}$, noting that $H_\emptyset[f_t(x)]$ is a constant (degree-0) vector, which we denote $H_\emptyset[f_t]$. Next, note that $H_{\{1\}}[f_t(x)]$ has only “ a ” variables (not “ b ” or “ c ”), which we denote by $H_{\{1\}}[f_t(x^{(\ell)})] = H_{\{1\}}[f_t(a = a^{(\ell)})]$. Each $H_{\{1\}}[f_t(a = a^{(\ell)})]$ lies in $\text{span}\{H_{\{1\}}[f_t(a = e^{(j)})] : j \in [n_1]\}$ where $e^{(j)}$ are unit basis vectors. This means $\text{rank}(M_{\{1\}}) \leq sn_1$ because $\text{colspan}(M_{\{1\}}) \subseteq \text{span}\{H_{\{1\}}[f_t(a = e^{(j)})] : t \in [s], j \in [n_1]\}$. Similarly, $\text{rank}(M_{\{2\}}) \leq sn_2$ and $\text{rank}(M_{\{3\}}) \leq sn_3$. Applying the same argument to the row span, we also have $\text{rank}(M_{\{1,2,3\}}) \leq s$, $\text{rank}(M_{\{2,3\}}) \leq sn_1$, $\text{rank}(M_{\{1,3\}}) \leq sn_2$, and $\text{rank}(M_{\{1,2\}}) \leq sn_3$. Putting it together,

$$\text{rank} \left(M \left(\sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)} \right) \right) \leq \sum_{S \subseteq [3]} \text{rank}(M_S) \leq 2s(n_1 + n_2 + n_3 + 1).$$

Additivity would imply that this rank is rs , from which we conclude $r \leq 2(n_1 + n_2 + n_3 + 1)$.

Remark 4.2. *The proof above is similar to that of Theorem 4.4 of [EGOW18], where they give a bound of $r \leq 8n$ for $n \times n \times n$ tensors. We now explain the source of our improvement from $8n$ to $6n + 2$. When decomposing M into 8 terms M_S for $S \subseteq [3]$, the original proof of [EGOW18] bounds the rank of each term by sn . We instead keep track of the specific structure of each term and see that two of them, namely M_\emptyset and $M_{\{1,2,3\}}$, admit a better bound of s .*

4.4 Proof of Theorem 2.12

Let s be the symbolic rank of

$$\tilde{M}(b^{(1,1)} \otimes \dots \otimes b^{(k,1)}, \dots, b^{(1,d)} \otimes \dots \otimes b^{(k,d)}) \tag{35}$$

in the variables $b = \{b^{(i,j)} : i \in [k], j \in [d]\}$; this is also the generic rank of $\tilde{M}(T^{(1)}, \dots, T^{(d)})$. The entries of (35) are degree- kd homogeneous polynomials in the b variables, so by Lemma 4.1, (35) admits the decomposition

$$\sum_{t=1}^s H_{kd}[f_t(b)g_t(b)^\top].$$

The variables b are naturally partitioned into $\bigsqcup_{i \in [k], j \in [d]} B_{ij}$ where $B_{ij} := \{b_m^{(i,j)} : m \in [n_i]\}$. For $S \subseteq [k] \times [d]$, let $H_S(\cdot)$ denote the terms that are degree 1 in B_{ij} for each $(i, j) \in S$ and degree 0 in B_{ij} for each $(i, j) \notin S$. Note that all terms in (35) are degree 1 in every B_{ij} , so we can rewrite the above decomposition as

$$\sum_{t=1}^s H_{[k] \times [d]}[f_t(b)g_t(b)^\top].$$

With $\bar{S} := ([k] \times [d]) \setminus S$, we can write

$$H_{[k] \times [d]}[f_t(b)g_t(b)^\top] = \sum_{S \subseteq [k] \times [d]} H_S[f_t(b)]H_{\bar{S}}[g_t(b)]^\top,$$

and so

$$\sum_{1 \leq \ell_1 < \dots < \ell_d \leq r} \tilde{M}(T^{(\ell_1)}, \dots, T^{(\ell_d)}) = \sum_{S \subseteq [k] \times [d]} M_S$$

where

$$M_S := \sum_{t=1}^s \sum_{1 \leq \ell_1 < \dots < \ell_d \leq r} H_S[f_t(b^{(i,j)} = a^{(i,\ell_j)})]H_{\bar{S}}[g_t(b^{(i,j)} = a^{(i,\ell_j)})]^\top.$$

Here, $H_S[f_t(b^{(i,j)} = a^{(i,\ell_j)})]$ is obtained from $H_S[f_t(b)]$ by making the substitution $b^{(i,j)} = a^{(i,\ell_j)}$ for each $(i, j) \in S$; note that the variables $b^{(i,j)}$ with $(i, j) \notin S$ do not appear. Similarly, $H_{\bar{S}}[g_t(b^{(i,j)} = a^{(i,\ell_j)})]$ is obtained from $H_{\bar{S}}[g_t(b)]$ by making the substitution $b^{(i,j)} = a^{(i,\ell_j)}$ for each $(i, j) \notin S$. Note that $H_S[f_t(b)]$ is a multilinear function of the d tensors $\{\otimes_{i:(i,j) \in S} b^{(i,j)} : j \in [d]\}$. After making the substitution $b^{(i,j)} = a^{(i,\ell_j)}$, we have

$$\otimes_{i:(i,j) \in S} b^{(i,j)} \subseteq \otimes_{i:(i,j) \in S} \mathbb{R}^{n_i},$$

where the right-hand side is a subspace of dimension $\prod_{i:(i,j) \in S} n_i$. We also have

$$\otimes_{i:(i,j) \in S} b^{(i,j)} \subseteq \text{span} \left\{ \otimes_{i:(i,j) \in S} a^{(i,\ell)} : \ell \in [r] \right\},$$

where the right-hand side is a subspace of dimension $\leq r$. For each S, j we will choose the better of these two bounds. We conclude, for any setting of the variables $a^{(i,\ell)}$,

$$\dim \text{colspan}(M_S) \leq s \prod_{j \in [d]} \min \left\{ \prod_{i:(i,j) \in S} n_i, r \right\}.$$

By the same argument applied to the row span,

$$\dim \text{rowspan}(M_S) \leq s \prod_{j \in [d]} \min \left\{ \prod_{i:(i,j) \in \bar{S}} n_i, r \right\}.$$

For each S we are free to choose the better of these two bounds to bound $\text{rank}(M_S)$. For each S , call an index $j \in [d]$ “good” if

$$\prod_{i:(i,j) \in S} n_i \leq n_* := \max_{U \subseteq [k]} \min \left\{ \prod_{i \in U} n_i, \prod_{i \in [k] \setminus U} n_i \right\}.$$

Each $j \in [d]$ must be good for either S or \bar{S} . Therefore, either S or \bar{S} has the property that at least $d/2$ indices j are good. If it's S then use the colspan bound for $\text{rank}(M_S)$, and otherwise use the rowspan bound. For “bad” indices, use the bound of r in the second term in $\min\{\dots\}$. We conclude

$$\text{rank}(M_S) \leq s(n_*)^{d/2} r^{d/2}$$

and so

$$\text{rank} \left(\sum_{1 \leq \ell_1 < \dots < \ell_d \leq r} \tilde{M}(T^{(\ell_1)}, \dots, T^{(\ell_d)}) \right) \leq \sum_{S \subseteq [k] \times [d]} \text{rank}(M_S) \leq 2^{kd} \cdot s(n_*)^{d/2} r^{d/2}.$$

Additivity would imply that this rank must equal $\binom{r}{d} s$, so we have

$$2^{kd} (n_*)^{d/2} r^{d/2} \geq \binom{r}{d} \geq \left(\frac{r}{d}\right)^d,$$

which gives $r \leq 4^k d^2 n_*$ as claimed.

A Trivial Flattenings

The following result shows that trivial flattenings achieve rank detection up to a particular rank, as claimed in Section 2.2.1.

Theorem A.1. *If $T = \sum_{\ell=1}^r a^{(1,\ell)} \otimes \dots \otimes a^{(k,\ell)}$ is $n_1 \times \dots \times n_k$ with the components generically chosen, and*

$$r \leq n_* := \max_{S \subseteq [k]} \min \left\{ \prod_{i \in S} n_i, \prod_{i \in [k] \setminus S} n_i \right\}, \quad (36)$$

then the matrix $M^{\text{triv}}(T; S)$ defined in (2), with S taken to be any maximizer of (36), has rank exactly r .

Proof. The trivial flattening takes the form

$$M^{\text{triv}}(T; S) = \sum_{\ell=1}^r b^{(\ell)} c^{(\ell)\top}$$

where

$$b^{(\ell)} := \bigotimes_{i \in S} a^{(i,\ell)} \quad \text{and} \quad c^{(\ell)} := \bigotimes_{i \in [k] \setminus S} a^{(i,\ell)}.$$

To show $\text{rank}(M^{\text{triv}}) = r$, it suffices to show the following: $\{b^{(\ell)} : \ell \in [r]\}$ are linearly independent and $\{c^{(\ell)} : \ell \in [r]\}$ are linearly independent. To see this: linear independence implies that the matrix B with columns $b^{(\ell)}$ has an $r \times r$ nonsingular submatrix, and similarly for C with rows $c^{(\ell)}$, and the product of these two submatrices appears as an $r \times r$ nonsingular submatrix of M^{triv} . We focus on the $b^{(\ell)}$ vectors, as the proof for the $c^{(\ell)}$ vectors is identical.

Since rank (of the B matrix) cannot exceed generic rank (see Section 2.1.2), it suffices to exhibit a single choice of the variables $a^{(i,\ell)}$ such that the $b^{(\ell)}$ vectors are linearly independent. Provided $r \leq \prod_{i \in S} n_i$, we can choose the variables $a^{(i,\ell)}$ so that the $b^{(\ell)}$ are distinct standard unit basis vectors. Formally, fix an injection $[r] \rightarrow \prod_{i \in S} [n_i]$ and write $\ell \mapsto (\phi_i(\ell) : i \in S)$ where $\phi_i(\ell) \in [n_i]$. Then for $i \in S$, set $a^{(i,\ell)}$ equal to the standard unit basis vector $e^{(\phi_i(\ell))} \in \mathbb{R}^{n_i}$, so that $b^{(\ell)} = \bigotimes_{i \in S} e^{(\phi_i(\ell))}$. Applying the same argument to the $c^{(\ell)}$ vectors, we also incur the requirement $r \leq \prod_{i \in [k] \setminus S} n_i$. \square

B Commuting Extensions

For $r \geq n$, a tuple of matrices (Z_1, \dots, Z_m) in $\mathbb{R}^{r \times r}$ is called a *commuting extension* of a tuple of matrices (A_1, \dots, A_m) in $\mathbb{R}^{n \times n}$ if the Z_i pairwise commute and the upper-left $n \times n$ submatrix of each Z_i is equal to A_i . The algorithmic task of finding a commuting extension is studied in [Koi24b], where the following planted model is proposed in Section 5: For $i \in [m]$, let $Z_i = R^{-1}D_iR$ where R is a generically chosen $r \times r$ matrix, and each D_i is a generically chosen $r \times r$ diagonal matrix. For some $n \leq r$, let A_i be the upper-left $n \times n$ submatrix of Z_i . Given (A_1, \dots, A_m) and r , the goal is to construct a commuting extension $(\tilde{Z}_1, \dots, \tilde{Z}_m)$ of dimension $r \times r$. The Z_i make up one “planted” solution, but it is not unique: see Eq. (2) of [Koi24b].

The algorithm of [Koi24b] solves the above task provided $m \geq 3$ and $r \leq 4n/3$. Building on the connection between commuting extensions and tensor decomposition described in [Koi24a], we will give an algorithm for finding commuting extensions using our tensor decomposition algorithm. The algorithm will succeed for $r \leq (2 - \epsilon)n$ for an arbitrary $\epsilon > 0$, provided m is a large enough constant (depending on ϵ) and n is large. The threshold $r = 2n$ is fundamental: as pointed out in [Koi20], any tuple (A_1, \dots, A_m) of $n \times n$ matrices admits the commuting extension

$$Z_i = \begin{pmatrix} A_i & -A_i \\ A_i & -A_i \end{pmatrix}$$

of dimension $r = 2n$.

The following lemma from [Koi24a] will be a subroutine in the algorithm.

Lemma B.1 ([Koi24a], Lemma 13). *Let $U \in \mathbb{R}^{n \times r}$ and $V \in \mathbb{R}^{r \times n}$ with $UV = I_n$. Then $r \geq n$, and it is possible to “extend” U, V to two $r \times r$ matrices that are inverses. More precisely, $\tilde{U} = (U^\top \mid A)^\top$ and $\tilde{V} = (V \mid B)$ where $A, B \in \mathbb{R}^{r \times (r-n)}$ and $\tilde{U}\tilde{V} = I_r$. Specifically, B can be any matrix for which $\text{Im}(B) = \ker(U)$, and A is the unique matrix such that $A^\top B = I_{r-n}$ and $A^\top U = 0$.*

Algorithm 2 (Finding a commuting extension).

- *Input: A tuple of $n \times n$ matrices (A_1, \dots, A_m) .*
 - *Input: A parameter $r \geq n$.*
 - *Input: A generic $n \times n$ matrix M (say, chosen at random).*
1. *Let $T_i = MA_i$ for $i \in [m]$, and let T be the $m \times n \times n$ tensor with slices T_i .*
 2. *Let q be the largest odd integer with $q \leq m$, and let $p = (q - 1)/2$.*
 3. *Run the tensor decomposition algorithm (Algorithm 1) on T with parameters p, q, r to produce a decomposition $T = \sum_{\ell=1}^r a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)}$.*
 4. *Let B be $n \times r$ with columns $b^{(\ell)}$, let C be $r \times n$ with rows $c^{(\ell)\top}$, and for $i \in [m]$ let $\tilde{D}_i = \text{diag}(a_i^{(\ell)} : \ell = 1, \dots, r)$.*
 5. *Solve for an $r \times r$ diagonal matrix \hat{D} satisfying the linear equations $M^{-1}B\hat{D}C = I_n$.*
 6. *Invoke Lemma B.1 to extend $M^{-1}B\hat{D}$ and C to $r \times r$ matrices \tilde{U} and \tilde{V} respectively.*
 7. *Output $(\tilde{Z}_1, \dots, \tilde{Z}_m)$ where $\tilde{Z}_i = \tilde{U}\hat{D}^{-1}\tilde{D}_i\tilde{V}$.*

The matrix M might be unnecessary, but it will help with the genericity analysis below.

Theorem B.2. *Let $r \geq n \geq 3$ and $m \geq 19$. The following holds for a generic choice of an $r \times r$ matrix R , an $n \times n$ matrix M , and m diagonal $r \times r$ matrices D_1, \dots, D_m (that is, the $r^2 + n^2 + mr$ variables $R_{ij}, M_{ij}, (D_i)_{jj}$ are chosen generically from $\mathbb{R}^{r^2+n^2+mr}$). For $i \in [m]$, let $Z_i = R^{-1}D_iR$ and let A_i be the upper-left $n \times n$ submatrix of Z_i . If*

$$r \leq 2n \left(1 - \frac{4}{m-1}\right) - \frac{m^3}{4}$$

then Algorithm 2 with input $(A_1, \dots, A_m), r, M$ outputs a valid commuting extension $(\tilde{Z}_1, \dots, \tilde{Z}_m)$, that is, the \tilde{Z}_i pairwise commute and for $i \in [m]$, the upper-left $n \times n$ submatrix of \tilde{Z}_i is A_i .

Proof. By construction, we know $A_i = UD_iV$ where U is $n \times r$ consisting of the first n rows of R^{-1} , and V is $r \times n$ consisting of the first n columns of R . This means $T_i = MUD_iV$, which gives a tensor decomposition $T = \sum_{\ell=1}^r w^{(\ell)} \otimes u^{(\ell)} \otimes v^{(\ell)}$ where $u^{(\ell)}$ are the columns of MU , $v^{(\ell)\top}$ are the rows of V , and $D_i = \text{diag}(w_i^{(1)}, \dots, w_i^{(r)})$. We will verify that this decomposition of T satisfies the conditions of our uniqueness theorem (Theorem 2.7), applied with the parameters p, q specified in Algorithm 2. The components $w^{(\ell)}, u^{(\ell)}, v^{(\ell)}$ are not generic in the usual sense, so we cannot apply Theorem 2.8 directly. Instead, since $R^{-1} = \text{adj}(R)/\det(R)$, the entries of the components $w^{(\ell)}, u^{(\ell)}, v^{(\ell)}$ can be expressed as rational functions in the underlying variables

$$x = \{R_{ij} : i, j \in [r]\} \sqcup \{(D_i)_{jj} : i \in [m], j \in [r]\} \sqcup \{M_{ij} : i, j \in [n]\}.$$

For generic x , our goal is to verify conditions (i)-(xi) from Theorem 2.7 for the components $w^{(\ell)}, u^{(\ell)}, v^{(\ell)}$. Inspecting each of these conditions and recalling that the components are rational functions, it suffices to exhibit a single choice of the underlying variables x (with R, M invertible) for which conditions (i)-(xi) hold; it then follows that they hold for generic x . From Theorem 2.8 we know that a generic choice of components will satisfy the conditions (we will verify the requirement of the theorem later); fix such a generic choice $\tilde{w}^{(\ell)}, \tilde{u}^{(\ell)}, \tilde{v}^{(\ell)}$. We will exhibit values for the variables $x = (R, D_i, M)$ such that $w^{(\ell)}(x) = \tilde{w}^{(\ell)}$, $u^{(\ell)}(x) = \tilde{u}^{(\ell)}$, and $v^{(\ell)}(x) = \tilde{v}^{(\ell)}$, implying that $w^{(\ell)}, u^{(\ell)}, v^{(\ell)}$ satisfy the conditions as desired.

First, set $D_i = \text{diag}(\tilde{w}_i^{(1)}, \dots, \tilde{w}_i^{(r)})$, which ensures $w^{(\ell)}(x) = \tilde{w}^{(\ell)}$. To ensure $u^{(\ell)}(x) = \tilde{u}^{(\ell)}$, we need to arrange $MU = \tilde{U}$ where \tilde{U} has columns $\tilde{u}^{(1)}, \dots, \tilde{u}^{(r)}$ and where U consists of the first n rows of R^{-1} . Equivalently, $U = M^{-1}\tilde{U}$. Similarly, to ensure $v^{(\ell)}(x) = \tilde{v}^{(\ell)}$, we need to arrange $V = \tilde{V}$ where \tilde{V} has rows $\tilde{v}^{(1)\top}, \dots, \tilde{v}^{(r)\top}$ and where V consists of the first n columns of R . Set $M = \tilde{U}\tilde{V}$, which is invertible due to the generic choice of $\tilde{u}^{(\ell)}, \tilde{v}^{(\ell)}$. For our desired values of $U = M^{-1}\tilde{U}$ and $V = \tilde{V}$, we have $UV = M^{-1}\tilde{U}\tilde{V} = I_n$. Therefore by Lemma B.1 there exists an $r \times r$ matrix R such that the first n rows of R^{-1} are $U = M^{-1}\tilde{U}$ and the first n columns of R are $V = \tilde{V}$. This completes the construction of $x = (R, D_i, M)$ and therefore completes the proof that the components $w^{(\ell)}, u^{(\ell)}, v^{(\ell)}$ satisfy the conditions of Theorem 2.7 when R, D_i, M are chosen generically.

Now by Theorem 2.7 we know T has a unique (in the sense of Definition 2.2) rank- r decomposition, and that the components $a^{(\ell)}, b^{(\ell)}, c^{(\ell)}$ recovered by our algorithm are related to $w^{(\ell)}, u^{(\ell)}, v^{(\ell)}$ as follows: there exists a permutation $\pi : [r] \rightarrow [r]$ such that $a^{(\ell)} \otimes b^{(\ell)} \otimes c^{(\ell)} = w^{(\pi(\ell))} \otimes u^{(\pi(\ell))} \otimes v^{(\pi(\ell))}$ for all $\ell \in [r]$. This means $b^{(\ell)} = \beta_\ell \cdot u^{(\pi(\ell))}$ for a scalar $\beta_\ell \neq 0$, which we can write in matrix form as $B = MU\Pi \cdot \text{diag}(\beta)$ where Π is the permutation matrix associated with π and B is defined in Algorithm 2. Similarly, $c^{(\ell)} = \gamma_\ell \cdot v^{(\pi(\ell))}$ for a scalar $\gamma_\ell \neq 0$, and $C = \text{diag}(\gamma) \cdot \Pi^\top V$. The algorithm's next step is to solve for a diagonal matrix \hat{D} such that $M^{-1}B\hat{D}C = I_n$. One solution to this linear system is $\hat{D} = \text{diag}(\beta_\ell^{-1}\gamma_\ell^{-1} : \ell = 1, \dots, r)$; to see this, note that $UV = I_n$ (since U, V are defined as particular submatrices of R^{-1}, R) and $\Pi \cdot \text{diag}(\beta) \cdot \hat{D} \cdot \text{diag}(\gamma) \cdot \Pi^\top = \text{diag}(\beta_{\pi^{-1}(\ell)}\hat{D}_{\pi^{-1}(\ell)}\gamma_{\pi^{-1}(\ell)} : \ell = 1, \dots, r)$.

$\ell = 1, \dots, r$). Later, we will show that generically, this is the only solution. Therefore the algorithm recovers $\hat{D} = \text{diag}(\beta_\ell^{-1} \gamma_\ell^{-1} : \ell = 1, \dots, r)$. The algorithm's next step is to extend $M^{-1}B\hat{D}$ and C to $r \times r$ matrices \tilde{U}, \tilde{V} with $\tilde{U}\tilde{V} = I_r$, where $M^{-1}B\hat{D}$ makes up the first n rows of \tilde{U} and C makes up the first n columns of \tilde{V} ; this uses Lemma B.1 along with the previous claim $M^{-1}B\hat{D}C = I_n$.

Recall that the algorithm has access to

$$\tilde{D}_i = \text{diag}(a_i^{(\ell)} : \ell = 1, \dots, r) = \text{diag}(\beta_\ell^{-1} \gamma_\ell^{-1} w_i^{(\pi(\ell))} : \ell = 1, \dots, r)$$

and outputs $\tilde{Z}_i = \tilde{U}\hat{D}^{-1}\tilde{D}_i\tilde{V}$ for $i \in [m]$. We need to show this is a valid commuting extension. Since \tilde{U}, \tilde{V} are inverses and $\hat{D}^{-1}\tilde{D}_i$ is diagonal, the Z_i pairwise commute. To verify that Z_i "extends" A_i : the upper-left $n \times n$ block of Z_i is

$$\begin{aligned} M^{-1}B\hat{D} \cdot \hat{D}^{-1}\tilde{D}_i \cdot C &= M^{-1}(MU\Pi \cdot \text{diag}(\beta))\tilde{D}_i(\text{diag}(\gamma) \cdot \Pi^\top V) \\ &= U\Pi \cdot \text{diag}(\beta) \cdot \tilde{D}_i \cdot \text{diag}(\gamma) \cdot \Pi^\top V \\ &= U\Pi \cdot \text{diag}(w_i^{(\pi(\ell))} : \ell = 1, \dots, r) \cdot \Pi^\top V \\ &= U \cdot \text{diag}(w_i^{(\ell)} : \ell = 1, \dots, r) \cdot V \\ &= UD_iV \\ &= A_i \end{aligned}$$

as desired.

Next we will prove the claim, made above, that the linear system of equations for \hat{D} has a unique solution. The equations are $M^{-1}B\hat{D}C = I_n$ where \hat{D} is diagonal, so it suffices to show that the matrices $M^{-1}b^{(\ell)}c^{(\ell)\top}$ for $\ell \in [r]$ are linearly independent. Since M^{-1} is invertible, it suffices to show that the matrices $b^{(\ell)}c^{(\ell)\top}$ for $\ell \in [r]$ are linearly independent. Since $b^{(\ell)}$ and $c^{(\ell)}$ are nonzero scalar multiples of $u^{(\pi(\ell))}$ and $v^{(\pi(\ell))}$ respectively, it suffices to show that $u^{(\ell)}v^{(\ell)\top}$ for $\ell \in [r]$ are linearly independent. Recalling that $u^{(\ell)}$ are the columns of MU and M is invertible, it suffices to show that $\hat{u}^{(\ell)}v^{(\ell)\top}$ for $\ell \in [r]$ are linearly independent, where $\hat{u}^{(\ell)}$ are the columns of U . Recalling that the entries of R^{-1} are rational functions in the entries of R , it suffices to exhibit a single (invertible) value of R for which $\hat{u}^{(\ell)}v^{(\ell)\top}$ are linearly independent, where $\hat{u}^{(\ell)}$ and $v^{(\ell)}$ are viewed as functions of R : namely, the first n rows of R^{-1} have columns $\hat{u}^{(\ell)}$ and the first n columns of R have rows $v^{(\ell)\top}$. Constructing such a value for R amounts to finding linearly independent matrices $\hat{u}^{(\ell)}v^{(\ell)\top}$ for $\ell \in [r]$ that sum to I_n , since the rest of R^{-1}, R can be filled in by Lemma B.1. This in turn is guaranteed by Lemma B.3 below, noting that our assumption $n \geq 3$ implies $n^2 - n + 1 \geq 2n$.

Finally, we need to check that our parameters p, q, r verify the requirements for Theorem 2.8. Note we are in the case $\alpha = 1$ so we need $q \geq 18$ and $r \leq 2n(1 - 4/q) - q^3/4$. These follow from our assumptions, and our choice $m - 1 \leq q \leq m$. \square

Lemma B.3. *For $n \geq 1$ and $n \leq r \leq n^2 - n + 1$, there exist r linearly independent rank-1 $n \times n$ matrices that sum to I_n .*

Proof. The rank-1 matrices used in our construction will each be supported on a rectangle, either with all nonzero values equal to 1 or all nonzero values equal to -1 . One possible construction is to take the rectangle $[n] \times [n]$ and subtract off a collection of rectangles that partition $([n] \times [n]) \setminus \{(1, 1), (2, 2), \dots, (n, n)\}$. The number of rectangles in the partition can range from $2(n - 1)$ to $n^2 - n$, so this covers all the r values ranging from $2(n - 1) + 1$ to $n^2 - n + 1$. An alternative construction is to, for some $k \in [n]$, take the rectangle $[k] \times [k]$ and subtract off a collection of $2(k - 1)$ rectangles that partition $([k] \times [k]) \setminus \{(1, 1), (2, 2), \dots, (k, k)\}$, and then add $n - k$ additional 1×1 rectangles on the diagonal. This uses $1 + 2(k - 1) + (n - k) = n + k - 1$ rectangles in total, so by varying k this covers all the r values ranging from n to $2n - 1$. \square

Acknowledgments

We thank Aravindan Vijayaraghavan for helpful discussions.

References

- [AFT11] Boris Alexeev, Michael A Forbes, and Jacob Tsimerman. Tensor rank: Some lower and upper bounds. In *26th Annual Conference on Computational Complexity*, pages 283–291. IEEE, 2011.
- [AGKM23] Omar Alrabiah, Venkatesan Guruswami, Pravesh K Kothari, and Peter Manohar. A near-cubic lower bound for 3-query locally decodable codes from semirandom CSP refutation. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1438–1448, 2023.
- [AOW15] Sarah R Allen, Ryan O’Donnell, and David Witmer. How to refute a random CSP. In *56th Annual Symposium on Foundations of Computer Science*, pages 689–708. IEEE, 2015.
- [BCM14a] 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. PMLR, 2014.
- [BCM14b] Aditya Bhaskara, Moses Charikar, Ankur Moitra, and Aravindan Vijayaraghavan. Smoothed analysis of tensor decompositions. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 594–603, 2014.
- [BCO14] Cristiano Bocci, Luca Chiantini, and Giorgio Ottaviani. Refined methods for the identifiability of tensors. *Annali di Matematica Pura ed Applicata (1923-)*, 193:1691–1702, 2014.
- [BESV24] Aditya Bhaskara, Eric Evert, Vaidehi Srinivas, and Aravindan Vijayaraghavan. New tools for smoothed analysis: Least singular value bounds for random matrices with dependent entries. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 375–386, 2024.
- [BM16] Boaz Barak and Ankur Moitra. Noisy tensor completion via the sum-of-squares hierarchy. In *Conference on Learning Theory*, pages 417–445. PMLR, 2016.
- [BVD⁺18] Martijn Boussé, Nico Vervliet, Ignat Domanov, Otto Debals, and Lieven De Lathauwer. Linear systems with a canonical polyadic decomposition constrained solution: Algorithms and applications. *Numerical Linear Algebra with Applications*, 25(6):e2190, 2018.
- [Car91] Jean-François Cardoso. Super-symmetric decomposition of the fourth-order cumulant tensor. Blind identification of more sources than sensors. In *ICASSP*, volume 91, pages 3109–3112, 1991.
- [CGL07] Amin Coja-Oghlan, Andreas Goerdt, and André Lanka. Strong refutation heuristics for random k-SAT. *Combinatorics, Probability and Computing*, 16(1):5–28, 2007.

- [CKW11] Xi Chen, Neeraj Kayal, and Avi Wigderson. Partial derivatives in arithmetic complexity and beyond. *Foundations and Trends® in Theoretical Computer Science*, 6(1–2):1–138, 2011.
- [CR22] Haolin Chen and Luis Rademacher. Overcomplete order-3 tensor decomposition, blind deconvolution, and gaussian mixture models. *SIAM Journal on Mathematics of Data Science*, 4(1):336–361, 2022.
- [DdL⁺22] Jingqiu Ding, Tommaso d’Orsi, Chih-Hung Liu, David Steurer, and Stefan Tiegel. Fast algorithm for overcomplete order-3 tensor decomposition. In *Conference on Learning Theory*, pages 3741–3799. PMLR, 2022.
- [DJL24] Harm Derksen, Nathaniel Johnston, and Benjamin Lovitz. \mathcal{X} -arability of mixed quantum states. *arXiv preprint arXiv:2409.18948*, 2024.
- [DL06] Lieven De Lathauwer. A link between the canonical decomposition in multilinear algebra and simultaneous matrix diagonalization. *SIAM journal on Matrix Analysis and Applications*, 28(3):642–666, 2006.
- [DLCC07] Lieven De Lathauwer, Josphine Castaing, and Jean-Francois Cardoso. Fourth-order cumulant-based blind identification of underdetermined mixtures. *IEEE Transactions on Signal Processing*, 55(6):2965–2973, 2007.
- [EGOW18] Klim Efremenko, Ankit Garg, Rafael Oliveira, and Avi Wigderson. Barriers for rank methods in arithmetic complexity. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [Far16] Cameron Farnsworth. Koszul–Young flattenings and symmetric border rank of the determinant. *Journal of Algebra*, 447:664–676, 2016.
- [GKM22] Venkatesan Guruswami, Pravesh K Kothari, and Peter Manohar. Algorithms and certificates for Boolean CSP refutation: smoothed is no harder than random. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 678–689, 2022.
- [GKS20] Ankit Garg, Neeraj Kayal, and Chandan Saha. Learning sums of powers of low-degree polynomials in the non-degenerate case. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 889–899. IEEE, 2020.
- [GM15] Rong Ge and Tengyu Ma. Decomposing overcomplete 3rd order tensors using sum-of-squares algorithms. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2015)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2015.
- [GMOW19] Ankit Garg, Visu Makam, Rafael Oliveira, and Avi Wigderson. More barriers for rank methods, via a “numeric to symbolic” transfer. In *60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 824–844. IEEE, 2019.
- [Gua17] Yonghui Guan. Flattenings and Koszul Young flattenings arising in complexity theory. *Communications in Algebra*, 45(9):4002–4017, 2017.

- [GVX14] Navin Goyal, Santosh Vempala, and Ying Xiao. Fourier PCA and robust tensor decomposition. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 584–593, 2014.
- [Har70] Richard A Harshman. Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multi-modal factor analysis. *UCLA working papers in phonetics*, 16(1):84, 1970.
- [Hås90] Johan Håstad. Tensor rank is NP-complete. *Journal of algorithms*, 11(4):644–654, 1990.
- [Has20] Matthew B Hastings. Classical and quantum algorithms for tensor principal component analysis. *Quantum*, 4:237, 2020.
- [HKM23] Jun-Ting Hsieh, Pravesh K Kothari, and Sidhanth Mohanty. A simple and sharper proof of the hypergraph Moore bound. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2324–2344. SIAM, 2023.
- [HL13] Christopher J Hillar and Lek-Heng Lim. Most tensor problems are NP-hard. *Journal of the ACM (JACM)*, 60(6):1–39, 2013.
- [HOOS19] Jonathan D Hauenstein, Luke Oeding, Giorgio Ottaviani, and Andrew J Sommese. Homotopy techniques for tensor decomposition and perfect identifiability. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2019(753):1–22, 2019.
- [HSS19] Samuel B Hopkins, Tselil Schramm, and Jonathan Shi. A robust spectral algorithm for overcomplete tensor decomposition. In *Conference on Learning Theory*, pages 1683–1722. PMLR, 2019.
- [HSSS16] Samuel B Hopkins, Tselil Schramm, Jonathan Shi, and David Steurer. Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 178–191, 2016.
- [JLV23] Nathaniel Johnston, Benjamin Lovitz, and Aravindan Vijayaraghavan. Computing linear sections of varieties: quantum entanglement, tensor decompositions and beyond. In *64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1316–1336. IEEE, 2023.
- [KB09] Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- [KM24a] Pravesh K Kothari and Peter Manohar. An exponential lower bound for linear 3-query locally correctable codes. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 776–787, 2024.
- [KM24b] Pravesh K Kothari and Peter Manohar. Superpolynomial lower bounds for smooth 3-LCCs and sharp bounds for designs. *arXiv preprint arXiv:2404.06513*, 2024.
- [Koi20] Pascal Koiran. On tensor rank and commuting matrices. *arXiv preprint arXiv:2006.02374*, 2020.

- [Koi24a] Pascal Koiran. An efficient uniqueness theorem for overcomplete tensor decomposition. *arXiv preprint arXiv:2404.07801*, 2024.
- [Koi24b] Pascal Koiran. On the uniqueness and computation of commuting extensions. *arXiv preprint arXiv:2401.01302*, 2024.
- [KP20] Bohdan Kivva and Aaron Potechin. Exact nuclear norm, completion and decomposition for random overcomplete tensors via degree-4 SOS. *arXiv preprint arXiv:2011.09416*, 2020.
- [Kru77] Joseph B Kruskal. Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear algebra and its applications*, 18(2):95–138, 1977.
- [KS23] Pascal Koiran and Subhayan Saha. Complete decomposition of symmetric tensors in linear time and polylogarithmic precision. In *International Conference on Algorithms and Complexity*, pages 308–322. Springer, 2023.
- [Lan11] Joseph M Landsberg. *Tensors: geometry and applications*, volume 128. American Mathematical Soc., 2011.
- [Lan15] JM Landsberg. Nontriviality of equations and explicit tensors in $\mathbb{C}^m \otimes \mathbb{C}^m \otimes \mathbb{C}^m$ of border rank at least $2m - 2$. *Journal of Pure and Applied Algebra*, 219(8):3677–3684, 2015.
- [LM17] Joseph M Landsberg and Mateusz Michałek. A $2n^2 - \log_2 n - 1$ lower bound for the border rank of matrix multiplication. *International Mathematics Research Notices*, 2018(15), 2017.
- [LM20] Allen Liu and Ankur Moitra. Tensor completion made practical. *Advances in Neural Information Processing Systems*, 33:18905–18916, 2020.
- [LO13] Joseph M Landsberg and Giorgio Ottaviani. Equations for secant varieties of veronese and other varieties. *Annali di Matematica Pura ed Applicata*, 192(4):569–606, 2013.
- [LO15] Joseph M Landsberg and Giorgio Ottaviani. New lower bounds for the border rank of matrix multiplication. *Theory of Computing*, 11(1):285–298, 2015.
- [LP23] Benjamin Lovitz and Fedor Petrov. A generalization of Kruskal’s theorem on tensor decomposition. In *Forum of Mathematics, Sigma*, volume 11, page e27. Cambridge University Press, 2023.
- [LRA93] Sue E Leurgans, Robert T Ross, and Rebecca B Abel. A decomposition for three-way arrays. *SIAM Journal on Matrix Analysis and Applications*, 14(4):1064–1083, 1993.
- [Moi18] Ankur Moitra. *Algorithmic aspects of machine learning*. Cambridge University Press, 2018.
- [MSS16] Tengyu Ma, Jonathan Shi, and David Steurer. Polynomial-time tensor decompositions with sum-of-squares. In *57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 438–446. IEEE, 2016.
- [OO13] Luke Oeding and Giorgio Ottaviani. Eigenvectors of tensors and algorithms for Waring decomposition. *Journal of Symbolic Computation*, 54:9–35, 2013.

- [Per18] Elena-Mădălina Persu. *Tensors, sparse problems and conditional hardness*. PhD thesis, Massachusetts Institute of Technology, 2018.
- [PWB⁺19] Amelia Perry, Jonathan Weed, Afonso S Bandeira, Philippe Rigollet, and Amit Singer. The sample complexity of multireference alignment. *SIAM Journal on Mathematics of Data Science*, 1(3):497–517, 2019.
- [Raz13] Ran Raz. Tensor-rank and lower bounds for arithmetic formulas. *Journal of the ACM (JACM)*, 60(6):1–15, 2013.
- [RSG17] Stephan Rabanser, Oleksandr Shchur, and Stephan Günnemann. Introduction to tensor decompositions and their applications in machine learning. *arXiv preprint arXiv:1711.10781*, 2017.
- [Swe18] Joseph Swernofsky. Tensor rank is hard to approximate. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [Vij20] Aravindan Vijayaraghavan. Efficient tensor decomposition. *arXiv preprint arXiv:2007.15589*, 2020.
- [Wei23] Alexander S Wein. Average-case complexity of tensor decomposition for low-degree polynomials. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1685–1698, 2023.
- [WEM19] Alexander S Wein, Ahmed El Alaoui, and Cristopher Moore. The Kikuchi hierarchy and tensor PCA. In *60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1446–1468. IEEE, 2019.