Discreteness of asymptotic tensor ranks

Jop Briët[1], Matthias Christandl[2], Itai Leigh[3], Amir Shpilka[3], and Jeroen Zuiddam[4]

[1]Centrum Wiskunde & Informatica, Amsterdam, Netherlands
[2]University of Copenhagen, Copenhagen, Denmark
[3]Tel Aviv University, Tel Aviv, Israel
[4]University of Amsterdam, Amsterdam, Netherlands

September 25, 2023

Abstract

Tensor parameters that are amortized or regularized over large tensor powers, often called "asymptotic" tensor parameters, play a central role in several areas including algebraic complexity theory (constructing fast matrix multiplication algorithms), quantum information (entanglement cost and distillable entanglement), and additive combinatorics (bounds on cap sets, sunflower-free sets, etc.). Examples are the asymptotic tensor rank, asymptotic slice rank and asymptotic subrank. Recent works (Costa–Dalai, Blatter–Draisma–Rupniewski, Christandl–Gesmundo–Zuiddam) have investigated notions of discreteness (no accumulation points) or "gaps" in the values of such tensor parameters.

We prove a general discreteness theorem for asymptotic tensor parameters of order-three tensors and use this to prove that (1) over any finite field (and in fact any finite set of coefficients in any field), the asymptotic subrank and the asymptotic slice rank have no accumulation points, and (2) over the complex numbers, the asymptotic slice rank has no accumulation points.

Central to our approach are two new general lower bounds on the asymptotic subrank of tensors, which measures how much a tensor can be diagonalized. The first lower bound says that the asymptotic subrank of any concise three-tensor is at least the cube-root of the smallest dimension. The second lower bound says that any concise three-tensor that is "narrow enough" (has one dimension much smaller than the other two) has maximal asymptotic subrank.

Our proofs rely on new lower bounds on the maximum rank in matrix subspaces that are obtained by slicing a three-tensor in the three different directions. We prove that for any concise tensor, the product of any two such maximum ranks must be large, and as a consequence there are always two distinct directions with large max-rank.

Contents

1   Introduction

Tensor parameters that are amortized or regularized over large tensor powers, often called "asymptotic" tensor parameters, play a central role in several areas of theoretical computer science including algebraic complexity theory (constructing fast matrix multiplication algorithms, and barriers for such constructions), quantum information (entanglement cost and distillable entanglement), and additive combinatorics (bounds on cap sets, sunflower-free sets, etc.). Examples are the asymptotic tensor rank (famous for its connection to the matrix multiplication exponent), the asymptotic subrank, and the asymptotic slice rank. These asymptotic tensor parameters are of the form $\underset{\sim}{F}(T) = \lim_{n\to\infty} F(T^{\boxtimes n})^{1/n}$ for some integer-valued function F (e.g. tensor rank, subrank or slice rank), where $\boxtimes$ denotes the Kronecker product of tensors. The computation of these parameters, which in some applications is the main goal and in others is done to bound other parameters of interest, has turned out to be very challenging, and many questions about them are open despite much interest.

The fundamental question whether the matrix-multiplication exponent ω equals 2 is closely related to the question whether asymptotic tensor rank is integral-valued. Contrary to matrix rank, some asymptotic tensor parameters may indeed take non-integral values. For instance, the asymptotic slice rank of the W-tensor (which appears in the study of sunflower-free sets) equals $2^{h(1/3)} \approx 1.88$, where h is the binary entropy function, and the asymptotic slice rank of the cap set tensor (which appears in the study of arithmetic progression-free sets or cap sets) is known to be the non-integral value ≈ 2.755 over the finite field $\mathbb{F}_3$.

This raises the fundamental question, for a given function F, what values $\underset{\sim}{F}(T)$ can take when varying T over all tensors of order three with arbitrary dimensions (over any fixed field). More generally, what is the structure (geometric, algebraic, topological, etc) of the set of values

$$\{\underset{\sim}{F}(T) : T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}, n_1, n_2, n_3 \in \mathbb{N}\}?$$

Does $\underset{\sim}{F}(T)$ have accumulation points, that is, are there non-trivial sequences of tensors $T_1, T_2, \ldots$ such that $\underset{\sim}{F}(T_i)$ converges? Or is it discrete? What "gaps" are there between the possible values? Even when $\mathbb{F}$ is a finite field, the answers to these questions are a priori not clear at all.

In this paper we prove *discreteness* of asymptotic tensor rank, asymptotic subrank and asymptotic slice rank in several regimes. This means that the values of each of these parameters have no accumulation points. In fact, the proof of discreteness of asymptotic tensor rank (over finite fields) follows from a simple argument. Using that same simple argument, together with several new results about tensors, we obtain discreteness for the other parameters. In particular, as a core ingredient, we prove a new result about diagonalizability of tensors. This comes in the form of a lower bound on the asymptotic subrank that relies only on the dimensions of the tensor (as opposed to the well-known laser method for fast matrix multiplication, which relies on much more information about the tensor). As another core ingredient we prove new results about matrix subspaces and their max-rank and min-rank. In particular, we prove that the max-ranks of the matrix spaces obtained by slicing a tensor in the three different directions are strongly related, in such a way that at least two of them must be large.

Our discreteness results show that there is a surprising rigidity in the asymptotic behaviour of tensors. The discreteness of the above parameters gives rise to the phenomenon of "rounding" or "boosting" (upper or lower) bounds on them to the next possible value

(although making this effective requires more knowledge of the possible values than just knowing discreteness). The discreteness of asymptotic tensor rank implies, for instance, that the asymptotic rank of the matrix multiplication tensor is bounded away from (or equal to) the asymptotic tensor rank of any other tensor. In particular, the matrix multiplication exponent ω is an isolated point among the exponents of all bilinear maps. If a tensor is "close enough" to the matrix multiplication tensor, its exponent must "snap" to ω. Similar statements hold for the other asymptotic parameters. In the context of combinatorial applications, this may moreover lead to limitations for the asymptotic slice rank to improve on existing results.

Before stating our results in detail, we will first discuss the various asymptotic ranks, their applications and context.

Matrix multiplication and asymptotic tensor rank

Determining the computational complexity of matrix multiplication is a fundamental problem in algebraic complexity theory. This complexity is controlled by the famous matrix multiplication exponent ω, which is defined as the infimum over nonnegative real numbers τ such that any two $n \times n$ matrices can be multiplied using $\mathcal{O}(n^\tau)$ arithmetic operations [BCS97, Blä13]. The naive matrix multiplication algorithm gives the upper bound $\omega \leq 3$. In 1969, Strassen proved that $\omega \leq 2.81$ [Str69]. Since then, using many different techniques, the best upper bound has been brought down to $\omega \leq 2.371552$ [LG14, AW21, DWZ23, WXXZ23]. There is a tantalizing possibility (and many have conjectured) that $\omega = 2$, and routes have been proposed that aim to prove this [CU03, CKSU05, CU13, BCC$^+$17, BCG$^+$22]. It is just as intriguing to consider the possibility that $\omega > 2$ and ω giving rise to a new fundamental constant, and there has been much work on this lower bound direction as well [BI11, LO13].

Not only can we currently not determine the value of ω, or decide whether $\omega = 2$ or $\omega > 2$, there is a much more relaxed problem that we cannot solve. Indeed ω is naturally described in terms of tensors as the logarithm of the asymptotic tensor rank of the matrix multiplication tensor $\langle 2, 2, 2 \rangle \in \mathbb{F}^4 \otimes \mathbb{F}^4 \otimes \mathbb{F}^4$, that is $\omega = \log \underset{\sim}{\mathrm{R}}(\langle 2, 2, 2 \rangle)$, and thus $\omega > 2$ is equivalent to $\underset{\sim}{\mathrm{R}}(\langle 2, 2, 2 \rangle) > 4$. The following much more relaxed problem is open:

Problem 1.1 ([BCS97, Open Problem 15.5]). Prove that there is a tensor $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ with $\underset{\sim}{\mathrm{R}}(T) > n$.

It is possible that for every concise tensor $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ we have $\underset{\sim}{\mathrm{R}}(T) = n$ (which would in particular imply $\omega = 2$) and that the image of $\underset{\sim}{\mathrm{R}}$ over all tensors is simply $\mathbb{N}$. This naturally leads to the (very general) question: What is the structure (geometric, algebraic, topological, etc) of the set

$$S = \{\underset{\sim}{\mathrm{R}}(T) : T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}, n_1, n_2, n_3 \in \mathbb{N}\}?$$

Is there anything we can prove about S without resolving Problem 1.1 or determining ω? Not much is known.

One known structural result is that S is closed under applying any univariate polynomial with non-negative integer coefficients [WZ22, Theorem 4.8]. This statement applies in fact more generally, and in particular also to asymptotic subrank and asymptotic slice rank. This thus says that S has "many" elements.

Our discreteness result says that S does not have "too many" elements, and for asymptotic tensor rank over a finite field[1] the proof is surprisingly simple. Here is a sketch: Let

---

[1]Or in fact over any finite set of coefficients coming from an arbitrary field.

$T_1, T_2, \ldots$ be any sequence of tensors such that $T_i \in \mathbb{F}^{a_i} \otimes \mathbb{F}^{b_i} \otimes \mathbb{F}^{c_i}$ and such that $\underaccent{\tilde}{\mathrm{R}}(T_i)$ takes infinitely many values. We may assume that every T_i is concise, meaning it does not fit into any smaller tensor space. Then $\underaccent{\tilde}{\mathrm{R}}(T_i) \geq \max\{a_i, b_i, c_i\}$. Since by assumption $\mathbb{F}$ is finite, there are only finitely many tensors per format $a_i \times b_i \times c_i$, so the set of triples $\{(a_i, b_i, c_i) : i \in \mathbb{N}\}$ is infinite. In particular, $\max\{a_i, b_i, c_i\}$ is unbounded, and so $\underaccent{\tilde}{\mathrm{R}}(T_i)$ is unbounded and cannot converge, which proves the claim.

While this argument is very simple, a much more subtle argument and new technical results will be needed to deal with the other tensor parameters that we consider.

Asymptotic subrank and asymptotic slice rank

Besides tensor rank, there are many other notions of rank of a tensor that play a role in applications, for instance the subrank, slice rank, analytic rank [GW11, Lov19, CM23], geometric rank [KMZ20], and G-stable rank [Der22]. We will focus here on the asymptotic subrank and asymptotic slice rank. The subrank was introduced by Strassen [Str87] in the study of matrix multiplication algorithms. The subrank $\mathrm{Q}(T)$ is the size of the largest diagonal tensor that can be obtained from T by taking linear combinations of the slices in the three different directions (i.e. "Gaussian elimination" for tensors). The slice rank was introduced by Tao [Tao16] to give a tensor proof of the cap set problem after the first proof of Gijswijt and Ellenberg [EG17]. The slice rank $\mathrm{SR}(T)$ is the smallest number of tensors with flattening rank one whose sum is T. Tao proved that $\mathrm{Q}(T) \leq \mathrm{SR}(T)$. Recent works have shown that analytic rank, geometric rank [CM21, CM23], and G-stable rank [Der22] are all equal to slice rank, up to a multiplicative constant. These results imply that the asymptotic versions of these parameters are all equal to the asymptotic slice rank, warranting our focus on it.

Slice rank method in combinatorics. The proof of the longstanding cap set problem [Tao16, ST16] (and other related results [NS17]) can be thought of as upper bounding the independence number of powers of a hypergraph, by constructing a tensor that "fits" on the hypergraph and then computing the slice rank of the powers of the tensor, that is, the asymptotic slice rank $\underaccent{\tilde}{\mathrm{SR}}(T)$. Knowing (the structure of) the set

$$S = \{\underaccent{\tilde}{\mathrm{SR}}(T) : T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}, n_1, n_2, n_3 \in \mathbb{N}\}$$

thus gives information on what bounds one can prove on the size of combinatorial objects using the slice rank method. One of our main results is that asymptotic slice rank is discrete, not only over every finite field, but even over the complex numbers. The latter crucially requires a result from [CVZ23] that characterizes asymptotic slice rank in terms of representation-theoretic objects called moment polytopes. In fact it is known by now that the four smallest value in S are $0, 1, \approx 1.88, 2, \approx 2.68$ (see next section) and our result says that also the larger values are discrete.

Matrix multiplication barriers. Besides the aforementioned combinatorial problems, the asymptotic subrank and asymptotic slice rank appear in several "barrier results" for matrix multiplication algorithms [Alm19, CVZ21, BCC+17]. Matrix multiplication algorithms are usually constructed by a reduction of matrix multiplication to another bilinear map, and these barrier results say what properties that intermediate bilinear map must have to obtain certain upper bounds on ω, or to reach $\omega = 2$. These properties can be phrased in terms of asymptotic subrank or asymptotic slice rank. In particular, the barrier of [CVZ21] states that to reach $\omega = 2$, an intermediate tensor T must have $\underaccent{\tilde}{\mathrm{R}}(T) = \mathrm{Q}(T)$, which has led to further research to find tensors with large asymptotic subrank [BL20]. Our discreteness result says that asymptotic subrank is discrete over every finite field. We do not get this

result over the complex numbers because the analogous representation-theoretic ingredient from above is missing here. Intriguingly, it is possible that

$$S = \{\underset{\sim}{Q}(T) : T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}, n_1, n_2, n_3 \in \mathbb{N}\}$$

equals the analogous set for asymptotic slice rank, that is, the following is open:

Problem 1.2. Prove that asymptotic slice rank equals asymptotic subrank.

We do as a side-result prove a new relation between asymptotic subrank and asymptotic slice rank (which we will discuss in more detail in the next section).

Previous work on discreteness of asymptotic ranks

Several works, among which some very recent ones, have investigated notions of discreteness in the values of tensor parameters.

Strassen [Str88, Lemma 3.7] proved that for any k-tensor T over any field, the asymptotic subrank (and, as a consequence of his method, also the asymptotic slice rank) of T is equal to 0, equal to 1, or at least $2^{2/k}$. This result established the first "gaps" in asymptotic tensor parameters. Costa and Dalai [CD21] proved that, for any k-tensor T over any field, the asymptotic slice rank of T is equal to 0, equal to 1 or at least $2^{h(1/k)}$ where h is the binary entropy function[2]. Christandl, Gesmundo and Zuiddam [CGZ22] extended the result of Costa and Dalai by proving that, for any k-tensor T over any field, the asymptotic subrank and asymptotic partition rank of T are equal to 0, equal to 1 or at least $2^{h(1/k)}$ (which is a tight bound). Additionally, they prove that for any 3-tensor T over any field, the asymptotic subrank and asymptotic slice rank of T are equal to 0, equal to 1, equal to $2^{h(1/3)} \approx 1.88$ or at least 2. Gesmundo and Zuiddam [GZ23] extended this result by proving that the next possible value after 2 is ≈ 2.68.

Blatter, Draisma and Rupniewski [BDR22b] proved that for any function on k-tensors, over any finite field, that is normalized and monotone the set of values that this function takes is well-ordered. The asymptotic subrank and asymptotic slice rank[3] are examples of such functions. This means that the values of any such function do not have accumulation points "from above", but leaves open the possibility that there are accumulation points "from below".

Christandl, Vrana and Zuiddam [CVZ23] proved that the asymptotic slice rank over the complex numbers takes only finitely many values on tensors of any fixed format, and thus only countably infinite many values in general. This is done by characterizing the asymptotic slice rank as an optimization over the moment polytope of the tensor and using the result that there are only finitely many such polytopes per tensor format.

Blatter, Draisma and Rupniewski [BDR22a] proved that for any "algebraic" tensor invariant over the complex numbers the related asymptotic parameter takes only countably many values. This implies in particular that the asymptotic subrank, asymptotic slice rank, asymptotic geometric rank, and asymptotic partition rank (all over the complex numbers) take only countably many values.

---

[2]The binary entropy function is defined for $p \in (0, 1)$ by $h(p) = -p \log_2 p - (1 - p) \log_2(1 - p)$ and $h(0) = h(1) = 0$.

[3]Over finite fields, since we do not know whether the limit $\lim_{n \to \infty} \mathrm{SR}(T^{\boxtimes n})^{1/n}$ exists in general, when we say asymptotic slice rank we will mean $\liminf_{n \to \infty} \mathrm{SR}(T^{\boxtimes n})^{1/n}$.

New results

In this paper:

- We prove two general lower bounds on the asymptotic subrank of concise tensors that depend only on the dimensions of the tensor. The first says that the asymptotic subrank of any concise tensor is at least the cube-root of its smallest dimension. The second says that the asymptotic subrank of any "narrow enough" tensor (meaning that one dimension is much smaller than the others) is maximal.

- We use those lower bounds to prove that over any finite set of coefficients in any field the asymptotic subrank has no accumulation points (i.e., is discrete). We moreover prove that over any finite field and over the complex numbers the asymptotic slice rank has no accumulation points. A much simpler argument gives that the asymptotic rank is discrete over any finite set of coefficients.

- As a core ingredient for the above, we prove optimal relations among the maximal rank of any matrix in the span of the slices of a tensor, when considering slicings in the three different directions.

- With similar techniques, we prove an upper bound on the difference between asymptotic subrank and asymptotic slice rank.

1.1 Discreteness of asymptotic tensor parameters

We will now discuss our results in more detail. We begin with some basic definitions (which we discuss in more detail in Section 1.4). Let $\mathbb{F}$ be any field. Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ be an order-three tensor over $\mathbb{F}$ with dimensions (n_1, n_2, n_3). The subrank of T, denoted by $Q(T)$, is the largest number r such that there are linear maps $L_i : \mathbb{F}^{n_i} \to \mathbb{F}^r$ such that $(L_1 \otimes L_2 \otimes L_3)T = \sum_{i=1}^{r} e_i \otimes e_i \otimes e_i$. In other words, the subrank measures how much a tensor can be diagonalized. The flattenings of T are the elements in $(\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}) \otimes \mathbb{F}^{n_3}$, $\mathbb{F}^{n_1} \otimes (\mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3})$ and $\mathbb{F}^{n_2} \otimes (\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_3})$, obtained by naturally grouping the tensor legs of T. We say T has slice rank one if at least one of the flattenings has rank one. The slice rank of T, denoted by $SR(T)$ is the smallest number r such that T is a sum of r tensors with slice rank one. The asymptotic subrank of T is defined as $\widetilde{Q}(T) = \lim_{n \to \infty} Q(T^{\boxtimes n})^{1/n}$ where $\boxtimes$ is the Kronecker product on tensors. The limit exists and equals the supremum by Fekete's lemma, since Q is super-multiplicative. The asymptotic slice rank of T we define as $\widetilde{SR}(T) = \liminf_{n \to \infty} SR(T^{\boxtimes n})^{1/n}$. (Over the complex numbers, it is known that the liminf can be replaced by a limit. Over other fields, however this is generally not known.)

1.1.1 Discreteness of asymptotic slice rank and asymptotic subrank

We prove discreteness (no accumulation points) for the asymptotic slice rank and asymptotic subrank, in several regimes, as follows.

Theorem 1.3. *Over any finite field, the asymptotic subrank and the asymptotic slice rank each have no accumulation points.*

Theorem 1.3 improves the result of Blatter, Draisma and Rupniewski [BDR22b] that the asymptotic subrank and asymptotic slice rank over any finite field have no accumulation points "from above", that is, have well-ordered sets of values. Indeed our result rules out all accumulation points (so also those "from below").

Theorem 1.4. *Over the complex numbers, the asymptotic slice rank has no accumulation points.*

Theorem 1.4 improves the result of Blatter, Draisma and Rupniewski [BDR22a] and Christandl, Vrana and Zuiddam [CVZ23] that the asymptotic slice rank over the complex numbers takes only countably many values. Our result is indeed strictly stronger, as countable sets may a priori have accumulation points.

Our results shed light on the recent results of Costa and Dalai [CD21] and Christandl, Gesmundo and Zuiddam [CGZ22] that found gaps between the smallest values of the asymptotic slice rank and asymptotic subrank, and answers positively (in some regimes) the question stated in [CGZ22] asking whether the values will always be "gapped".

Besides the above, we prove discreteness theorems over arbitrary fields for two subclasses of all tensors. Namely we consider the class of oblique tensors, which are the tensors whose support in some basis is an antichain, and the tight tensors, whose support in some basis can be characterized by algebraic equation (details in Section 4). The set of tight tensors is a strict subset of the set of oblique tensors, which is a strict subset of all tensors. Both classes originate in the work of Strassen [Str87, Str88, Str91]. Examples of tight tensors include the well-known matrix multiplication tensors. Tight tensors also play a central role in the laser method of Strassen to construct fast matrix multiplication algorithms. Our discreteness theorems in these regimes are as follows.

Theorem 1.5. *Over any field, on tight tensors, the asymptotic subrank and asymptotic slice rank (which are equal) have no accumulation points.*

Theorem 1.6. *Over any field, on oblique tensors, the asymptotic slice rank has no accumulation points.*

1.1.2 General discreteness theorem

We prove the above discreteness theorems as an application of a general discreteness theorem that we discuss now. This general theorem gives discreteness for real-valued tensor parameters that satisfy several conditions. To describe these conditions we need the notion of equivalence of tensors. We say two tensors $S \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ and $T \in \mathbb{F}^{m_1} \otimes \mathbb{F}^{m_2} \otimes \mathbb{F}^{m_3}$ are *equivalent* if there are linear maps $A_i : \mathbb{F}^{n_i} \to \mathbb{F}^{m_i}$ such that $(A_1 \otimes A_2 \otimes A_3)S = T$ and there are linear maps $B_i : \mathbb{F}^{m_i} \to \mathbb{F}^{n_i}$ such that $(B_1 \otimes B_2 \otimes B_3)S = T$. A tensor $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ is called concise if the three flattenings obtained by grouping two of the three tensor legs together, in the three possible ways, each have maximal rank. This means essentially that T cannot be embedded into a smaller tensor space.

Theorem 1.7. *Let $\mathbb{F}$ be any field. Let $\mathcal{C}$ be any subset of tensors over $\mathbb{F}$ such that for every $T \in \mathcal{C}$ there is an $S \in \mathcal{C}$ that is concise and equivalent to T. Let $f : \mathcal{C} \to \mathbb{R}_{\geq 0}$ be any function such that*

(i) *f does not change under equivalence of tensors*

(ii) *$f(T) \geq \underset{\sim}{Q}(T)$ for every $T \in \mathcal{C}$*

(iii) *For every $n_1, n_2, n_3 \in \mathbb{N}$, f takes finitely many values on $\mathcal{C} \cap (\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3})$*

(iv) *For every $n_1, n_2, n_3 \in \mathbb{N}$, $f(T) \leq \min_i n_i$ for every $T \in \mathcal{C} \cap (\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3})$.*

Then the set of values that f takes on $\mathcal{C}$ has no accumulation points.

We make some remarks on condition (iii) of Theorem 1.7. When applying Theorem 1.7 to a given real-valued function f on tensors, it will depend very much on the regime we are working in whether condition (iii) is non-trivial or not. In particular, over finite fields, condition (iii) is trivial, because then $\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ contains only finitely many elements. Another trivial situation is when f takes only integral values, as there are only finitely many integers between 0 and $\min_i n_i$ (in this case the conclusion of the theorem is also trivial). However, when f is not integral (say f is the asymptotic slice rank or asymptotic subrank) and the field is not finite (say $\mathbb{F}$ is the complex numbers) condition (iii) can be very non-trivial. For instance, our application of Theorem 1.7 to asymptotic slice rank over the complex numbers (leading to Theorem 1.4) relies on the representation-theoretic characterization of this parameter in terms of moment polytopes [CVZ23] and a result from invariant theory that there are only finitely many such polytopes per choice of (n_1, n_2, n_3).

Our proof of Theorem 1.7 relies mainly on new lower bounds on the asymptotic subrank $\underset{\sim}{Q}(T)$ of concise tensors T, which we will discuss next. Intuitively, these lower bounds will ensure (using condition (ii)) that for any infinite sequence of tensors T_i the value of $f(T_i)$ gets "pushed up" so much that it either cannot converge, or eventually becomes constant (when $\min_i n_i$ is bounded).

1.2 Lower bounds on asymptotic subrank

Having discussed our discreteness theorems, we will now discuss two results that are central in the proof of the discreteness theorems, and of independent interest. These results are about lower bounds on the asymptotic subrank.

The general goal of these results, in the context of the proof of the general discreteness theorem, is to establish that if we have a sequence of tensors $T_i \in \mathbb{F}^{a_i} \otimes \mathbb{F}^{b_i} \otimes \mathbb{F}^{c_i}$ for $i \in \mathbb{N}$ such that the set of triples $\{(a_i, b_i, c_i) : i \in \mathbb{N}\}$ is infinite, then the asymptotic subrank of these tensors must be either unbounded, or eventually constant and integral. We will discuss this in detail in the proof of the general discreteness theorem.

1.2.1 Previous work

Strassen [Str91] building on the work of Coppersmith and Winograd [CW90] introduced a method to prove (optimal) asymptotic subrank lower bounds for a subclass of structured tensors called "tight" tensors. This method formed an integral part of the laser method for constructing fast matrix multiplication algorithms, and similar ideas have also been applied in the context of additive combinatorics [KSS18].

1.2.2 Concise tensors have large asymptotic subrank

For concise tensors T we prove a cube-root lower bound on the asymptotic subrank $\underset{\sim}{Q}(T)$ in terms of the smallest dimension of the tensor.

Theorem 1.8. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ be concise. Then $\underset{\sim}{Q}(T) \geq (\min_i n_i)^{1/3}$.*

We emphasize that Theorem 1.8 does not rely on any special structure of the tensor T, unlike previous methods like the laser method that rely on T being "tight".

We do not know whether the lower bound $(\min_i n_i)^{1/3}$ in Theorem 1.8 is optimal. The best upper bound we know is from an example of a concise tensor $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ such that $\underset{\sim}{Q}(T) = 2\sqrt{n-1}$ (Example 2.6).

Alternatively, Theorem 1.8 can be phrased without the conciseness condition if we replace the dimension n_i by the flattening rank $R(T_i)$, as follows: Let T be any tensor. Then $\underset{\sim}{Q}(T) \geq (\min_i R_i(T))^{1/3}$.

For symmetric tensors we can prove the following stronger bound. (In fact, we can prove this stronger bound for a larger class of tensors which we call "pivot–matched", as we will explain in Section 5.3) We recall that a tensor $T = \sum_{i,j,k} T_{i,j,k} e_i \otimes e_j \otimes e_k \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ is called symmetric if permuting the three tensor legs does not change the tensor, i.e. for every $(i,j,k) \in [n]^3$ and every permutation $\sigma : \{i,j,k\} \to \{i,j,k\}$ it holds that $T_{i,j,k} = T_{\sigma(i),\sigma(j),\sigma(k)}$.

Theorem 1.9. *Let $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ be concise and symmetric. Then $\underset{\sim}{\mathrm{Q}}(T) \geq n^{1/2}$.*

We do not know whether the lower bound $n^{1/2}$ in Theorem 1.9 is optimal.

Again, alternatively, Theorem 1.9 can be phrased without conciseness if the dimension n is replaced by the flattening rank $\mathrm{R}_1(T)$ (for symmetric tensors the three flattening ranks are equal): Let $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ be symmetric. Then $\underset{\sim}{\mathrm{Q}}(T) \geq \mathrm{R}_1(T)^{1/2}$.

1.2.3 Narrow tensors have maximal asymptotic subrank

Theorem 1.8 implies that if n_1, n_2 and n_3 all grow, and $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ is any concise tensor, then $\underset{\sim}{\mathrm{Q}}(T)$ must also grow. This leaves open what happens in the regime where one of the n_i is constant. We will consider the "narrow" regime where $n_3 = c$ is constant, and one of the dimensions n_1, n_2 is large enough. Here we prove that the asymptotic subrank is maximal, that is, matches the upper bound c.

Theorem 1.10. *For every integer $c \geq 2$ there is an $N(c) \in \mathbb{N}$ such that for every n_1, n_2 with $\max\{n_1, n_2\} > N(c)$ and every concise tensor $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^c$ we have $\underset{\sim}{\mathrm{Q}}(T) = c$.*

Moreover, for the case $c = 2$ we prove with a direct construction that $N(2) = 2$ and that asymptotic subrank can be replaced by subrank.

Theorem 1.11. *Let $n_1, n_2 > 2$. Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^2$ be concise. Then $\mathrm{Q}(T) = 2$.*

1.2.4 Lower bound on asymptotic subrank in terms of slice rank

Besides the aforementioned bounds on the asymptotic subrank, we use some of the same methods to prove a lower bound on the asymptotic subrank in terms of the asymptotic slice rank.

Slice rank was introduced by Tao [Tao16]. He proved that for every tensor T we have $\mathrm{SR}(T) \geq \mathrm{Q}(T)$. The gap between $\mathrm{SR}(T)$ and $\mathrm{Q}(T)$ can be large, namely for generic tensors $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ (over algebraically closed fields $\mathbb{F}$) it is known that $\mathrm{SR}(T) = n$ while $\mathrm{Q}(T) = \Theta(\sqrt{n})$ [DMZ22]. It is, however, not known whether there can be a large gap between $\mathrm{SR}(T^{\boxtimes m})$ and $\mathrm{Q}(T^{\boxtimes m})$ when m is large. In particular, it is possible that $\lim_{n \to \infty} \mathrm{Q}(T^{\boxtimes n})^{1/n} = \lim_{n \to \infty} \mathrm{SR}(T^{\boxtimes n})^{1/n}$. Strassen's work implies this equality for the subset of tight tensors [Str91]. Over the complex numbers, the limit $\lim_{n \to \infty} \mathrm{SR}(T^{\boxtimes n})^{1/n}$ is also known to exist and has a characterization in terms of moment polytopes [CVZ23].

We prove the following:

Theorem 1.12. *Suppose the limit $\underset{\sim}{\mathrm{SR}}(T) = \lim_{n \to \infty} \mathrm{SR}(T^{\boxtimes m})^{1/m}$ exists. Then*

$$\underset{\sim}{\mathrm{Q}}(T) \geq \underset{\sim}{\mathrm{SR}}(T)^{2/3}.$$

Our proof of Theorem 1.12 consists of proving that the border subrank of the third power of T is bounded from below in terms of the slice rank of T, and applying a field-agnostic Flanders-type lower bound on the max-rank of Haramaty and Shpilka [HS10].

10

1.3 Max-rank and min-rank of slice spans

Our lower bounds on asymptotic subrank discussed in the previous subsection rely on results we prove about the ranks of elements in the span of the slices of a tensor. These may be of independent interest and we discuss some of them here.

1.3.1 Max-ranks of slice spans

Our proof of Theorem 1.8 relies (among other ingredients) on the notion of the max-rank of matrix subspaces, and the relation between max-ranks of matrix subspaces obtained by slicing a tensor in the three possible directions.

For any matrix subspace $\mathcal{A} \subseteq \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$, we let max-rank($\mathcal{A}$) denote the largest matrix rank of any element of $\mathcal{A}$. To any tensor $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ we may associate three matrix subspaces $\mathcal{A}_1 \subseteq \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$, $\mathcal{A}_2 \subseteq \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ and $\mathcal{A}_3 \subseteq \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$, obtained by taking the span of the slices of T in each of the three possible directions. We denote the max-ranks of these spaces by $Q_i(T) = \text{max-rank}(\mathcal{A}_i)$, for $i \in [3]$. We prove that for a concise tensor T the max-ranks $Q_i(T)$ cannot all be small, in the following sense.

Theorem 1.13. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ be concise. Let $i, j, k \in [3]$ be distinct. Suppose that $|\mathbb{F}| > n_k$. Then $Q_i(T) Q_j(T) \geq n_k$.*

We have explicit examples of families of tensors (provided later) that show how Theorem 1.13 is essentially optimal:

- For every n that is a square, there is a concise tensor $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ such that for all i we have $\sqrt{n} \leq Q_i(T) \leq 2\sqrt{n}$ (Example 2.8). In particular, for all $i \neq j$ we have $n \leq Q_i(T) Q_j(T) \leq 4n$.

- For every n, there is a concise tensor $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ such that $Q_1(T) = Q_3(T) = n$ and $Q_2(T) = 2$, so that $Q_2(T) Q_3(T) = 2n$ (Example 2.6).

- For every c and for every n that is a multiple of c, there is a concise tensor $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ such that $Q_1(T) = n$, $Q_2(T) \leq c+1$ and $Q_3(T) \leq n/c+1$ (Example 2.7). In particular, $Q_2(T) Q_3(T) \leq \frac{(c+1)}{c}n + c + 1$.

It follows from Theorem 1.13 (by a straightforward argument) that $Q_i(T)$ must be large for at least two directions i, in the following sense:

Corollary 1.14. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ be concise. There are distinct $\ell_1, \ell_2 \in \{1, 2, 3\}$ such that $Q_{\ell_1}(T) \geq (\max_i n_i)^{1/2}$ and $Q_{\ell_2}(T) \geq (\min_i n_i)^{1/2}$.*

From Corollary 1.14 we can prove a preliminary asymptotic subrank lower bound $\underset{\sim}{Q}(T) \geq (\min_i n_i)^{1/4}$, using the (easy to prove) fact that $\underset{\sim}{Q}(T)^2 \geq Q_i(T) Q_j(T)$ for any distinct $i, j \in [3]$. Proving our stronger cube-root lower bound $\underset{\sim}{Q}(T) \geq (\min_i n_i)^{1/3}$ of Theorem 1.8 requires slightly more work.

Work on max-rank (and its relation to dimension) goes back to Dieudonné [Die49], Meshulam [Mes85] and [Fla62]. Another relevant line of work here is on commutative and non-commutative rank, which has established strong connections between max-rank and non-commutative rank [Coh95, FR04, DM18].

1.3.2 Min-ranks of slice spans

Our proof of Theorem 1.10 relies on a careful analysis of the min-rank of matrix subspaces, the relation to subrank and the behaviour of min-rank under powering.

For any matrix subspace $\mathcal{A} \subseteq \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$, we let min-rank($\mathcal{A}$) denote the smallest matrix rank of any nonzero element of $\mathcal{A}$. We prove several properties of the min-rank, of which we give a rough outline here (the precise description we defer to later):

- If the slices of a tensor have large min-rank, then the tensor has large subrank (Lemma 3.3).

- Any concise tensor has a slice of large rank (Lemma 3.6).

- If a matrix subspace has large max-rank, then we can transform the subspace in a natural fashion such that it has large min-rank and all elements are diagonal (Lemma 3.5).

- Min-rank is super-multiplicative under tensor product, as long as at least one of the matrix subspaces is diagonal (Lemma 3.14).

A careful combination of the above ingredients leads to a proof of Theorem 1.10.

The min-rank has been investigated before in several different contexts. Amitsur [Ami65] used min-rank to characterize properties of rings of operators. Meshulam and Semrl [MS04] used the min-rank to study properties of operator spaces. In quantum information the rank is a measure of entanglement. Spaces of bipartite states which are all entangled are spaces of matrices over the complex numbers with min-rank strictly greater than 1. They were investigated in [Wal02] and [Par04]. In [HLW06] a slightly different angle was taken, analysing random subspaces. It was shown that most random subspaces have almost maximal min-rank, which was used for superdense coding in [AHSW06], and was summarised in [Hay04]. Generalising both of these lines of work, in [CMW08] the question of dimension versus min-rank was addressed as number of qubits versus guaranteed entanglement in a subspace of states, and their construction is used in a follow-up paper [CHL+08] to show there are counterexamples to the additivity of p-Rényi entropies, for all $p \leq p_0$ for some small constant $p_0 < 1$, utilising the fact that the 0-Rényi entropy is the min-rank.

1.4 Tensor preliminaries

In this section we discuss some standard tensor preliminaries and terminology that we will use throughout the paper, following the papers of Strassen [Str87, Str88, Str91] and the book of Bürgisser, Clausen and Shokrollahi [BCS97].

Tensors

Let $\mathbb{F}$ be any field. For any $n \in \mathbb{N}$ we denote by $e_1, \ldots, e_n$ the standard basis elements of the vector space $\mathbb{F}^n$. For any $n_1, n_2, n_3 \in \mathbb{N}$, we denote by $\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ the vector space of order-three tensors $T = \sum_{i \in [n_1]} \sum_{j \in [n_2]} \sum_{k \in [n_3]} T_{i,j,k}\, e_i \otimes e_j \otimes e_k$ with coefficients $T_{i,j,k} \in \mathbb{F}$. In most of this paper we will be dealing only with order-three tensors, so instead of order-three tensor we will just say tensor.

Restriction and equivalence

We will be using a natural preordering on tensors called restriction, which is defined as follows. For any two tensors $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ and $S \in \mathbb{F}^{m_1} \otimes \mathbb{F}^{m_2} \otimes \mathbb{F}^{m_3}$, we say

T *restricts to* S and write $T \geq S$ if there are linear maps $L_i : \mathbb{F}^{n_i} \to \mathbb{F}^{m_i}$ such that $(L_1 \otimes L_2 \otimes L_3)T = S$. We say S and T are *equivalent* if $T \geq S$ and $S \geq T$. Most tensor properties we will study are invariant under equivalence and monotone under restriction (see Lemma 1.17).

Flattenings and flattening ranks

Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$. Grouping together two of the three tensor legs gives an element $T_i \in \mathbb{F}^{n_i} \otimes (\mathbb{F}^{n_j} \otimes \mathbb{F}^{n_k})$ (for any choice of distinct $i, j, k \in [3]$). We may think of T_i as an $n_i \times (n_j \cdot n_k)$ matrix. We call the T_i the *flattenings* of T. For every $i \in [3]$ we define $\mathrm{R}_i(T)$ as the matrix rank of T_i. We call the $\mathrm{R}_i(T)$ the *flattening ranks* of T. The triple $(\mathrm{R}_1(T), \mathrm{R}_2(T), \mathrm{R}_3(T))$ is often called the *multilinear rank* of T.

Conciseness

Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$. From the definition of flattening rank it follows that $\mathrm{R}_i(T) \leq n_i$ for every $i \in [3]$. We call T *concise* if $\mathrm{R}_i(T) = n_i$ for every $i \in [3]$. Conciseness essentially says that the tensor cannot fit in a smaller tensor space. The property of being concise is very mild in the following sense. We call a tensor S a *subtensor* of $T = \sum_{i,j,k} T_{i,j,k} e_i \otimes e_j \otimes e_k$ if S is of the form $S = \sum_{i \in I, j \in J, k \in K} T_{i,j,k} e_i \otimes e_j \otimes e_k$ for some subsets I, J, K.

Lemma 1.15. *Every tensor is equivalent to a concise tensor, which is unique up to equivalence. Concretely, every tensor is equivalent to a subtensor that is concise.*

Proof. The first statement is a standard fact [BCS97, Section 14.3]. For the second statement, let $T = \sum_{i,j,k} T_{i,j,k}\, e_i \otimes e_j \otimes e_k \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$. Let $A_i = (T_{i,j,k})_{j,k} \in \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ for $i \in [n_1]$ be the 1-slices of T. Then one verifies that the dimension of the span of the A_i equals the flattening rank $r_1 = \mathrm{R}_1(T)$. Suppose $A_1, \ldots, A_r$ are linearly independent, and let S be the tensor with these matrices as 1-slices. Then S is a subtensor of T so clearly $T \geq S$. Also, since the matrices $A_{r+1}, \ldots, A_{n_1}$ are in the span of the matrices $A_1, \ldots, A_r$ we have that $S \geq T$. Since S and T are thus equivalent, their flattening ranks are equal. Repeating this construction for the other two directions gives the claimed subtensor. $\qquad \square$

Not every space $\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ contains concise tensors, as can be seen from the following implication (which is in fact known to be an equivalence [CK11, Theorem 2]):

Lemma 1.16. *For any $n_1, n_2, n_3 \in \mathbb{N}$, if there is a concise tensor in $\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$, then $n_1 \leq n_2 \cdot n_3$, $n_2 \leq n_1 \cdot n_3$, and $n_3 \leq n_1 \cdot n_2$.*

Proof. Suppose $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ is concise. Then $n_i = \mathrm{rank}(T_i)$ where T_i is the previously defined flattening of T. Since T_i is an $n_i \times (n_j \cdot n_k)$ matrix, we have $\mathrm{R}(T_i) \leq n_j \cdot n_k$. Thus $n_i \leq n_j \cdot n_k$ for every choice of distinct $i, j, k \in [3]$. $\qquad \square$

Unit tensors, tensor rank, subrank, and slice rank

For every $r \in \mathbb{N}$ we define the tensor $\langle r \rangle = \sum_{i=1}^{r} e_i \otimes e_i \otimes e_i \in \mathbb{F}^r \otimes \mathbb{F}^r \otimes \mathbb{F}^r$. We call $\langle r \rangle$ the *unit tensor* of rank r. We say a tensor has *tensor rank one* if it is of the form $u \otimes v \otimes w$ for nonzero vectors u, v, w. The tensor rank of a tensor T is the smallest number r such that T can be written as a sum of r tensors with tensor rank one. We denote tensor rank by $\mathrm{R}(T)$. Equivalently, the tensor rank $\mathrm{R}(T)$ is the smallest number r such that $T \leq \langle r \rangle$. The *subrank* of T is the largest number r such that $\langle r \rangle \leq T$. We denote subrank by $\mathrm{Q}(T)$. We say a tensor has *slice rank one* if any of its flattening ranks equals one. The slice rank of a tensor T is the smallest number r such that T can be written as a sum of

r tensors with slice rank one. Subrank, slice rank, the flattening ranks and tensor rank satisfy $Q(T) \leq SR(T) \leq R_i(T) \leq R(T)$ for every $i \in [3]$.

It is also easy to see that these measures are monotone under restrictions:

Lemma 1.17. *For any two tensors T and S, if $T \geq S$, then we have $R(T) \geq R(S)$, $SR(T) \geq SR(S)$, $Q(T) \geq Q(S)$, and for every i, $Q_i(T) \geq Q_i(S)$ and $R_i(T) \geq R_i(S)$.*

Kronecker product and asymptotic tensor parameters

Asymptotic tensor parameters are defined by "regularizing" a tensor parameter over large powers under the Kronecker product. For any two tensors

$$T = \sum_{i,j,k} T_{i,j,k}\, e_i \otimes e_j \otimes e_k \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$$

$$S = \sum_{a,b,c} S_{a,b,c}\, e_a \otimes e_b \otimes e_c \in \mathbb{F}^{m_1} \otimes \mathbb{F}^{m_2} \otimes \mathbb{F}^{m_3}$$

the Kronecker product $T \boxtimes S \in \mathbb{F}^{n_1 m_1} \otimes \mathbb{F}^{n_2 m_2} \otimes \mathbb{F}^{n_3 m_3}$ is defined as

$$T \boxtimes S = \sum_{i,j,k} \sum_{a,b,c} T_{i,j,k} S_{a,b,c}\, (e_i \otimes e_a) \otimes (e_j \otimes e_b) \otimes (e_k \otimes e_c).$$

We define $T^{\boxtimes n}$ as the product of n copies of T. The asymptotic tensor rank is defined as $\underset{\sim}{R}(T) = \lim_{n \to \infty} R(T^{\boxtimes n})^{1/n}$, which, because of sub-multiplicativity of R, equals $\inf_{n \to \infty} R(T^{\boxtimes n})^{1/n}$ (by a standard fact called Fekete's lemma). The asymptotic subrank is defined as $\underset{\sim}{Q}(T) = \lim_{n \to \infty} Q(T^{\boxtimes n})^{1/n}$, which, because of super-multiplicativity of the subrank Q, equals $\sup_{n \to \infty} Q(T^{\boxtimes n})^{1/n}$. Over the complex numbers, the asymptotic slice rank is defined as $\underset{\sim}{SR}(T) = \lim_{n \to \infty} SR(T^{\boxtimes n})^{1/n}$. By the characterization in [CVZ23] this limit exists. Over other fields, we generally do not know whether the limit $\lim_{n \to \infty} SR(T^{\boxtimes n})^{1/n}$ exists, except for oblique tensors. When we do not know whether the limit exists we will instead consider the liminf and refer to this as the "asymptotic slice rank".[4]

2 Max-rank of slice spans and asymptotic subrank lower bound

In this section we will prove a lower bound on the asymptotic subrank. The main goal is to obtain a lower bound on the asymptotic subrank of concise tensors that only depends on the dimensions of the tensor and that moreover grows as the dimensions grow. Unlike previous lower bound methods (like the laser method mentioned in the introduction) our lower bound does not rely on any special structure of the tensor. Indeed the property of being concise is very mild as any tensor can be made concise by embedding it in the smallest tensor space that it fits in.

Concretely we will prove that the asymptotic subrank of a concise tensor is at least the cube-root of the smallest dimension of the tensor. We will give two proofs of this in this paper, one in this section and another one in Section 5.2 (which uses different ideas which we then use to get stronger results for symmetric tensors).

The proof of the asymptotic subrank lower bound in this section relies on the notion of max-rank of a matrix subspace, which is simply the largest rank of any matrix in the subspace (Section 2.1). Any tensor naturally defines three matrix subspaces by slicing

---

[4]The results we have over finite fields for asymptotic slice rank hold equally well using liminf or limsup.

the tensor in the three different directions and taking spans. We prove that there is a strong relation between the max-ranks of such subspaces. Namely, the product of every pair of them is large (Section 2.2). This result, combined with known facts about matrix multiplication tensors, leads to the asymptotic subrank bound (Section 2.4).

2.1   Max-rank of matrix subspaces and slice spans of a tensor

We denote by $\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ the space of $n_1 \times n_2$ matrices over $\mathbb{F}$. We define the *max-rank* of any matrix subspace $\mathcal{A} \subseteq \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ as

$$\text{max-rank}(\mathcal{A}) := \max\{\text{rank}(A) : A \in \mathcal{A}\}.$$

For any tensor $T = \sum_{i,j,k} T_{i,j,k}\, e_i \otimes e_j \otimes e_k \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ we define the matrices

$$T_i^{(1)} = \sum_{j,k} T_{i,j,k}\, e_j \otimes e_k \quad (i \in [n_1])$$

$$T_j^{(2)} = \sum_{i,k} T_{i,j,k}\, e_i \otimes e_k \quad (j \in [n_2])$$

$$T_k^{(3)} = \sum_{i,j} T_{i,j,k}\, e_i \otimes e_j \quad (k \in [n_3]).$$

For every $\ell \in [3]$, we call $\{T_i^{(\ell)} : i \in [n_\ell]\}$ the *ℓ-slices* of T. We define the ℓ-slice span of T as the matrix subspace $\mathcal{A}_\ell = \text{span}\{T_i^{(\ell)} : i \in [n_\ell]\}$ and denote its max-rank by

$$Q_\ell(T) = \text{max-rank}(\mathcal{A}_\ell).$$

We note that the flattening rank $R_\ell(T)$ defined earlier, is equal to the dimension of the subspace $\mathcal{A}_\ell$.

2.2   Lower bound on products of max-ranks of slice spans

In this section we prove a fundamental property of the max-ranks of the slice spans of a tensor, showing that they cannot all be small. This will be a crucial ingredient for proving the asymptotic subrank lower bound afterwards.

Theorem 2.1. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ be concise. Let $i, j, k \in [3]$ be distinct. Suppose that $|\mathbb{F}| > n_k$. Then*
$$Q_i(T)\, Q_j(T) \geq n_k.$$

A corollary of this is that there are two distinct directions with large max-rank:

Corollary 2.2. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ be concise. Suppose $n_1 \leq n_2 \leq n_3$. Suppose that $|\mathbb{F}| > n_3$. Then there are distinct $i, j \in [3]$ such that*
$$Q_i(T) \geq \sqrt{n_3},$$
$$Q_j(T) \geq \sqrt{n_1}.$$

Proof. We have by Theorem 2.1 that

$$Q_1(T)\, Q_2(T) \geq n_3$$
$$Q_2(T)\, Q_3(T) \geq n_1$$
$$Q_1(T)\, Q_3(T) \geq n_2 \geq n_1.$$

Thus one of $Q_1(T)$ and $Q_2(T)$ is at least $\sqrt{n_3}$, and either $Q_3(T)$ is at least $\sqrt{n_1}$ or both $Q_1(T)$ and $Q_2(T)$ are at least $\sqrt{n_1}$. $\qquad\square$

We will now work towards the proof of Theorem 2.1.

For matrices $A \in \mathbb{F}^{n \times m}$ and $B \in \mathbb{F}^{n \times k}$ we denote by $[A; B] \in \mathbb{F}^{n \times (m+k)}$ the concatenation of A and B. For any matrix A we denote by $A|_{[d]}$ the matrix consisting of the first d columns of A. We denote by $\mathrm{colspan}(A)$ the span of the columns of A.

Lemma 2.3. *Let $A \in \mathbb{F}^{n \times m}$ and $B \in \mathbb{F}^{n \times k}$. There is a nonzero k^2-variable polynomial $p \in \mathbb{F}[x_{1,1}, x_{1,2}, \ldots, x_{k,k}]$ of degree $s = \mathrm{rank}([A; B]) - \mathrm{rank}(A)$ such that for any matrix $U \in \mathbb{F}^{k \times k}$ satisfying $p(U) \neq 0$, we have that*

$$\mathrm{colspan}([A; B]) = \mathrm{colspan}([A; (BU)|_{[s]}]).$$

Proof. Let $r = \mathrm{rank}([A; B])$. Then, there is an $r \times r$ submatrix of $[A; B]$ with rank r induced by $r - s$ columns of A and s columns of B. Let $R \subseteq [n]$ be the set of row indices of this submatrix and $C_A \subseteq [m], C_B \subseteq \{m+1, \ldots, m+k\}$ be the respective column indices of A and B in this submatrix. Let $p(X)$ be the the determinant of the submatrix of $[A; BX]$ induced by the rows in R and columns in $C_A \cup \{m+1, \ldots, m+s\}$, where $X = (x_{i,j})_{i,j=1}^{k}$ is a matrix of variables.

We claim that p is nonzero. There is a permutation matrix $V \in \mathbb{F}^{k \times k}$ that swaps the columns in C_B with the columns in $\{m+1, \ldots, m+s\}$. Since the submatrix of $[A; BV]$ induced by the rows in R and columns in $C_A \cup \{m+1, \ldots, m+s\}$ is then of full rank, it follows that $p(V) \neq 0$.

Let $U \in \mathbb{F}^{k \times k}$ be any matrix such that $p(U) \neq 0$. Then $\mathrm{rank}([A; (BU)|_{[s]}]) \geq r$. From $\mathrm{colspan}([A; (BU)|_{[s]}]) \subseteq \mathrm{colspan}([A; B])$ and $\mathrm{rank}([A; B]) = r$ it follows that $\mathrm{colspan}([A; B]) = \mathrm{colspan}([A; (BU)|_{[s]}])$. $\square$

Lemma 2.4. *Let $A_1, \ldots, A_\ell \in \mathbb{F}^{n \times k}$ be matrices. Suppose $|\mathbb{F}| > \mathrm{rank}([A_1; \cdots; A_\ell])$. Let*

$$s_i = \mathrm{rank}([A_1; \cdots; A_i]) - \mathrm{rank}([A_1; \cdots; A_{i-1}]).$$

There is a matrix $U \in \mathbb{F}^{k \times k}$ such that for all $i \in [\ell]$ we have

$$\mathrm{colspan}([A_1; A_2; \cdots; A_{i-1}; (A_i U)|_{[s_i]}]) = \mathrm{colspan}([A_1; \cdots; A_i]).$$

Proof. From Lemma 2.3, we get that for each $i \in [\ell]$, there exists a nonzero k^2-variable polynomial $p_i(X) \in \mathbb{F}[x_{1,1}, x_{1,2}, \ldots, x_{k,k}]$ of degree s_i such that for any $U \in \mathbb{F}^{k \times k}$ satisfying $p_i(U) \neq 0$,

$$\mathrm{colspan}([A_1; A_2; \cdots; A_{i-1}; (A_i U)|_{[s_i]}]) = \mathrm{colspan}([A_1; \cdots; A_i]).$$

Let $q(X) = p_1(X) \cdots p_\ell(X)$. This is a nonzero polynomial of degree $s_1 + \cdots + s_\ell = \mathrm{rank}([A_1; \cdots; A_\ell]) < |\mathbb{F}|$. It follows from the Schwartz–Zippel lemma [DL78, Sch80, Zip79] that there is a $U \in \mathbb{F}^{k \times k}$ such that $q(U) \neq 0$. This is to say that for all $i \in [\ell]$ simultaneously, $p_i(U) \neq 0$, which implies the result. $\square$

Lemma 2.5. *Let $A_1, \ldots, A_\ell \in \mathbb{F}^{n \times k}$ be matrices. Suppose $|\mathbb{F}| > \mathrm{rank}([A_1; \cdots; A_\ell])$. Let*

$$s_i = \mathrm{rank}([A_1; \cdots; A_i]) - \mathrm{rank}([A_1; \cdots; A_{i-1}]).$$

There is a matrix $U \in \mathbb{F}^{k \times k}$ such that

$$\mathrm{colspan}([(A_1 U)|_{[s_1]}; (A_2 U)|_{[s_2]}; \cdots; (A_\ell U)|_{[s_\ell]}]) = \mathrm{colspan}([A_1; \cdots; A_\ell]).$$

Proof. Let $U \in \mathbb{F}^{k \times k}$ be a matrix as in Lemma 2.4. We give a proof by induction. The base case is $\text{colspan}((A_1 U)|_{[s_1]}) = \text{colspan}(A_1)$, which follows directly from Lemma 2.4. For the induction step, we assume

$$\text{colspan}([A_1; A_2; \cdots ; A_i]) = \text{colspan}([(A_1 U)|_{[s_1]}; \cdots ; (A_i U)|_{[s_i]}]). \tag{1}$$

We now use that if $\text{colspan}(A) = \text{colspan}(B)$, then $\text{colspan}([A;C]) = \text{colspan}([B;C])$ holds for any matrices A, B, C with the same number of rows. Applying this to (1) gives

$$\begin{aligned}
\text{colspan}([A_1; A_2; \cdots ; A_i; (A_{i+1} U)|_{[s_{i+1}]}]) & \\
&= \text{colspan}([(A_1 U)|_{[s_1]}; \cdots ; (A_i U)|_{[s_i]}; (A_{i+1} U)|_{[s_{i+1}]}]). \tag{2}
\end{aligned}$$

From Lemma 2.4 we know that

$$\text{colspan}([A_1; \cdots ; A_{i+1}]) = \text{colspan}([A_1; A_2; \cdots ; A_i; (A_{i+1} U)|_{[s_{i+1}]}]). \tag{3}$$

Then (2) and (3) together give

$$\text{colspan}([A_1; A_2; \cdots ; A_i; A_{i+1}]) = \text{colspan}([(A_1 U)|_{[s_1]}; \cdots ; (A_i U)|_{[s_i]}; (A_{i+1} U)|_{[s_{i+1}]}]).$$

This proves the claim by induction. $\square$

Proof of Theorem 2.1. Let $A_1, \ldots, A_{n_3} \in \mathbb{F}^{n_1 \times n_2}$ be the 3-slices of T. Let U be as in Lemma 2.5. Let

$$s_i = \text{rank}([A_1; \cdots ; A_i]) - \text{rank}([A_1; \cdots ; A_{i-1}]).$$

Conciseness of T gives that $s_1 + \cdots + s_{n_3} = n_1$. It follows from this and Lemma 2.5 that the columns of the submatrices $(A_i U)|_{[s_i]}$ are linearly independent. This implies that for each $i \in [n_3]$, there is a 3-slice of rank at least s_i, and in particular $Q_3(T) \geq \max_i s_i$. Let S be the tensor with 3-slices $A_1 U, A_2 U, \ldots, A_{n_3} U$. Then $Q_2(S) \leq Q_2(T)$ since the 2-slices of S are linear combinations of the 2-slices of T. The matrix $M = [(A_1 U)|_{[1]}; \cdots ; (A_{n_3} U)|_{[1]}]$ consisting of every first column of the 3-slices of S equals the first 2-slice of S. Since the linearly independent columns in the 3-slices of S are flushed to the left, we find that M has at least $|\{i \in [n_3] : s_i \neq 0\}|$ linearly independent columns, and thus $\text{rank}(M) \geq |\{i \in [n_3] : s_i \neq 0\}|$, which gives a lower bound on $Q_2(S)$ and thus on $Q_2(T)$. Since

$$n_1 = s_1 + \cdots + s_{n_3} \leq |\{i \in [n_3] : s_i \neq 0\}| \max_i s_i,$$

the result follows. $\square$

2.3 Examples of optimality of max-rank bounds

The following examples show that the bound in Theorem 2.1 is tight. In particular, the right hand side cannot be super-linear in n_k, and in fact the leading coefficient 1 is optimal.

Example 2.6 (Null-algebra [Str91, p. 168])**.** For every $n \in \mathbb{N}$, there is a concise tensor $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ such that $Q_1(T) = Q_3(T) = n$ and $Q_2(T) = 2$. In particular, $Q_2(T) Q_3(T) = 2n$. To define this tensor, let

$$T := e_1 \otimes e_1 \otimes e_1 + \sum_{i=2}^n e_1 \otimes e_i \otimes e_i + \sum_{i=2}^n e_i \otimes e_i \otimes e_1.$$

We see directly that $Q_1(T) = Q_3(T) = n$. Because the non-zero elements of the 2-slices are confined to the first column and the first row, we have $Q_2(T) = 2$. We also mention that for all $n \geq 5$ the tensor T has asymptotic subrank $\underset{\sim}{Q}(T) = 2\sqrt{n-1}$ [Str91, p. 169].

Example 2.7. This example extends Example 2.6 (by taking $c = 1$). Let c be a constant and let n be any multiple of c. There is a concise tensor $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ such that $Q_1(T) = n$, $Q_2(T) \leq c+1$ and $Q_3(T) \leq n/c+1$. In particular,

$$Q_2(T) Q_3(T) \leq \tfrac{(c+1)}{c} n + c + 1.$$

To define this tensor, let

$$T := \sum_{i=1}^{n} e_1 \otimes e_i \otimes e_i + \sum_{k=1}^{c} \sum_{j=1}^{n/c} e_{(k-1)\frac{n}{c}+j} \otimes e_j \otimes e_k$$

$$= \sum_{i=1}^{n} e_1 \otimes e_i \otimes e_i + \sum_{i=0}^{n-1} e_{i+1} \otimes e_{(i \bmod \frac{n}{c})+1} \otimes e_{\lfloor i/(\frac{n}{c})\rfloor+1}.$$

We will prove that $Q_2(T) \leq c+1$ and $Q_3(T) \leq n/c+1$. Looking at direction 2 (interpreting direction 1 as enumerating the rows and 3 as enumerating the columns), we can see all slices are restricted to the first row (first sum) or the first c columns (the other sum) for non-zero entries, so each matrix in their span, is of rank $Q_2(T) \leq c+1$. Similarly, all slices in direction 3 (interpreting direction 2 as enumerating the rows) are restricted to the first n/c rows or the first column, meaning $Q_3(T) \leq n/c+1$.

Example 2.8. Let n be a square. There is a concise tensor $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ such that for all i we have $\sqrt{n} \leq Q_i(T) \leq 2\sqrt{n}$. In particular, for all $i \neq j$ we have $n \leq Q_i(T) Q_j(T) \leq 4n$.

To define this tensor, let $f : [n] \to [\sqrt{n}]$ and $g : [n] \to [\sqrt{n}]$ be any two functions such the function $(f, g) : [n] \to [\sqrt{n}] \times [\sqrt{n}] : i \mapsto (f(i), g(i))$ is injective and such that for every $i \in [\sqrt{n}]$ we have $f(i) = g(i) = i$. Let

$$T := \sum_{i=1}^{\sqrt{n}} e_i \otimes e_i \otimes e_i + \sum_{i=\sqrt{n}+1}^{n} e_i \otimes e_{f(i)} \otimes e_{g(i)} + e_{f(i)} \otimes e_i \otimes e_{g(i)} + e_{f(i)} \otimes e_{g(i)} \otimes e_i$$

Then T has a diagonal subtensor of size $\sqrt{n}$, so $\sqrt{n} \leq Q(T) \leq Q_i(T)$. On the other hand, $Q_i(T) \leq SR_i(T) \leq 2\sqrt{n}$ as the support of all its slices is confined to the first $\sqrt{n}$ rows and the first $\sqrt{n}$ columns. From injectivity of (f, g) it follows that T is concise.

2.4 Lower bound on the asymptotic subrank via max-ranks and matrix multiplication tensors

We will now prove a lower bound on the asymptotic subrank of concise tensors in terms of the dimensions of the tensor. We will use the max-rank inequality (Theorem 2.1) proven before and basic properties of the matrix multiplication tensors that we will discuss now. The main result here is the cube-root lower bound $Q(T) \geq (\min_i n_i)^{1/3}$ for concise tensors $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ (Theorem 2.12). With a simpler argument we can already obtain the bound $\tilde{Q}(T) \geq (\min_i n_i)^{1/4}$ (Theorem 2.11), which we will prove first. This fourth-root lower bound is already sufficient for our use in the proof of the discreteness theorem. In Section 5 we will prove an even better square-root lower bound under a symmetry assumption.

For $a, b, c \in \mathbb{N}$ we denote by $\langle a, b, c \rangle \in \mathbb{F}^{ab} \otimes \mathbb{F}^{bc} \otimes \mathbb{F}^{ca}$ the matrix multiplication tensor $\langle a, b, c \rangle = \sum_{i=1}^{a} \sum_{j=1}^{b} \sum_{k=1}^{c} e_{i,j} \otimes e_{j,k} \otimes e_{k,i}$. These tensors have the multiplicative property $\langle a, b, c \rangle \otimes \langle x, y, z \rangle = \langle ax, by, cz \rangle$. They are naturally related to the parameters Q_i as follows:

Lemma 2.9. *Let T be a tensor and $r \in \mathbb{N}$. The following holds:*

- $Q_1(T) \geq r$ if and only if $T \geq \langle 1, 1, r \rangle$.

- $Q_2(T) \geq r$ if and only if $T \geq \langle r, 1, 1 \rangle$.

- $Q_3(T) \geq r$ if and only if $T \geq \langle 1, r, 1 \rangle$.

Lemma 2.10. *If there are distinct i, j such that $Q_i(T), Q_j(T) \geq r$, then $Q(T^{\boxtimes 2}) \geq r$.*

Proof. Suppose $Q_1(T), Q_2(T) \geq r$ (the proof for the other cases goes the same). Then we have $T \geq \langle 1, 1, r \rangle = \sum_{i=1}^{r} e_1 \otimes e_i \otimes e_i$ and $T \geq \langle r, 1, 1 \rangle = \sum_{i=1}^{r} e_i \otimes e_1 \otimes e_i$. We multiply these inequalities to get $T^{\boxtimes 2} \geq \langle r, 1, r \rangle = \sum_{i,j=1}^{r} e_i \otimes e_j \otimes (e_i \otimes e_j)$. We apply the projection $(e_i \otimes e_j) \mapsto \delta_{i=j} e_i$ to the third tensor leg to get $T^{\boxtimes 2} \geq \sum_{i=1}^{r} e_i \otimes e_i \otimes e_i$ which proves the claim. $\qquad\square$

We are now ready to prove the fourth-root lower bound on the asymptotic subrank. After that we will prove the stronger cube-root lower bound (Theorem 2.12).

Theorem 2.11. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ be concise. Then $\underset{\sim}{Q}(T) \geq (\min_i n_i)^{1/4}$.*

Proof. The proof is an application of Theorem 2.1. That theorem only holds over large enough fields however. Since the asymptotic subrank does not change under field extension [Str88, Theorem 3.10], we may assume that our base field $\mathbb{F}$ is infinite (by working over the algebraic closure of the original field) so that the field size condition of Theorem 2.1 is satisfied. We use that there are distinct $\ell_1, \ell_2 \in [3]$ such that $Q_{\ell_1}(T), Q_{\ell_2}(T) \geq \min_i n_i$ (Corollary 2.2). Then $\underset{\sim}{Q}(T^{\boxtimes 2}) \geq \min_i n_i$ (Lemma 2.10) which gives the claim. $\qquad\square$

It turns out that the above construction is not optimal. We will now give a slightly more elaborate construction that leads to the better cube-root bound:

Theorem 2.12. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ be concise. Then $\underset{\sim}{Q}(T) \geq (\min_i n_i)^{1/3}$.*

The proof of Theorem 2.12 makes use of basic properties of matrix multiplication tensors which we discuss in the following two lemmas. Recall that we use the notation $\langle a, b, c \rangle$ for $a, b, c \in \mathbb{N}$ to denote the matrix multiplication tensors.

Lemma 2.13. $T^{\boxtimes 3} \geq \langle Q_2(T), Q_3(T), Q_1(T) \rangle$.

Proof. From Lemma 2.9 we have three inequalities $T \geq \langle Q_2(T), 1, 1 \rangle$, $T \geq \langle 1, Q_3(T), 1 \rangle$, and $T \geq \langle 1, 1, Q_1(T) \rangle$. We multiply these inequalities using the Kronecker product to get the claim. $\qquad\square$

Lemma 2.14. $\underset{\sim}{Q}(T^{\boxtimes 3}) \geq \min_{i,j} Q_i(T) Q_j(T)$.

Proof. It is known that for every $a, b, c \in \mathbb{N}$ the asymptotic subrank of the matrix multiplication tensor $\langle a, b, c \rangle$ is given by $\underset{\sim}{Q}(\langle a, b, c \rangle) = \min\{ab, bc, ac\}$ [Str88, Equation 4.6]. We combine this with Lemma 2.13 to get the claim. $\qquad\square$

Proof of Theorem 2.12. As in the proof of Theorem 2.11 we may assume that the base field $\mathbb{F}$ is infinite. We use $\underset{\sim}{Q}(T^{\boxtimes 3}) \geq \min_{i,j} Q_i(T) Q_j(T)$ (Lemma 2.14) and we use that for every distinct $i, j, k \in [3]$ we have $Q_i(T) Q_j(T) \geq n_k$ (Theorem 2.1) to obtain $\underset{\sim}{Q}(T^{\boxtimes 3}) \geq \min_{i,j} Q_i(T) Q_j(T) \geq \min_k n_k$. $\qquad\square$

19

3 Min-rank of slice spans and asymptotic subrank of narrow tensors

In Section 2 we lower bounded the asymptotic subrank of a concise tensor by a growing function in the smallest local dimension (Theorem 2.12). Namely, we showed that for every concise $n_1 \times n_2 \times n_3$ tensor the asymptotic subrank is at least $(\min_i n_i)^{1/3}$. We will use that result as an ingredient to prove the discreteness theorem in Section 4. However, that bound will not be enough as we must also understand how asymptotic subrank behaves when the smallest dimension $\min_i n_i$ is fixed (we elaborate on this point in Section 4). This is the topic of this section.

To briefly explain what is missing, we note that it is a priori possible that there is a sequence of concise $c \times n_1 \times n_2$ tensors with c constant and at least one of n_1, n_2 growing[5] such that the asymptotic subrank (or asymptotic slice rank) has an accumulation point. Our result will rule this out in the appropriate sense. In particular, we prove (Theorem 3.18) that for every $c \geq 2$ there is a number $N(c)$ such that for every n_1, n_2 with $\max\{n_1, n_2\} > N(c)$, for every concise $c \times n_1 \times n_2$ tensor, the asymptotic subrank equals c. This will be a crucial ingredient in our proof of the discreteness theorem in Section 4.

Whereas the proof of the asymptotic subrank lower bound in Section 2 relied for a large part on the max-rank of matrix subspaces, here we will need the related notion of min-rank of a matrix subspace. In Section 3.1 we introduce the notion of min-rank, which we will use throughout. In Section 3.1.1 we prove that if the slices of a concise tensor have large min-rank, then the tensor has large subrank. In Section 3.1.2 we see how from matrices with large max-rank we can construct matrices with large min-rank that are moreover diagonal. In Section 3.1.3 we then prove that min-rank is super-multiplicative (when at least one of the matrix subspaces is diagonal), and finally we generalize the super-multiplicativity property to make it work with the output of Section 3.1.2.

In Section 3.2 we combine the lemmas on min-rank to prove the main result (Theorem 3.18). Finally, as an extra, in Section 3.3, we determine precisely what happens in the smallest case $c = 2$, and in particular determine that $N(2) = 2$.

To avoid potential confusion we remark that throughout we will call any matrix $A \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ (that is not necessarily square) *diagonal* if its support, which is defined as $\mathrm{supp}(A) := \{(i,j) : A_{i,j} \neq 0\}$, is a subset of $\{(i,i) : 1 \leq i \leq \min\{n_1, n_2\}\}$.

3.1 Min-rank of matrix subspaces and its properties

The min-rank of a matrix subspace is simply the smallest rank of any non-zero matrix in the space.

Definition 3.1. We define the *min-rank* of a matrix subspace $\mathcal{A} \subseteq \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ as

$$\text{min-rank}(\mathcal{A}) := \min\{\mathrm{rank}(A) : A \in \mathcal{A} \setminus \{0\}\}.$$

For any collection of matrices $A_1, \ldots, A_c \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ we define

$$\text{min-rank}(A_1, \ldots, A_c) := \text{min-rank}(\mathrm{span}\{A_1, \ldots, A_c\}).$$

We will use the following straightforward equivalence throughout this section. Given matrices $A_1, \ldots, A_c \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ we call a linear combination $\alpha_1 A_1 + \cdots + \alpha_c A_c$ with coefficients $\alpha_i \in \mathbb{F}$ *nontrivial* if at least one of the coefficients α_i is nonzero.

---

[5]We note that in this context, since we consider only concise tensors, if c is constant and at least one of n_1, n_2 is growing, then both n_1 and n_2 must be growing by Lemma 1.16.

Lemma 3.2. *Let $A_1, \ldots, A_c \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ be matrices. For every $r \geq 1$, the following two statements are equivalent:*

- *Every nontrivial linear combination of the $A_1, \ldots, A_c$ is of rank at least r.*

- *The matrices $A_1, \ldots, A_c$ are linearly independent and $\text{min-rank}(A_1, \ldots, A_c) \geq r$.*

3.1.1 Large min-rank implies large subrank

In this section we prove that if the slices of a tensor have large min-rank, then the tensor has large subrank, in the following sense:

Lemma 3.3. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^c$ be a tensor with 3-slices $A_1, \ldots, A_c \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$. If $A_1, \ldots, A_c$ are linearly independent and $\text{min-rank}(A_1, \ldots, A_c) \geq 10c^2$, then $\text{Q}(T) \geq c$.*

Remark 3.4. It is immediate from Lemma 3.3 that for any tensor $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$, if any c linearly independent 3-slices $A_1, \ldots, A_c$ of T satisfy $\text{min-rank}(A_1, \ldots, A_c) \geq 10c^2$, then $\text{Q}(T) \geq c$.

Proof. The proof is by induction on c. For the base case $c = 1$, note that A_1 is nonzero by linear independence, and so T is a nonzero tensor, which implies $\text{Q}(T) \geq 1$.

For the induction step, let $c > 1$. Consider A_1, which (again) is nonzero by linear independence. Diagonalize it with invertible row and column operations (Gaussian elimination) such that $(A_1)_{1,1} = 1$. Specifically note that besides the $(1,1)$ entry, all other entries in the first column and first row of A_1 are zero. We can also make sure that A_1 is the only slice with a nonzero $(1,1)$ entry by subtracting a multiple of it from the other slices.

Next, for every $i > 1$, we would like to make the first column of A_i zero. As we will explain, we will not be able to achieve quite this, but we will get close: every such first column will be zero except for a fixed set of coordinates, of size at most $c - 1$. These will correspond to rows which we will remove from the tensor by restriction.

We already have that $(A_2)_{1,1} = 0$. If there is a row $i_2 > 1$ such that $(A_2)_{i_2,1} \neq 0$, then by doing invertible row operations that involve subtracting multiples of row i_2 from the other rows (except the first one, which already has $(A_2)_{1,1} = 0$), we can make sure that all entries in the first column of A_2 are zero, except the $(i_2, 1)$ entry. Note that these row operations did not affect any column of A_1 except column i_2, because for every $j \neq i_2$ we have $(A_1)_{i_2,j} = 0$ since A_1 is diagonal. The row operations also did not affect the first row of any A_i. We let $I_R = \{i_2\}$.

Repeat this procedure for $A_3, \ldots, A_c$, ignoring the indices in I_R, i.e. those rows that we already picked to zero-out the first column of the previous slices. Each such step adds at most one row-index to I_R.

We next repeat a similar procedure to make rows zero (the first row of each of $A_2, \ldots, A_c$). We ignore those indices already in I_R. Similarly, this results in a set of column indices, I_C, of size at most $c - 1$.

Thus, there is a set $I = I_R \cup I_C$, of size $|I| \leq 2(c - 1)$, such that $1 \notin I$ and such that if we restrict our matrices to coordinates $([n_1] \setminus I) \times ([n_2] \setminus I)$ then A_1 is diagonal and has 1 in its $(1,1)$ entry, and each A_i, $i > 1$, has 0 in the first row and column. Consider any non-trivial linear combination of $\{A_2, \ldots, A_c\}$. By assumption and Lemma 3.2 we know its rank is at least $10c^2$. Note that as $|I| \leq 2(c - 1)$, restricted to $([n_1] \setminus I) \times ([n_2] \setminus I)$ the rank of this linear combination must be at least $10c^2 - 4(c - 1) = 10c^2 - 4c + 4 \geq 10c^2 - 20c + 10 = 10(c - 1)^2$. (Indeed, thinking of the rank as the dimension of the row-span shows deleting $2(c - 1)$ rows can decrease the rank by at most $2(c - 1)$, and

21

similarly with the columns.) Denote the restriction of the A_i matrix to these coordinates by $A_i' := A_i|_{([n_1]\backslash I) \times ([n_2]\backslash I)}$. Then $A_2', \ldots, A_c'$ are linearly independent (since any non-trivial linear combination of them has nonzero rank) and min-rank$(A_1, \ldots, A_c) > 10(c-1)^2$, so $\{A_2', \ldots, A_c'\}$ satisfies the condition for $c-1$. By the induction hypothesis there is a restriction of tensors $(A_2', \ldots, A_c') \geq \langle c-1 \rangle$ (the tensor which has $A_i', i > 1$ as its 1-slices). All these matrices are zero in their first row and column, so using the fact that A_1' is diagonal with a nonzero $(1,1)$-entry, we get a restriction $T \geq (A_1, \ldots, A_c) \geq (A_1', \ldots, A_c') \geq \langle c \rangle$, and so $Q(T) \geq c$. $\qquad \square$

3.1.2 From large max-rank to large min-rank and diagonal

In this section we discuss how from any collection of matrices that contains *at least one* matrix of large rank (large max-rank) we can produce a collection of slightly smaller matrices such that *all* nonzero matrices in their span have large rank (large min-rank). Moreover, the resulting matrices are all diagonal, which will be crucial in our application.

For any matrix $A \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ and subsets $I \subseteq [n_1]$, $J \subseteq [n_2]$ we denote by $A_{I \times J}$ the submatrix of A with rows indexed by I and columns indexed by J. For any subset $X \subseteq \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ we denote by $X_{I \times J}$ the set $\{A_{I \times J} : A \in X\}$.

Lemma 3.5. *For every $c \in \mathbb{N}$, there is an $\varepsilon(c) > 0$ such that the following holds. Let $\mathcal{A} \subseteq \mathbb{F}^{n_1 \times n_2}$ be a matrix subspace of dimension c. Then, there exist $U \in \mathrm{GL}(\mathbb{F}^{n_1})$, $V \in \mathrm{GL}(\mathbb{F}^{n_2})$, a subset $J \subseteq [\min_i n_i]$, a basis $B_1, \ldots, B_c$ for the space $U\mathcal{A}V$ and $b \in [c]$ such that:*

- *$(B_1)_{J \times J}, \ldots, (B_b)_{J \times J}$ are diagonal and linearly independent;*

- *min-rank$\big((B_1)_{J \times J}, \ldots, (B_b)_{J \times J}\big) \geq \varepsilon(c) \, \mathrm{max\text{-}rank}(\mathcal{A})$;*

- *$(B_{b+1})_{J \times J} = \cdots = (B_c)_{J \times J} = 0$.*

Before discussing the proof in detail, as a motivation and ingredient for later, we observe using the pigeonhole principle that any concise tensor has at least one slice of large rank.

Lemma 3.6. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^c$ be a concise tensor. Then T has at least one 3-slice of rank at least $\max\{n_1, n_2\}/c$.*

Proof. Suppose $n_1 \geq n_2$. Concatenate the 3-slices $A_1, \ldots, A_c$ of T to an $n_1 \times cn_2$ matrix. Conciseness of T implies that the rank of this matrix is n_1 so there are n_1 linearly independent columns. By the pigeonhole principle, there is an $i \in [c]$ such that A_i contains at least n_1/c linearly independent columns. $\qquad \square$

Diagonalization. The proof of Lemma 3.5 has three parts. In this first part, we show in Lemma 3.9 that any collection of matrices can be made simultaneously diagonal on a large principal submatrix in such a way that the property of having large max-rank is preserved.

Lemma 3.7. *Let $A \in \mathbb{F}^{n \times n}$. Then there are $U, V \in \mathrm{GL}(\mathbb{F}^n)$ and $i, j \in \{2, \ldots, n\}$ such that for $I = [n] \backslash \{i, j\}$,*

(i) $(UAV)_{I \times I} = \alpha_{1,1} E_{1,1} + \sum_{a,b \geq 2} \alpha_{a,b} E_{a,b}$ for some $\alpha_{a,b} \in \mathbb{F}$;

(ii) for every diagonal matrix $B \in \mathbb{F}^{n \times n}$, we have that $(UBV)_{I \times I} = B_{I \times I}$.

Proof. Let $X = \{\ell \in \{2, \ldots, n\} : A_{1,\ell} \neq 0\}$. If $X = \emptyset$, then we let $V = \mathrm{Id}$ and let j be any element of $[n]$. Otherwise, we let j be any element of X, and we let V be the matrix that, when multiplied with A from the right, adds multiples of the jth column of A to the ℓth column of A such that the $(1, \ell)$ entry becomes zero, for every $\ell \in X \setminus \{i\}$. In other words, we let

$$V = \mathrm{Id} - \sum_{\ell \in X \setminus \{j\}} A_{1,j}^{-1} A_{1,\ell} E_{j,\ell}.$$

Let $Y = \{\ell \in \{2, \ldots, n\} : A_{\ell,1} \neq 0\}$. If $Y = \emptyset$ or $Y = \{j\}$, then we let $U = \mathrm{Id}$ and we let i be any element of $[n] \setminus \{j\}$. Otherwise, let i be any element of $Y \setminus \{j\}$, and we let U be the matrix that, when multiplied with A from the left, add multiples of the ith row of A to the ℓth row of A such that the $(\ell, 1)$ entry becomes zero, for every $\ell \in Y \setminus \{i, j\}$. In other words,

$$U = \mathrm{Id} - \sum_{\ell \in Y \setminus \{i,j\}} A_{i,1}^{-1} A_{\ell,1} E_{\ell,i}.$$

Then claim (i) is satisfied. To see that (ii) is satisfied we note that $U = \mathrm{Id} + U_0$ and $V = \mathrm{Id} + V_0$, where U_0 supported by the ith column and of V_0 is supported by the jth row. Then, for any diagonal matrix $B \in \mathbb{F}^{n \times n}$,

$$UBV = B + U_0 B + B V_0 + U_0 B V_0.$$

Since $i \neq j$, it follows that $U_0 B V_0 = 0$. Moreover, $U_0 B$ is supported by the ith column and $B V_0$ is supported by the jth row, so $(U_0 B)_{I \times I} = 0$ and $(B V_0)_{I \times I} = 0$, and so $(UBV)_{I \times I} = B_{I \times I}$. □

Lemma 3.8. *Let $A \in \mathbb{F}^{n \times n}$. Then there are $U, V \in \mathrm{GL}(\mathbb{F}^n)$ and $I \subseteq [n]$ such that $|I| \geq n/3$ and*

(i) $(UAV)_{I \times I}$ is diagonal;

(ii) for every diagonal matrix $B \in \mathbb{F}^{n \times n}$, we have that $(UBV)_{I \times I} = B_{I \times I}$.

Proof. The proof is by repeatedly applying Lemma 3.7. Applying Lemma 3.7 to A gives U, V, I such that $(UAV)_{I \times I}$ is of the form

$$\alpha_{1,1} E_{1,1} + \sum_{a,b \geq 2} \alpha_{a,b} E_{a,b}$$

and such that for any diagonal B we have $(UBV)_{I \times I} = B_{I \times I}$. Suppose we are given $A' \in \mathbb{F}^{m \times m}$ of the form

$$\sum_{i=1}^{k} \alpha_{i,i} E_{i,i} + \sum_{a,b \geq k+1} \alpha_{a,b} E_{a,b}.$$

Applying Lemma 3.7 to $A'' = (A')_{\{k+1,\ldots m\} \times \{k+1,\ldots m\}}$ gives U', V', I' such that $(U'A''V')_{I' \times I'}$ is of the form $\alpha'_{1,1} E_{1,1} + \sum_{a,b \geq 2} \alpha'_{a,b} E_{a,b}$. Then

$$((U' \oplus \mathrm{Id}) A' (V' \oplus \mathrm{Id}))_{I' \times I'}$$

is of the form

$$\sum_{i=1}^{k+1} \alpha_{i,i} E_{i,i} + \sum_{a,b \geq k+2} \alpha_{a,b} E_{a,b}$$

and for any diagonal B we have $((U' \oplus \mathrm{Id}) UBV (V' \oplus \mathrm{Id}))_{I' \times I'} = B_{I' \times I'}$. In every step of the process k goes up by one, and $|I|$ goes down by two, so that in the end we obtain the claim with $I \geq n/3$. □

Lemma 3.9. *Let $A_1 = \mathrm{Id}_n, A_2, \ldots, A_c \in \mathbb{F}^{n \times n}$. Then there are $U, V \in \mathrm{GL}(\mathbb{F}^n)$ and $I \subseteq [n]$ such that $|I| \geq 3^{-(c-1)}n$ and*

(i) $(U \mathrm{Id}_n V)_{I \times I} = (\mathrm{Id}_n)_{I \times I}$

(ii) $(U A_i V)_{I \times I}$ *is diagonal for all $i \in [c]$.*

Proof. The proof is by induction over c using Lemma 3.8. The base case $c = 1$ is clear. For the induction step, we are given $A_1, \ldots, A_m$, U, V and $I \subseteq [n]$ with $|I| \geq 3^{-(m-1)}n$ such that for $B_i = (U A_i V)_{I \times I}$ we have $B_1 = \mathrm{Id}_I$ and every B_i is diagonal. Let $B_{m+1} = (U A_{m+1} V)_{I \times I}$. We apply Lemma 3.8 to $A = B_{m+1}$. This gives $U', V' \in \mathrm{GL}(\mathbb{F}^I)$ and $J \subseteq I$ such that $|J| \geq \frac{1}{3}|I|$, $(U' B_{m+1} V')_{J \times J}$ is diagonal, and $(U' B_i V')_{J \times J} = (B_i)_{J \times J}$ for every $i \in [m]$. This proves the claim. $\qquad\square$

From large max-rank to large min-rank. In this second part of the proof of Lemma 3.5, we show in Lemma 3.12 that any low-dimensional matrix subspace $\mathcal{A}$ of diagonal matrices with large max-rank contains a principal-matrix subspace with large min-rank. Moreover, in Lemma 3.13, we show that any matrix subspace has a basis that, when restricted to a particular principal submatrix, consists of a set of linearly independent matrices and all-zero matrices.

For any vector $w \in \mathbb{F}^n$ let $\mathrm{supp}(w) = \{i \in [n] : w_i \neq 0\}$ and let $|w| = |\mathrm{supp}(w)|$. For any subspace $W \subseteq \mathbb{F}^n$ let

$$\text{max-supp}(W) = \max\{|w| : w \in W\},$$
$$\text{min-supp}(W) = \min\{|w| : w \in W \setminus \{0\}\}.$$

Lemma 3.10. *Let $V \subseteq \mathbb{F}^n$ be a subspace such that $\dim(V) \leq c$. There is an $I \subseteq [n]$ such that V_I is not zero and*
$$\text{min-supp}(V_I) \geq \tfrac{1}{c} \text{max-supp}(V).$$

Proof. We will use the general fact that for arbitrary vectors $w_1, \ldots, w_\ell \in W$, if for every $j \in \{0, \ldots, \ell - 1\}$ we have $\cup_{i \leq j} \mathrm{supp}(w_i) \neq \cup_{i \leq j+1} \mathrm{supp}(w_i)$, then $w_1, \ldots, w_\ell$ are linearly independent.

Let $k = \text{max-supp}(V)$ and $d = \dim(V)$. Before constructing I, we first construct a set $X \subseteq V$ by the following process, starting with X empty. If there is a nonzero element $v_1 \in V$ such that $|v_1| < \frac{1}{c}k$, then add v_1 to X. If $X = \{v_1, \ldots, v_\ell\}$ and there is an element $v_{\ell+1} \in V$ such that $0 < |\cup_{i \leq \ell+1} \mathrm{supp}(v_i) \setminus \cup_{i \leq \ell} \mathrm{supp}(v_i)| < \frac{1}{c}k$, then add $v_{\ell+1}$ to X. By the aforementioned general fact, X is linearly independent and thus X is finite, namely $|X| \leq d$.

We let $J = \cup_i \mathrm{supp}(v_i)$ and $I = [n] \setminus J$.

First we claim that V_I is not zero. We have $|J| < d\frac{1}{c}k \leq k$. Then for any element $v \in V$ with $|v| = k$ we have $|v_I| \geq k - |J| > 1$ and thus $v_I \neq 0$.

Second, we claim $\text{min-supp}(V_I) \geq \frac{1}{c}k$. Suppose there is a nonzero element $w \in V_I$ such that $|w| < \frac{1}{c}k$, then this contradicts the way we constructed X. $\qquad\square$

Lemma 3.11. *Let $V \subseteq \mathbb{F}^n$ be a subspace of dimension c, and let $I \subseteq [n]$. Then there is a basis $v_1, \ldots, v_c$ for V such that*

- $(v_1)_I, \ldots, (v_b)_I$ *are linearly independent*

- $(v_{b+1})_I = \cdots = (v_c)_I = 0$

for $b = \dim(V_I)$.

Proof. Choose any basis $w_1, \ldots, w_c$ for V. Let M be the matrix with these basis vectors as columns. The submatrix $M_{I \times [n]}$ has rank b so there is a subset $J \subseteq [n]$ of size b such that $M_{I \times J}$ has rank b. Let $\{v_1, \ldots, v_b\} = \{w_j : j \in J\}$ and $\{v_{b+1}, \ldots, v_c\} = \{w_j : j \notin J\}$. Then $(v_1)_I, \ldots, (v_b)_I$ form a basis of V_I. By subtracting multiples of $v_1, \ldots, v_b$ from the vectors $v_{b+1}, \ldots, v_c$ we can ensure that $(v_{b+1})_I = \cdots = (v_c)_I = 0$. $\square$

Lemma 3.12. *Let $\mathcal{A} \subseteq \mathbb{F}^{n \times n}$ be a subspace of diagonal matrices such that $\dim(\mathcal{A}) \leq c$. There is an $I \subseteq [n]$ such that*

$$\text{min-rank}(\mathcal{A}_{I \times I}) \geq \tfrac{1}{c} \max\text{-rank}(\mathcal{A}).$$

Proof. This follows from Lemma 3.10 by identifying every diagonal matrix $A \in \mathcal{A}$ with the vector of its diagonal coefficients $v = (A_{1,1}, \ldots, A_{n,n})$ and $\text{rank}(A) = |\text{supp}(v)|$. $\square$

Lemma 3.13. *Let $\mathcal{A} \subseteq \mathbb{F}^{n_1 \times n_2}$ be matrix subspace of dimension c, and $J \subseteq [\min_i n_i]$. Then there is a basis $B_1, \ldots, B_c$ for $\mathcal{A}$ such that*

- *$(B_1)_{J \times J}, \ldots, (B_b)_{J \times J}$ are linearly independent*

- *$(B_{b+1})_{J \times J} = \cdots = (B_c)_{J \times J} = 0$*

for $b = \dim(\mathcal{A}_{J \times J})$.

Proof. This follows from applying Lemma 3.11 to $V = \mathcal{A}$ and $I = J \times J$. $\square$

Combining the above. We now combine Lemma 3.9, Lemma 3.12 and Lemma 3.13 to prove Lemma 3.5.

Proof of Lemma 3.5. Let $k = \max\text{-rank}(\mathcal{A})$. By Lemma 3.6, we have that $k \geq n_2/c$. After a basis transformation, we may assume that $\mathcal{A}$ has a basis $A_1 = \text{Id}_k, A_2, \ldots, A_c$. Applying Lemma 3.9 to the submatrices $(A_i)_{[k] \times [k]}$, we obtain matrices $U, V \in \text{GL}(\mathbb{F}^k)$ and a set $I \subseteq [k]$ of size at least $3^{-(c-1)}k$ such that

(i) $A'_1 = \left((U \oplus \text{Id}_{n_1-k})A_1(V \oplus \text{Id}_{n_2-k})\right)_{I \times I} = (\text{Id}_k)_{I \times I}$;

(ii) $A'_i = \left((U \oplus \text{Id}_{n_1-k})A_i(V \oplus \text{Id}_{n_2-k})\right)_{I \times I}$ is diagonal for all $i \in [c]$.

By Lemma 3.12 there is a set $J \subseteq I$ such that

$$\text{min-rank}\left((A'_1)_{J \times J}, \ldots, (A'_c)_{J \times J}\right) \geq \tfrac{1}{c}3^{-(c-1)}k.$$

The claim follows from applying Lemma 3.13 with $B_i = A'_i$. $\square$

3.1.3 Super-multiplicativity of min-rank

For subspaces of matrices $\mathcal{A}$ and $\mathcal{B}$ their tensor product $\mathcal{A} \otimes \mathcal{B}$ is the subspace of matrices spanned by the products $A \otimes B$ for all $A \in \mathcal{A}$ and $B \in \mathcal{B}$. In this section we will prove that the min-rank is super-multiplicative under the tensor product, under the assumption that at least one of the matrix subspaces is *diagonal*. (This assumption is necessary.) We then apply this to the min-rank of powers of any diagonal subspace of matrices.

Lemma 3.14. *Let $\mathcal{A}$ and $\mathcal{B}$ be subspaces of matrices. Suppose that all matrices in $\mathcal{A}$ are diagonal. Then $\text{min-rank}(\mathcal{A} \otimes \mathcal{B}) \geq \text{min-rank}(\mathcal{A}) \cdot \text{min-rank}(\mathcal{B})$.*

Proof. Let $C \in \mathcal{A} \otimes \mathcal{B}$ be a nonzero element. Then C is of the form $C = \sum_{i=1}^{\ell} A_i \otimes B_i$ for some $A_i \in \mathcal{A}$, $B_i \in \mathcal{B}$ and $\ell \geq 1$. We may assume $\{A_i\}_{i=1}^{\ell}$ and $\{B_i\}_{i=1}^{\ell}$ are both linearly independent sets of matrices.

Since every A_i is diagonal, the matrix C is block-diagonal with blocks D_j given by $D_j = \sum_{i=1}^{\ell} (A_i)_{j,j} B_i$. Thus $\mathrm{rank}(C) = \sum_j \mathrm{rank}(D_j)$. For any j such that $(A_1)_{j,j} \neq 0$, the matrix D_j is a nontrivial linear combination of the B_i and thus $\mathrm{rank}(D_j) \geq \mathrm{min\text{-}rank}(\mathcal{B})$. The number of j such that $(A_1)_{j,j} \neq 0$ is equal to $\mathrm{rank}(A_1)$ since A_1 is diagonal. Thus

$$\mathrm{rank}(C) = \sum_j \mathrm{rank}(D_j) \geq \sum_{j:(A_1)_{j,j} \neq 0} \mathrm{rank}(D_j) \geq \mathrm{rank}(A_1) \, \mathrm{min\text{-}rank}(\mathcal{B})$$

$$\geq \mathrm{min\text{-}rank}(\mathcal{A}) \, \mathrm{min\text{-}rank}(\mathcal{B}),$$

which proves the claim. $\qquad\square$

Example 3.15. The condition in Lemma 3.14 that at least one of the matrix subspaces consists of diagonal matrices is necessary. To see this we may take the subspace of matrices over the reals

$$\mathcal{A} = \left\{ \begin{pmatrix} a & b \\ -b & a \end{pmatrix} : a, b \in \mathbb{R} \right\}.$$

Then $\mathrm{min\text{-}rank}(\mathcal{A}) = 2$. However, $\mathcal{A}^{\otimes 2}$ contains the element

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}^{\otimes 2} + \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}^{\otimes 2} = \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & -1 & 0 \\ 0 & -1 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix}$$

which has rank $2 < 2^2$, so $\mathrm{min\text{-}rank}(\mathcal{A}^{\otimes 2}) < \mathrm{min\text{-}rank}(\mathcal{A})^2$.

Lemma 3.16. *Let $\mathcal{A}$ be a subspace of matrices. Suppose that all matrices in $\mathcal{A}$ are diagonal. Then for all $k \in \mathbb{N}$ we have $\mathrm{min\text{-}rank}(\mathcal{A}^{\otimes k}) \geq \mathrm{min\text{-}rank}(\mathcal{A})^k$.*

Proof. Repeatedly apply Lemma 3.14. $\qquad\square$

In the rest of this section we prove a slightly more general version of Lemma 3.16. The reason that we need to do this is that we want to take the output of Lemma 3.5 (a collection of linearly independent matrices, such that restricted to specified rows and columns the matrices are diagonal and have large min-rank, or are zero) and use the super-multiplicativity property of Lemma 3.16 to construct a large collection of matrices (that are tensor products of the original matrices) with large min-rank. However, we cannot do this directly as some of the matrices after restriction are zero. The lemma we prove here deals with that.

Recall that for any matrix $A \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ and subsets $I \subseteq [n_1]$, $J \subseteq [n_2]$ we denote by $A_{I \times J}$ the submatrix of A indexed by rows I and columns J. For any subset $X \subseteq \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ we denote by $X_{I \times J}$ the set $\{A_{I \times J} : A \in X\}$.

Lemma 3.17. *Let $X = \{B_1, \ldots, B_b, C_{b+1} \ldots, C_c\} \subseteq \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ be a set of linearly independent matrices. Let $I \subseteq [n_1]$, $J \subseteq [n_2]$, $|I| = |J|$, such that*

- *the matrices $(B_i)_{I \times J}$, for $i = 1, \ldots, b$, are diagonal and linearly independent,*

- *the matrices $(C_i)_{I \times J}$, for $i = b+1, \ldots, c$, are zero.*

For any $m \geq \ell \geq 1$, let Y be the set of all order-m Kronecker products of elements of X such that at least ℓ of the factors are from $B_1, \ldots, B_b$. Then

$$\text{min-rank}(Y) \geq \text{min-rank}(X_{I \times J})^\ell.$$

Proof. Let A be any nontrivial linear combination of elements of Y. Then

$$A = \sum_{t=1}^r a_t \, D_{i_{t,1}} \otimes \cdots \otimes D_{i_{t,m}}$$

for some nonzero $a_t \in \mathbb{F}$ and $D_{i_{t,j}} \in X$. Every $D_{i_{t,j}}$ is either in $\{B_1, \ldots, B_b\}$, in which case we will call it a B-matrix, or in $\{C_{b+1}, \ldots, C_c\}$, in which case we will call it a C-matrix. Let k be the maximal number of B-matrices that appear as factors in any of the summands. Then $k \geq \ell$. Consider any summand with the maximal number of B-matrices in its factors. Suppose it is the first summand. For simplicity of notation, we assume the first k factors of this summand are B-matrices and the last $m - k$ factors are C-matrices:

$$D_{i_{1,1}} \otimes \cdots \otimes D_{i_{1,m}} = B_{i_{1,1}} \otimes \cdots \otimes B_{i_{1,k}} \otimes C_{i_{1,k+1}} \otimes \cdots \otimes C_{i_{1,m}}.$$

We restrict the first k factors of A to $I \times J$ to get

$$A' = \sum_{t=1}^r a_t \, (D_{i_{t,1}})_{I \times J} \otimes \cdots \otimes (D_{i_{t,k}})_{I \times J} \otimes D_{i_{t,k+1}} \otimes \cdots \otimes D_{i_{t,m}}.$$

If there is a C-matrix among $D_{i_{t,1}}, \ldots, D_{i_{t,k}}$, then $(D_{i_{t,1}})_{I \times J} \otimes \cdots \otimes (D_{i_{t,k}})_{I \times J}$ is zero. Moreover, nonzero summands cannot have any B-matrix among $D_{i_{t,k+1}}, \ldots, D_{i_{t,m}}$, as this would contradict the maximality of k. We know that the first k factors of the first summand of A are B-matrices, so A' is nonzero by the linear-independence of $\{(B_i)_{I \times J}\}^{\otimes k} \otimes \{B_i, C_i\}^{\otimes m-k}$. Using the notation $B_i' = (B_i)_{I \times J}$, we may rewrite A' as

$$A' = \sum_{j \in [b]^k} B_{j_1}' \otimes \cdots \otimes B_{j_k}' \otimes \left(\sum_{t=1}^r \delta_{i_{t,1}=j_1} \cdots \delta_{i_{t,k}=j_k} \cdot a_t \, C_{i_{t,k+1}} \otimes \cdots \otimes C_{i_{t,m}} \right).$$

Let $\tilde{j} = (i_{1,1}, \ldots, i_{1,k}) \in [b]^k$ be the first k indices of the first summand in A. We have $a_1 \neq 0$ and the set $\{C_i\}_{i \in \{b+1, \ldots, c\}}^{\otimes(m-k)}$ is linearly independent, thus

$$\sum_{t=1}^r \delta_{i_{t,1}=i_{1,1}} \cdots \delta_{i_{t,k}=i_{1,k}} \cdot a_t \, C_{i_{t,k+1}} \otimes \cdots \otimes C_{i_{t,m}} \tag{4}$$

is nonzero.

We may think of A' as an $n_1^{m-k} \times n_2^{m-k}$ block matrix where each block is an $|I|^k \times |J|^k$ matrix. That is, the "outer" blocks correspond to linear combinations of products of C_i's, and the "inner" blocks correspond to linear combinations of tensor products of B_i''s. Consider the nonzero linear combination Equation 4 corresponding to $\tilde{j}$ and denote it by S. Assume without loss of generality that the $(1,1)$ entry of S is not zero ($S_{11} \neq 0$). Block $(1,1)$ in A' is a linear combination of the order-k tensor products of the B_i's with a nonzero coefficient (S_{11}) for $B_{i_{1,1}} \otimes \cdots \otimes B_{i_{1,k}}$, meaning this linear combination is nontrivial. Lemma 3.16 then implies that this block has rank at least $\text{min-rank}(X_{I \times J})^k$. The matrix A' is block-diagonal, so the rank of A' is at least $\text{min-rank}(X_{I \times J})^k \geq \text{min-rank}(X_{I \times J})^\ell$. The rank of A is at least the rank of A', which proves the claim. $\qquad \square$

3.2 Narrow tensors have maximal asymptotic subrank

We will now prove the main theorem of this section, which says that for every constant c, and large enough n_1, n_2, all tensors in $\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^c$ have maximal asymptotic subrank.

Theorem 3.18. *For every $c \geq 2$ there is $N(c) \in \mathbb{N}$ such that for any $n_1, n_2 \in \mathbb{N}$ such that $\max\{n_1, n_2\} > N(c)$, every concise tensor $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^c$ satisfies $\underset{\sim}{Q}(T) = c$.*

The proof of Theorem 3.18 relies on min-rank lemmas that we proved in Section 3.1. For convenience of the reader we briefly restate here the ones we need.

Lemma 3.6. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^c$ be a concise tensor. Then T has at least one 3-slice of rank at least $\max\{n_1, n_2\}/c$.*

Lemma 3.5. *For every $c \in \mathbb{N}$, there is an $\varepsilon(c) > 0$ such that the following holds. Let $\mathcal{A} \subseteq \mathbb{F}^{n_1 \times n_2}$ be a matrix subspace of dimension c. Then, there exist $U \in \mathrm{GL}(\mathbb{F}^{n_1})$, $V \in \mathrm{GL}(\mathbb{F}^{n_2})$, a subset $J \subseteq [\min_i n_i]$, a basis $B_1, \ldots, B_c$ for the space $U\mathcal{A}V$ and $b \in [c]$ such that:*

- *$(B_1)_{J \times J}, \ldots, (B_b)_{J \times J}$ are diagonal and linearly independent;*

- *$\mathrm{min}\text{-}\mathrm{rank}\big((B_1)_{J \times J}, \ldots, (B_b)_{J \times J}\big) \geq \varepsilon(c) \max\text{-}\mathrm{rank}(\mathcal{A})$;*

- *$(B_{b+1})_{J \times J} = \cdots = (B_c)_{J \times J} = 0$.*

Lemma 3.17. *Let $X = \{B_1, \ldots, B_b, C_{b+1} \ldots, C_c\} \subseteq \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ be a set of linearly independent matrices. Let $I \subseteq [n_1]$, $J \subseteq [n_2]$, $|I| = |J|$, such that*

- *the matrices $(B_i)_{I \times J}$, for $i = 1, \ldots, b$, are diagonal and linearly independent,*

- *the matrices $(C_i)_{I \times J}$, for $i = b+1, \ldots, c$, are zero.*

For any $m \geq \ell \geq 1$, let Y be the set of all order-m Kronecker products of elements of X such that at least ℓ of the factors are from $B_1, \ldots, B_b$. Then

$$\mathrm{min}\text{-}\mathrm{rank}(Y) \geq \mathrm{min}\text{-}\mathrm{rank}(X_{I \times J})^{\ell}.$$

Lemma 3.3. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^c$ be a tensor with 3-slices $A_1, \ldots, A_c \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$. If $A_1, \ldots, A_c$ are linearly independent and $\mathrm{min}\text{-}\mathrm{rank}(A_1, \ldots, A_c) \geq 10c^2$, then $Q(T) \geq c$.*

Proof of Theorem 3.18. Suppose $n_1 \geq n_2$. Suppose $n_1 > N(c)$ for an $N(c)$ to be determined in the proof. Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^c$ be concise. Let $\mathcal{A} \subseteq \mathbb{F}^{n_1 \times n_2}$ be the 3-slice span of T. By Lemma 3.6 we have $\max\text{-}\mathrm{rank}(\mathcal{A}) \geq \frac{1}{c} n_1$. We apply Lemma 3.5 to $\mathcal{A}$, giving matrices $B_1, \ldots, B_b$, $B_{b+1}, \ldots, B_c$ and subset a $J \subseteq [\min_i n_i]$. To these, we apply Lemma 3.17 for $m \geq \ell \geq 1$ (to be specified later) to obtain a set Y of slices from $T^{\boxtimes m}$ satisfying

$$\mathrm{min}\text{-}\mathrm{rank}(Y) \geq \big(\varepsilon(c) \max\text{-}\mathrm{rank}(\mathcal{A})\big)^{\ell} \geq (\varepsilon(c) \tfrac{1}{c} n_1)^{\ell}. \tag{5}$$

The size of Y equals the number of m-tuples in $[c]^m$ containing at least ℓ elements from $[b]$. Setting $m = 8c^3$ and $\ell = \frac{1}{2c} m$, it then follows from the Chernoff bound that $|Y| \geq \frac{1}{2} c^m$.

Since T is concise, so is $T^{\boxtimes m}$ and therefore, the set Y is linearly independent. Let T' be the tensor with 3-slices given by the elements in Y. Then $T^{\boxtimes m} \geq T'$. Provided

$$\mathrm{min}\text{-}\mathrm{rank}(Y) \geq 10(\tfrac{1}{2} c^m)^2, \tag{6}$$

28

it follows from Lemma 3.3 that

$$Q(T^{\boxtimes m}) \geq Q(T') \geq \tfrac{1}{2}c^m. \tag{7}$$

By our choice of ℓ and (5), we have that (6) holds if

$$\left(\varepsilon(c)\tfrac{1}{c}n_1\right)^{\frac{1}{2c}m} \geq 10(\tfrac{1}{2}c^m)^2.$$

It is clear that there is an $N(c) \in \mathbb{N}$ such that this is satisfied for any $n_1 > N(c)$ and every m. Letting m go to infinity we then get from (7) that

$$\underset{\sim}{Q}(T) = \lim_{m \to \infty} Q(T^{\boxtimes m})^{1/m} \geq c. \qquad \square$$

Remark 3.19. To conclude this section, let us note that a similar approach fails for k-tensors, for $k > 3$. The motivating question mentioned at the beginning of Section 3, whether there exist accumulation points in the image of asymptotic slicerank (over $\mathbb{C}$ and finite fields), makes sense for k-tensors for any k. In order to show there cannot exist such points in the same way for $k > 3$ we would need to show there are only finitely many formats that allow asymptotic sliceranks strictly smaller than the maximal constant c. For $k = 3$ we showed it for asymptotic subrank, which lower bounds asymptotic slicerank, but this is false for asymptotic subrank and any $k > 3$. For $k = 3$ conciseness implies the formats that allow for a constant asymptotic subrank are only (n_1, n_2, c) for a constant c and n_1, n_2 growing. But for $k = 4$ and above there can be formats of the form $(n_1, n_2, \ldots, n_{k-2}, c_1, c_2)$ for c_1, c_2 constants and $n_1, n_2, \ldots, n_{k-2}$ growing. For these formats the straightforward generalisation of Theorem 3.18, asking to show that for every c_1, c_2 constants there would be N for which any concise $(n_1, n_2, \ldots, n_{k-2}, c_1, c_2)$-tensor with $n_1, n_2, \ldots, n_{k-2} > N$ has asymptotic subrank $\min\{c_1, c_2\}$ does not hold for k-tensors, for any $k > 3$:

The 4-tensor $\mathrm{Id}_n \otimes \mathrm{Id}_2 = \sum_{i=1}^{n} e_i \otimes e_i \otimes e_1 \otimes e_1 + \sum_{i=1}^{n} e_i \otimes e_i \otimes e_2 \otimes e_2$ is a concise 4-tensor of format $(n, n, 2, 2)$ (for each $i \in [4]$, the set of i-slices is linearly independent) but it is a tensor-product as a 2-tensor (for the bipartition $\{\{1, 2\}, \{3, 4\}\}$) so its subrank and its asymptotic subrank are 1 as a 2-tensor and thus as a 4-tensor as well.

This example can be generalised to any $k > 3$ and any $c_1 = c_2 =: c$: $\langle n \rangle_{k-2} \otimes \mathrm{Id}_c$ where $\langle n \rangle_{k-2} = \sum_{i=1}^{n} e_i \otimes \cdots \otimes e_i$ is the diagonal (unit) $(k-2)$-tensor of size n.

3.3 Concrete case: smallest dimension equals two

In the previous section we proved for any constant c that for large enough n_1, n_2 all concise tensors in $\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^c$ have asymptotic subrank c. In the special case $c = 2$ we can prove the following stronger statement in which we make the constant $N(2)$ explicit, namely $N(2) = 2$, and we are able to replace asymptotic subrank by subrank.

Theorem 3.20. *Let $n_1, n_2 > 2$ and let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^2$ be a concise tensor. Then*

$$Q(T) = 2.$$

Example 3.21. In Theorem 3.20 it is required that at least one of the dimensions n_1, n_2 is strictly larger than 2. Namely, there are concise tensors $T \in \mathbb{F}^2 \otimes \mathbb{F}^2 \otimes \mathbb{F}^2$ such that $Q(T) = 1$, for instance the tensor $T = e_1 \otimes e_2 \otimes e_2 + e_2 \otimes e_1 \otimes e_2 + e_2 \otimes e_2 \otimes e_1$ that is often called the "W-tensor".

Remark 3.22. Theorem 3.20 leaves open what happens for $n_1 = 2$ and $n_2 > 2$. Over the complex numbers the classification of Miyake and Verstraete [MV04] shows that indeed for

every $n > 2$ and every concise tensor $T \in \mathbb{F}^n \otimes \mathbb{F}^2 \otimes \mathbb{F}^2$, $Q(T) = 2$. Over any algebraically closed field, Tobler [Tob97, Section 4.2] showed that for every $n_2 > n_1 \geq 2$ and every concise tensor $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^2$, $\underline{Q}(T) = 2$. This in particular implies that the asymptotic subrank is 2 (Lemma 5.12).

Lemma 3.23. *For any $s, t, u \in \mathbb{F}$, the tensor with slices*

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, B = \begin{pmatrix} t & u \\ s & 1 \end{pmatrix}$$

has subrank 2.

Proof. We first subtract s times the second column from the first column and u times the second row from the first row to get

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, B = \begin{pmatrix} t' & 0 \\ 0 & 1 \end{pmatrix}$$

for some t'. We then subtract $t' \cdot A$ from B to get

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, B = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}.$$

This is $\langle 2 \rangle$, the diagonal tensor of rank 2. $\qquad\square$

Proof of Theorem 3.20. Let A, B be the two $n_1 \times n_2$ slices. We diagonalize the $n_1 \times 2n_2$ matrix $(A|B)$ by diagonalizing A and doing the same row and column operations on B. Observe that $\mathrm{rank}(A|B) = n_1$ as T is concise. Assume w.l.o.g. that the top $r \times r$ minor of A is the identity and the rest is 0. Specifically, the last $n_1 - r$ rows in it are 0. We can now continue and diagonalize the last $n_1 - r$ rows of B, noting that this will not change the structure of A (we are only doing row operations now). We will get a subset I of B's columns such that restricted to some of the last $n_1 - r$ rows, $B|_{\{r+1,\dots,r+\ell\} \times I}$ is the identity. By doing more row operations we can assume that the first r rows in the columns I in B are 0 (the $n_1 - r$ rows in A are zeros so we are free to subtract scaled versions of these rows from above rows to zero out the desired cells in B). Here's an illustration of our two $n_1 \times n_2$ matrices (the dashed lines separate the first r rows and columns from the rest and the I indices of columns are indicated beneath matrix B. Asterisks signify arbitrary values that we cannot assume anything about). In particular, a 0 above the dashed line in B represents an $r \times 1$ columns of zeros:

$$A = \left(\begin{array}{c:c} \mathrm{Id}_r & 0 \\ \hdashline 0 & 0 \end{array} \right) \quad B = \left(\begin{array}{ccccc:c} \ast & 0 & \ast & 0 & \cdots & 0 \\ \hdashline \ast & 1 & \ast & 0 & \cdots & 0 & \ast \\ \ast & 0 & \ast & 1 & \cdots & 0 & \ast \\ \ast & \vdots & \ast & \vdots & \ddots & \vdots & \vdots \\ \ast & 0 & \ast & 0 & \cdots & 1 & \ast \\ \ast & 0 & \ast & 0 & \cdots & 0 & \ast \end{array} \right)$$
$$ \quad I_1 \quad I_2 \quad \cdots \quad I_\ell$$

We now have three cases to consider:

Case 1: $1 < r < \min\{n_1, n_2\}$. Let $i \in I$ be arbitrary but different than r. Note that $B_{r+i,i} = 1$ (it is part of the identity matrix on columns I). Consider the restriction of our tensor to rows $r, r+i$ and columns r and i. This gives the tensor with slices

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, B = \begin{pmatrix} t & 0 \\ s & 1 \end{pmatrix}$$

for some s and t. From Lemma 3.23 we know that this tensor has subrank 2.

Case 2: $r = \min\{n_1, n_2\}$. We consider two subcases:

- Case 2a: The matrix B has a nonzero nondiagonal coefficient. Assume w.l.o.g. that $B_{1,2} = t \neq 0$. Restrict our tensor to rows 1,3 and columns 2,3 ($n_1, n_2 \geq 3$). We get slices

$$A = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, B = \begin{pmatrix} t & * \\ * & * \end{pmatrix}.$$

  Using Lemma 3.23, as $t \neq 0$, this tensor has subrank 2.

- Case 2b: All nondiagonal coefficients of B are 0 ($\forall i \neq j : B_{i,j} = 0$). As the tensor is concise, the matrix B is not a multiple of the matrix A. The diagonal coefficients of A are all 1. Hence, w.l.o.g., $B_{1,1} \neq B_{2,2}$. Restricting to rows and columns $\{1, 2\}$ we get

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, B = \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}.$$

  Taking appropriate linear combinations of A and B we get

$$A = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, B = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

Case 3: $r = 1$. In this case we repeat the argument above with B in place of A. As $\operatorname{rank}(A) = 1$, $n_1, n_2 > 2$, and the tensor is concise we must have $\operatorname{rank}(B) > 1$. $\qquad\square$

Remark 3.24. The proof of Theorem 3.20 only uses conciseness with respect to two directions, namely that the matrices A and B are linearly independent (Case 2b) and that the matrix $(A|B)$ has full rank (all cases). An example (that is not concise) is

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, B = \begin{pmatrix} 1 & 1 \\ 2 & 2 \end{pmatrix}.$$

4 Discreteness of asymptotic tensor ranks

In this section we prove a general sufficient condition for a tensor parameter to be discrete (using the lower bounds on the asymptotic subrank of concise tensors that we proved in the previous sections, Section 2 and Section 3). Then we show that this condition is satisfied (and thus obtained discreteness) for several natural tensor parameters and regimes.

4.1 General sufficient condition for discreteness

We recall that two tensors S and T are called equivalent if $S \geq T$ and $T \geq S$. Every tensor is equivalent to a concise tensor (Lemma 1.15). We will be interested in functions f on tensors that are invariant under equivalence, meaning that if two tensors S and T are equivalent, then $f(S) = f(T)$. This is a very mild condition, in particular any function that is monotone under restriction is also invariant under equivalence.

Theorem 4.1. *Let $\mathbb{F}$ be any field. Let $\mathcal{C}$ be any subset of order-three tensors over $\mathbb{F}$, such that for any $T \in \mathcal{C}$ there is an $S \in \mathcal{C}$ that is concise and equivalent to T. Let $f : \mathcal{C} \to \mathbb{R}_{\geq 0}$ be any function such that*

  (i) f does not change under equivalence of tensors

(ii) $f(T) \geq \underset{\sim}{\mathrm{Q}}(T)$ for every $T \in \mathcal{C}$

(iii) For every $n_1, n_2, n_3 \in \mathbb{N}$, f takes finitely many values on $\mathcal{C} \cap (\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3})$

(iv) For every $n_1, n_2, n_3 \in \mathbb{N}$, $f(T) \leq \min_i n_i$ for every $T \in \mathcal{C} \cap (\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3})$.

Then the set of values that f takes on $\mathcal{C}$ has no accumulation points.

Proof. Suppose we have, for every $i \in \mathbb{N}$, three numbers $a_i, b_i, c_i \in \mathbb{N}$ and a tensor $T_i \in \mathcal{C} \cap (\mathbb{F}^{a_i} \otimes \mathbb{F}^{b_i} \otimes \mathbb{F}^{c_i})$ in such a way that the set $\{f(T_i) : i \in \mathbb{N}\}$ is infinite. We will show that the $f(T_i)$ cannot converge.

We may assume that every tensor T_i is concise by condition (i) and Lemma 1.15. By condition (iii), f takes only finitely many values on $\mathcal{C} \cap (\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3})$ for any choice of n_1, n_2, n_3, so the set of triples $\{(a_i, b_i, c_i) : i \in \mathbb{N}\}$ must be infinite.

We consider two cases. First, suppose $\{\min\{a_i, b_i, c_i\}\}_i$ is bounded. Suppose for simplicity of notation that $a_i = \min\{a_i, b_i, c_i\}$ for every $i \in \mathbb{N}$. Then a_i takes only finitely many values. Take any infinite subsequence of the T_i such that a_i is constant. Then both b_i and c_i are unbounded, because $b_i \leq a_i c_i$ and $c_i \leq a_i b_i$ by Lemma 1.16. When b_i and c_i are large enough we have $\underset{\sim}{\mathrm{Q}}(T_i) = a_i$ by Theorem 3.18, and hence $f(T_i) = a_i$ by condition (ii) and (iv). We conclude that $f(T_i)$ takes only finitely many values, which contradicts our assumption that $\{f(T_i) : i \in \mathbb{N}\}$ is infinite.

Second, suppose $\min\{a_i, b_i, c_i\}$ is unbounded. We have $f(T_i) \geq \underset{\sim}{\mathrm{Q}}(T_i)$ by condition (ii). We have $\underset{\sim}{\mathrm{Q}}(T_i) \geq \min\{a_i, b_i, c_i\}^{1/3}$ by Theorem 2.12. It then follows that the values $f(T_i)$ do not converge. $\square$

4.2 Asymptotic subrank and asymptotic slice rank

We will now use Theorem 4.1 to prove a discreteness theorem for asymptotic slice rank and asymptotic subrank.

Before stating our theorem, we recall the notions of tight and oblique tensors, which are used in one part of the theorem. Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$. We recall that the support of T is the set $\mathrm{supp}(T) = \{(i, j, k) : T_{i,j,k} \neq 0\}$. The tensor T is called *tight* if (after an appropriate basis transformation) there exist injective maps $f_i : [n_i] \to \mathbb{Z}$ such that for every $(i_1, i_2, i_3) \in \mathrm{supp}(T)$ we have $f_1(i_1) + f_2(i_2) + f_3(i_3) = 0$. The tensor T is called *oblique* if (after an appropriate basis transformation) $\mathrm{supp}(T)$ is an antichain in the product order (i.e. pointwise order) on $[n_1] \times [n_2] \times [n_3]$. Tight tensors are oblique.

Theorem 4.2.

(i) *Let $\mathbb{F}$ be any field and $S \subseteq \mathbb{F}$ any finite subset. Then the asymptotic subrank and asymptotic slice rank on order-three tensors over S have no accumulation points.*

(ii) *The asymptotic slice rank on order-three tensors over $\mathbb{C}$ has no accumulation points.*

(iii) *Let $\mathbb{F}$ be any field.*

- *The asymptotic subrank and asymptotic slice rank on tight order-three tensors over $\mathbb{F}$ have no accumulation points.*

- *The asymptotic slice rank on oblique order-three tensors over $\mathbb{F}$ has no accumulation points.*

Proof. The proof of each case follows from Theorem 4.1. As slice rank is lower bounded by subrank, $f(T) \geq \underset{\sim}{Q}(T)$ for every $T \in \mathcal{C}$, over any field $\mathbb{F}$. Thus, conditions (i), (ii) and (iv) of Theorem 4.1 are fulfilled. We next show that condition (iii) of Theorem 4.1 holds for each of the cases in Theorem 4.2, i.e. that in each of these cases f takes only finitely many values on $\mathcal{C} \cap (\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3})$.

(i) Let f be the asymptotic subrank or the asymptotic slice rank and let $\mathcal{C}$ be the set of all order-three tensors over the finite set S. Indeed $\mathcal{C}$ is closed under taking the concise version of any element by Lemma 1.15, since any tensor is equivalent to a concise subtensor (which in particular again has coefficients in S). Then f takes only finitely many values on $\mathcal{C} \cap (\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3})$ simply because the set of tensors in $\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ with coefficients in S is finite.

(ii) Let f be the asymptotic slice rank and let $\mathcal{C}$ be the set of all order-three tensors over $\mathbb{C}$. Fix dimensions n_1, n_2 and n_3. By the characterization of asymptotic slice rank in terms of entanglement (moment) polytopes (via the quantum functionals) [CVZ23, Corollary 5.7] and the fact that for fixed n_1, n_2, n_3 there are finitely many such polytopes [WDGC13, Theorem 2], it holds that f takes only finitely many values on $\mathcal{C} \cap (\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3})$.

(iii) Let f be the asymptotic subrank or the asymptotic slice rank and let $\mathcal{C}$ be the set of all tight order-three tensors over arbitrary $\mathbb{F}$. Fix dimensions n_1, n_2 and n_3. By the characterization of asymptotic subrank in terms of the tight support (via the support functionals) [Str91, Proposition 5.4] and [CVZ23, Theorem 5.9], and the fact there are only a finite number of possible supports, it holds that f takes only finitely many values on $\mathcal{C} \cap (\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3})$.

Let f be the asymptotic slice rank and let $\mathcal{C}$ be the set of all oblique order-three tensors over arbitrary $\mathbb{F}$. By the characterization of asymptotic slice rank in terms of the oblique support [ST16, Proposition 4] it follows that f takes only finitely many values on $\mathcal{C} \cap (\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3})$. $\qquad\square$

Remark 4.3. Theorem 4.2 gives several regimes in which the asymptotic subrank and asymptotic slice rank are discrete, that is, between every two consecutive values there is a "gap". Using Example 2.6 we can say something about the size of these gaps. Namely, Example 2.6 gives for every large enough n (namely $n \geq 5$) a concise tensor $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ such that the asymptotic subrank and asymptotic slice rank of T equal $2\sqrt{n-1}$. This means that the gap between two consecutive values is at most $2\sqrt{n-1} - 2\sqrt{n-2}$, which is at most $O(1/\sqrt{n})$. This shows in particular that the consecutive gaps of the asymptotic subrank and asymptotic slice rank tend to zero as n grows.

With similar ideas as above (but much simpler) we can prove the analogous statement for asymptotic rank over finite fields. Recall that the asymptotic rank of a tensor T is defined as $\underset{\sim}{R}(T) = \lim_{n\to\infty} R(T^{\boxtimes n})^{1/n}$.

Theorem 4.4. *Let $k \in \mathbb{N}$. Let $\mathbb{F}$ be any finite field. The asymptotic rank on k-tensors over $\mathbb{F}$ has no accumulation points.*

Proof. Over any finite field $\mathbb{F}$, for tensors in $\mathbb{F}^{n_1} \otimes \cdots \otimes \mathbb{F}^{n_k}$ for a fixed format $(n_1, \ldots, n_k)$, there are only finitely many possible values of the asymptotic rank (simply because there are only finitely many tensors of that format). Thus, to make a sequence of concise tensors T_i such that the asymptotic rank $\underset{\sim}{R}(T_i)$ takes infinitely many different values, we need infinitely many different formats $(n_1, \ldots, n_k)$. In particular, $\max_i n_i$ needs to be unbounded. Since $\underset{\sim}{R}(T_i) \geq \max_i n_i$, the sequence $\underset{\sim}{R}(T_i)$ cannot converge. $\qquad\square$

5 Pivots of tensors and better lower bounds on the asymptotic subrank

In Section 2 we proved a cube-root lower bound on the asymptotic subrank of concise tensors, using a result on max-rank and matrix multiplication tensors.

In this section we develop a different strategy, using pivots, to first reprove the cube-root lower bound on asymptotic subrank of Section 2 and then prove a better lower bound under a symmetry condition that we introduce (pivot-matched).

Namely we show that the asymptotic subrank of a concise pivot-matched tensor is at least the square-root of the smallest dimension. In particular, this holds for symmetric concise tensors.

5.1 Pivots of tensors

We will first discuss what a pivot of a matrix is and what the pivots of a matrix subspace are, then introduce the pivots of a tensor, then relate this notion to the border subrank and the asymptotic subrank.

5.1.1 Pivot of a matrix and pivots of matrix subspaces

Following Meshulam [Mes85], we recall the notion of pivots of a matrix subspace and associated covering and packing parameters ρ and σ, which are known to be equal by an application of Kőnig's theorem. In the next section we will use these notions in the context of tensors.

Definition 5.1. For any $n_1 \times n_2$ matrix A we call the smallest element $(i, j) \in [n_1] \times [n_2]$ in the lexicographic ordering on $[n_1] \times [n_2]$ for which $A_{i,j} \neq 0$ the *pivot* of A, and we denote this element by $p(A) = (f(A), g(A))$.

Lemma 5.2. *Let $\mathcal{A}$ be a subspace of $n_1 \times n_2$ matrices. There is a basis for $\mathcal{A}$, $\{A_1, \ldots, A_{n_3}\}$, with distinct pivots $\forall i \neq j : p(A_i) \neq p(A_j)$. Additionally, for any such basis, the set of its matrices' pivots $\{p(A_1), \ldots, p(A_{n_3})\}$ is the same.*

To see this, vectorise the matrices row-by-row and stack them as the rows of a matrix M. Changing the basis amounts to invertible row operations on M and having distinct pivots amounts to getting M into its row-echelon form (up to the order of the rows). Uniqueness of the row-echelon form gives the uniqueness of the set of pivots.

By this lemma we may define the set of pivots as a parameter of the linear subspace of matrices itself:

Definition 5.3. For any subspace of $n_1 \times n_2$ matrices $\mathcal{A}$, let $A_1, \ldots, A_{n_3}$ be any basis of $\mathcal{A}$ such that the pivots of the A_i are distinct. We call the set of all pivots $\{p(A_1), \ldots, p(A_{n_3})\}$ the *pivots* of $\mathcal{A}$, and we denote this set by $p(\mathcal{A})$.

Remark 5.4. Let T be a tensor with distinct pivots of its 1-slices, then this set of pivots is a special case of a linearly independent set of cell-locations for T in direction 1.

Definition 5.5. For any subspace of $n_1 \times n_2$ matrices $\mathcal{A}$, we let $\rho(\mathcal{A})$ be the smallest number r such that there are subsets $S_1 \subseteq [n_1], S_2 \subseteq [n_2]$ of size $|S_1| + |S_2| = r$ for which $p(\mathcal{A}) \subseteq (S_1 \times [n_2]) \cup ([n_1] \times S_2)$.

Definition 5.6. For any subspace of $n_1 \times n_2$ matrices $\mathcal{A}$, we let $\sigma(\mathcal{A})$ be the largest number s such that there are matrices $A_1, \ldots, A_s \in \mathcal{A}$ for which for every $i \neq j$ we have $f(A_i) \neq f(A_j)$ and $g(A_i) \neq g(A_j)$.

The following is an application of Kőnig's theorem on the equivalence of maximum matchings and minimal vertex covers in bipartite graphs

Lemma 5.7 ([Mes85]). *For any subspace of $n_1 \times n_2$ matrices $\mathcal{A}$ we have $\rho(\mathcal{A}) = \sigma(\mathcal{A})$.*

5.1.2 Pivots of a tensor

Having discussed the notion of pivots of a matrix subspace, we will now associate to any tensor matrix subspaces obtained by taking the span of the slices of the tensor. We have considered slice spans before when defining the parameters Q_i, but here we will need to be more precise and distinguish six matrix subspaces instead of three (because the property we consider is not invariant under taking transpose). We then consider the pivots of these subspaces. We will relate these to the parameters Q_i and to the border subrank. Combining these ideas we then give another proof of the cube-root lower bound on the asymptotic subrank.

Definition 5.8. Let $T = \sum_{i,j,k} T_{i,j,k} \, e_i \otimes e_j \otimes e_k \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ be a tensor. Define the subspaces

$$\mathcal{A}_{1,2} = \text{span}\{A_1, \ldots, A_{n_3}\} \subseteq \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}, \quad A_k = \sum_{i,j} T_{i,j,k} \, e_i \otimes e_j$$

$$\mathcal{A}_{2,1} = \text{span}\{A_1, \ldots, A_{n_3}\} \subseteq \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_1}, \quad A_k = \sum_{i,j} T_{i,j,k} \, e_j \otimes e_i$$

$$\mathcal{A}_{1,3} = \text{span}\{A_1, \ldots, A_{n_2}\} \subseteq \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_3}, \quad A_j = \sum_{i,k} T_{i,j,k} \, e_i \otimes e_k$$

$$\mathcal{A}_{3,1} = \text{span}\{A_1, \ldots, A_{n_2}\} \subseteq \mathbb{F}^{n_3} \otimes \mathbb{F}^{n_1}, \quad A_j = \sum_{i,k} T_{i,j,k} \, e_k \otimes e_i$$

$$\mathcal{A}_{2,3} = \text{span}\{A_1, \ldots, A_{n_1}\} \subseteq \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}, \quad A_i = \sum_{j,k} T_{i,j,k} \, e_j \otimes e_k$$

$$\mathcal{A}_{3,2} = \text{span}\{A_1, \ldots, A_{n_1}\} \subseteq \mathbb{F}^{n_3} \otimes \mathbb{F}^{n_2}, \quad A_i = \sum_{j,k} T_{i,j,k} \, e_k \otimes e_j.$$

Let $\rho_{i,j}(T) = \rho(\mathcal{A}_{i,j})$.

Remark 5.9. Note that $\rho_{i,j}(T)$ can depend on the order of i, j, which defines which direction enumerates the rows and which the columns. This is a direct consequence of using the lexicographic order in the definition of the pivots. For example, the tensor T with

$$\mathcal{A}_{1,2} = \left\{ \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \right\},$$

has

$$\mathcal{A}_{2,1} = \left\{ \begin{pmatrix} 0 & 1 \\ 0 & 0 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 0 \end{pmatrix} \right\}.$$

Then $\rho_{1,2}(T) = 1 \neq 2 = \rho_{2,1}(T)$.

5.2 Alternative proof of cube-root lower bound via pivots

In this section we will give another proof of the cube-root lower bound on asymptotic subrank (Theorem 2.12) using pivots.

Theorem 5.10. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$ be concise. For every distinct i, j, k we have*

$$\rho_{i,j}(T) \max\{Q_i(T), Q_j(T)\} \geq n_k.$$

Proof. We give the proof for $i = 2, j = 3, k = 1$ (the proof for the other cases is the same). By an invertible linear transformation of direction 1 (linear combinations of the 1-slices) we may assume that for every $i \in [n_1]$, the i'th 1-slice is the only one with a non-zero entry in its pivot $p(i) = (f(i), g(i))$. Now we will prove that in every row there are at most $Q_2(T)$ pivots and in every column at most $Q_3(T)$. Let x in the image of f, and project the second leg of T to e_x. Restricted in the third leg to the columns with pivots in row x $\{k \in [n_3] : \exists i \in [n_1], (f(i), g(i)) = (x, k)\}$ this gives the follwoing diagonal submatrix of the x'th 2-slice

$$\sum_{i:f(i)=x} e_i \otimes e_x \otimes e_{g(i)}$$

which implies $Q_2(T) \geq |\{i : f(i) = x\}|$ for every x in the image of f. Similarly, $Q_3(T) \geq |\{i : g(i) = y\}|$ for every y in the image of g. By Definition 5.5, $p(\mathcal{A}_{2,3}) \subseteq (S_2 \times [n_3]) \cup ([n_2] \times S_3)$ for some subsets $S_i \subseteq [n_i]$ such that $|S_2| + |S_3| = \rho_{2,3}(T)$. Both of these combine to:

$$
\begin{aligned}
n_1 &= |\{(f(i), g(i)) : i \in [n_1]\}| \\
&\leq |\{(f(i), g(i)) : i \in [n_1], f(i) \in S_2\}| + |\{(f(i), g(i)) : i \in [n_1], g(i) \in S_3\}| \\
&\leq |S_2| \, Q_2(T) + |S_3| \, Q_3(T) \\
&\leq (|S_2| + |S_3|) \max\{Q_2(T), Q_3(T)\} \\
&= \rho_{2,3}(T) \max\{Q_2(T), Q_3(T)\}.
\end{aligned}
$$

This proves the claim. $\qquad\square$

Definition 5.11. The *border subrank* of a tensor $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$, denoted by $\underline{Q}(T)$, is the largest number r such that there are matrices $A(\varepsilon), B(\varepsilon), C(\varepsilon)$ whose coefficients are Laurent polynomials in the formal variable ε, such that

$$(A(\varepsilon) \otimes B(\varepsilon) \otimes C(\varepsilon))T = \langle r \rangle + \varepsilon S_1 + \varepsilon^2 S_2 + \cdots + \varepsilon^t S_t$$

where $\langle r \rangle$ is the $r \times r \times r$ diagonal tensor and the S_i are arbitrary tensors (with coefficients in $\mathbb{F}$).

The following known relation between border subrank and asymptotic subrank can be found in [Str88, Proposition 5.10] and [BCS97, Proposition 15.26].

Lemma 5.12. *For any tensor T we have $\underline{Q}(T) \leq \widetilde{Q}(T)$.*

Theorem 5.13. *For any tensor T and any $i \neq j$ we have $\rho_{i,j}(T) \leq \underline{Q}(T)$.*

Proof. Let $r = \rho_{i,j}(T)$. Then by Lemma 5.7, after taking appropriate linear combinations of the slices of T, T has an $r \times r \times r$ subtensor with slices $A_1, \ldots, A_r$ whose pivots $p(A_i)$ are pairwise disjoint in both coordinates. We may assume that $p(A_i) = (i, i)$ and that the ith row of A_i is zero except for the pivot coordinate (i, i). Thus A_i is of the form

$$A_i = e_i \otimes e_i + \sum_{j>i} \sum_k (A_i)_{j,k} \, e_j \otimes e_k.$$

So T restricts to the tensor

$$S = \sum_i e_i \otimes e_i \otimes e_i + \sum_{j>i} \sum_k (A_i)_{j,k} \, e_i \otimes e_j \otimes e_k.$$

We define the matrices $A(\varepsilon), B(\varepsilon), C(\varepsilon)$ as the ones that map

$$A(\varepsilon) : e_i \mapsto \varepsilon^i e_i$$
$$B(\varepsilon) : e_i \mapsto \varepsilon^{-i} e_i$$
$$C(\varepsilon) : e_i \mapsto e_i.$$

Then

$$(A(\varepsilon) \otimes B(\varepsilon) \otimes C(\varepsilon))S = \sum_i e_i \otimes e_i \otimes e_i + \sum_{j>i} \varepsilon^{j-i} \sum_k (A_i)_{j,k}\, e_i \otimes e_j \otimes e_k$$

so $\underline{Q}(T) \geq r$. $\qquad\square$

Theorem 5.14. *Let $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ be concise. Then $\underset{\sim}{Q}(T) \geq n^{1/3}$.*

Proof. We consider two cases. For the first case, suppose there are i, j such that $\rho_{i,j}(T) \geq n^{1/3}$. Then $\underset{\sim}{Q}(T) \geq n^{1/3}$, because generally we have $\underset{\sim}{Q}(T) \geq \underline{Q}(T) \geq \rho_{i,j}(T)$ (Theorem 5.13). For the second case, suppose for all $i \neq j$ we have $\rho_{i,j}(T) < n^{1/3}$. Then $\max\{Q_i(T), Q_j(T)\} \geq n^{2/3}$ for all distinct i, j (Theorem 5.10). From this follows that there are distinct i, j such that $Q_i(T) \geq n^{2/3}$ and $Q_j(T) \geq n^{2/3}$. Then $Q(T^{\boxtimes 2}) \geq n^{2/3}$ (Lemma 2.10) so $\underset{\sim}{Q}(T) \geq n^{1/3}$. $\qquad\square$

We finish this subsection with a relation between the border subrank and the parameters Q_i that may be of independent interest.

Theorem 5.15. *Let T be any tensor. Suppose $|\mathbb{F}| > \underline{Q}(T) + 1$. For every i we have*

$$\underline{Q}(T) \leq Q_i(T).$$

Proof. We will give the proof for $i = 3$. Let $q = \underline{Q}(T)$. Then there are matrices $A(\varepsilon)$, $B(\varepsilon)$, $C(\varepsilon)$ whose coefficients are Laurent polynomials over $\mathbb{F}$ in the variable ε such that

$$(A(\varepsilon) \otimes B(\varepsilon) \otimes C(\varepsilon))(T) = \langle q \rangle + \varepsilon S$$

for some tensor S whose coefficients are polynomials over $\mathbb{F}$ in ε. Taking the uniform linear combination over its 3-slices is acting by yet another local linear transformation on the direction 3, so we can absorb it into the third transformation to get a new matrix $C'(\varepsilon)$ such that $(A(\varepsilon) \otimes B(\varepsilon) \otimes C'(\varepsilon))(T) = \mathrm{Id}_q + \varepsilon N(\varepsilon)$ for some matrix $N(\varepsilon) \in M_q(\mathbb{F}[\varepsilon])$. The determinant of the matrix $\mathrm{Id}_q + \varepsilon N(\varepsilon)$ is a nonzero polynomial in ε of degree at most q, as it is evaluated to a nonzero element for $\varepsilon = 0$. Thus, it has at most q zeros and we have another evaluation point for which it is non-zero, as long as $|\mathbb{F}| > q + 1$. Let $x \in \mathbb{F}$ be such an evaluation point. Then $A(x), B(x), C'(x)$ are linear transformations over $\mathbb{F}$ for which we get a 3-slice $(A(x) \otimes B(x) \otimes C'(x))(T) = \mathrm{Id}_q + x N(x) \in M_q(\mathbb{F})$ of rank q, showing $Q_3(T) \geq q = \underline{Q}(T)$. $\qquad\square$

Combining Theorem 5.13 and Theorem 5.15, assuming $|\mathbb{F}| > \underline{Q}(T) + 1$, we find for any directions $i, j, k \in [3]$ such that $i \neq j$ the relation $\rho_{i,j}(T) \leq \underline{Q}(T) \leq Q_k(T)$.

5.3 Square-root lower bound for symmetric tensors

For symmetric tensors we can improve the cube-root lower bound of Theorem 5.14 to a square-root lower bound, using the notion of pivots.

Theorem 5.16. *Let $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ be symmetric and concise. Then $\underset{\sim}{\mathrm{Q}}(T) \geq \sqrt{n}$.*

In fact, we will prove a stronger theorem which implies Theorem 5.16. For this we will define the notion of being pivot-matched.

First, for any collection of matrices $A_1, \ldots, A_n \in \mathbb{F}^{m \times k}$ we define the pivot map

$$p = (f, g) : [n] \to [m] \times [k]$$

by mapping $\ell \in [n]$ to the pivot $p(A_\ell) = (f(A_\ell), g(A_\ell))$ of A_ℓ.

Definition 5.17 (Pivot-matched)**.** We call $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ *pivot-matched* if there are tensors T_A and T_B isomorphic to T such that, letting $A_i = \sum_{j,k} (T_A)_{i,j,k} \, e_j \otimes e_k \in \mathbb{F}^n \otimes \mathbb{F}^n$ for $i \in [n]$ be the 1-slices of T_A, and $p_A = (f_A, g_A)$ the pivot map of these matrices, and letting $B_k = \sum_{i,j} (T_B)_{i,j,k} \, e_i \otimes e_j \in \mathbb{F}^n \otimes \mathbb{F}^n$ for $k \in [n]$ be the 3-slices of T_B, and $p_B = (f_B, g_B)$ the pivot map of these matrices, we have that p_A and p_B are injective and for all $\ell \in [n]$ we have $f_A(\ell) = f_B(\ell)$.

The condition that p_A and p_B are injective can always be satisfied by Lemma 5.2.

We note that in the above definition, since pivots are defined using the lexicographic ordering, to make our proofs work it is crucial to define the slices A_i as stated (and not as their transpose) and the same for the B_i.

Remark 5.18. By reordering the 1-slices of T_A and the 3-slices of T_B, the definition is equivalent to requiring a seemingly weaker requirement of having the same number of slices with pivots in each row for the two sets of slices, i.e. $\forall \ell \in [n] : |f_A^{-1}(\ell)| = |f_B^{-1}(\ell)|$.

Theorem 5.19. *Let $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ be a concise tensor that is pivot-matched. Then*

$$\underset{\sim}{\mathrm{Q}}(T) \geq \sqrt{n}.$$

Remark 5.20. We do not know whether Theorem 5.16 can be used to prove that $\underset{\sim}{\mathrm{Q}}(T) \geq \sqrt{n}$ for general non-symmetric concise $T \in \mathbb{F}^{n \times n \times n}$. A natural approach is to symmetrize T by taking $\mathrm{sym}(T) = \boxtimes_{\pi \in S_3} \pi T$ (where we let the symmetric group S_3 act by permuting the three tensor legs) and apply Theorem 5.16 to $\mathrm{sym}(T)$. We note, however, that the asymptotic subrank of a tensor generally does not behave nicely with respect to symmetrization (in contrast with matrix rank). In particular, if we let $T = \langle n, 1, 1 \rangle$, which has asymptotic subrank 1, then the product of the cyclic permutations of T, $T \boxtimes (1,2,3)T \boxtimes (1,2,3)^2 T$, equals the matrix multiplication tensor $\langle n, n, n \rangle \in \mathbb{F}^{n^2 \times n^2 \times n^2}$ which has asymptotic subrank n^2. In other words, tensoring T with permutations of T can drastically increase the asymptotic subrank.

Lemma 5.21. *Suppose $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ is a concise tensor that is pivot-matched. Then $T \boxtimes T$ restricts to a tensor of the form*

$$\sum_{i=1}^{n} \left[e_i \otimes e_{f(i)} \right] \otimes \left[e_{f(i)} \otimes e_{g(i)} \right] \otimes \left[e_{h(i)} \otimes e_i \right]$$

$$+ \sum_{i=1}^{n} \sum_{j<k} \sum_{l,u,v=1}^{n} c_{ijkluv} \left[e_i \otimes e_j \right] \otimes \left[e_k \otimes e_l \right] \otimes \left[e_u \otimes e_v \right].$$

for $f, g, h : [n] \to [n]$ functions such that $i \mapsto (f(i), g(i))$ is injective.

Proof. Let $A_1, \ldots, A_n$ be the 1-slices of T_A as in Definition 5.17, and $B_1, \ldots, B_n$ the 3-slices of T_B. Without loss of generality, we may assume that

$$(A_i)_{p_A(j)} = (B_i)_{p_B(j)} = \delta_{ij} \tag{8}$$

where $p_A = (f_A, g_A)$ is the pivot map for the 1-slices of T_A (rows enumerated by direction 2) and $p_B = (f_B, g_B)$ is for the 3-slices of T_B (rows by direction 1), otherwise we can take invertible linear combinations of the slices to guarantee it.

Using these two slicings, we have the decomposition

$$T^{\boxtimes 2} \geq T_A \boxtimes T_B = \sum_{i,u=1}^{n} \sum_{j,v=1}^{n} \sum_{k,w=1}^{n} (A_i)_{jk}(B_w)_{uv}(e_i \otimes e_u) \otimes (e_j \otimes e_v) \otimes (e_k \otimes e_w).$$

We apply the following three linear maps to the respective legs of this tensor:

$$U = \sum_{i=1}^{n} E_{i,i} \otimes E_{f_A(i),f_A(i)}$$

$$V = \sum_{v=1}^{n} E_{f_B(v),f_B(v)} \otimes E_{g_B(v),g_B(v)}$$

$$W = \sum_{w=1}^{n} E_{g_A(w),g_A(w)} \otimes E_{w,w}.$$

where $E_{i,i}$ is the projector onto e_i. A small calculation then gives

$$T^{\boxtimes 2} \geq (U \otimes V \otimes W)(T_A \boxtimes T_B) =$$
$$\sum_{i,v,w=1}^{n} (A_i)_{f_B(v)\,g_A(w)} (B_w)_{f_A(i)\,g_B(v)} (e_i \otimes e_{f_A(i)}) \otimes (e_{f_B(v)} \otimes e_{g_B(v)}) \otimes (e_{g_A(w)} \otimes e_w).$$

This is the final tensor. We now proceed to show how it fits the desired form.

By definition of f_A as giving the first non-zero rows of the A slices, we get non-zero contributions only from the (i, v)-pairs satisfying $f_B(v) \geq f_A(i)$ from the first factor after the triple sum. This already shows that there is an upper-triangular structure present, by looking at the second component of the first leg and the first component of the second leg.

To finish showing $(U \otimes V \otimes W)(T_A \boxtimes T_B)$ is in the desired form we focus on the pairs (i, v) satisfying $f_A(i) = f_B(v)$, aiming to show that this fits the first sum. Let S denote the set of such pairs. Restricted to S, the triple sum becomes

$$\sum_{(i,v) \in S} \sum_{w=1}^{n} (A_i)_{f_A(i)\,g_A(w)} (B_w)_{f_B(v)\,g_B(v)} (e_i \otimes e_{f_A(i)}) \otimes (e_{f_B(v)} \otimes e_{g_B(v)}) \otimes (e_{g_A(w)} \otimes e_w).$$

By (8), we have that

$$(B_w)_{f_B(v)\,g_B(v)} = (B_w)_{p_B(v)} = \delta_{vw}.$$

Hence, the sum over w collapses to a single term and we are left with the tensor

$$\sum_{(i,v) \in S} (A_i)_{f_A(i)\,g_A(v)} (e_i \otimes e_{f_A(i)}) \otimes (e_{f_B(v)} \otimes e_{g_B(v)}) \otimes (e_{g_A(v)} \otimes e_v).$$

Fix i, j such that $f_A(i) = j$. Then v ranges over the set $f_B^{-1}(j)$, so all the B-slices whose pivot lies on row j. But by assumption, these are precisely the A-slices whose pivot lies

39

on row j. This implies that the coordinate $(j, g_A(v))$ is the pivot of some A-slice. Again by (8) we get that
$$(A_i)_{f_A(i) \, g_A(v)} = \delta_{iv}.$$
Thus our tensor looks like
$$\sum_{i=1}^{n} (e_i \otimes e_{f_A(i)}) \otimes (e_{f_B(i)} \otimes e_{g_B(i)}) \otimes (e_{g_A(i)} \otimes e_i).$$

Note that the maps $i \mapsto (i, f_A(i))$, $i \mapsto (f_B(i), g_B(i))$ and $i \to (g_A(i), i)$ are injective, as required (i.e. this first sum is a diagonal of length n). $\square$

Lemma 5.22. *Suppose $T \in \mathbb{F}^n \otimes \mathbb{F}^n \otimes \mathbb{F}^n$ restricts to a tensor S of the form*

$$\sum_{i=1}^{n} \left[e_i \otimes e_{f(i)} \right] \otimes \left[e_{f(i)} \otimes e_{g(i)} \right] \otimes \left[e_{h(i)} \otimes e_i \right]$$

$$+ \sum_{i=1}^{n} \sum_{j<k} \sum_{l,u,v=1}^{n} c_{ijkluv} \left[e_i \otimes e_j \right] \otimes \left[e_k \otimes e_l \right] \otimes \left[e_u \otimes e_v \right]$$

for $f, g, h : [n] \to [n]$ functions such that $i \mapsto (f(i), g(i))$ is injective, then

$$\underline{\mathrm{Q}}(T) \geq n.$$

Proof. We define the maps
$$A(\varepsilon) : [e_i \otimes e_j] \mapsto \varepsilon^{-j} [e_i \otimes e_j]$$
$$B(\varepsilon) : [e_i \otimes e_j] \mapsto \varepsilon^{i} [e_i \otimes e_j]$$
$$C(\varepsilon) : [e_i \otimes e_j] \mapsto [e_i \otimes e_j].$$

Then

$$(A(\varepsilon) \otimes B(\varepsilon) \otimes C(\varepsilon)) S = \sum_{i=1}^{n} \left[e_i \otimes e_{f(i)} \right] \otimes \left[e_{f(i)} \otimes e_{g(i)} \right] \otimes \left[e_{h(i)} \otimes e_i \right]$$

$$+ \sum_{i=1}^{n} \sum_{j<k} \sum_{l,u,v=1}^{n} \varepsilon^{k-j} c_{ijkluv} \left[e_i \otimes e_j \right] \otimes \left[e_k \otimes e_l \right] \otimes \left[e_u \otimes e_v \right].$$

while the first sum is a diagonal of length n. Therefore, $\underline{\mathrm{Q}}(T) \geq n$. $\square$

Proof of Theorem 5.19. By Lemma 5.21, $T^{\boxtimes 2}$ restricts to the upper-triangular form that is the hypothesis for Lemma 5.22, which in turn states that it degenerates to a diagonal tensor of size n. By Lemma 5.12 this lower bounds its asymptotic subrank $\underset{\sim}{\mathrm{Q}}(T^{\boxtimes 2}) \geq \underline{\mathrm{Q}}(T^{\boxtimes 2}) \geq n$. So $\underset{\sim}{\mathrm{Q}}(T) \geq \sqrt{n}$. $\square$

6 Asymptotic slice rank versus asymptotic subrank

In this section we prove a lower bound on the asymptotic subrank in terms of the asymptotic slice rank, thus bounding how different the asymptotic subrank and asymptotic slice rank can be. This requires some of the methods developed so far, and some extra ingredients. These extra ingredients are the notion of the minimal covering number of a matrix subspace, which is also known as non-commutative rank in the literature, and a Flanders-type relation to the max-rank.

6.1 Minimal covering number of matrix subspaces and slice spans

For any matrix subspace $\mathcal{A} \subseteq \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$ we define the *minimal covering number* min-cov$(\mathcal{A})$ as the smallest number $m_1 + m_2$ such that there are subspaces $V_1 \subseteq \mathbb{F}^{n_1}$ and $V_2 \subseteq \mathbb{F}^{n_2}$ with $\dim V_1 = m_1$ and $\dim V_2 = m_2$ such that $\mathcal{A} \subseteq V_1 \otimes \mathbb{F}^{n_2} + \mathbb{F}^{n_1} \otimes V_2$. In the literature, the parameter min-cov is also known as the non-commutative rank [FR04], and a matrix subspace $\mathcal{A}$ satisfying min-cov$(\mathcal{A}) \leq r$ is sometimes called r-decomposable [Mes85].

Lemma 6.1. *For any matrix subspace $\mathcal{A}$,* max-rank$(\mathcal{A}) \leq$ min-cov$(\mathcal{A})$.

Proof. There are subspaces $V_1 \subseteq \mathbb{F}^{n_1}$ and $V_2 \subseteq \mathbb{F}^{n_2}$ with $\dim(V_1) + \dim(V_2) =$ min-cov$(\mathcal{A})$ and such that $\mathcal{A} \subseteq V_1 \otimes \mathbb{F}^{n_2} + \mathbb{F}^{n_1} \otimes V_2$. Any matrix in $\mathcal{A}$ can then be written as $A + B$ for some $A \in V_1 \otimes \mathbb{F}^{n_2}$ and $B \in \mathbb{F}^{n_1} \otimes V_2$. Then rank$(A) \leq \dim V_1$ and rank$(B) \leq \dim(V_2)$. This gives rank$(A + B) \leq$ rank$(A) +$ rank$(B) \leq$ min-cov$(\mathcal{A})$, as required. $\square$

In Section 6.2 we discuss reverse versions of Lemma 6.2.

For any tensor $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$, let $\mathcal{A}_i$ denote the slice spans of T for $i \in [3]$. We define $\mathrm{SR}_i(T) =$ min-cov$(\mathcal{A}_i)$ for $i \in [3]$. Lemma 6.1 directly gives:

Lemma 6.2. *For any tensor T and every $i \in [3]$, we have* $\mathrm{Q}_i(T) \leq \mathrm{SR}_i(T)$.

Finally, it is easily seen how Q_i is a relaxation of Q, and similarly how SR is a relaxation of SR_i, leading to the following inequalities:

Lemma 6.3. *For any tensor T and every $i \in [3]$, we have* $\mathrm{Q}(T) \leq \mathrm{Q}_i(T)$.

Proof. Let $r = \mathrm{Q}(T)$. Then $T \geq \langle r \rangle$ and specifically there are matrices A, B, C such that $(A \otimes B \otimes C)T$ is a matrix of rank r, and a linear combination of the i-slices acted on by a matrix from the left and a matrix on the right. So the linear combination itself is of rank at least r. $\square$

Lemma 6.4. *For any tensor T and every $i \in [3]$, we have* $\mathrm{SR}(T) \leq \mathrm{SR}_i(T)$.

Proof. This inequality is immediate by noting that $\mathrm{SR}_i(T)$ is the minimal size of a slicerank decomposition with the added restriction that the used slices are only j-slices with $j \neq i$. $\square$

6.2 Flanders-type bounds on the max-rank

We know for any matrix subspace $\mathcal{A}$ that max-rank$(\mathcal{A}) \leq$ min-cov$(\mathcal{A})$ (Lemma 6.2). The goal of this section is to discuss inequalities in the reverse direction. While these are not new, we provide self-contained proofs for the convenience of the reader. Since the tensor parameters Q_i and SR_i are defined in terms of max-rank and min-cov we immediately get similar inequalities for those. We start with the most basic reverse inequality, which is due to Flanders [Fla62]:

Lemma 6.5. *Let $\mathcal{A}$ be a matrix subspace over the field $\mathbb{F}$. Suppose $|\mathbb{F}| >$ max-rank$(\mathcal{A})$. Then* min-cov$(\mathcal{A}) \leq 2$ max-rank$(\mathcal{A})$.

Lemma 6.5 requires the base field $\mathbb{F}$ to be large enough. For arbitrary fields, using a slightly larger constant on the right hand side, a similar inequality is still known to be true by a result of Haramaty and Shpilka [HS10, Lemma 3.5] (Lemma 3.7 in the Arxiv version).

Lemma 6.6. *Let $\mathcal{A}$ be any matrix subspace. Then* min-cov$(\mathcal{A}) \leq 4$ max-rank$(\mathcal{A})$.

The result in [HS10, Lemma 3.5] is stated in the more general setting of quadratic polynomials. We follow their proof closely, but present it for the special case of matrices (bilinear forms).

Proof. Throughout the proof we will consider matrices as vectors in the tensor product $\mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2}$. Denote $r = \text{max-rank}(\mathcal{A})$. We will find subspaces $U_1 \subseteq \mathbb{F}^{n_1}, V_1 \subseteq \mathbb{F}^{n_2}$, both of dimension r, such that restricting $\mathcal{A}$ to some complement subspaces for them $U_1^\perp, V_1^\perp$ we get maximum rank in $\mathcal{A}|_{U_1^\perp \times V_1^\perp}$ at most $\frac{r}{2}$ (as we are work over arbitrary fields, we are not assuming an inner product. $U_1^\perp, V_1^\perp$ are complements w.r.t. an arbitrary direct sum decomposition of $\mathbb{F}^{n_1}, \mathbb{F}^{n_2}$ as explained in detail below). Repeating this process, we get by induction subspaces $U_1, U_2 \ldots, U_m$ and $V_1, V_2, \ldots, V_m$ such that restricting $\mathcal{A}$ to complements of them all we have the trivial space of only the 0 matrix. The dimension of both U_i and V_i is at most $\frac{r}{2^{i-1}}$ so the dimension of both $\sum_{i=1}^m U_i$ and $\sum_{i=1}^m V_i$ is at most $\sum_{i=1}^m \frac{r}{2^{i-1}} \leq 2r$. Then $\mathcal{A} \subseteq (\sum_{i=1}^m U_i) \otimes \mathbb{F}^{n_2} + \mathbb{F}^{n_1} \otimes (\sum_{i=1}^m V_i)$ shows the desired min-cov$(\mathcal{A}) \leq 2r + 2r = 4r$.

We have left to show that indeed we can find U_1, V_1 as described above. Let $A \in \mathcal{A}$ be a maximal-rank element, so $\text{rank}(A) = r$. Let $A = \sum_{i=1}^r u_i \otimes v_i$ be a rank-decomposition of A. Define $U_1 := \text{span}\{u_1, \ldots, u_r\}$ and $V_1 := \text{span}\{v_1, \ldots, v_r\}$. By the fact that it is a rank decomposition we have that $\{u_1, \ldots, u_r\}$ and $\{v_1, \ldots, v_r\}$ are both linearly independent sets of vectors, showing $\dim(U_1) = \dim(V_1) = r$. Complete them to bases $u_1, \ldots, u_{n_1}$ and $v_1, \ldots, v_{n_2}$ of $\mathbb{F}^{n_1}$ and $\mathbb{F}^{n_2}$ respectively. Define $U_1^\perp := \text{span}\{u_{r+1}, \ldots, u_{n_1}\}$ and $V_1^\perp := \text{span}\{v_{r+1}, \ldots, v_{n_2}\}$ and denote the projectors onto them with respect to these bases by $\Pi_{U_1^\perp}$ and $\Pi_{V_1^\perp}$ respectively. Define the restricted subspace $\mathcal{A}|_{U_1^\perp \times V_1^\perp} := \{(\Pi_{U_1^\perp} \otimes \Pi_{V_1^\perp})(B) : B \in \mathcal{A}\}$. Let $B \in \mathcal{A}$ be any element in $\mathcal{A}$. Denote its restriction by $B|_{U_1^\perp \times V_1^\perp}$ and denote the rank of this restriction by $s := \text{rank}(B|_{U_1^\perp \times V_1^\perp})$. We want to show $s \leq \frac{r}{2}$. Write the restriction as a rank-decomposition $B|_{U_1^\perp \times V_1^\perp} = \sum_{i=1}^s x_i \otimes y_i$ for elements $x_i \in U_1^\perp$ and $y_i \in V_1^\perp$. Then $\{u_1, \ldots, u_r, x_1, \ldots, x_s\}$ and $\{v_1, \ldots, v_r, y_1, \ldots, y_s\}$ are both linearly independent. Thus, to write the unrestricted B we only need to introduce tensor products of either elements from U_1 with $\ell_i^{(2)} \in \mathbb{F}^{n_2}$ or tensor products of $\ell_i^{(1)} \in \mathbb{F}^{n_1}$ with elements from V_1, that is: $B = \sum_{i=1}^s x_i \otimes y_i + \sum_{i=1}^r (u_i \otimes \ell_i^{(2)} + \ell_i^{(1)} \otimes v_i)$. Decompose $\ell_i^{(2)} = \tilde{v}_i + \tilde{y}_i + \tilde{\ell}_i^{(2)}$ where $\tilde{v}_i \in V_1$, $\tilde{y}_i \in \text{span}\{y_1, \ldots, y_s\}$ and $\tilde{\ell}_i^{(2)}$ is independent of both subspaces (or 0). Do the same for the $\ell_i^{(1)}$'s. Denote $\tilde{y}_i = \sum_{j=1}^s b_{i,j} y_j$ and $\tilde{x}_i = \sum_{j=1}^s a_{i,j} x_j$ for $a_{i,j}, b_{i,j} \in \mathbb{F}$, and define $\ell_j'^{(1)} := \sum_{i=1}^r b_{i,j} u_i \in U_1$ and similarly $\ell_j'^{(2)} := \sum_{i=1}^r a_{i,j} v_i \in V_1$. Then

$$B = \underbrace{\sum_{i=1}^s (x_i + \ell_i'^{(1)}) \otimes (y_i + \ell_i'^{(2)})}_{=:C}$$

$$+ \underbrace{\sum_{i=1}^r (u_i \otimes (\tilde{v}_i + \tilde{\ell}_i^{(2)}) + (\tilde{u}_i + \tilde{\ell}_i^{(1)}) \otimes v_i) - \sum_{i=1}^s (\ell_i'^{(1)} \otimes \ell_i'^{(2)})}_{=:D}.$$

Note that $\{x_i + \ell_i'^{(1)}\}_{i=1}^s \cup \{u_i\}_{i=1}^r \cup \{\tilde{\ell}_i^{(1)} : i \in [r], \tilde{\ell}_i^{(1)} \neq 0\}$ is linearly independent and so is $\{y_i + \ell_i'^{(2)}\}_{i=1}^s \cup \{v_i\}_{i=1}^r \cup \{\tilde{\ell}_i^{(2)} : i \in [r], \tilde{\ell}_i^{(2)} \neq 0\}$. In particular we have $\text{rank}(C) = s$ and $B = C \oplus D$. Then $s + \text{rank}(D) = \text{rank}(C) + \text{rank}(D) = \text{rank}(B) \leq r$ so $\text{rank}(D) \leq r - s$. Recall that $A \in U_1 \otimes V_1$ so $C \in \text{span}\{x_i + \ell_i'^{(1)}\}_{i=1}^s \otimes \text{span}\{y_i + \ell_i'^{(2)}\}_{i=1}^s$ and $D + A$ are supported on linearly independent vectors (in both factors), implying the rank is additive on them: $r \geq \text{rank}(B + A) = \text{rank}(C + D + A) = \text{rank}(C) + \text{rank}(D + A) \geq$

42

$s + (\operatorname{rank}(A) - \operatorname{rank}(D)) \geq s + (r - (r-s)) = 2s$ where the first inequality is by $A + B \in \mathcal{A}$, the second is generally true for any two matrices and the third is the rank of A and the inequality for the rank of D we established earlier. $\qquad\square$

From the above follows immediately the corresponding inequalities for the tensor parameters Q_i and SR_i.

Lemma 6.7. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$. Suppose $|\mathbb{F}| > \mathrm{Q}_i(T)$. Then $\mathrm{SR}_i(T) \leq 2\,\mathrm{Q}_i(T)$ for every $i \in [3]$.*

Lemma 6.8. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$. Then $\mathrm{SR}_i(T) \leq 4\,\mathrm{Q}_i(T)$ for every $i \in [3]$.*

6.3   Lower bound on asymptotic subrank in terms of slice rank

In previous sections we have proved lower bounds on the asymptotic subrank of concise tensors that only depend on the dimensions of the tensor. Here we shift gears, and use some of the methods that we have developed so far, including the Flanders-type bounds of Section 6.2, to prove a lower bound on the asymptotic subrank in terms of the slice rank. In fact, we will prove a lower bound on the border subrank in terms of slice rank. As a consequence, we obtain a bound on the gap between asymptotic subrank and asymptotic slice rank.

Lemma 6.9. *Let $T \in \mathbb{F}^{n_1} \otimes \mathbb{F}^{n_2} \otimes \mathbb{F}^{n_3}$. Then $\underline{\mathrm{Q}}(T^{\boxtimes 3}) \geq \frac{3}{64}\,\mathrm{SR}(T)^2$.*

Proof. By Lemma 2.13 we have $T^{\boxtimes 3} \geq \langle \mathrm{Q}_1(T), \mathrm{Q}_2(T), \mathrm{Q}_3(T)\rangle$. Strassen [Str87, Theorem 6.6] proved for any $e \leq h \leq \ell$ that

$$\underline{\mathrm{Q}}(\langle e, h, \ell\rangle) \geq \begin{cases} eh - \lfloor \frac{(e+h-\ell)^2}{4} \rfloor & e + h \geq \ell \\ eh & \text{otherwise.} \end{cases}$$

In both cases we have $\underline{\mathrm{Q}}(\langle e, h, \ell\rangle) \geq \frac{3}{4}eh$ (the worst case being $\ell = h$). Combining these we get

$$\underline{\mathrm{Q}}(T^{\boxtimes 3}) \geq \tfrac{3}{4}\,\mathrm{Q}_1(T) \cdot \mathrm{Q}_2(T) \geq \tfrac{3}{4}(\tfrac{1}{4}\,\mathrm{SR}_1(T)) \cdot (\tfrac{1}{4}\,\mathrm{SR}_2(T)) \geq \tfrac{3}{4}(\tfrac{1}{4}\,\mathrm{SR}(T))^2$$

for which we used Lemma 6.8 and Lemma 6.3. $\qquad\square$

Theorem 6.10. *Suppose the limit $\widetilde{\mathrm{SR}}(T) = \lim_{n\to\infty} \mathrm{SR}(T^{\boxtimes m})^{1/m}$ exists. Then*

$$\widetilde{\mathrm{SR}}(T) \geq \widetilde{\mathrm{Q}}(T) \geq \widetilde{\mathrm{SR}}(T)^{2/3}.$$

Proof. The first inequality follows directly from the fact that slice rank upper bounds subrank. For the second inequality, we have from Lemma 5.12 and Lemma 6.9 that $\mathrm{Q}(T) \geq \underline{\mathrm{Q}}(T^{\boxtimes 3})^{1/3} \geq (\frac{3}{64}\,\mathrm{SR}(T)^2)^{1/3}$. Now replace T by $T^{\boxtimes m}$, take the m'th root on both sides and let m go to infinity to get the claim. $\qquad\square$

Remark 6.11. The constant in Lemma 6.9 can be slightly improved if we replace the border subrank by the asymptotic subrank. Namely, for any 3-tensor T, we have $\widetilde{\mathrm{Q}}(T) \geq (\frac{1}{4}\,\mathrm{SR}(T))^{2/3}$. The proof is as follows. We have $\underline{\mathrm{Q}}(T^{\boxtimes 3}) \geq \min_{i,j} \mathrm{Q}_i(T)\,\mathrm{Q}_j(T)$ (Lemma 2.14). Combining with $\mathrm{Q}_i(T) \geq \frac{1}{4}\,\mathrm{SR}_i(T)$ (Lemma 6.8) and $\mathrm{SR}_i(T) \geq \mathrm{SR}(T)$ (Lemma 6.3) gives $\widetilde{\mathrm{Q}}(T)^3 = \widetilde{\mathrm{Q}}(T^{\boxtimes 3}) \geq \min_{i,j}(\frac{1}{4}\,\mathrm{SR}_i(T))(\frac{1}{4}\,\mathrm{SR}_j(T)) \geq (\frac{1}{4}\,\mathrm{SR}(T))^2$. Taking the cube root on both sides gives the claim. Another situation in which we can improve the constant is when the base field $\mathbb{F}$ is large enough, which allows using $\mathrm{Q}_i(T) \geq \frac{1}{2}\,\mathrm{SR}_i(T)$ (Lemma 6.7) instead of $\mathrm{Q}_i(T) \geq \frac{1}{4}\,\mathrm{SR}_i(T)$ (Lemma 6.8) in the above proofs, leading to a stronger bound.

References

[AHSW06] Anura Abeyesinghe, Patrick Hayden, Graeme Smith, and Andreas J. Winter. Optimal superdense coding of entangled states. *IEEE Trans. Inform. Theory*, 52(8):3635–3641, 2006. `doi:10.1109/TIT.2006.878174`. 12

[Alm19] Josh Alman. Limits on the universal method for matrix multiplication. In *Proceedings of the 34th Computational Complexity Conference (CCC 2019)*, pages 12:1–12:24, 2019. `arXiv:1812.08731`, `doi:10.4230/LIPIcs.CCC.2019.12`. 5

[Ami65] Shimshon Avraham Amitsur. Generalized polynomial identities and pivotal monomials. *Trans. Amer. Math. Soc.*, 114:210–226, 1965. `doi:10.2307/1993998`. 12

[AW21] Josh Alman and Virginia Vassilevska Williams. A refined laser method and faster matrix multiplication. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA 2021)*, pages 522–539. SIAM, 2021. `doi:10.1137/1.9781611976465.32`. 4

[BCC+17] Jonah Blasiak, Thomas Church, Henry Cohn, Joshua A. Grochow, Eric Naslund, William F. Sawin, and Chris Umans. On cap sets and the group-theoretic approach to matrix multiplication. *Discrete Anal.*, 2017. `arXiv:1605.06702`, `doi:10.19086/da.1245`. 4, 5

[BCG+22] Jonah Blasiak, Henry Cohn, Joshua A. Grochow, Kevin Pratt, and Chris Umans. Matrix multiplication via matrix groups, 2022. `arXiv:2204.03826`. 4

[BCS97] Peter Bürgisser, Michael Clausen, and Mohammad Amin Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1997. `doi:10.1007/978-3-662-03338-8`. 4, 12, 13, 36

[BDR22a] Andreas Blatter, Jan Draisma, and Filip Rupniewski. Countably many asymptotic tensor ranks, 2022. `doi:10.48550/ARXIV.2212.12219`. 6, 8

[BDR22b] Andreas Blatter, Jan Draisma, and Filip Rupniewski. A tensor restriction theorem over finite fields, 2022. `doi:10.48550/ARXIV.2211.12319`. 6, 7

[BI11]   Peter Bürgisser and Christian Ikenmeyer. Geometric complexity the-
           ory and tensor rank. In *Proceedings of the 43rd ACM Symposium
           on Theory of Computing (STOC 2011)*, pages 509–518. ACM, 2011.
           `doi:10.1145/1993636.1993704`. 4

[BL20]   Markus Bläser and Vladimir Lysikov. Slice rank of block tensors and irre-
           versibility of structure tensors of algebras. In *45th International Symposium
           on Mathematical Foundations of Computer Science (MFCS 2020)*, pages 17:1–
           17:15, 2020. `doi:10.4230/LIPIcs.MFCS.2020.17`. 5

[Blä13]   Markus Bläser. *Fast Matrix Multiplication*. Number 5 in Graduate Surveys.
           Theory of Computing Library, 2013. `doi:10.4086/toc.gs.2013.005`. 4

[CD21]   Simone Costa and Marco Dalai. A gap in the slice rank of k-tensors. *J. Comb.
           Theory, Ser. A*, 177:105335, 2021. `doi:10.1016/j.jcta.2020.105335`. 6, 8

[CGZ22]   Matthias Christandl, Fulvio Gesmundo, and Jeroen Zuiddam. A gap in the
           subrank of tensors, 2022. `doi:10.48550/ARXIV.2212.01668`. 6, 8

[CHL+08]   Toby Cubitt, Aram Wettroth Harrow, Debbie Leung, Ashley Montanaro,
           and Andreas Winter. Counterexamples to additivity of minimum output p-
           Rényi entropy for p close to 0. *Comm. Math. Phys.*, 284(1):281–290, 2008.
           `doi:10.1007/s00220-008-0625-z`. 12

[CK11]   Enrico Carlini and Johannes Kleppe. Ranks derived from mul-
           tilinear maps. *J. Pure Appl. Algebra*, 215(8):1999–2004, 2011.
           `doi:10.1016/j.jpaa.2010.11.010`. 13

[CKSU05]   Henry Cohn, Robert Kleinberg, Balazs Szegedy, and Christopher Umans.
           Group-theoretic algorithms for matrix multiplication. In *46th Annual IEEE
           Symposium on Foundations of Computer Science (FOCS 2005)*, pages 379–388.
           IEEE, 2005. `doi:10.1109/SFCS.2005.39`. 4

[CM21]   Alex Cohen and Guy Moshkovitz. Structure vs. randomness for bilinear maps.
           In *53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC
           2021)*, pages 800–808. ACM, 2021. `doi:10.1145/3406325.3451007`. 5

[CM23]   Alex Cohen and Guy Moshkovitz. Partition and analytic rank are equiv-
           alent over large fields. *Duke Mathematical Journal*, pages 1–38, 2023.
           `doi:10.1215/00127094-2022-0086`. 5

[CMW08]   Toby Cubitt, Ashley Montanaro, and Andreas Winter. On the dimension of
           subspaces with bounded Schmidt rank. *J. Math. Phys.*, 49(2):022107, 6, 2008.
           `doi:10.1063/1.2862998`. 12

[Coh95]   Paul M. Cohn. *Skew fields*, volume 57 of *Encyclopedia of Mathematics and
           its Applications*. Cambridge University Press, Cambridge, 1995. Theory of
           general division rings. `doi:10.1017/CBO9781139087193`. 11

[CU03]   Henry Cohn and Christopher Umans. A group-theoretic approach to fast
           matrix multiplication. In *Proceedings of the 44th Annual IEEE Symposium on
           Foundations of Computer Science (FOCS 2003)*, pages 438–449. IEEE, 2003.
           `doi:10.1109/SFCS.2003.1238217`. 4

[CU13]   Henry Cohn and Christopher Umans. Fast matrix multiplication using coherent configurations. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2013)*, pages 1074–1086. SIAM, 2013. 4

[CVZ21]   Matthias Christandl, Péter Vrana, and Jeroen Zuiddam. Barriers for fast matrix multiplication from irreversibility. *Theory of Computing*, 17(2):1–32, 2021. URL: http://www.theoryofcomputing.org/articles/v017a002, doi:10.4086/toc.2021.v017a002. 5

[CVZ23]   Matthias Christandl, Péter Vrana, and Jeroen Zuiddam. Universal points in the asymptotic spectrum of tensors. *J. Amer. Math. Soc.*, 36(1):31–79, 2023. doi:10.1090/jams/996. 5, 6, 8, 9, 10, 14, 33

[CW90]   Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *J. Symbolic Comput.*, 9(3):251–280, 1990. doi:10.1016/S0747-7171(08)80013-2. 9

[Der22]   Harm Derksen. The G-stable rank for tensors and the cap set problem. *Algebra Number Theory*, 16(5):1071–1097, 2022. doi:10.2140/ant.2022.16.1071. 5

[Die49]   Jean Dieudonné. Sur une généralisation du groupe orthogonal à quatre variables. *Arch. Math.*, 1:282–287, 1949. doi:10.1007/BF02038756. 11

[DL78]   Richard A. Demillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193–195, 1978. doi:10.1016/0020-0190(78)90067-4. 16

[DM18]   Harm Derksen and Visu Makam. On non-commutative rank and tensor rank. *Linear and Multilinear Algebra*, 66(6):1069–1084, 2018. doi:10.1080/03081087.2017.1337058. 11

[DMZ22]   Harm Derksen, Visu Makam, and Jeroen Zuiddam. Subrank and Optimal Reduction of Scalar Multiplications to Generic Tensors. In *37th Computational Complexity Conference (CCC 2022)*, page 9:1–9:23, 2022. doi:10.4230/LIPIcs.CCC.2022.9. 10

[DWZ23]   Ran Duan, Hongxun Wu, and Renfei Zhou. Faster matrix multiplication via asymmetric hashing, 2023. arXiv:2210.10173. 4

[EG17]   Jordan S. Ellenberg and Dion Gijswijt. On large subsets of with no three-term arithmetic progression. *Annals of Mathematics*, page 339–343, 2017. doi:10.4007/annals.2017.185.1.8. 5

[Fla62]   Harley Flanders. On spaces of linear transformations with bounded rank. *J. London Math. Soc.*, 37:10–16, 1962. doi:10.1112/jlms/s1-37.1.10. 11, 41

[FR04]   Marc Fortin and Christophe Reutenauer. Commutative/noncommutative rank of linear matrices and subspaces of matrices of low rank. *Sém. Lothar. Combin.*, 52:Art. B52f, 12, 2004. URL: http://eudml.org/doc/125000. 11, 41

[GW11]   W. T. Gowers and J. Wolf. Linear forms and higher-degree uniformity for functions on $\mathbf{F}_p^n$. *Geom. Funct. Anal.*, 21(1):36–69, 2011. doi:10.1007/s00039-010-0106-3. 5

[GZ23]   Fulvio Gesmundo and Jeroen Zuiddam. The next gap in the subrank of 3-tensors, 2023. `arXiv:2307.06115`. 6

[Hay04]   Patrick Hayden. Entanglement in random subspaces. In *Quantum communication, measurement and computing*, volume 734 of *AIP Conf. Proc.*, pages 226–229. Amer. Inst. Phys., Melville, NY, 2004. `doi:10.1063/1.1834421`. 12

[HLW06]   Patrick Hayden, Debbie W. Leung, and Andreas Winter. Aspects of generic entanglement. *Comm. Math. Phys.*, 265(1):95–117, 2006. `doi:10.1007/s00220-006-1535-6`. 12

[HS10]   Elad Haramaty and Amir Shpilka. On the structure of cubic and quartic polynomials. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC 2010)*, pages 331–340, 2010. `arXiv:0908.2853`, `doi:10.1145/1806689.1806736`. 10, 41, 42

[KMZ20]   Swastik Kopparty, Guy Moshkovitz, and Jeroen Zuiddam. Geometric Rank of Tensors and Subrank of Matrix Multiplication. In *35th Computational Complexity Conference (CCC 2020)*, 2020. `doi:10.4230/LIPIcs.CCC.2020.35`. 5

[KSS18]   Robert Kleinberg, Will Sawin, and David Speyer. The growth rate of tri-colored sum-free sets. *Discrete Analysis*, 2018. `doi:10.19086/da.3734`. 9

[LG14]   François Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation (ISSAC 2014)*, pages 296–303. ACM, 2014. `doi:10.1145/2608628.2608664`. 4

[LO13]   J. M. Landsberg and Giorgio Ottaviani. Equations for secant varieties of Veronese and other varieties. *Ann. Mat. Pura Appl. (4)*, 192(4):569–606, 2013. `doi:10.1007/s10231-011-0238-6`. 4

[Lov19]   Shachar Lovett. The analytic rank of tensors and its applications. *Discrete Anal.*, pages Paper No. 7, 10, 2019. `arXiv:1806.09179`, `doi:10.19086/da.8654`. 5

[Mes85]   Roy Meshulam. On the maximal rank in a subspace of matrices. *Quart. J. Math. Oxford Ser. (2)*, 36(142):225–229, 1985. `doi:10.1093/qmath/36.2.225`. 11, 34, 35, 41

[MS04]   Roy Meshulam and Peter Semrl. Locally linearly dependent operators and reflexivity of operator spaces. *Linear Algebra Appl.*, 383:143–150, 2004. `doi:10.1016/j.laa.2003.12.024`. 12

[MV04]   Akimasa Miyake and Frank Verstraete. Multipartite entanglement in $2 \times 2 \times n$ quantum systems. *Phys. Rev. A (3)*, 69(1):012101, 9, 2004. `doi:10.1103/PhysRevA.69.012101`. 29

[NS17]   Eric Naslund and Will Sawin. Upper bounds for sunflower-free sets. *Forum of Mathematics, Sigma*, 5:e15, 2017. `doi:10.1017/fms.2017.12`. 5

[Par04]   Kalyanapuram Rangachari Parthasarathy. On the maximal dimension of a completely entangled subspace for finite level quantum systems. *Proc. Indian Acad. Sci. Math. Sci.*, 114(4):365–374, 2004. `doi:10.1007/BF02829441`. 12

[Sch80]   Jacob T. Schwartz.   Fast probabilistic algorithms for verification of polynomial identities.   *Journal of the ACM*, 27(4):701–717, 1980. `doi:10.1145/322217.322225`. 16

[ST16]   Will Sawin and Terence Tao. Notes on the "slice rank" of tensors, 2016. URL: `https://terrytao.wordpress.com/2016/08/24/`. 5, 33

[Str69]   Volker Strassen.   Gaussian elimination is not optimal.   *Numer. Math.*, 13(4):354–356, 1969. `doi:10.1007/BF02165411`. 4

[Str87]   Volker Strassen.   Relative bilinear complexity and matrix multiplication.   *J. Reine Angew. Math.*, 375/376:406–443, 1987. `doi:10.1515/crll.1987.375-376.406`. 5, 8, 12, 43

[Str88]   Volker Strassen. The asymptotic spectrum of tensors. *J. Reine Angew. Math.*, 384:102–152, 1988. `doi:10.1515/crll.1988.384.102`. 6, 8, 12, 19, 36

[Str91]   Volker Strassen.   Degeneration and complexity of bilinear maps: some asymptotic spectra.   *J. Reine Angew. Math.*, 413:127–180, 1991. `doi:10.1515/crll.1991.413.127`. 8, 9, 10, 12, 17, 33

[Tao16]   Terence Tao.   A symmetric formulation of the Croot-Lev-Pach-Ellenberg-Gijswijt capset bound, 2016. URL: `https://terrytao.wordpress.com/2016/05/18/`. 5, 10

[Tob97]   Verena Tobler.   Spezialisierung und degeneration von tensoren.   Konstanzer Schriften in Mathematik und Informatik; 47, 1997. URL: `http://nbn-resolving.de/urn:nbn:de:bsz:352-opus-20324`. 30

[Wal02]   Nolan Russell Wallach. An unentangled Gleason's theorem. In *Quantum computation and information*, volume 305 of *Contemp. Math.*, pages 291–298. Amer. Math. Soc., Providence, RI, 2002. `doi:10.1090/conm/305/05226`. 12

[WDGC13]   Michael Walter, Brent Doran, David Gross, and Matthias Christandl. Entanglement polytopes: multiparticle entanglement from single-particle information. *Science*, 340(6137):1205–1208, 2013. `doi:10.1126/science.1232957`. 33

[WXXZ23]   Virginia Vassilevska Williams, Yinzhan Xu, Zixuan Xu, and Renfei Zhou. New bounds for matrix multiplication: from alpha to omega, 2023. `arXiv:2307.07970`. 4

[WZ22]   Avi Wigderson and Jeroen Zuiddam.   Asymptotic spectra: Theory, applications and extensions, 2022. URL: `https://staff.fnwi.uva.nl/j.zuiddam/papers/convexity.pdf`. 4

[Zip79]   Richard Zippel. *Probabilistic algorithms for sparse polynomials*. Springer, 1979. `doi:10.1007/3-540-09519-5_73`. 16