The Log-Rank Conjecture: New Equivalent Formulations

Lianna Hambardzumyan [*§] Shachar Lovett [†] Morgan Shirley [‡§]

April 9, 2026

Abstract

The log-rank conjecture is a longstanding open problem with multiple equivalent formulations in complexity theory and mathematics. In its linear-algebraic form, it asserts that the rank and partitioning number of a Boolean matrix are quasi-polynomially related.

We propose a relaxed but still equivalent version of the conjecture based on a new matrix parameter, signed rectangle rank: the minimum number of all-1 rectangles needed to express the Boolean matrix as a ± 1-sum. Signed rectangle rank lies between rank and partition number, and our main result shows that it is in fact equivalent to rank up to a logarithmic factor. Additionally, we extend the main result to tensors. This reframes the log-rank conjecture as: can every signed decomposition of a Boolean matrix be made positive with only quasi-polynomial blowup?

As an application, we prove an equivalence between the log-rank conjecture and a conjecture of Lovett and Singer–Sudan on cross-intersecting set systems.

1 Introduction

The log-rank conjecture is one of the most well-known problems in complexity theory, and despite extensive work it remains unsolved. It asserts that for a Boolean matrix its communication complexity and the logarithm of its matrix rank over the reals are polynomially related. An equivalent linear-algebraic formulation of the conjecture is that for Boolean matrices, the matrix rank and partitioning number (sometimes called the binary rank) are quasi-polynomially related.

More formally, for a Boolean matrix M let $\mathrm{r}(M)$ denote its rank over the reals, and let $\mathrm{p}(M)$ denote its partitioning number: the minimum number of all-1 submatrices that partition the 1-entries of M. Equivalently, the partitioning number is the minimum number p such that $M = \sum_{i=1}^{p} R_i$, where each R_i is a *primitive* matrix—an all-1 submatrix, possibly after adding all-zero rows and columns. If we relax the decomposition to allow general rank-1 matrices instead of primitive matrices, we recover the standard matrix rank. That is, $\mathrm{r}(M)$ is the minimum number r such that $M = \sum_{i=1}^{r} M_i$, with each M_i of rank-1. In this sense, matrix rank can be seen as a relaxation of the partitioning number, and trivially $\mathrm{r}(M) \le \mathrm{p}(M)$. The log-rank conjecture asks how well this relaxation estimates the partitioning number:

[*]The University of Copenhagen, Denmark. lianna.hambardzumyan@gmail.com. This work is funded by the European Research Council (ERC) under grant agreement no. 101125652 (ALBA).

[†]University of California San Diego, CA, USA. slovett@ucsd.edu. Research supported by Simons Investigator Award #929894 and NSF award CCF-2425349.

[‡]Lund University, Sweden. morgan.shirley@cs.lth.se. Supported by Knut and Alice Wallenberg grant KAW 2023.0116.

[§]Most of the work was done while the first and third authors were postdoctoral researchers at University of Victoria, Canada, and funded by NSERC.

Conjecture 1.1 (Log-rank conjecture [LS88])**.** *For any Boolean matrix M,*

$$\log \mathrm{p}(M) \leq (\log \mathrm{r}(M))^{O(1)}.$$

The log-rank conjecture was first posed by Lovász and Saks [LS88] in the context of communication complexity: is the deterministic communication complexity of a Boolean matrix polynomially related to the logarithm of its rank over the reals? A closely related question had appeared even earlier in graph theory. There, the log-rank conjecture is equivalent to asking whether the logarithm of a graph's chromatic number is quasi-polynomially related to the rank of its adjacency matrix [Nuf76, Faj88, LS88].

Regarding the state of the art, the best known upper bound is due to Sudakov and Tomon [ST25], improving on previous work of Lovett [Lov16], and shows that

$$\log \mathrm{p}(M) \leq O\left(\sqrt{\mathrm{r}(M)}\right).$$

The largest known separation follows from [BBG+21, Corollary 3], which yields a matrix M satisfying $\log \mathrm{p}(M) \geq \Omega(\log^2 \mathrm{r}(M))$. For a more detailed overview of the log-rank conjecture and its equivalent formulations, we refer the reader to the survey by Lee and Shraibman [LS23], as well as the textbooks by Jukna [Juk12] and Rao and Yehudayoff [RY20].

In this paper, we consider a "gradual relaxation" from partitioning number to rank via an intermediate complexity measure: the *signed rectangle rank*. This notion allows decomposition of the matrix into *signed* primitive matrices. Formally, let the *signed rectangle rank* of M, denoted by $\mathrm{srr}(A)$, be the minimum number t such that $M = \sum_{i=1}^{t} \varepsilon_i R_i$, where $\varepsilon_i \in \{1, -1\}$ and each R_i is a primitive matrix. Trivially, we have

$$\mathrm{r}(M) \leq \mathrm{srr}(M) \leq \mathrm{p}(M).$$

A promising approach to the log-rank conjecture is to study how "close" is signed rectangle rank to rank and partitioning number. To resolve the conjecture, it would suffice to either prove both of the following statements or disprove one of them:

- $\mathrm{srr}(M)$ is quasi-polynomially related to $\mathrm{r}(M)$, and
- $\mathrm{srr}(M)$ is quasi-polynomially relate to $\mathrm{p}(M)$.

The main result of the paper is that rank and signed rectangle rank are tightly related.

Theorem 1.2 (Main Theorem)**.** *Every Boolean matrix M of rank r can be written as a ± 1-linear-combination of at most $O(r \log r)$ primitive matrices, that is,*

$$\mathrm{srr}(M) \leq O(r \log r).$$

From this we immediately get an equivalent formulation of the log-rank conjecture:

Conjecture 1.3. *For every Boolean matrix M, $\log \mathrm{p}(M) \leq (\log \mathrm{srr}(M))^{O(1)}$.*

Corollary 1.4 (from Theorem 1.2)**.** *The log-rank conjecture (Conjecture 1.1) is equivalent to Conjecture 1.3.*

A natural question left open here is whether the bound in Theorem 1.2 can be improved, or if it is already tight:

Question 1.5. *Is it true that for every Boolean matrix M, $\mathrm{srr}(M) \leq O(\mathrm{r}(M))$?*

A result of Lindström [Lin65] shows that the bound is tight for our proof technique (see the discussion at the end of Section 2). Therefore, any improvements on the bound require different ideas.

Power of cancellations vs. monotonicity. The log-rank conjecture compares two ways of measuring the complexity of a matrix: rank and partitioning number. These measures differ in two fundamental ways:

- **Combinatorial structure**: Rank expresses a matrix as a summation of rank-1 matrices with arbitrary real coefficients and no apparent combinatorial structure. In contrast, partitioning number only allows primitive matrices, which are purely *combinatorial* objects.
- **Cancellations**: Rank allows for *cancellations* in the summation. Partitioning, on the other hand, is a *monotone* model: all-1 submatrices (also known as *rectangles*) must be combined positively and disjointly.

The log-rank conjecture asks whether these two measures are nevertheless quasi-polynomially related.

The measure we study, *signed rectangle rank*, is an intermediate notion that allows cancellations, but only at the level of rectangles. In this way, it relaxes partitioning by permitting signed combinations of rectangles, while still restricting to simple combinatorial building blocks. Intuitively, our main theorem shows that allowing cancellation between rectangles already captures essentially all of the linear structure present in a Boolean matrix.

An analogous perspective arises from studying another measure of a matrix, the *non-negative rank*. If one forbids cancellations but allows non-negative real values, without imposing any combinatorial structure on the decomposition, the resulting measure is the non-negative rank. Formally, non-negative rank of a $n \times n$ matrix M, denoted by $\mathrm{nnr}(M)$ is the minimum number r such that $M = \sum_{i=1}^{r} u_i v_i^T$, where $u_i, v_i \in \mathbb{R}^n$ are non-negative column vectors. Trivially, we have

$$\mathrm{r}(M) \leq \mathrm{nnr}(M) \leq \mathrm{p}(M).$$

Analogous to our main theorem, it is known that for Boolean matrices this relaxation of partitioning number does not increase its power too much (this follows from a result of Lovász and Saks [LS93]; see e.g. Watson [Wat16]):

$$\log \mathrm{p}(M) \leq O(\log^2 \mathrm{nnr}(M)).$$

Consequently, the log rank conjecture is equivalent to asking whether there is a quasi-polynomial relationship between rank and non-negative rank, namely, $\log \mathrm{nnr}(M) \leq (\log \mathrm{r}(M))^{O(1)}$. Comparing this with Corollary 1.4, our results imply that the log-rank conjecture also reduces to converting a signed rectangle decomposition into a fully monotone one, with at most a quasi-polynomial increase in the number of primitive matrices. See Figure 1 for an illustration of the relationships between these measures.

Together, these paint a picture of what is left to understand about the log-rank conjecture. The fact that rank allows arbitrary real coefficients and partitioning number is combinatorial is a red herring: the only relevant question is whether allowing cancellations is significantly stronger than monotonicity.

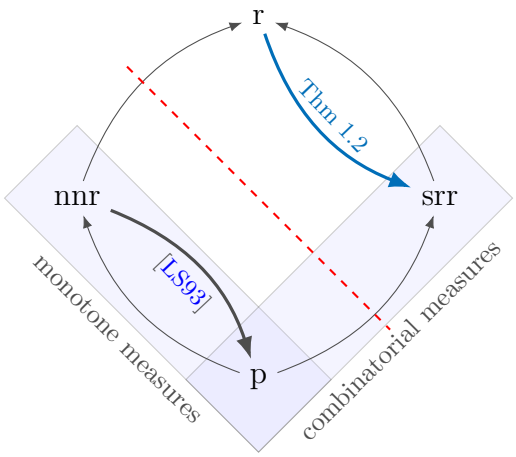


Figure 1: Relationships between the matrix measures discussed; an arrow $a \to b$ indicates that $\log b \leq \log^{O(1)} a$. The thin arrows indicate bounds which are trivial from the definitions. The red dashed line indicates that proving the log-rank conjecture is equivalent to showing that *any* measure below the line is quasi-polynomially related to *any* measure above the line.

Equivalent conjecture on cross-intersecting set systems. Let $\mathcal{S}, \mathcal{T} \subseteq 2^{[d]}$ be two set families. We say the pair $(\mathcal{S}, \mathcal{T})$ is *L-cross-intersecting* for some $L \subseteq \{0, \dots, d\}$ if for all $S \in \mathcal{S}$ and $T \in \mathcal{T}$, we have $|S \cap T| \in L$; that is, the size of every pairwise intersection belongs to L. Cross-intersecting set families have been widely studied in combinatorics, with much of the work focusing on their extremal properties [FR87, Sga99, KS05, HMST24].

The following conjecture about $\{a, b\}$-cross-intersecting set systems was independently proposed by Lovett [Lov21] and Singer–Sudan [SS22], who both observed that it is implied by the log-rank conjecture.

Conjecture 1.6. *Let $\mathcal{S} = \{S_1, \dots, S_m\}$ and $\mathcal{T} = \{T_1, \dots, T_n\}$ be an $\{a, b\}$-cross-intersecting pair of families from $2^{[d]}$, where $a, b \in \{0, \dots, d\}$. Then there exist subfamilies $\mathcal{A} \subseteq \mathcal{S}$ and $\mathcal{B} \subseteq \mathcal{T}$ such that $(\mathcal{A}, \mathcal{B})$ is either $\{a\}$- or $\{b\}$-cross-intersecting, and*

$$|\mathcal{A}|, |\mathcal{B}| \geq 2^{-\operatorname{polylog}(d)} \cdot |\mathcal{S}||\mathcal{T}|.$$

As an application of Theorem 1.2, we show that the log-rank conjecture is equivalent to this conjecture.[1]

Theorem 1.7. *The log-rank conjecture (Conjecture 1.1) is equivalent to Conjecture 1.6.*

[1] This equivalence was claimed in [SS22] as a parenthetical remark. However, this was a typo. The remark was meant to claim the implication from the log-rank conjecture [SS25].

Tensors. In addition to the matrix case we extend Theorem 1.2 to Boolean tensors, showing analogous relation between the tensor rank and signed rectangle rank of a tensor. This has implications for the generalization of the log-rank conjecture to *number-in-hand* communication complexity (asked explicitly in [DKW11]).

Organization. We give the proof of Theorem 1.2 in Section 2 and discuss the open question of how tight the bound is. In Section 3, we prove the equivalence of log-rank conjecture to Conjecture 1.6. Finally, prove the extension of the main theorem to tensors in Section 4.

2 Proof of the Main Theorem

Let $M = (m_{i,j})$ be an $m \times n$ Boolean matrix of rank r. For a subset S of the columns, define its *column-sum* to be the column vector c such that for $i \in [m]$, $c_i = \sum_{j \in S} m_{i,j}$. That is, c is the entrywise sum of all the columns of S along all the rows. Denote the column-sum of S as $\mathrm{sum}(S)$. Call a subset S of columns *independent* if all the subsets of S have distinct column-sums.

Claim 2.1. *Let S be an independent set of columns of M. Then $|S| \leq O(r \log r)$, where r is the rank of M.*

Proof. Let A_S be the matrix whose columns are the column-sums of S, so A_S has $2^{|S|}$ columns. The entries of A_S are in $\{0, \ldots, |S|\}$ as they are sums of at most $|S|$ entries of M, which take value $\{0, 1\}$. The rank of A_S is at most r because its columns are in the span of the columns of M, which itself has rank r. Thus, there is a set of rows R with $|R| = r$ such that, for any column c of A_S, the values $(c_i)_{i \in R}$ are sufficient to determine every entry of c. Therefore, there are at most $(|S| + 1)^r$ unique columns of A_S, and since every column of A_S is unique, we have that the subsets of S have at most $(|S| + 1)^r$ column-sums. Therefore, $2^{|S|} \leq (|S| + 1)^r$, which implies $|S| \leq O(r \log r)$. □

Claim 2.2. *Let S be a maximal independent set of columns of M. Then every column of M can be expressed as a ± 1-linear combination of columns in S.*

Proof. This trivially holds for the columns in S. Fix a column $c \notin S$. Since S is a maximal independent set, $S \cup \{c\}$ is not independent. This means that there are two subsets A and B of $S \cup \{c\}$ with the same column-sum. We can assume that these subsets are disjoint; if they are not, then removing the common columns still results in equal column-sums. Note A and B cannot both exclude c, as this would contradict S being independent. Assume $c \in B$, and let $B' = B \setminus \{c\}$. Combining this with $\mathrm{sum}(A) = \mathrm{sum}(B)$, we get $\mathrm{sum}(A) = \mathrm{sum}(B') + c$. Noting that $A \cap B' = \emptyset$ and $A, B' \subseteq S$ concludes that $c = \mathrm{sum}(A) - \mathrm{sum}(B')$ is the desired linear combination. □

Combining these two claims concludes the proof of Theorem 1.2 as follows. Let S be the largest independent set of columns of M. By Claim 2.2 every column y can be expressed as a ± 1-linear combination of columns in S. Thus, M can be written as

$$M(x, y) = \sum_{c \in S} \alpha_c(y) c(x),$$

where each coefficient $\alpha_c(y) \in \{-1, 0, 1\}$. Decompose $\alpha_c(y) = \alpha_c^+(y) - \alpha_c^-(y)$, where $\alpha_c^+(y), \alpha_c^-(y) \in \{0, 1\}$. For each column $c \in S$, define,

$$R_c^+(x, y) = \alpha_c^+(y) c(x) \qquad \text{and} \qquad R_c^-(x, y) = \alpha_c^-(y) c(x)$$

Each of R_c^+ and R_c^- is either an all-zeroes matrix or a primitive matrix, since they are outer products of $\{0,1\}$-valued vectors. Hence, $M = \sum_{c \in S}(R_c^+ - R_c^-)$, which expresses M as ± 1-sum of at most $2|S|$ primitive matrices. Finally, applying the bound on $|S|$ from Claim 2.1, we conclude that $\mathrm{srr}(M) \leq O(r \log r)$.

We remind the reader of the following open question:

Question 1.5. *Is it true that for every Boolean matrix M, $\mathrm{srr}(M) \leq O(\mathrm{r}(M))$?*

An approach one might hope to take is to improve the statement of Claim 2.1, that is, to prove a better upper bound on the size of an independent set of columns in a low-rank Boolean matrix. However, a result of Lindström [Lin65, Theorem 1] shows that Claim 2.1 is tight. Concretely, Lindström gives a construction of a set of $\Theta(r \log r)$ vectors in $\{0,1\}^r$ which are independent (that is, all their column sums are distinct).

3 Equivalence to the cross-intersecting set systems conjecture

In order to prove Theorem 1.7, we will use the equivalent version of the log-rank conjecture by Nisan and Wigderson [NW95]. A submatrix of a matrix is called a *monochromatic rectangle* if all of its entries have the same value.

Conjecture 3.1 ([NW95])**.** *Any Boolean matrix M has a monochromatic rectangle of density $2^{-\operatorname{polylog}(\mathrm{r}(M))}$.*

Sgall [Sga99] considered Conjecture 1.6 in the case of $\{a, a+1\}$-cross-intersecting set systems and noted that it would follow from the log-rank conjecture. Sgall noted that any $\{a,b\}$-cross-intersecting set pair of families $(\mathcal{S}, \mathcal{T})$ from $2^{[d]}$ can be represented as an $m \times n$ matrix $M_{\mathcal{S},\mathcal{T}}$ over $\{a,b\}$ with rank at most d, where $M_{\mathcal{S},\mathcal{T}}[i,j] := |S_i \cap T_j|$. One can write $M_{\mathcal{S},\mathcal{T}}$ as a sum of d primitive matrices R_k indexed by elements of $[d]$, where $R_k[i,j] = 1$ iff $k \in S_i \cap T_j$. Each R_k has rank 1, so the total rank of $M_{\mathcal{S},\mathcal{T}}$ is at most d. Conversely, any $\{a,b\}$-valued matrix that is a sum of d primitive matrices corresponds to an $\{a,b\}$-cross-intersecting set system over a universe of size d.

Theorem 1.7. *The log-rank conjecture (Conjecture 1.1) is equivalent to Conjecture 1.6.*

Proof. As mentioned above, the fact that Conjecture 1.6 is implied by the log-rank conjecture was proved by Sgall [Sga99]. We include the proof here for completeness. For an $\{a,b\}$-cross-intersecting family pair $(\mathcal{S}, \mathcal{T})$ with $a \leq b$, we can write $M_{\mathcal{S},\mathcal{T}} = (b-a)B + aJ$, where B is some Boolean matrix and J is the all-ones matrix. Then, $\mathrm{r}(B) - 1 \leq \mathrm{r}(M_{\mathcal{S},\mathcal{T}}) \leq \mathrm{r}(B) + 1$. By Conjecture 3.1, B has a monochromatic rectangle of density $2^{-\operatorname{polylog}(\mathrm{r}(B))}$, which yields a monochromatic rectangle in $M_{\mathcal{S},\mathcal{T}}$ of density $2^{-\operatorname{polylog}(\mathrm{r}(M))} \geq 2^{-\operatorname{polylog}(d)}$ from the discussion above.

For the reverse direction, assume Conjecture 1.6 holds. Let A be a Boolean $m \times n$ matrix with signed rectangle rank u. By Theorem 1.2 and Conjecture 3.1, it suffices to find a monochromatic rectangle in A of density at least $2^{-\operatorname{polylog}(u)}$. We reduce this to the problem of finding a large $\{a\}$- or $\{b\}$-cross-intersecting subfamily in an $\{a,b\}$-cross-intersecting set system. Let a, b be integers chosen later. Define a matrix $A' \in \{a,b\}^{m \times n}$ as follows:

$$A'[i,j] := \begin{cases} a & \text{if } A[i,j] = 1, \\ b & \text{if } A[i,j] = 0. \end{cases}$$

Our goal is to show that A' corresponds to an $\{a, b\}$-cross-intersecting set system over a universe of size $d = \Theta(u)$.

Let $A = \sum_{i=1}^{u} \varepsilon_i R_i$ be the signed rectangle rank decomposition for A. We now construct a set of $2u$ primitive matrices $\{R'_i\}$ such that $A' = \sum_{i=1}^{2u} R'_i$, by replacing each R_i with a pair of primitive matrices R'_{2i-1}, R'_{2i} depending on the sign of ε_i.

- If $\varepsilon_i = 1$, let $R'_{2i-1} = J$ (the all-ones matrix), and $R'_{2i} = R_i$. Then, $R'_{2i-1} + R'_{2i} = J + R_i$
- If $\varepsilon_i = -1$, let $A_i \subseteq [m]$, $B_i \subseteq [n]$ be such that $R_i = 1$ on $A_i \times B_i$ and zero elsewhere. Define R'_{2i-1} to be 1 on $([m] \setminus A_i) \times [n]$ and 0 elsewhere, and R'_{2i} to be 1 on $A_i \times ([n] \setminus B_i)$. Then $R'_{2i-1} + R'_{2i} = J - R_i$.

In either case, the pair (R'_{2i-1}, R'_{2i}) replaces $\varepsilon_i R_i$ by $J + \varepsilon_i R_i$. Thus, each entry of A' satisfies $A'[i, j] = A[i, j] + u$, and therefore A' corresponds to a $\{u, u + 1\}$-cross-intersecting set family over $[d]$ for $d = 2u$. Applying Conjecture 1.6, we obtain a large monochromatic rectangle in A', and hence in A, of density at least $2^{-\operatorname{polylog}(u)}$. □

Remark 3.2. The above proof shows that the log-rank conjecture is equivalent to a special case of Conjecture 1.6 for $\{a, a + 1\}$-cross-intersecting set systems.

4 Generalizing the Main Theorem to tensors

The log-rank conjecture extends naturally to Boolean tensors. Motivated by this perspective, we show that Theorem 1.2 continues to hold for the tensor rank of constant-order Boolean tensors, giving an equivalent formulation of the log-rank conjecture for tensors.

Let T be an order-ℓ Boolean tensor: a multilinear map $[n_1] \times \ldots \times [n_\ell] \to \{0, 1\}$ (for some natural numbers $n_1, \ldots, n_\ell$). In other words, T can be expressed as an ℓ-dimensional array whose entries are all in $\{0, 1\}$.

We say T has tensor rank 1 if there are maps $v_i : [n_i] \to \mathbb{R}$ for $i \in [\ell]$ such that

$$T(x_1, \ldots, x_\ell) = v_1(x_1) \cdot \ldots \cdot v_\ell(x_\ell).$$

The *tensor rank* of a tensor T, denoted by $\operatorname{tr}(T)$ is the minimum number r such that T can be expressed as the sum of r rank-1 tensors.

A *primitive tensor* is the natural generalization of a primitive matrix: it is a rank-1 tensor where v_i's map to $\{0, 1\}$. That is, a primitive tensor can be expressed as a multidimensional array which is all-1 on some product set $Q_1 \times \ldots \times Q_\ell$ for $Q_i \subseteq [n_i]$ and is all-0 elsewhere. Having this, the *partitioning number* and *signed rectangle rank* of a tensor are defined analogously as the minimum integers p and r, respectively, such that $T = \sum_{i=1}^{p} S_i$ and $T = \sum_{i=1}^{r} \varepsilon_i S_i$, where S_i are primitive tensors and $\varepsilon_i \in \{\pm 1\}$.

The log-rank conjecture for tensors can be formulated as follows.

Conjecture 4.1 (Log-rank conjecture for tensors)**.** *For any constant-order Boolean tensor T,*

$$\log \mathrm{p}(T) \leq (\log \operatorname{tr}(T))^{O(1)}.$$

In the communication complexity language, the log-rank conjecture for tensors states that the *multiparty number-in-hand communication complexity* of a tensor T is polynomially related to the logarithm of the tensor rank of T (see [DKW11]). Note, it is known that the log-rank conjecture does not hold for tensors of super-polynomial order [Li11].

We prove Theorem 1.2 generalizes to tensors.

Theorem 4.2. *Let $T : [n_1] \times \ldots \times [n_\ell] \to \{0,1\}$ be an order-ℓ Boolean tensor with $\ell \geq 2$. If the tensor rank of T is r, then T can be written as a ± 1-linear-combination of at most $(cr \log r)^{(\ell-1)}$ primitive tensors, where c is an absolute constant.*

A direct corollary of this theorem is an equivalent formulation of the log-rank conjecture for tensors, analogous to Conjecture 1.3 for matrices.

Conjecture 4.3. *Every constant-order Boolean tensor T satisfies $\log \mathrm{p}(T) \leq (\log \mathrm{srr}(T))^{O(1)}$.*

Finally, let us set up the proof of Theorem 4.2. A *slice* of an order-ℓ tensor is the order-$(\ell-1)$ tensor obtained by taking a coordinate and setting it to a fixed value. For $\lambda \in [\ell]$, a *λ-slice* is a slice where we specify that λ is the coordinate to be fixed. Similar to the proof of Theorem 1.2, we define a slice-sum to be the entrywise sum of slices, and an independent set of slices is one where all of its subsets have distinct slice-sums.

Claim 4.4. *Let S be an independent set of λ-slices of T. Then $|S| \leq O(r \log r)$, where r is the tensor rank of T.*

Proof. Let $m = \prod_{i \neq \lambda} n_i$. Consider the flattening of T to a matrix $M_T \in \{0,1\}^{m \times n_\lambda}$. Then M_T has rank at most r and, following the bijection between λ-slices of T and columns of M_T, there is an independent set of columns S' in M_T with $|S'| = |S|$. Apply Claim 2.1 to obtain $|S'| \leq O(r \log r)$. □

Claim 4.5. *Let S be a maximal independent set of λ-slices of T. Then every λ-slice of T can be expressed as a ± 1-linear combination of λ-slices in S.*

We omit the proof of Claim 4.5 as it is analogous to Claim 2.2.

Proof of Theorem 4.2. The proof is by induction. The base case is $\ell = 2$, which is proven by Theorem 1.2. For $\ell > 2$, arbitrarily choose a coordinate λ and use Claim 4.5 to express T as a ± 1-linear combination of a maximal independent set of λ-slices S. Each λ-slice of S has tensor rank at most r, and so by the inductive hypothesis, each λ-slice in S can be expressed as a ± 1-linear combination of $(cr \log r)^{(\ell-2)}$ order-$(\ell-1)$ primitive tensors.

In a similar fashion to the proof of Theorem 1.2, we therefore can write T as a ± 1-sum of $|S| \cdot (cr \log r)^{(\ell-2)}$ order-ℓ primitive tensors. Conclude by substituting $|S| = O(r \log r)$ using Claim 4.4. □

Remark 4.6. We note that the proof of Theorem 4.2 only uses the tensor rank as an upper bound for the *flattening rank* of a tensor, which is the maximal rank of any flattening of it as an $n_\lambda \times \left(\prod_{i \neq \lambda} n_i\right)$ matrix over $\lambda \in [\ell]$. If this flattening rank is r, then the conclusion of Theorem 4.2 still holds.

Moreover, in this formulation, the quantitative bound of $O((r \log r)^{\ell-1})$ primitive tensors is near optimal. Indeed, consider all tensors $T : [r]^\ell \to \{0,1\}$. Any such tensor has flattening rank

$\leq r$. On the other hand, a simple counting argument shows that most such tensors need $\Omega(r^{\ell-1}/\ell)$ primitive tensors in their decomposition.

If we return however to the original formulation of Theorem 4.2 using tensor rank, it is unclear if the bound is close to tight. In fact, we suspect that it can be significantly improved.

Question 4.7. *Can the bound in Theorem 4.2 be improved to $o(r^{\ell-1})$?*

Acknowledgments

We are grateful to Adi Shraibman and Amir Yehudayoff for enlightening discussions.

References

[BBG+21] Kaspars Balodis, Shalev Ben-David, Mika Göös, Siddhartha Jain, and Robin Kothari. Unambiguous DNFs and Alon–Saks–Seymour. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 116–124. IEEE Computer Society, 2021.

[DKW11] Jan Draisma, Eyal Kushilevitz, and Enav Weinreb. Partition arguments in multiparty communication complexity. *Theoretical Computer Science*, 412(24):2611–2622, May 2011.

[Faj88] Siemion Fajtlowicz. On conjectures of Graffiti. In *Proceedings of the First Japan Conference on Graph Theory and Applications (Hakone, 1986)*, volume 72, pages 113–118, 1988.

[FR87] Peter Frankl and Vojtěch Rödl. Forbidden intersections. *Trans. Amer. Math. Soc.*, 300(1):259–286, 1987.

[HMST24] Zach Hunter, Aleksa Milojević, Benny Sudakov, and István Tomon. Disjoint pairs in set systems and combinatorics of low rank matrices. *arXiv preprint arXiv:2411.13510*, 2024.

[Juk12] Stasys Jukna. *Boolean function complexity: advances and frontiers*, volume 27. Springer, 2012.

[KS05] Peter Keevash and Benny Sudakov. Set systems with restricted cross-intersections and the minimum rank of inclusion matrices. *SIAM J. Discrete Math.*, 18(4):713–727, 2005.

[Li11] Yang D Li. Applications of monotone rank to complexity theory. *arXiv preprint arXiv:1102.2932*, 2011.

[Lin65] Bernt Lindström. On a combinatorial problem in number theory. *Canadian Mathematical Bulletin*, 8(4):477–490, 1965.

[Lov16] Shachar Lovett. Communication is bounded by root of rank. *J. ACM*, 63(1):Art. 1, 9, 2016.

[Lov21] Shachar Lovett. Unpublished note, 2021.

[LS88] László Lovász and Michael Saks. Lattices, Möbius functions and communications complexity. In *Proceedings of 29th Annual Symposium on Foundations of Computer Science*, pages 81–90. IEEE Computer Society, 1988.

[LS93] László Lovász and Michael Saks. Communication complexity and combinatorial lattice theory. *Journal of Computer and System Sciences*, 47(2):322–349, October 1993.

[LS23] Troy Lee and Adi Shraibman. Around the log-rank conjecture. *Israel Journal of Mathematics*, 256(2):441–477, 2023.

[Nuf76] Cyriel van Nuffelen. Research problems: A bound for the chromatic number of a graph. *Amer. Math. Monthly*, 83(4):265–266, 1976.

[NW95] Noam Nisan and Avi Wigderson. On rank vs. communication complexity. *Combinatorica*, 15(4):557–565, 1995.

[RY20] Anup Rao and Amir Yehudayoff. *Communication complexity and applications*. Cambridge University Press, Cambridge, 2020.

[Sga99] Jiří Sgall. Bounds on pairs of families with restricted intersections. *Combinatorica*, 19(4):555–566, 1999.

[SS22] Noah Singer and Madhu Sudan. Point-hyperplane incidence geometry and the log-rank conjecture. *ACM Trans. Comput. Theory*, 14(2):Art. 7, 16, 2022.

[SS25] Madhu Sudan and Noah Singer. Personal communication, 2025.

[ST25] Benny Sudakov and István Tomon. Matrix discrepancy and the log-rank conjecture. *Math. Program.*, 212(1-2):567–579, 2025.

[Wat16] Thomas Watson. Nonnegative rank vs. binary rank. *Chicago Journal of Theoretical Computer Science*, 2016(2), 2016.