

Tensor Rank is Hard to Approximate

Joseph Swernofsky - KTH Royal Institute of Technology*

April 20, 2018

Abstract

We prove that approximating the rank of a 3-tensor to within a factor of $1 + 1/1852 - \delta$, for any $\delta > 0$, is NP-hard over any finite field. We do this via reduction from bounded occurrence 2-SAT.

1 Introduction

The rank of a matrix is well understood and easy to compute. The student first introduced to the notion of rank often learns that they can perform Gaussian elimination on a matrix and the number of nonzero rows they obtain is its rank. They may even learn that this is equivalent to the number of linearly independent rows of the matrix. A rank 1 matrix can be written as an outer product of two vectors u and v , meaning its ij^{th} entry is $u_i \cdot v_j$. The rank of M is also equivalent to the number of outer products (rank 1 matrices) that must be added to form M .

This notion of rank has been generalized to describe tensors. A 3-tensor T is a grid of numbers with 3 indices. A tensor is rank 1 if it is the outer product of 3 vectors, and the rank of T is the minimum number of rank 1 tensors that must be added to form it.

Rank of 3-tensors is, apart from its inherent mathematical appeal, a natural tool for reasoning about bilinear circuit complexity. The inputs and outputs of a bilinear circuit correspond to indices in the 3 coordinates of a 3-tensor, and the rank is equal to the number of multiplication nodes needed in the circuit [7]. Strassen's [15] famous algorithm for 2×2 matrix multiplication can be viewed in this light. The multiplication is a circuit computing the entries of the output matrix from the entries of the input matrices. Writing this as a $2^2 \times 2^2 \times 2^2$ tensor, his basic construction corresponds to a rank 7 decomposition. He then applied this recursively to create the first algorithm for $n \times n$ matrix multiplication that was faster than the naive algorithm.

Unfortunately, while matrix rank is easy to compute, Håstad [5] showed that tensor rank is NP-hard to calculate. More recently, Shitov [13] and Schaefer and

*Stockholm S100 44, Sweden (JosephSw@kth.se) Supported by the Knut and Alice Wallenberg Foundation

Štefankovič [12] showed that rank over a field \mathbb{F} is complete for the existential theory of \mathbb{F} and is also uncomputable over \mathbb{Z} . This presents a roadblock for researchers hoping to analyze the rank of tensors for specific applications. Much work has been done since Strassen’s initial paper to analyze larger tensors modeling matrix multiplication and many improvements have been obtained. However, even very simple tensors stubbornly resist analysis, such as the $3^2 \times 3^2 \times 3^2$ tensor for 3×3 matrix multiplication. Its rank is only known to be between 19 [3] and 23 [8].

Another application for tensor rank is demonstrated by Raz [11], who shows how strong circuit lower bounds for more general circuits may be possible with a better understanding of rank. Tensor rank is also a practical tool in a variety of fields, such as signal processing and computer vision [6]. Our knowledge is very limited though, because despite knowing that there are 3-tensors with rank at least $n^2/3$ (which can be seen by dimension counting), the highest rank known for an explicit family of 3-tensors is $3n - o(n)$ [14]. Also, while we know rank is hard to compute exactly, we do not know if it can be approximated. Alexeev et al. [1] and Bläser [4] both mention this open question. In this paper we prove that:

Theorem 1. *It is NP-hard to approximate 3-tensor rank over a finite field \mathbb{F} within a factor of $1 + 1/1852 - \delta$, for any $\delta > 0$.*

Our construction is a re-analysis and simplification of the reduction by Håstad [5]. We show that if bounded occurrence SAT is used as the starting point for the reduction then significant extra rank can be guaranteed from unsatisfied clauses over the rank in the satisfiable case.

A natural follow-up question to ask is whether tensor rank is hard to approximate within any constant, or within a specific unbounded function. It would also be interesting to have any nontrivial approximation algorithm. We discuss these questions briefly in Section 4.

2 Preliminaries

We work over an arbitrary finite field \mathbb{F} throughout this paper and we are given an explicit representation of \mathbb{F} .

A **d -tensor** is a function from d -tuples of natural numbers to \mathbb{F} . The entries in the tuple represent coordinates in a d -dimensional space. A given d -tensor has size n_i in coordinate i and is defined at exactly those s where $s_i \in [n_i]$ for all i . It is helpful to think of this as a grid with numbers written in the cells. We typically write T_s to denote $T(s)$, or even $T_{s_1 s_2 \dots s_d}$. For $v_i \in \mathbb{F}^{n_i}$, an outer product $\otimes_{i \in [d]} v_i$ of d vectors is a rank 1 tensor. The **rank** $\mathbf{rk}(T)$ is the minimum r so that $T = \sum_{j \in [r]} \otimes_{i \in [d]} v_i^j$.

An minimization problem P consists of a set of valid strings L together with a score function $s : L \rightarrow \mathbb{R}^+$. An **approximation algorithm** for P is an algorithm A that takes strings in L and outputs numbers in \mathbb{R}^+ with

$\forall x \in L. A(x) \geq s(x)$. Let L_n be the strings in L of length n . For a function $c: \mathbb{N} \rightarrow \mathbb{R}^+$, A is a c -approximation algorithm if $\max_{x \in L_n} \frac{A(x)}{s(x)} \leq c(n)$.

2.1 Slices

A **slice** of a tensor is the grid obtained from fixing one of the coordinates. The slice $T_{i:x}$ of a d -tensor T is a $(d-1)$ -tensor with $T_{i:x}(s) := T(s_1, s_2, \dots, s_{i-1}, x, s_i, \dots, s_{d-1})$. That is, when indexing into $T_{i:x}$ one is indexing into T but omitting the coordinate i . For example, if $d = 3$ then $T_{3:k}$ is the matrix M where $M_{ij} = T_{ijk}$. Even more concretely, if we draw the first coordinate as the row and the second coordinate as the column:

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}, A_{2:1} = \begin{bmatrix} a \\ c \end{bmatrix}$$

We want to recover some of our intuition about matrix rank and use row reduction when computing tensor rank.

Lemma 1. *If a slice is scaled by a nonzero constant or added to another slice in the same coordinate, the rank does not change.*

Proof. Suppose we have a rank r decomposition $T = \sum_{h \in [r]} \otimes_{j \in [d]} v_j^h$. If slice $T_{i:x}$ is scaled by λ , simply replace $v_i^h(x)$ with $\lambda v_i^h(x)$. This shows that the rank does not increase. Given a rank r' decomposition of the new tensor, scale by $1/\lambda$ to see that $r' \geq r$.

To add $T_{i:x}$ to $T_{i:y}$, simply replace $v_i^h(y)$ with $v_i^h(y) + v_i^h(x)$. Again, this operation can be inverted by replacing $v_i^h(y)$ with $v_i^h(y) - v_i^h(x)$, so the ranks before and after adding $T_{i:x}$ to $T_{i:y}$ are the same. \square

Corollary 1. *If a slice is a linear combination of other slices in the same coordinate, setting every entry to 0 or removing it does not change the rank.*

We call the process of iteratively removing dependent slices until all remaining slices are independent **slice elimination**. When performing row reduction on a matrix, one can add any row u to any other row v without changing the rank. Gaussian elimination turns this into a simple strategy for calculating rank where a row u is added, scaled appropriately, to each other row according to a simple rule. One can think of this as choosing some vector c to extend u by, forming the outer product $u \otimes c$, and adding this to the rank decomposition of the matrix.

Unfortunately, slice elimination does not give an algorithm for tensor rank for two reasons. First, unlike in a matrix, the leftover slices can have rank greater than 1 and it will be unclear what the overall rank of the tensor is. Second, even choosing what multiples to use when adding one slice to the others is NP-hard, and this is the property exploited in showing NP-hardness of tensor rank.

2.2 Substitution

While slice elimination does not give an algorithm for tensor rank, it can be used to analyze the rank of some tensors. A rank 1 slice T' of T can be written uniquely as an outer product, so it is natural to assume T' appears in some minimum rank decomposition of T . Indeed, we show this can be assumed in the following lemma. A form of this for polynomials comes from Pan [10] and goes by many names, such as “slice reduction”, “layer reduction”, and “substitution”. We slightly generalize the proof by Håstad [5]. More general versions can be found in papers by Alexeev et al. [1] and Landsberg and Michałek [9]:

Lemma 2 (Substitution). *Given a d -tensor T of rank r , suppose the 1-slices $T_{1:j}$ for $j \in [k]$ are rank 1 and linearly independent as vectors. Then there is an expansion $T = \sum_{j=1}^r v^j \otimes M^j$ with $\text{rk}(M^j) = 1$ and $M^j = T_{1:j}$ for $j \in [k]$.*

Note that the lemma is stated using 1-slices and using coordinates $j \in [k]$ for $T_{1:j}$ for simplicity. It also holds when taking slices in an arbitrary coordinate i and considering an arbitrary list $S \subset [n_i]$ of i -coordinates.

Proof. Suppose $k = 1$ and write $T_{1:1}$ as a linear combination of M^j . We pick one of the M^j and rearrange the equation to write it in terms of $T_{1:1}$ and the other $M^{j'}$. Then we substitute this into the equations for the other slices, thereby eliminating the use of M^j and introducing $T_{1:1}$.

Explicitly, the rank decomposition gives us

$$T_{1:x} = \sum_{j \in [r]} v^j(x) \otimes M^j$$

Assume without loss of generality that $v^1(1) \neq 0$. We use this equation to replace all occurrences of M^1 with $T_{1:1}$. First we rearrange the equation for $T_{1:1}$ to get

$$M^1 = \frac{1}{v^1(1)} T_{1:1} - \sum_{j=2}^r \frac{v^j(1)}{v^1(1)} M^j$$

Now we can substitute this in the remaining equations and simplify to get

$$\begin{aligned} T_{1:x} &= v^1(x) \left(\frac{1}{v^1(1)} T_{1:1} - \sum_{j=2}^r \frac{v^j(1)}{v^1(1)} M^j \right) + \left(\sum_{j=2}^r v^j(x) M^j \right) \\ &= \frac{v^1(x)}{v^1(1)} T_{1:1} + \sum_{j=2}^r \left(v^j(x) - \frac{v^j(1)}{v^1(1)} \right) M^j \end{aligned}$$

For example

$$\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \text{ becomes } \begin{bmatrix} 1 & b & c \\ \frac{d}{a} & e - \frac{d}{a}b & f - \frac{d}{a}c \\ \frac{g}{a} & h - \frac{g}{a}b & i - \frac{g}{a}c \end{bmatrix}$$

To find an expansion with $M^j = T_{1:j}$ for subsequent j , we simply repeat this procedure. We must be careful that we do not replace an earlier M^j . Fortunately the $T_{1:j}$ for $j \leq k$ are independent, so each one must use some M^j for $j > k$ when it is reached. \square

We would like to extend this proof to delete rank 1 slices and simplify the tensor. With the same notation as in Lemma 2, write $T = \tilde{T} + \sum_{h=1}^k v^h \otimes M^h$ and note that $\text{rk}(T) = \text{rk}(\tilde{T}) + k$. We get the following consequence of Lemma 2.

Corollary 2. *Given a d -tensor T of rank r , suppose the 1-slices $T_{1:j}$ for $j \in [k]$ are rank 1 and linearly independent as vectors. Then*

$$\text{rk}(T) = k + \min(\text{rk}(T - \sum_{j=1}^k v^j \otimes T_{1:j})),$$

where the minimum is taken over the choice of $\{v^j\}_{j \in [k]} \subseteq \mathbb{F}^{n_1}$. Furthermore, in the $T - \sum_{j=1}^k v^j \otimes T_{1:j}$ of minimum rank, we can assume slice $1 : j$ has all zero entries.

Later, we will use T' to refer to the minimum rank tensor in Corollary 2 with slices $1 : j$ removed, for $j \in [k]$.

3 Hardness

To establish NP-hardness of approximation we adapt the proof of NP-hardness by Håstad [5] but reduce from bounded occurrence SAT. Håstad started with a 3-SAT instance ϕ with n variables and m clauses and created a $(2 + n + 2m) \times 3n \times (3n + m)$ tensor T . If ϕ was satisfiable then $\text{rk}(T) = 4n + 2m$, and otherwise $\text{rk}(T) > 4n + 2m$. We follow the same general approach but by starting with a bounded occurrence SAT instance we are able to more tightly relate the rank of the resulting tensor to the minimal number of falsified clauses.

3.1 Reduction

Let us give an overview of our reduction. Given a SAT formula ϕ with n variables, m clauses, and v literal occurrences, we create a $(1 + n + v - m) \times 2n \times (n + m)$ tensor T with n variable slices and m clause slices. Though the reduction works for a varying number of literals per clause, we reduce from MAX-E2-SAT, where each clause is of length exactly 2. We represent each literal as a vector as follows: $v_{x_i} \in \mathbb{F}^{2n}$ contains a 1 in position $2i - 1$ and 0 everywhere else. Vector $v_{\bar{x}_i} \in \mathbb{F}^{2n}$ contains a 1 in positions $2i - 1$ and $2i$, and 0 everywhere else. The 3-slices are then defined as follows:

- For $i \leq n$, $T_{3:i}$ represents variable x_i . It has a 1 in positions $(1, 2i - 1)$ and $(i + 1, 2i)$, and is otherwise 0.
- Suppose the h^{th} clause is $C_h = \ell_1 \vee \ell_2$, where ℓ_i is a literal. Then

- $(T_{3:(n+h)})_{1:1} = v_{y_1}$
- $(T_{3:(n+h)})_{1:h+n+1} = v_{\ell_1} - v_{\ell_2}$
- $T_{3:(n+h)}$ is 0 everywhere else

Here are the slices for the formula $(x \vee y) \wedge (\bar{x} \vee \bar{y})$. We use “.” to represent 0 for readability:

$$\begin{bmatrix} 1 & . & . & . \\ . & 1 & . & . \\ . & . & . & . \\ . & . & . & . \\ . & . & . & . \end{bmatrix}, \begin{bmatrix} . & . & 1 & . \\ . & . & . & . \\ . & . & . & 1 \\ . & . & . & . \\ . & . & . & . \end{bmatrix}, \begin{bmatrix} 1 & . & . & . \\ . & . & . & . \\ . & . & . & . \\ 1 & . & -1 & . \\ . & . & . & . \end{bmatrix}, \begin{bmatrix} 1 & 1 & . & . \\ . & . & . & . \\ . & . & . & . \\ . & . & . & . \\ 1 & 1 & -1 & -1 \end{bmatrix}$$

This tensor is a subset of the tensor from Håstad’s paper [5], with all auxiliary 3-slices and the third copy of the variable columns removed. We use Corollary 2 to reduce T to a matrix M' , which consists of:

- For each $i \in [n]$, a row v'_i , representing variable x_i . This row has a 1 in position $2i - 1$, a variable a_i in position $2i$, and 0 everywhere else.
- For each $h \in [m]$, a row c'_h , representing clause $C_h = \ell_1 \vee \ell_2$. This row is nonzero only in the columns for variables appearing in the clause. It depends on a variable b_h and equals $(1 - b_h)v_{\ell_1} + b_h v_{\ell_2}$.

For example, for the formula $(x \vee y) \wedge (\bar{x} \vee \bar{y})$ we obtain

$$\begin{bmatrix} 1 & a_1 & . & . \\ . & . & 1 & a_2 \\ 1 - b_1 & . & b_1 & . \\ 1 - b_2 & 1 - b_2 & b_2 & b_2 \end{bmatrix}, \text{ which might yield } \begin{bmatrix} 1 & . & . & . \\ . & . & 1 & 1 \\ 1 & . & . & . \\ . & . & 1 & 1 \end{bmatrix}$$

for a specific assignment to the variables. We use M to represent M' under some assignment to its variables, and call it a **realization** of M' . We use c_h to refer to c'_h and u_i to refer to v'_i under an assignment to the variables of M' . Then we have the following:

Lemma 3. *If ϕ has v variable occurrences then $\text{rk}(T) = \min(\text{rk}(M)) + n + m$, where the minimum is taken over all realizations M of M' .*

As one might expect, the variables in M' correspond to the choices taken in Corollary 2 of how to subtract slices.

Proof. We perform slice elimination on the 1-slices of T . We think of the third coordinate as being perpendicular to the page, so we can think of 1-slices as cutting horizontally across the page and also through the page. For $i \geq 2$ each 1-slice $T_{1:i}$ is a matrix with a single nonzero row. Only the 3-slice $T_{3:i-1}$ has nonzero values on $T_{1:i}$. Slice $T_{3:i-1}$ is the 3-slice for variable x_{i-1} if $i \leq n$ or clause C_{i-n-1} if $i > n$.

Each of these 1-slices thus has nonzero entries only in positions $T_{1:i}(x, i - 1)$ for some x . The 1-slices are thus independent as vectors. We apply Corollary 2 and call the simplified tensor T' . There is a variable a_i in T' obtained from subtracting some multiple of $T_{1:i+1}$ from $T_{1:1}$ for $i \leq n$. There is a variable b_j obtained from subtracting some multiple of $T_{1:n+j+1}$ from $T_{1:1}$. Finally, T' has only size 1 in its 1st coordinate, so we view it as the matrix $T'_{3:1} = M'$ and it has the same rank as T' . \square

3.2 Bounding the rank

Now we use hardness properties of SAT. We define an *a-good* instance to be a formula ϕ with m clauses where some assignment leaves at most am clauses unsatisfied and an *a-bad* instance ψ to be a formula where every assignment leaves at least am clauses unsatisfied. We work with E3-OCC-MAX-2SAT, which is SAT where each clause is of size exactly 2 and each variable occurs exactly 3 times. We use the following theorem from Berman and Karpinski [2]:

Theorem 2. *It is NP-hard to distinguish between $(4/792 + \epsilon)$ -good and $(5/792 - \epsilon)$ -bad instances of E3-OCC-MAX-2SAT, for any $\epsilon > 0$.*

We relate the rank of M to the fraction of clauses that ϕ can satisfy to obtain our main theorem.

Lemma 4 (Completeness). *If ϕ is a-good then for some realization M of M' , we have $\text{rk}(M) \leq n + am$.*

Proof. There is some assignment ρ so $\rho(C)$ is true for the maximum number of clauses $C \in \phi$. Set $a_i = 1 - \rho(x_i)$. Suppose clause $C_h = \ell_1 \vee \ell_2$, that ℓ_1 is a literal of x_i , and that ℓ_2 is a literal of x_j . If $\rho(\ell_1)$ is true then set $b_h = 0$ so that $c_h = v_{\ell_1} = u_i$. If instead $\rho(\ell_2)$ is true, then set $b_h = 1$ so that $c_h = v_{\ell_2} = u_j$. Thus no satisfied clause contributes to the rank of M . There are only am remaining rows, so $\text{rk}(M) \leq n + am$. \square

So we see that the variable rows contribute n to the rank and satisfied clause rows contribute nothing. Unsatisfied clause rows should contribute to the rank as well, but we have to be careful. In general it is hard to get a precise bound on the contribution of clauses to the rank of M . When many clauses share the same variable, we can only guarantee that the rank increases by some fraction of the number of unsatisfied clauses. However, for the very simple formulas we study, we get the sharp relationship in Lemma 5.

Lemma 5 (Soundness). *If ϕ is an a-bad instance of E3-OCC-MAX-2SAT then for every realization M of M' , we have $\text{rk}(M) \geq n + am$.*

Proof. Fix M , a realization of M' of minimum rank. We build a boolean assignment ρ that falsifies at most $\text{rk}(M) - n$ clauses in ϕ . Say a row **uses** i if its last nonzero entry is in column $2i - 1$ or $2i$. Sort the rows of M by the index used.

One row using i is u_i . There are at most 3 more rows using i , which we call r, s , and t , if they exist. Call the corresponding clauses C_r, C_s , and C_t . If $\text{rk}(M_i) \setminus \text{rk}(M_{i-1}) = 1$ then they share the same literal of x_i . Call this literal ℓ_i . Then setting $\rho(\ell_i)$ true satisfies all these clauses. Otherwise, set $\rho(\ell_i)$ to satisfy the majority of the clauses C_r, C_s , and C_t .

Because there are only 3 clauses that contain x_i , this leaves at most 1 clause unsatisfied. There are at most $\text{rk}(M) - n$ indices $i \in [n]$ where $\text{rk}(M_i) \setminus \text{rk}(M_{i-1}) > 1$, so we end up with at most $\text{rk}(M) - n$ falsified clauses. Since every assignment to ϕ falsifies at least am clauses, we conclude that $\text{rk}(M) \geq n + am$. \square

3.3 Inapproximability

Now we get our main theorem.

Theorem 3 (Theorem 1). *It is NP-hard to approximate 3-tensor rank over any finite field \mathbb{F} within a factor of $1 + 1/1852 - \delta$, for any $\delta > 0$.*

Proof. We combine Theorem 2, Lemma 4, and Lemma 5. We know $\text{rk}(T) = \min(\text{rk}(M)) + n + m$. If ϕ is a -good then $\min(\text{rk}(M)) \leq n + am$. If ϕ is a -bad then $\min(\text{rk}(M)) \geq n + am$. By Theorem 2 it is NP-hard to distinguish $(4/792 + \epsilon)$ -good and $(5/792 - \epsilon)$ -bad instances. Hence, it is NP-hard to distinguish whether $\text{rk}(T) \leq 2n + (1 + 4/792 + \epsilon)m$ or $\text{rk}(T) \geq 2n + (1 + 5/792 - \epsilon)m$.

Using $m = 3n/2$ and given $\delta > 0$ we can choose $\epsilon > 0$ so that we have an inapproximability ratio of $1 + 1/1852 - \delta$. \square

4 Discussion

It seems likely that much better inapproximability results are possible for tensor rank. By counting, we know there are tensors of shape $n \times n \times n$ with rank at least $n^2/3$. Say that a family $T(n)$ of tensors of increasing size n is “explicit” if there is a polynomial time algorithm that takes input n written in unary and prints $T(n)$. Despite the existence of $n \times n \times n$ tensors of quadratic rank, the only “explicit” such tensors have rank at most $3n - o(n)$ [14]. In fact, no explicit 3-tensors T of any shape are known where $\text{rk}(T)$ is greater than the sum of the coordinates of T .

Alexeev et al. [1] claim that “any gap-preserving reduction from NP to tensor rank would automatically yield lower bounds for explicit tensors”. While it is not obvious to us how to turn this into a technically precise statement, we sketch the main idea. Suppose we reduce SAT to 3-tensor rank and obtain a superconstant hardness of approximation for tensor rank. If the reduction always outputs a tensor T with largest dimension n , and where the slices in every direction are independent, then the tensor has rank at least n . If, for some $c > 3$, the reduction demonstrates c -hardness of approximation then it must sometimes output tensors of rank at least cn . If this happens on unsatisfiable formulas, we can apply the reduction to get an explicit high rank tensor.

While it is possible that a randomized reduction could avoid this barrier, we are still motivated to find high rank tensors. One approach to finding higher rank explicit 3-tensors or to improving the gap in a reduction is to Kronecker multiply together smaller 3-tensors. It is known [16] that the product is not multiplicative. For the curious reader, let us give a simple counterexample:

Example 1. Take a 3-tensor T with slices

$$\begin{bmatrix} 1 & \cdot \\ \cdot & 1 \end{bmatrix} \begin{bmatrix} \cdot & 1 \\ \cdot & \cdot \end{bmatrix}$$

Kronecker multiplying this with itself, we get a 3-tensor with slices

$$\begin{bmatrix} 1 & \cdot & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & 1 \end{bmatrix} \begin{bmatrix} \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \begin{bmatrix} \cdot & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \begin{bmatrix} \cdot & \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix}$$

The former has rank 3, but the latter set can be formed as a linear combination of 8 rank 1 matrices:

$$\begin{bmatrix} 1 & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \begin{bmatrix} \cdot & -1 & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \begin{bmatrix} \cdot & 1 & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & -1 & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & -1 \\ \cdot & \cdot & \cdot & -1 \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix}$$

$$\begin{bmatrix} \cdot & \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 \end{bmatrix}$$

T has rank 3, but when Kronecker multiplying with itself it has rank at most 8, not 9. This is true over any field \mathbb{F} .

Perhaps it is still true that when k copies of a 3-tensor T of shape $n \times n \times n$ with $\text{rk}(T) > n$ are multiplied together, the rank grows as r^k for some $\text{rk}(T) \geq r > n$. If true, this would immediately give high rank 3-tensors.

Alternatively, consider summing tensors together as an approach to constructing high rank explicit tensors. The *direct sum* of two d -tensors, S and T , is written $S \oplus T$. If the i^{th} coordinate of S has size n_i^S and the i^{th} coordinate of T has size n_i^T then the i^{th} coordinate of $S \oplus T$ has size $n_i^S + n_i^T$. Define t by $t_i = s_i - n_i^S$. Then $S \oplus T$ is given by

$$(S \oplus T)(s) = \begin{cases} S(s) & \text{if } \forall i. s_i \leq n_i^S \\ T(t) & \text{if } \forall i. s_i > n_i^S \\ 0 & \text{otherwise} \end{cases}$$

Strassen's additivity conjecture states that, given tensors T and R of the same order, the tensor $T \oplus R$ has rank $\text{rk}(T) + \text{rk}(R)$. This has been proved for a

few special cases, but in each case the tensors involved have rank at most the sum of their coordinates. Were this conjecture to be true in full generality, we immediately get slightly superlinear 3-tensors.

Lemma 6. *Strassen's additivity conjecture implies explicit 3-tensors of size $n \times n \times n$ and rank $\Omega(n(\log(n))^{1/3})$ over any finite field \mathbb{F} .*

Proof. Let $s = (\log(n))^{1/3}$ and exhaustively search over $s \times s \times s$ tensors over field \mathbb{F} for a high rank tensor. To do so, create a table *Rank* that lists an estimate of the rank all possible $s \times s \times s$ tensors. Initially, set every entry in *Rank* to s^2 (the maximum possible rank). Iterate over all $|\mathbb{F}|^{O(s)}$ rank 1 tensors A . For each one, iterate over all $s \times s \times s$ tensors T in and replace $\text{Rank}(A + T)$ with $\min(\text{Rank}(A + T), \text{Rank}(T) + 1)$. After all rank 1 tensors have been iterated over, *Rank* correctly lists the rank of every tensor in the table.

The above process takes $|\mathbb{F}|^{O(s)}|\mathbb{F}|^{s^3} = |\mathbb{F}|^{O(s^3)}$ time. It can be seen by counting that there is at least one rank $\frac{s^2}{3}$ tensor of size $s \times s \times s$ so iterate over all such tensors T and pick the first where $\text{Rank}(T) \geq \frac{s^2}{3}$. Direct sum $\frac{n}{s}$ copies of T together to obtain T' . By Strassen's additivity conjecture, $\text{rk}(T') = \frac{n}{s} \cdot \text{rk}(T) \geq \frac{n}{s} \cdot \frac{s^2}{3} = \frac{ns}{3}$. \square

Acknowledgements

I thank Johan Håstad for guidance during research and for thoroughly reviewing and giving advice on the writing of this paper. I thank Per Austrin for discussions and for improving the inapproximability constant.

References

- [1] B. Alexeev, M. Forbes, and J. Tsimmerman, *Tensor rank: Some lower and upper bounds*, Conference on Computational Complexity (2011), pp. 283-291.
- [2] P. Berman and M. Karpinski, *Efficient amplifiers and bounded degree optimization*, Inst. fr Informatik (2001).
- [3] M. Bläser, *On the complexity of the multiplication of matrices of small formats*, Journal of Complexity 19.1 (2003), pp. 43-60
- [4] M. Bläser, *Explicit tensors*, Perspectives in Computational Complexity (2014), pp. 117-130.
- [5] J. Håstad, *Tensor rank is NP-Complete*, International Colloquium on Automata, Languages, and Programming (1989), pp. 451-460.
- [6] T. Kolda and B. Bader, *Tensor decompositions and applications*, SIAM review 51.3 (2009), pp. 455-500.

- [7] J. Kruskal, *Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics*, Linear algebra and its applications 18.2 (1977), pp. 95-138.
- [8] J. Laderman, *A noncommutative algorithm for multiplying 3 3 matrices using 23 multiplications*, Am. Math. Soc. Vol. 82. (1976), pp. 126-128
- [9] J. Landsberg and M. Michałek, *Abelian tensors*, Journal de Mathématiques Pures et Appliquées (2016).
- [10] V. Pan, *Methods of computing values of polynomials*, Russian Mathematical Surveys 21.1 (1966), pp. 105-136.
- [11] R. Raz, *Tensor-rank and lower bounds for arithmetic formulas*, Journal of the ACM 60.6 (2013), page 40.
- [12] M. Schaefer and D. Štefankovič. *The complexity of tensor rank*, arXiv preprint arXiv:1612.04338 (2016).
- [13] Y. Shitov, *How hard is the tensor rank?* arXiv preprint arXiv:1611.01559 (2016).
- [14] A. Shpilka, *Lower bounds for matrix product*, SIAM Journal on Computing 32.5 (2003), pp. 1185-1200.
- [15] V. Strassen, *Gaussian elimination is not optimal*, Numerische mathematik 13.4 (1969), pp. 354-356.
- [16] N. Yu, E. Chitamber, C. Guo, and R. Duan, *The tensor rank of the tripartite state $|W\rangle^{\otimes n}$* , Physical Review A 81.1 (2010).