

NOTIONS OF TENSOR RANK

MANDAR JUVEKAR

ABSTRACT. Tensors, or multi-linear forms, are important objects in a variety of areas from analytics, to combinatorics, to computational complexity theory. Notions of tensor rank aim to quantify the “complexity” of these forms, and are thus also important. While there is one single definition of rank that completely captures the complexity of matrices (and thus linear transformations), there is no definitive analog for tensors. Rather, many notions of tensor rank have been defined over the years, each with their own set of uses. In this paper we survey the popular notions of tensor rank. We give a brief history of their introduction, motivating their existence, and discuss some of their applications in computer science. We also give proof sketches of recent results by Lovett, and Cohen and Moshkovitz, which prove asymptotic equivalence between three key notions of tensor rank over finite fields with at least three elements.

1. INTRODUCTION

We first come across the notion of rank in a course on linear algebra. If A is an $m \times n$ matrix over a field \mathbb{F} , the rank of A is the dimension of the space spanned by its rows (or columns). To help generalize this definition, it will be useful to reinterpret the matrix A as a bilinear form $T : \mathbb{F}^m \times \mathbb{F}^n \rightarrow \mathbb{F}$ by the natural correspondence

$$(x, y) \in \mathbb{F}^m \times \mathbb{F}^n \mapsto T(x, y) = \sum_{i,j} A_{i,j} x^i y^j.$$

As a notational convention, here and elsewhere in this paper we will use superscripts to index coordinates of a vector, and subscripts to index over different vectors. For example, if x_1, x_2, \dots are vectors, x_5^7 would be the seventh coordinate of the fifth vector. Getting back to ranks, using this correspondence, we can formulate an alternative definition of rank where the rank of T is the minimal natural number r such that T can be written as a sum of r bilinear forms of “lowest” complexity, or *rank 1 matrices*. The natural objects of lowest complexity are the linear 1-forms, i.e., the dot product with a fixed vector. Products of 1-forms $T_1(x)T_2(y)$ are then bilinear forms. Since every bilinear form can be written as a finite sum of forms of this type, we set them to be our rank 1 bilinear forms (matrices). Putting everything together, our alternative definition then says that the rank of a matrix A is the minimum natural number r such that the bilinear

Date: May 2022.

form T corresponding to A can be written as the sum of r bilinear forms each of the type $T_1(x)T_2(y)$ where T_1 and T_2 are 1-forms. This alternative definition agrees with the usual linear algebra definition of rank (see for instance [Hal58, Sections 32 and 51]).

In this paper we focus on generalizing the notion of rank to higher-dimensional analogs of matrices. What are these higher-dimensional matrices? The higher-dimensional analogs we will look at are the d -tensors, which are defined as follows.

Definition 1.1 (Tensor). Let \mathbb{F} be a field. A d -tensor T on $\mathbb{F}^{n_1} \times \dots \times \mathbb{F}^{n_d}$ is a multilinear map $T : \mathbb{F}^{n_1} \times \dots \times \mathbb{F}^{n_d} \rightarrow \mathbb{F}$.

Note that a d -tensor $T : \mathbb{F}^{n_1} \times \dots \times \mathbb{F}^{n_d} \rightarrow \mathbb{F}$ can naturally be identified with a d -dimensional array M such that

$$T(x_1, \dots, x_d) = \sum_{i_1 \in [n_1], \dots, i_d \in [n_d]} M_{i_1, \dots, i_d} x_1^{i_1} \dots x_d^{i_d}.$$

So, following the convention for matrices, we will often denote the space of d -tensors on $\mathbb{F}^{n_1} \times \dots \times \mathbb{F}^{n_d}$ by $\mathbb{F}^{n_1 \times \dots \times n_d}$.¹

The definition of rank given above generalizes well to arbitrary d -tensors. Setting the rank 1 d -tensors to products of d linear 1-forms (or, since linear 1-forms are the same as 1-tensors, d 1-tensors) leads to the notion of rank that is traditionally associated with tensors. We call it *traditional rank*, or TR for short.

Definition 1.2 (Traditional Rank). A d -tensor T has *traditional rank 1* if we can write

$$T(x_1, \dots, x_d) = T_1(x_1)T_2(x_2) \cdots T_d(x_d),$$

where each T_j is a 1-tensor.

The *traditional rank* of an arbitrary d -tensor T , $\text{TR}(T)$, is the minimum number r such that we can write

$$T(x) = \sum_{i=1}^r T_i(x)$$

where each T_i is a d -tensor with traditional rank 1.

While this definition does generalize matrix rank it turns out that there are also other, non-equivalent generalizations of matrix rank to arbitrary tensors that characterize the combinatorial, analytic, and geometric properties of tensors. Each of these notions is useful in its own way. Finding tight relationships between these notions remains an open research question.

One disadvantage of the traditional notion of tensor rank is that it is prohibitively hard to compute in general. Since Håstad's work in 1989 [Hås89] it has been known that computing traditional tensor rank over finite fields is NP-complete, and over \mathbb{Q} is NP-hard. Later on, Hillar and Lim [HL13]

¹Here and elsewhere in the paper we will use $[n]$ to denote the set $\{1, 2, \dots, n\}$.

showed that Håstad’s proof could be modified to show that traditional rank is NP-hard over \mathbb{R} and \mathbb{C} as well. Traditional rank also turns out to be NP-hard to approximate with arbitrarily small error bounds [Swe18]. The non-traditional notions we will discuss are less “stringent” than traditional rank, so it is possible that they are easier to compute. Showing (exact or approximate) hardness results for those notions, however, is still an open problem.

In this paper we introduce the landscape of notions of tensor rank, motivating their existence (Section 2). We then give examples of applications of these notions in computational complexity theory (Section 3). In Section 4 we give a rundown of the trivial relationships between the notions that come from their definitions. We then give sketches of results by Lovett [Lov19] (Section 5) and Cohen and Moshkovitz [CM21] (Section 6) which, put together, prove asymptotic equivalence between three key notions of rank over finite fields with three or more elements.

2. TENSOR RANK THROUGH THE AGES

The oldest non-traditional notion of tensor rank is the so-called *analytic rank*.

Definition 2.1 (Analytic Rank). Let \mathbb{F} be a finite field. The *bias* of a d -tensor $T \in \mathbb{F}^{n_1 \times \dots \times n_d}$ is given by

$$\text{bias}(T) = \mathbb{E}_{(x_1, \dots, x_d) \in \mathbb{F}^{n_1} \times \dots \times \mathbb{F}^{n_d}} \chi(T(x_1, \dots, x_d)),$$

where χ is a nontrivial additive character which, in the case that $\mathbb{F} = \mathbb{F}_p$, can be taken to be $\chi(x) = e^{2\pi i x/p}$.

The *analytic rank* of T , denoted by $\text{AR}(T)$, is then given by

$$\text{AR}(T) = -\log_{|\mathbb{F}|} \text{bias}(T).$$

This measure of the complexity of a tensor was first introduced by Gowers and Wolf in the context of higher-order Fourier analysis [GW11]. In Section 5 we will see the connection between the analytic rank and the other combinatorial and geometric notions described below.

Going back to the traditional definition of tensor rank, there is no reason why one cannot use other kinds of tensors as rank 1 tensors, as long as the new notion agrees with the old in the case of matrices (in order to be a true generalization). Doing so, it turns out, is useful for a whole host of applications. The first such “alternative” notion was introduced in 2016 in the context of the so-called “capset problem.”

A *capset* in \mathbb{F}_3 is a subset of \mathbb{F}_3 with no non-trivial three-term arithmetic progressions. That is, a capset is a set $A \subseteq \mathbb{F}_3$ that does not contain $\{x, x+r, x+2r\}$ for any $x, r \in \mathbb{F}_3$ with $r \neq 0$. The capset problem asks what the maximum size of a capset in \mathbb{F}_3 can be. In the spring of 2016, Croot, Lev, and Pach [CLP17] proved a breakthrough result for a similar problem, but in the additive group $\mathbb{Z}/4\mathbb{Z}$. They showed that if $A \subseteq (\mathbb{Z}/4\mathbb{Z})^n$ contains no

non-trivial three-term arithmetic progressions, then $|A| \leq 3.60172^n$. Soon after, Ellenberg and Gijswijt [EG17] generalized the Croot-Lev-Pach argument to show that any progression-free set $A \subseteq (\mathbb{Z}/p\mathbb{Z})^n$, where p is a prime, satisfies $|A| \leq (J(p)p)^n$ where $J(p)$ is an explicit constant less than 1. This exponentially improved the known bound for the capset problem, bringing it down from $O(3^n/n^{1+c})$ (where c is some absolute constant) to $O(2.756^n)$. Tao [Tao16] reformulated the Ellenberg-Gijswijt argument in terms of tensors, introducing what is now known as the *slice rank* of a tensor.

Definition 2.2 (Slice Rank). A d -tensor T has *slice rank 1* if we can write

$$T(x_1, \dots, x_d) = T_1(x_i)T_2(x_j : j \neq i),$$

where $i \in [d]$, T_1 is a 1-tensor, and T_2 is a $(d-1)$ -tensor.

The *slice rank* of an arbitrary d -tensor T , $\text{SR}(T)$, is the minimum number r such that we can write

$$T(x) = \sum_{i=1}^r T_i(x)$$

where each T_i is a d -tensor with slice rank 1.

Notice that in essence this definition just modifies Definition 1.2 so that our rank 1 tensors go from being products of d 1-tensors to being the product of two tensors: one of order 1 and one of order $d-1$. Also notice that in the case of matrices (which are 2-tensors), Definitions 2.2 and 1.2 are exactly the same and agree with the usual definition of matrix rank.

Slice rank, as we defined above, defines its rank 1 tensors by “slicing” off one coordinate and multiplying a 1-tensor applied to that coordinate with a $(d-1)$ -tensor applied to the remaining coordinates. A natural extension to this is to, instead of slicing off one coordinate, allow arbitrary partitions of the coordinates into two parts. This was done in 2017 (preprint; published in 2020) by Naslund [Nas20], who called this new notion of rank *partition rank*. We formally define partition rank as follows.

Definition 2.3 (Partition Rank). A d -tensor T has *partition rank 1* if we can write

$$T(x_1, \dots, x_d) = T_1(x_i : i \in S)T_2(x_j : j \notin S),$$

where $S \subset [d]$ with $1 \leq |S| < d$, T_1 is a $|S|$ -tensor, and T_2 is a $(d-|S|)$ -tensor.

The *partition rank* of an arbitrary d -tensor T , $\text{PR}(T)$, is the minimum number r such that we can write

$$T(x) = \sum_{i=1}^r T_i(x)$$

where each T_i is a d -tensor with partition rank 1.

Using this new notion of rank, Naslund showed that any set $A \subseteq \mathbb{F}_q^n$ with size bigger than $\binom{n+(k-1)q}{(k-1)(q-1)}$ must contain distinct vectors x_1, x_2, \dots, x_{k+1}

such that the vectors $x_1 - x_{k+1}, x_2 - x_{k+1}, \dots, x_k - x_{k+1}$ are mutually orthogonal (he called a collection of such x_1, \dots, x_{k+1} a *k-right corner*).

Our last, most recent notion of rank is motivated by an age-old, fundamental question in computer science: asymptotically how many scalar additions and multiplications are necessary to multiply two $n \times n$ matrices? Tensors are naturally connected to this problem because the operation of matrix multiplication itself can be thought of as a tensor.

Recall that matrices M (2-tensors) can simultaneously be thought of as bilinear forms $T : \mathbb{F}^m \times \mathbb{F}^n \rightarrow \mathbb{F}$ that map $(x, y) \mapsto x^T M y$ and linear maps $T' : \mathbb{F}^m \rightarrow \mathbb{F}^n$ that map $x \mapsto Mx$. Similarly, we can also think of d -tensors

$$T(x_1, \dots, x_d) = \sum_{i_1 \in [n_1], \dots, i_d \in [n_d]} T_{i_1, \dots, i_d} x_1^{i_1} \cdots x_d^{i_d}$$

as $(d-1)$ -linear maps $T' : \mathbb{F}^{n_1} \times \cdots \times \mathbb{F}^{n_{d-1}} \rightarrow \mathbb{F}^{n_d}$ given by

$$[T'(x_1, \dots, x_{d-1})]_k = \sum_{i_1 \in [n_1], \dots, i_{d-1} \in [n_{d-1}]} T'_{i_1, \dots, i_{d-1}, k} x_1^{i_1} \cdots x_{d-1}^{i_{d-1}}.$$

By the usual formula for matrix multiplication, if M and N are $n \times n$ matrices over \mathbb{F} , then MN is an $n \times n$ matrix with

$$(MN)_{ij} = \sum_{\ell=1}^n M_{i\ell} N_{\ell j}.$$

Identifying the space of $n \times n$ matrices with \mathbb{F}^{n^2} , the operation of matrix multiplication takes two elements of \mathbb{F}^{n^2} to another element of \mathbb{F}^{n^2} where each entry is a bilinear form applied to M and N . Thus matrix multiplication is a 3-tensor over the space of $n \times n$ matrices. As we discuss in Section 3 the rank of this tensor is intimately related to the computational complexity of matrix multiplication.

Motivated by this application (as well as a few others), Kopparty, Moshkovitz, and Zuiddam [KMZ20] recently introduced the notion of *geometric rank*. Unlike the other types of ranks discussed, this notion of rank does not aim to capture the combinatorial or analytic properties of the tensor. Rather, it looks at the geometric properties of the tensor, defining rank as the codimension of an algebraic variety. Kopparty, Moshkovitz, and Zuiddam used this new notion of rank to prove tight bounds about the subrank of the matrix multiplication tensor (the subrank is a quantity related to the rank that is useful in the computational complexity analysis for matrix multiplication).

The geometric rank of a tensor is formally defined as follows.

Definition 2.4 (Geometric Rank). Let $T \in \mathbb{F}^{n_1 \times \cdots \times n_d}$ be a d -tensor with $d \geq 2$. The *geometric rank* of T is

$$\begin{aligned} \text{GR}(T) = \text{codim}\{ & (x_1, \dots, x_{d-1}) \in \mathbb{F}^{n_1} \times \cdots \times \mathbb{F}^{n_{d-1}} \\ & : \forall z \in \mathbb{F}^{n_d} : T(x_1, \dots, x_{d-1}, z) = 0\}. \end{aligned}$$

Here we use the usual definition of the codimension of an algebraic variety. If $V \subseteq \mathbb{F}^n$ is an algebraic variety (that is possibly reducible), the dimension of V , $\dim(V)$, is defined to be the length of a maximal chain of irreducible subvarieties of V . The codimension, $\text{codim}(V)$, is then defined to be $n - \dim(V)$. For more detailed explanations of these concepts we recommend looking at [Har13].

Geometric rank—as Kopparty, Moshkovitz, and Zuiddam mention without proof—coincides with the linear algebra definition of matrix rank when $d = 2$. We give a quick proof of that here.

Proposition 2.5. *Let $M \in \mathbb{F}^{m \times n}$ be a matrix (a 2-tensor), and let r be the rank of M as per the linear algebra definition. Then $r = \text{GR}(M)$.*

Proof. By definition of geometric rank,

$$\begin{aligned} \text{GR}(M) &= \text{codim}\{x \in \mathbb{F}^m : \forall y \in \mathbb{F}^n x^T M y = 0\} \\ &= \text{codim}\{x \in \mathbb{F}^m : x^T M = \mathbf{0}^T\} \\ &= m - \dim\{x \in \mathbb{F}^m : x^T M = \mathbf{0}^T\}, \end{aligned}$$

where $\mathbf{0}$ is the zero vector in \mathbb{F}^m . Notice that $\dim\{x \in \mathbb{F}^m : x^T M \equiv \mathbf{0}^T\}$ is just the nullity of M , and so by the rank nullity theorem we are done. \square

There is a slight caveat to Definition 2.4, which is that it implicitly assumes \mathbb{F} to be algebraically closed. This, however, is not a problem because in a general field we can extend this definition via the embedding of the field into its algebraic closure.

3. APPLICATIONS IN COMPUTER SCIENCE

Uttering the words “tensor rank” to a computer scientist will, with high probability, elicit a response that contains the phrase “matrix multiplication.” And for good reason; notions of tensor rank happen to be intimately related to the complexity and efficient computation of matrix multiplication.

Multiplying two matrices is a fundamental computation with applications in nearly every field computers are used in. It is thus not surprising that the following question is of great interest.

Question 3.1 (Arithmetic Complexity of Matrix Multiplication). *Given two $n \times n$ matrices A and B over a field \mathbb{F} , what is the minimum number of addition and scalar multiplication operations needed to compute the $n \times n$ product matrix AB where*

$$(1) \quad (AB)_{ij} = \sum_{\ell=1}^n A_{i\ell} B_{\ell j}?$$

Can one find an algorithm that actually computes the product with that number of operations?

The number of operations in Question 3.1 is known as the *arithmetic complexity* of $n \times n$ matrix multiplication. Let us call this number $\mathcal{A}(n)$. An exact determination of $\mathcal{A}(n)$ seems to be outside the range of methods available at the present time, so much of the work around this has been focused on getting asymptotic bounds on arithmetic complexity. It is then useful to define the following quantity.

Definition 3.2 (Exponent of Matrix Multiplication). The *exponent of matrix multiplication*, ω , is defined as

$$\omega = \inf\{\tau \in \mathbb{R} : \mathcal{A}(n) = O(n^\tau)\}.$$

Equivalently,

$$\omega = \inf\{\tau : \exists \text{ an algorithm to compute } n \times n \text{ MM with } O(n^\tau) \text{ operations}\}.$$

It is easy to see that $\omega \in [2, 3]$. The upper bound comes from the fact that the naïve algorithm of computing each entry of the product using Equation 1 is $O(n^3)$. The lower bound holds because any algorithm for matrix multiplication would have to make at least n^2 computations since the output is an $n \times n$ matrix. A long-standing conjecture in algebraic complexity theory is that ω is in fact equal to 2.

As mentioned in Section 2, the operation of matrix multiplication is itself a 3-tensor. It is convention to denote the tensor to multiply a $k \times m$ matrix and an $m \times n$ matrix with $\langle k, m, n \rangle$. Since we are only focusing on square matrices, we will focus on the tensors $\langle n, n, n \rangle$, which we will often shorten to $\langle n \rangle$. Seeing as rank is supposed to encode the complexity of a tensor, it is not surprising that it is closely connected to $\mathcal{A}(n)$. The following result is known to hold.

Theorem 3.3. *Over any field, we have*

$$\omega = \inf\{\tau \in \mathbb{R} : \text{TR}(\langle n \rangle) = O(n^\tau)\}.$$

Proof. See [BCS97, Proposition 15.1] □

An important application of Theorem 3.3 is in showing that the exponent of matrix multiplication does not change under another important, related measure of complexity called the *multiplicative complexity* of matrix multiplication. The multiplicative complexity of matrix multiplication is defined as the minimum number of scalar multiplication operations required for $n \times n$ matrix multiplication (the algorithm is of course also allowed additions). Let us call this quantity $\mathcal{M}(n)$. It has been shown that \mathcal{M} relates to traditional tensor rank in the following way (see, for instance, [BCS97, Section 14.1]).

$$\frac{1}{2} \text{TR}(\langle n \rangle) \leq \mathcal{M}(n) \leq \text{TR}(\langle n \rangle).$$

From this inequality it follows that $\text{TR}(\langle n \rangle) = O(n^\tau)$ implies $\mathcal{M}(n) = O(n^\tau)$ and $\mathcal{M}(n) = O(n^\tau)$ implies $\text{TR}(\langle n \rangle) = O(n^\tau)$. Combining this with Theorem 3.3 gives us $\omega = \inf\{\tau \in \mathbb{R} : \mathcal{M}(n) = O(n^\tau)\}$.

The connection in Theorem 3.3 also implies that upper and lower bounds on TR directly translate to upper and lower bounds on the exponent of matrix multiplication. The following result makes this explicit.

Theorem 3.4 ([Blä13]). *If $\text{TR}(\langle k, m, n \rangle) \leq r$, then $\omega \leq \log_{kmn} r$.*

In his seminal work, Strassen [Str69] showed that square matrices of size 2 can be multiplied with seven multiplication operations (as opposed to eight using the naïve method) at the cost of a few more addition operations. It follows from his work that $\text{TR}(\langle 2 \rangle) \leq 7$. Plugging this into Theorem 3.4 tells us that $\omega \leq 3 \log_8 7 = \log_2 7 = 2.801\dots \leq 2.81$. Hopcroft and Kerr [HK71] and Winograd [Win71] later (independently) showed that 2×2 multiplication is not possible with just six multiplications, which implies that $\text{TR}(\langle 2 \rangle)$ is in fact equal to 7. Thus $\log_2 7$ is the sharpest bound one can get from Theorem 3.4 using $\langle 2 \rangle$. It turns out that $\text{TR}(\langle 70 \rangle) \leq 143.640$ [Pan80], which gives a slightly better bound of 2.80. Whether other tensors can lead to sharper bounds remains an open question.

Strassen [Str86], using more traditional-rank-related techniques, was able to prove that $\omega < 2.48$. Coppersmith and Winograd [CW87], using a construction of arithmetic-progression-free sets, showed that $\omega \leq 2.375\dots$. Starting in 2010, by analyzing higher-order variants of the Coppersmith-Winograd construction, Stothers [Sto10], then Vassilevska Williams [Wil12], and then Le Gall [LG14] made incremental improvements. This led to the current best upper bound of $\omega < 2.372\dots$

The connections between matrix multiplication and tensor rank do not end here. As mentioned in the previous paragraph, Coppersmith and Winograd’s proof involved a construction of arithmetic-progression-free sets. This suggests that the complexity of matrix multiplication might be related to the capset problem and Tao’s slice rank. And indeed, that did turn out to be the case. In 2003, Cohn and Umans [CU03] described a framework for proving upper bounds on ω that involves reducing matrix multiplication to group algebra multiplication. In 2012, Alon, Shpilka, and Umans [ASU12] proved relations between various then-open conjectures in combinatorics and bounds on ω . The resolution of the capset problem in 2016 settled some of the conjectures involved in the 2012 work. That led to the 2017 work of Blasiak et al. [BC⁺17], which used the resolution of the capset problem to rule out obtaining an $\omega = 2$ using a subclass of Cohn-Umans-style constructions. In doing so, they extended the capset result, making extensive use of the notion of slice ranks.

Finally, the introduction of the geometric rank was in part motivated by the complexity of matrix multiplication. In particular, Kopparty, Moshkovitz, and Zuiddam [KMZ20] used the notion of geometric rank to prove a tight upper bound on the so-called *border subrank* of the matrix multiplication tensor, which matched a known lower bound. While the exact definitions of subrank are beyond the scope of this survey, we recommend reading [KMZ20] for further details.

4. SOME TRIVIAL RELATIONSHIPS

Before getting to the deeper connections between the different notions of rank, we state some consequences obvious from the definitions. Any rank 1 form used in Definition 2.2,

$$T_1(x_1)T_2(x_j : j \neq i),$$

is in particular a rank 1 form of the form in Definition 2.3,

$$(2) \quad T_1(x_i : i \in S)T_2(x_j : j \notin S).$$

A tensor of the form in (2) is in turn of the form of the rank 1 form in Definition 1.2,

$$T_1(x_1) \cdots T_d(x_d).$$

This implies the following result.

Proposition 4.1. *For any tensor T ,*

$$\text{PR}(T) \leq \text{SR}(T) \leq \text{TR}(T).$$

5. RELATING ANALYTIC RANK TO THE OTHERS

We now discuss some deeper relationships which will show a cycle of relationships between AR, SR, and TR. We start with an elegant result due to Lovett relating the analytic and partition rank. For ease of notation we will look at tensors $T : V^d \rightarrow \mathbb{F}$ where $V = \mathbb{F}^n$ for some n . It is easy to see that this proof can be generalized to arbitrary $T \in \mathbb{F}^{n_1 \times \cdots \times n_d}$.

Theorem 5.1 ([Lov19]). *Let $T : V^d \rightarrow \mathbb{F}$ be a d -tensor. Then $\text{AR}(T) \leq \text{PR}(T)$.*

To prove this theorem we will need a lemma as well as another theorem about the arithmetic rank. Both of these are due to Lovett.

Lemma 5.2 ([Lov19]). *For each $I \subseteq [d]$ let $R_I : V^{|I|} \rightarrow \mathbb{F}$ be an $|I|$ -tensor. Let*

$$R(x) = \sum_{I \subseteq [d]} R_I(x_j : j \in I).$$

Then $|\text{bias}(R)| \leq \text{bias}(R_{[d]})$.

Proof. First, let W_0, \dots, W_n be arbitrary functions from \mathbb{F}^m to \mathbb{F} . Let A and B be functions from $\mathbb{F}^n \times \mathbb{F}^m$ to \mathbb{F} defined by

$$A(x, y) = \sum_{i=1}^n x^i W_i(y), \text{ and}$$

$$B(x, y) = A(x, y) + W_0(y).$$

From the definition of bias we get

$$\begin{aligned}
\text{bias}(B) &= \mathbb{E}_{x,y} \chi(B(x,y)) \\
&= \mathbb{E}_{x,y} \chi(A(x,y) + W_0(y)) \\
&= \mathbb{E}_{x,y} [\chi(A(x,y)) \cdot \chi(W_0(y))] \\
&= \mathbb{E}_y [\chi(W_0(y)) \cdot \mathbb{E}_x \chi(\sum_{i=1}^n x^i W_i(y))] \\
&= \mathbb{E}_y [1_{W_1(y)=\dots=W_n(y)=0} \chi(W_0(y))],
\end{aligned}$$

where the last equality holds because the expectation over x is 0 whenever any of the W_j with $1 \leq j \leq n$ are non-zero. Taking absolute values and using the triangle inequality gives us

$$(3) \quad |\text{bias}(B)| \leq \mathbb{E}_y [1_{W_1(y)=\dots=W_n(y)=0}] = \text{bias}(A).$$

With this smaller result, we prove the lemma by applying it repeatedly. Fixing $i \in [d]$, we can break up the summation in the definition of R as follows.

$$R(x) = \sum_{I \subseteq [d]; i \in I} R_I(x_j : j \in I) + \sum_{I \subseteq [d]; i \notin I} R_I(x_j : j \in I).$$

Notice that each $R_I(x_j : j \in I)$ where $i \in I$ is a tensor depending on x_i , and thus can be written as $x_i^j W_j(x_k : k \neq i)$. Thus the first sum is of the form $\sum_j x_i^j W_j(y)$, which matches the form of A above. The second sum does not depend on x_i , and so is of the form of B . Thus, applying Equation 3 gives us

$$|\text{bias}(R)| \leq \text{bias} \left(\sum_{I \subseteq [d]; i \in I} R_I(x_j : j \in I) \right).$$

We use this inequality iteratively. First, applying it for $i = d$ we have

$$|\text{bias}(R)| \leq \text{bias} \left(\sum_{I \subseteq [d]; d \in I} R_I(x_j : j \in I) \right).$$

Using $\sum_{I \subseteq [d]; d \in I} R_I(x_j : x_j \in I)$ instead of the tensor R in our inequality with $i = d - 1$ gives us

$$\left| \text{bias} \left(\sum_{I \subseteq [d]; d \in I} R_I(x_j : j \in I) \right) \right| \leq \text{bias} \left(\sum_{I \subseteq [d-1]; d-1 \in I; d \in I} R_I(x_k : k \in I) \right).$$

Continuing this process and applying the inequality for $d - 2, \dots, 1$, and then chaining the inequalities gives us

$$|\text{bias}(R)| \leq \text{bias}(R_{[d]})$$

as desired. \square

Using this lemma, Lovett proves that the analytic rank is sub-additive.

Theorem 5.3 ([Lov19]). *Let $T, S : V^d \rightarrow \mathbb{F}$ be d -tensors. Then*

$$\text{AR}(T + S) \leq \text{AR}(T) + \text{AR}(S).$$

Proof Sketch. For this proof, Lovett defines functions T_I and S_I such that for any x, y , $T(x, y) = \sum_{I \subseteq [d]} T_I(x_I, y_{[d]-I})$ (similarly for S) where $x_I = (x_i : i \in I)$ (similarly for y). Using this decomposition and performing some algebra leads to

$$\text{bias}(T)\text{bias}(S) \leq \left| \text{bias} \left((T + S)(x) + \sum_{I \subseteq [d]} S_I(x_I, b_{[d]-I}) \right) \right|,$$

where b is a fixed choice for y .

Applying Lemma 5.2 to the functions $R_{[d]} = (T + S)(x)$ and $R_I = S_I(x_I, b_{[d]-I})$ shows that the right hand side of the inequality is less than or equal to $\text{bias}(R_{[d]})$, which in our case equals $\text{bias}(T + S)$. Putting everything together we get

$$\text{bias}(T + S) \geq \text{bias}(T)\text{bias}(S),$$

and so by the definition of analytic rank,

$$\text{AR}(T + S) \leq \text{AR}(T)\text{AR}(S).$$

□

Using these two results, we can now prove Theorem 5.1.

Proof of Theorem 5.1. Given Theorem 5.3, it suffices to prove Theorem 5.1 for tensors of partition rank 1. So suppose $T : V^d \rightarrow \mathbb{F}$ has partition rank 1. Then we can find a partition $A \sqcup B = [d]$ with $|A|, |B| \geq 1$ such that

$$T(x) = T_1(x_i : i \in A)T_2(x_j : j \in B).$$

For convenience, let us denote $(x_i : i \in A)$ by x_A and $(x_j : j \in B)$ by x_B .

Since the partition rank of T is 1, we need to show that $\text{AR}(T) \leq 1$. To do so, it suffices to show (by definition of analytic rank) that $\text{bias}(T) \geq |\mathbb{F}|^{-1}$. For any $a, b \in \mathbb{F}$ define the function

$$F_{a,b}(x) = (T_1(x_A) + a)(T_2(x_B) + b).$$

Expanding it out we can write

$$\begin{aligned} F_{a,b}(x) &= T_1(x_A)T_2(x_B) + T_1(x_A)b + T_2(x_B)a + ab \\ &= T(x) + T_1(x_A)b + T_2(x_B)a + ab. \end{aligned}$$

Letting $R_A(x) = T_1(x_A)b$, $R_B(x) = T_2(x_B)a$, $R_\emptyset = ab$, and $R_{[d]} = T$, we can apply Lemma 5.2 to get

$$|\text{bias}(F_{a,b})| \leq \text{bias}(T).$$

On the other hand, if a and b are chosen uniformly, we can take the expectation of the bias over a and b to get

$$\begin{aligned} \mathbb{E}_{a,b} \text{ bias}(F_{a,b}) &= \mathbb{E}_{a,b \in \mathbb{F}; x \in V^d} [\chi((T_1(x_A) + a)(T_2(x_B) + b))] \\ &= \mathbb{E}_{a,b \in \mathbb{F}} [\chi(ab)] \\ &= \Pr_{b \in \mathbb{F}} [b = 0] \\ &= |\mathbb{F}|^{-1}. \end{aligned}$$

This proves our theorem. \square

Combining Theorem 5.1 with Proposition 4.1 gives us the following corollary.

Corollary 5.4. *Let $T : V^d \rightarrow \mathbb{F}$ be a d -tensor. Then $\text{AR}(T) \leq \text{SR}(T)$.*

6. CLOSING THE LOOP FOR 3-TENSORS: COHEN AND MOSHKOVITZ'S ARGUMENT

Cohen and Moshkovitz [CM21] proved the following two theorems which, for 3-tensors over finite fields with at least 3 elements, showed asymptotic equivalence between AR, GR, and SR when combined with Lovett's result.

Theorem 6.1 ([CM21]). *For any 3-tensor T over a perfect field \mathbb{F} ,*

$$\text{SR}(T) \leq 3\text{GR}(T)$$

Theorem 6.2 ([CM21]). *For any 3-tensor T over any finite field \mathbb{F} ,*

$$\text{AR}(T) \geq (1 - \log_{|\mathbb{F}|} 2)\text{GR}(T).$$

Our goal will be to give an intuitive sketch of their argument. Techniques from algebraic geometry will be essential, so we give some background.

6.3. Background from Algebraic Geometry. From now on, let \mathbb{K} be a field. Recall that a *variety cut out by a (finite) number of polynomials* over \mathbb{K}^n is the subspace of \mathbb{K}^n where all the polynomials vanish. The ideal $I(V)$ of a variety V is defined as follows.

Definition 6.4 (Ideal of a Variety). Let $V \subset \mathbb{K}^n$ be a variety. The *ideal* of V is

$$I(V) = \{f \in \mathbb{K}[x] : f(p) = 0 \text{ for each } p \in V\}.$$

Now we define the tangent space $T_p V$ to a variety V at p .

Definition 6.5 (Tangent Space). Let $V \subset \mathbb{K}^n$ be a variety. The *tangent space* $T_p V$ to a variety V at p is

$$T_p V = \left\{ \mathbf{v} \in \mathbb{K}^n : \frac{\partial g}{\partial \mathbf{v}}(p) = 0, \text{ for each } g \in I(V) \right\}.$$

One can check that any variety can be written uniquely as the union of *irreducible varieties*. An irreducible variety is one which cannot be written as the union of strictly contained subvarieties. Then the *dimension* of a variety V can be defined as follows.

Definition 6.6 (Dimension of a Variety). The *dimension* of a variety $V \subset \mathbb{K}^n$, written $\dim V$, is the maximal length of a chain of irreducible varieties such that

$$\emptyset \neq V_1 \subsetneq V_2 \subsetneq \cdots \subsetneq V_k \subsetneq V.$$

The *codimension* of a variety $V \subset \mathbb{K}^n$ is

$$\text{codim}V = n - \dim V.$$

6.7. Proof Sketch of Theorem 6.1. For the full argument see [CM21]. We will give a broader overview of the proof with the aim of conveying the intuition behind it.

The proof of Theorem 6.1 hinges on the following previously proved rephrasing of geometric rank.

Lemma 6.8. *For a 3-tensor over any field \mathbb{K} ,*

$$\text{GR}(T) = \min_r \{r + \text{codim}\{x : \text{rank } T(x, \cdot, \cdot) = r\}\}.$$

Proof. See [KMZ20, Theorem 3.1]. □

Let r achieve the minimum in Lemma 6.8. Defining

$$X_r = \{x \in \overline{\mathbb{F}}^{n_1} : \text{rank}T(x, \cdot, \cdot) \leq r\},$$

one can use Lemma 6.8 to show that $\text{GR}(T) = r + \text{codim}X_r$. A 3-tensor T with encoding array $(T_{i,j,k})$ can be reinterpreted as a linear space $L \subset \overline{\mathbb{F}}^{n_1 \times n_2}$ spanned by $\{M_1, \dots, M_{n_3}\}$, where $M_k = (T_{i,j,k})_{i,j}$. One also has an association going the other way. Given a matrix space L , choose a basis M_1, \dots, M_k , and let T be the tensor with encoding array (M_1, \dots, M_k) . Thus the notions of rank for tensors can be transferred to matrix spaces.

Now let

$$L_r = L \cap M_r,$$

where M_r is the space of $n_1 \times n_2$ matrices of rank at most r . Note that M_r , being the space cut out by the $(r+1) \times (r+1)$ minors, is itself a variety. It turns out that the variety X_r is isomorphic to $L_r \times \overline{\mathbb{F}}^{n_1-d}$. Thus $\text{codim}X_r = \text{codim}_L L_r$.² We have reduced the problem to analyzing these “slices” of the matrix space L . Above, $\overline{\mathbb{F}}$ is the algebraic closure of \mathbb{F} , which will be easier to work over in the arguments to follow. We let $\overline{\text{SR}}(T)$ be as in Definition 2.2, except the coefficients are allowed to come from $\overline{\mathbb{F}}$. Clearly $\overline{\text{SR}}(T) \leq \text{SR}(T)$. One also has $2\text{SR}(T) \leq 3\overline{\text{SR}}(T)$ which is proved in [CM21].

We would be done if we had the following.

Lemma 6.9. *Let $L \subset \mathbb{K}^{m \times n}$ be a matrix space over any algebraically closed field \mathbb{K} . For any $r \in \mathbb{N}$,*

$$\text{SR}(L) \leq 2r + \text{codim}_L L_r.$$

²Here $\text{codim}_L X$ is the codimension of a variety X in a linear space L , i.e. $\text{codim}_L X = \dim L - \dim X$.

To see this, observe that

$$\begin{aligned}\overline{\text{SR}}(T) &= \text{SR}(L) \leq 2r + \text{codim}_L L_r \\ &= 2r + \text{codim} X_r \\ &= 2\text{GR}(T) - \text{codim} X_r \\ &\leq 2\text{GR}(T).\end{aligned}$$

Applying $2\text{SR}(T) \leq 3\overline{\text{SR}}(T)$ gives $\text{SR}(T) \leq 3\text{GR}(T)$.

Proof sketch of Lemma 6.9. We proceed by induction. The base case is $\text{SR}(L) \leq \dim L$, which is not difficult to show. Now consider the inductive step. Let V be an irreducible component of L_r with $\dim V = \dim L_r$, and $A \in V \setminus M_{r-1}$. If $V \setminus M_{r-1} = \emptyset$ then we have $V \subseteq L_{r-1}$, and the result follows from induction. So we assume $V \setminus M_{r-1} \neq \emptyset$. The trick now is to decompose the matrix space L into subspaces $P = L \cap T_A M_r$ and P^\perp . This particular decomposition is useful because one can run an explicit calculation to find that

$$\text{SR}(T_A M_r) \leq 2r.$$

We need the following simple properties of slice rank and tangent spaces to conclude the result from that calculation.

Lemma 6.10.

- (1) $\text{SR}(L) \leq \dim L$.
- (2) $\text{SR}(L') \leq \text{SR}(L)$ if $L' \subset L$.
- (3) $\text{SR}(L + L') \leq \text{SR}(L) + \text{SR}(L')$.
- (4) $T_p(V \cap W) \subset T_p V \cap T_p W$.
- (5) If V is irreducible and $p \in V$ then $\dim V \leq \dim T_p V$.

Using Lemma 6.10 and the main estimate $\text{SR}(T_A M_r) \leq 2r$,

$$\text{SR}(P) = \text{SR}(L \cap T_A M_r) \leq \text{SR}(T_A M_r) \leq 2r.$$

We also have

$$\begin{aligned}\dim L_r &\leq \dim T_A V \\ &\leq \dim T_A L_r \\ &\leq \dim(T_A L \cap T_A M_r) \\ &\leq \dim(L \cap T_A M_r) = \dim P.\end{aligned}$$

Thus we have $\text{codim}_L P \leq \text{codim}_L L_r$. Using Lemma 6.10 we get

$$\text{SR}(P^\perp) \leq \text{codim}_L P \leq \text{codim}_L L_r.$$

Hence

$$\begin{aligned}\text{SR}(L) &= \text{SR}(P + P^\perp) \\ &\leq \text{SR}(P) + \text{SR}(P^\perp) \\ &\leq 2r + \text{codim}_L L_r.\end{aligned}$$

□

6.11. **Proof Sketch of Theorem 6.2.** The key ingredient of the proof is the following generalization of the Schwarz-Zippel lemma. Define $V(\mathbb{F}) = V \cap \mathbb{F}^n$ for a variety $V \subset \overline{\mathbb{F}}^n$ defined over \mathbb{F} .

Lemma 6.12. *For any variety $V \subset \overline{\mathbb{F}}^n$ defined over any finite field \mathbb{F} cut out by polynomials of degree at most d ,*

$$\frac{|V(\mathbb{F})|}{|\mathbb{F}|^n} \leq \left(\frac{d}{|\mathbb{F}|} \right)^{\text{codim}V}.$$

We use Lemma 6.12 to prove Theorem 6.2. Consider $T \in \mathbb{F}^{n_1 \times n_2 \times n_3}$, where we interpret T as its defining array. Set $V = \ker(T) \subset \overline{\mathbb{F}}^N$ with $N = n_1 + n_2$, where here we use the bilinear map formulation of T . By Lemma 6.12,

$$\frac{|V(\mathbb{F})|}{|\mathbb{F}|^n} \leq \left(\frac{d}{|\mathbb{F}|} \right)^{\text{codim}V},$$

since T is bilinear. Thus

$$\begin{aligned} \text{AR}(T) &= -\log_{|\mathbb{F}|} \frac{|V(\mathbb{F})|}{|\mathbb{F}|^N} \\ \text{(Using Lemma 6.12)} \quad &\geq \text{codim}V \cdot (1 - \log_{|\mathbb{F}|} 2) \\ &= \text{GR}(T)(1 - \log_{|\mathbb{F}|} 2). \end{aligned}$$

The first equality above is true since

$$\begin{aligned} \text{bias}(T) &= \mathbb{E}_{(x_1, x_2) \in \mathbb{F}^{n_1} \times \mathbb{F}^{n_2}} [\mathbb{E}_{x_3 \in \mathbb{F}^{n_3}} \chi(T(x_1, x_2, x_3))] \\ (4) \quad &= \Pr_{(x_1, x_2) \in \mathbb{F}^{n_1} \times \mathbb{F}^{n_2}} [T(x_1, x_2, \cdot) \equiv 0] \\ &= \frac{|V(\mathbb{F})|}{|\mathbb{F}|^N}. \end{aligned}$$

Equation 4 holds because the bias of a linear form is 0 unless the linear form is identically 0, in which case its bias is 1. \square

ACKNOWLEDGEMENTS

I would like to thank Alex Iosevich for being a brilliant mentor and friend throughout my undergraduate career. I would like to thank Kaave Hosseini for advising this thesis. This thesis is based on work done in a joint independent study with Arian Nadjizadah, and I would like to thank him for being a fantastic collaborator and a dear friend. I would also like to thank the UR mathematics faculty and staff: it is because of them that the department is not just a place I went to school, but is a place I am going to miss. Finally, I would like to thank my friends and my family, who not only kept me sane over the last four years but also made my time in college a blast.

REFERENCES

- [ASU12] Noga Alon, Amir Shpilka, and Christopher Umans. On sunflowers and matrix multiplication. In *2012 IEEE 27th Conference on Computational Complexity*, pages 214–223. IEEE, 2012.
- [BC⁺17] Jonah Blasiak, , Thomas Church, Henry Cohn, Joshua A. Grochow, Eric Naslund, William F. Sawin, and Chris Umans and. On cap sets and the group-theoretic approach to matrix multiplication. *Discrete Analysis*, January 2017.
- [BCS97] Peter Bürgisser, Michael Clausen, and M. Amin Shokrallahi. *Algebraic complexity theory*. Springer-Verlag Berlin Heidelberg, 1997.
- [Blä13] Markus Bläser. Fast matrix multiplication. *Theory of Computing*, pages 1–60, 2013.
- [CLP17] Ernie Croot, Vsevolod Lev, and Péter Pach. Progression-free sets in \mathbb{Z}_4^n are exponentially small. *Annals of Mathematics*, 185(1):331 – 337, 2017.
- [CM21] Alex Cohen and Guy Moshkovitz. Structure vs. randomness for bilinear maps. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 800–808, 2021.
- [CU03] Henry Cohn and Christopher Umans. A group-theoretic approach to fast matrix multiplication. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings*. IEEE Computer. Soc, 2003.
- [CW87] Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing, STOC '87*, page 1–6, New York, NY, USA, 1987. Association for Computing Machinery.
- [EG17] Jordan Ellenberg and Dion Gijswijt. On large subsets of \mathbb{F}_q^n with no three-term arithmetic progression. *Annals of Mathematics*, 185(1):339 – 343, 2017.
- [GW11] William Gowers and Julia Wolf. Linear forms and higher-degree uniformity for functions on \mathbb{F}_p^n . *Geometric and Functional Analysis*, 21(1):36–69, 2011.
- [Hal58] Paul Halmos. *Finite-dimensional vector spaces*. Van Nostrand, Princeton, NJ, 2nd ed. edition, 1958.
- [Har13] Joe Harris. *Algebraic geometry: a first course*, volume 133. Springer Science & Business Media, 2013.
- [Hås89] Johan Håstad. Tensor rank is np-complete. In *International Colloquium on Automata, Languages, and Programming*, pages 451–460. Springer, 1989.
- [HK71] John Hopcroft and Leslie Kerr. On minimizing the number of multiplications necessary for matrix multiplication. *SIAM Journal on Applied Mathematics*, 20(1):30–36, 1971.

- [HL13] Christopher J. Hillar and Lek-Heng Lim. Most tensor problems are np-hard. *J. ACM*, 60(6), nov 2013.
- [KMZ20] Swastik Kopparty, Guy Moshkovitz, and Jeroen Zuiddam. Geometric rank of tensors and subrank of matrix multiplication. In *Proceedings of the 35th Computational Complexity Conference (CCC 2020)*, Leibniz International Proceedings in Informatics (LIPIcs), pages 35:1–35:21, 2020.
- [LG14] François Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*, ISSAC '14, page 296–303, New York, NY, USA, 2014. Association for Computing Machinery.
- [Lov19] Shachar Lovett. The analytic rank of tensors and its applications. *Discrete Analysis*, 7:1–10, 2019.
- [Nas20] Eric Naslund. The partition rank of a tensor and k-right corners in fqn. *Journal of Combinatorial Theory, Series A*, 174(105190), 2020.
- [Pan80] V. Ya Pan. New fast algorithms for matrix operations. *SIAM Journal on Computing*, 9(2):321–342, 1980.
- [Sto10] Andrew Stothers. *On the complexity of matrix multiplication*. PhD thesis, University of Edinburgh, 2010.
- [Str69] Volker Strassen. Gaussian elimination is not optimal. *Numerische mathematik*, 13(4):354–356, 1969.
- [Str86] Volker Strassen. The asymptotic spectrum of tensors and the exponent of matrix multiplication. In *27th Annual Symposium on Foundations of Computer Science (sfcs 1986)*, pages 49–54, 1986.
- [Swe18] Joseph Swernofsky. Tensor Rank is Hard to Approximate. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018)*, volume 116 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 26:1–26:9, Dagstuhl, Germany, 2018.
- [Tao16] Terence Tao. A symmetric formulation of the Croot-Lev-Pach-Ellenberg-Gijswijt capset bound. <https://terrytao.wordpress.com/2016/05/18/a-symmetric-formulation-of-the-croot-lev-pach-ellenberg-gijswijt-capset-bound/>, May 2016.
- [Wil12] Virginia Vassilevska Williams. Multiplying matrices faster than coppersmith-winograd. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, STOC '12, page 887–898, New York, NY, USA, 2012. Association for Computing Machinery.
- [Win71] Shmuel Winograd. On multiplication of 2×2 matrices. *Linear Algebra and its Applications*, 4(4):381–388, 1971.