Asymptotic Tensor-Rank Sublevels Need Not Be Irreducible

A counterexample in order four and the growth of components

Wenqi Zhu* Qiyuan Chen† Ping Nie‡

November 2026

Abstract

The set of complex tensors of asymptotic rank at most $\sqrt[3]{62}$ in format $(3,3,3,2)$ is reducible in the Zariski topology. This gives a negative answer to the irreducibility question when quantified over all tensor orders at least three; the order-three question remains open. A finite Fourier degeneration and Strassen's asymptotic-spectrum and tight-support theorems give the bound $\widetilde{\mathrm{R}}(P)^3 \le 74 - 3\sqrt[3]{81} < 62$ for the three-mode permutation tensor P. The scalar centroid of P supplies an ambient open neighborhood in which the balanced flattening constraints force the fourth factor to have rank one. A second rank-two tensor then exhibits two proper relatively closed subsets covering the entire sublevel. More generally, for every $n \ge 3$ and every order $k \ge 4$, we exhibit thresholds $n < B_n < n+1$ such that, for $B_n \le a < n+1$, the complex asymptotic-rank sublevel in format $(d_1, \dots, d_k)$ has at least $\binom{t_n}{3} + 1$ irreducible components, where $t_n = \#\{i : d_i \ge n\}$. The finite identities and witness geometry admit exact arithmetic checks and a Lean certificate; the asymptotic and Zariski arguments are given in the text.

1 The main theorem and its context

For a tensor $T \in V_1 \otimes \cdots \otimes V_k$, let $\mathrm{R}(T)$ be its rank with respect to the indicated k factors. We use the grouped tensor power

$$T^{\boxtimes m} \in V_1^{\otimes m} \otimes \cdots \otimes V_k^{\otimes m}$$

and the asymptotic rank

$$\widetilde{\mathrm{R}}(T) = \lim_{m\to\infty} \mathrm{R}(T^{\boxtimes m})^{1/m}, \qquad \widetilde{\mathrm{R}}(0) = 0. \tag{1.1}$$

For a nonzero tensor the limit exists by submultiplicativity of rank. Throughout the paper the ground field is $\mathbb{C}$.

Theorem 1.1. *In the* 54*-dimensional vector space*

$$V = \mathbb{C}^3 \otimes \mathbb{C}^3 \otimes \mathbb{C}^3 \otimes \mathbb{C}^2, \qquad a = \sqrt[3]{62},$$

the asymptotic-rank sublevel

$$Y = \{T \in V : \widetilde{\mathrm{R}}(T) \le a\} \tag{1.2}$$

is reducible in the Zariski topology. In particular, the assertion that all asymptotic-rank sublevels are irreducible, for every algebraically closed field, every tensor order $k \ge 3$, every tuple of positive factor dimensions, and every real threshold $a \ge 0$, is false.

The counterexample has order four. The theorem does not settle the question restricted to order three. Known closedness of asymptotic-rank sublevels identifies Y as a reducible algebraic variety; the proper relatively closed cover constructed below proves reducibility directly.

*Corresponding author. Mathematical Institute, University of Oxford, Andrew Wiles Building, Woodstock Road, Oxford OX2 6GG, United Kingdom. wenqi.zhu@maths.ox.ac.uk

†Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing 100190, China.

‡David R. Cheriton School of Computer Science, University of Waterloo, 200 University Avenue West, Waterloo, ON N2L 3G1, Canada. ping.nie@uwaterloo.ca

1.1 The question and the classical ingredients

Christandl, Hoeberechts, Nieuwboer, Vrana and Zuiddam proved that, for every field, every fixed tensor format of order at least three, and every real threshold, the corresponding asymptotic-rank sublevel is Zariski closed [5, Theorem 1.2]. They then asked whether these sublevels are irreducible over algebraically closed fields [5, Section 5]. The source explicitly distinguishes the familiar irreducible determinantal varieties of matrices from the open question for orders $k \geq 3$. The public catalogue records the latter with its universal quantifiers [7]. The present example answers that all-orders question negatively in order four. It leaves the separately restricted three-mode question open.

The asymptotic-rank estimate uses Strassen's asymptotic spectrum and his entropy formula for tight three-mode supports [8, 9]. We use the formulations in Christandl, Vrana and Zuiddam [4, Theorem 1.1, Proposition 1.6, Definition 4.1, Theorem 4.4 and Corollary 4.5]. Spectral points turn direct sums and grouped products into addition and multiplication. The tight-support theorem supplies a lower bound on an auxiliary tensor's asymptotic subrank, which can then be subtracted from a direct-sum spectral inequality. Both classical theorems are applied to three-mode tensors. Appending the fourth factor in the geometric example uses elementary rank and flattening arguments.

The six-permutation tensor P defined in (1.3) is the variant denoted cw'_2 in the Coppersmith–Winograd literature [1, 6]. Alman and Li developed quantitative asymptotic-rank speedups using degenerations with nontrivial direct sums on both sides [1, Sections 4 and 7.1]. Their Corollary 7.2 in the pinned arXiv version, equivalently Corollary 34 in the conference version, applies to this exact tensor and already establishes $\widetilde{R}(P) < 4$. They also reproduce the classical $(q+2)$-term Laurent degeneration of the small Coppersmith–Winograd tensor, crediting Coppersmith and Winograd.

For the constants needed here, we give a direct finite Fourier exchange and obtain the exact estimate

$$\widetilde{R}(P)^3 \leq 74 - 3\sqrt[3]{81} < 62.$$

The exchange is related to the direct-sum speedup method: for example, Alman and Li's Proposition 7.1 in the arXiv version (Proposition 32 in the conference version) exchanges auxiliary matrix tensors with dimensions 18 and 8 when $q = 2$ and $n = 3$. The identity proved below instead displays the auxiliary paired Coppersmith–Winograd tensors and all Fourier coefficients explicitly. We use the estimate in an ambient open-set argument to produce two proper relatively closed subsets covering the full asymptotic-rank sublevel. The comparisons here are confined to the cited sources; no priority claim is made for the tensor constructions or rank estimates.

1.2 The two witnesses and the structure of the proof

Let e_0, e_1, e_2 be the standard basis of $\mathbb{C}^3$, and let d_0, d_1 be the standard basis of $\mathbb{C}^2$. Define

$$\begin{aligned} P &= \sum_{(i,j,k) \text{ a permutation of } (0,1,2)} e_i \otimes e_j \otimes e_k, \qquad x = P \otimes d_0, \\ y &= e_0 \otimes e_0 \otimes e_0 \otimes d_0 + e_1 \otimes e_1 \otimes e_1 \otimes d_1. \end{aligned} \tag{1.3}$$

Section 3 proves

$$\widetilde{\mathrm{R}}(P)^3 \leq 74 - 3\sqrt[3]{81} < 62, \tag{1.4}$$

so $x \in Y$. The displayed decomposition and a flattening give $\widetilde{\mathrm{R}}(y) = 2$, hence $y \in Y$ as well.

Let $X \subseteq V$ be the closed locus of fourth-mode flattening rank at most one. We construct a polynomial open set $U \subseteq V$ such that $x \in U$ and $Y \cap U \subseteq X$. Its defining condition is the persistence of a one-dimensional kernel of a matrix of appended minors. Since $y \notin X$, the identity

$$Y = (Y \cap X) \ \cup \ \big(Y \cap (V \setminus U)\big) \tag{1.5}$$

is a cover by two proper relatively closed subsets. In particular, the open set is taken in V, without imposing a fixed linear slice on the sublevel problem.

2 The asymptotic spectrum and tight supports

2.1 Restrictions and the full three-mode spectrum

A *restriction* $S \leq T$ means that S is obtained from T by a linear map on each factor. The direct sum $S \oplus T$ uses the direct sum of the respective factor spaces; the product $S \boxtimes T$ groups corresponding factors. These operations make the equivalence classes under mutual restriction of all complex three-mode tensors into a preordered semiring. Put

$$\Delta_r = \sum_{j=1}^{r} e_j \otimes e_j \otimes e_j.$$

A spectral point on this semiring is a map ϕ to the nonnegative real numbers that is additive under $\oplus$, multiplicative under $\boxtimes$, normalized by $\phi(\Delta_r) = r$, and monotone under restriction. The spectrum in the next theorem is the spectrum of *all* complex three-mode tensors.

Theorem 2.1 (Strassen's spectral characterization)**.** *For every complex three-mode tensor T,*

$$\widetilde{\mathrm{R}}(T) = \max_{\phi} \phi(T). \tag{2.1}$$

We use the formulation and tensor-semiring application in [4, Theorem 1.1 and Proposition 1.6]. For later use, the subrank $\mathrm{Q}(T)$ is the largest integer r for which $\Delta_r \leq T$, and

$$\widetilde{\mathrm{Q}}(T) = \lim_{m\to\infty} \mathrm{Q}(T^{\boxtimes m})^{1/m}.$$

The limit follows from supermultiplicativity for nonzero tensors. Every spectral point satisfies

$$\widetilde{\mathrm{Q}}(T) \leq \phi(T) \leq \mathrm{R}(T). \tag{2.2}$$

Indeed, a rank decomposition realizes T as a restriction of $\Delta_{\mathrm{R}(T)}$. Conversely, $\mathrm{Q}(T^{\boxtimes m}) \leq \phi(T)^m$ for every m; taking mth roots gives the lower bound. Invertible changes of basis and nonzero scalar multiplication preserve all spectral values, ranks and asymptotic ranks, since they give restrictions in both directions.

2.2 Polynomial degenerations and interpolation

We will only use explicitly given degenerations. The following elementary argument explains why each spectral point is monotone under them.

Lemma 2.2. *Suppose that for every nonzero complex t, linear maps on the three factors of K produce the tensor*

$$F(t) = T + tT_1 + \cdots + t^d T_d.$$

The maps may have Laurent-polynomial entries. Then $\phi(T) \leq \phi(K)$ for every spectral point ϕ.

Proof. For each positive integer m, the polynomial $F(t)^{\boxtimes m}$ has degree at most dm and constant coefficient $T^{\boxtimes m}$. Choose $dm+1$ distinct nonzero complex values of t. Polynomial interpolation expresses that constant coefficient as a linear combination of the values at those points. Each value is a restriction of $K^{\boxtimes m}$. A linear combination of $dm+1$ such restrictions is a restriction of their direct sum: on each direct-sum block use its three maps, absorbing the scalar coefficient into the first map. Consequently

$$\phi(T)^m \leq (dm+1)\phi(K)^m.$$

Taking mth roots and letting m tend to infinity proves the claim. If $\phi(K) = 0$, the same displayed inequality already implies $\phi(T) = 0$. Laurent entries cause no difficulty because the evaluations avoid zero and the *combined output* is polynomial. □

2.3 The three-mode tight-support theorem

Fix bases indexed by finite sets I_1, I_2, I_3. A support $\Phi \subseteq I_1 \times I_2 \times I_3$ is *tight* if there are injective maps $\alpha_j : I_j \to \mathbb{Z}$ such that

$$\alpha_1(i_1) + \alpha_2(i_2) + \alpha_3(i_3) = 0 \quad \text{for every } (i_1, i_2, i_3) \in \Phi. \tag{2.3}$$

For a probability distribution p on Φ, write p_j for its marginal on I_j, and use natural logarithms in the entropy $H(p_j) = -\sum_i p_j(i) \log p_j(i)$, with $0 \log 0 = 0$.

Theorem 2.3 (Strassen's tight-support theorem)**.** *Let T be a complex three-mode tensor whose support in some bases is tight. For every probability distribution p on that support,*

$$\widetilde{\mathrm{Q}}(T) \geq \min_{1 \leq j \leq 3} \exp H(p_j). \tag{2.4}$$

This is the needed consequence of [4, Definition 4.1, Theorem 4.4 and Corollary 4.5]. The hypothesis concerns the support and injective integer weights; there is no extra equality or positivity hypothesis on its nonzero coefficients. The theorem is a three-mode result. We apply it to a three-mode auxiliary tensor, before adjoining the fourth factor in Section 5.

3 A strict asymptotic-rank bound for the permutation tensor

3.1 Paired Coppersmith–Winograd tensors

For an even positive integer q, let J be a set of size q with a fixed-point-free involution $r \mapsto r'$. In a space with basis z_0 and z_r for $r \in J$, define

$$C_q = \sum_{r \in J} \big(z_0 \otimes z_r \otimes z_{r'} + z_r \otimes z_0 \otimes z_{r'} + z_r \otimes z_{r'} \otimes z_0 \big). \tag{3.1}$$

Different labelings of the pairs give equivalent tensors. The usual small Coppersmith–Winograd tensor is

$$\mathrm{cw}_q = \sum_{j=1}^{q} \big(e_0 \otimes e_j \otimes e_j + e_j \otimes e_0 \otimes e_j + e_j \otimes e_j \otimes e_0 \big). \tag{3.2}$$

Lemma 3.1. *Over $\mathbb{C}$, the tensors C_q and cw_q are equivalent under invertible changes of basis. Every spectral point satisfies $\phi(C_q) \leq q + 2$.*

Proof. Label the paired vectors u_h, v_h, for $1 \leq h \leq q/2$. In each of the three modes apply the map

$$z_0 \longmapsto \tfrac{1}{2} e_0, \qquad u_h \longmapsto e_{2h-1} + \mathrm{i} e_{2h}, \qquad v_h \longmapsto e_{2h-1} - \mathrm{i} e_{2h}.$$

The two-by-two block on a pair has determinant $-2\mathrm{i}$, and the zero coordinate is scaled by $1/2$, so the map is invertible. For each pair,

$$u_h \otimes v_h + v_h \otimes u_h \longmapsto 2(e_{2h-1} \otimes e_{2h-1} + e_{2h} \otimes e_{2h}).$$

The factor two is canceled by the zero-coordinate scaling, in each of its three possible positions. Thus the image is exactly cw_q.

For the bound, set $s = \sum_{j=1}^{q} e_j$. For $t \neq 0$ the expression

$$t^{-2} \sum_{j=1}^{q} (e_0 + t e_j)^{\otimes 3} - t^{-3} (e_0 + t^2 s)^{\otimes 3} + (t^{-3} - q t^{-2}) e_0^{\otimes 3} \tag{3.3}$$

is a sum of at most $q + 2$ rank-one tensors. Expanding it gives

$$\mathrm{cw}_q + t \left(\sum_{j=1}^{q} e_j^{\otimes 3} - e_0 \otimes s \otimes s - s \otimes e_0 \otimes s - s \otimes s \otimes e_0 \right) - t^3 s^{\otimes 3}. \tag{3.4}$$

The coefficients of $e_0^{\otimes 3}$ cancel, and the degree -1 terms with two zero coordinates cancel. The remaining degree-zero terms are exactly (3.2). Applying Lemma 2.2 with source Δ_{q+2} shows $\phi(\mathrm{cw}_q) \leq q + 2$. Basis invariance proves the same bound for C_q. □

3.2 The group tensor and its Fourier decomposition

Identify $\mathbb{F}_2^2$ with the symbols $0, 1, 2, 3$ by their two-bit expansions, so addition is bitwise exclusive-or. Set

$$\begin{aligned} G &= (\mathbb{F}_2^2)^3, \qquad S = \{1,2,3\}^3, \qquad b = (1,1,1), \\ I &= S \cap (S+b), \qquad E = G \setminus \big(S \cup (S+b)\big). \end{aligned} \tag{3.5}$$

Here $I = \{2,3\}^3$. Since $|G| = 64$ and $|S| = 27$, inclusion-exclusion gives

$$|I| = 8, \qquad |E| = 64 - (27 + 27 - 8) = 18. \tag{3.6}$$

Both I and E are invariant under $r \mapsto r + b$. This is a fixed-point-free involution because $b \neq 0$. Also $0 \notin S \cup E$: indeed $b \in S$ implies $0 \in S + b$.

Let f_g, for $g \in G$, denote a basis and define the group tensor

$$K_G = \sum_{u+v+w=0} f_u \otimes f_v \otimes f_w. \tag{3.7}$$

Its Walsh decomposition is

$$K_G = \frac{1}{64} \sum_{\chi \in \widehat{G}} \left(\sum_{g \in G} \chi(g) f_g \right)^{\otimes 3}, \tag{3.8}$$

where $\widehat{G}$ is the set of the 64 characters with values in $\{1, -1\}$. For a coefficient indexed by (u, v, w) the right side is $64^{-1} \sum_\chi \chi(u+v+w)$. This equals one if the sum is zero. Otherwise, choosing a character that is negative on that nonzero group element pairs the terms with their negatives. This proves (3.8) and $\mathrm{R}(K_G) \leq 64$.

3.3 The finite exchange

In each target factor use the disjoint basis vectors

$$\{x_g : g \in S\} \;\sqcup\; \{z_0\} \;\sqcup\; \{z_g : g \in E\}.$$

Apply in all three modes the Laurent map

$$f_g \longmapsto \mathbf{1}_{g \in S} x_g + t^{-2} \mathbf{1}_{g=b} z_0 + t \mathbf{1}_{g \in E} z_g. \tag{3.9}$$

In particular, f_b has both an x_b component and a z_0 component; the formula is a sum of components, not a partition of the source basis. The image of f_0 is zero.

Lemma 3.2. *There are Laurent-polynomial restrictions of $\Delta_{64} \oplus C_8$ whose combined output is a polynomial in t of degree at most three with constant coefficient*

$$P^{\boxtimes 3} \oplus C_{18}. \tag{3.10}$$

Consequently every spectral point satisfies

$$\phi(P)^3 + \phi(C_{18}) \leq 64 + \phi(C_8). \tag{3.11}$$

Proof. First apply (3.9) to K_G. A term using only x components has degree zero. For each of the three coordinates of G, three elements of $\{1,2,3\}$ sum to zero exactly when they are the three distinct nonzero elements of $\mathbb{F}_2^2$. Identify $1, 2, 3$ with the basis labels $0, 1, 2$ of P. The x-only part is then $P^{\boxtimes 3}$, with $6^3 = 216$ nonzero coefficients, all equal to one.

Consider a term with exactly one z_0 component. Its corresponding group index is b, and the other two group indices must sum to b. There are three possibilities for their target types:

1. Two x components have degree -2. Their indices lie in I, paired by $r \mapsto r + b$. Taking all three positions of z_0 gives exactly $t^{-2} C_8(z_0, x_i : i \in I)$.

2. One x component and one z_r, for $r \in E$, would have degree -1, but cannot occur. If the first index is in S, the second must be in $S + b$, which is disjoint from E.

3. Two E components have degree zero. Their indices are the 18 elements of E, paired by the same involution. These terms give exactly C_{18}, with $3 \cdot 18 = 54$ support coefficients.

Two z_0 components would require the third group index to be zero, whose image is zero. Three would require $b+b+b=0$, contrary to $b \neq 0$. Finally, terms with no z_0 and at least one E component have positive degree, at most three. This exhausts every possible term, including the two components of the image of f_b.

Now adjoin a disjoint copy of the paired tensor C_8 to the source. Label its nonzero basis vectors by I with the pairing $i \mapsto i+b$. In each mode send its zero-index basis vector to $-t^{-2}z_0$ and its other basis vectors to x_i. Every term of that tensor has exactly one zero-index basis vector, so its image is precisely $-t^{-2}C_8(z_0, x_i : i \in I)$. This cancels all negative degrees. The resulting output is polynomial of degree at most three, with the constant coefficient in (3.10). Its two summands occupy disjoint coordinates in every factor.

Equation (3.8) realizes K_G as a restriction of Δ_{64}, so the same construction has source $\Delta_{64} \oplus C_8$. Lemma 2.2, together with additivity and multiplicativity of ϕ, proves (3.11). □

3.4 The tight auxiliary tensor and the exact constant

Choose labels u_h, v_h, for $1 \leq h \leq 9$, for the pairs of C_{18}. Assign weight 0 to z_0, weight h to u_h and weight $-h$ to v_h, in each of the three modes. These weights are injective within a mode and sum to zero on every support triple. Thus C_{18} has tight support.

Give each of its 54 support triples probability $1/54$. Every marginal has mass $1/3$ at the zero coordinate and mass $2/54 = 1/27$ at each of its 18 other coordinates. Its entropy therefore satisfies

$$\begin{aligned} \exp H &= \exp\left(-\tfrac{1}{3}\log\tfrac{1}{3} - 18 \cdot \tfrac{1}{27}\log\tfrac{1}{27}\right) \\ &= 3^{7/3} = 3\sqrt[3]{81}. \end{aligned} \tag{3.12}$$

Theorem 2.3 and (2.2) give $\phi(C_{18}) \geq \widetilde{\mathrm{Q}}(C_{18}) \geq 3\sqrt[3]{81}$ for every spectral point. Lemma 3.1 gives $\phi(C_8) \leq 10$. Substituting both inequalities into (3.11) yields the uniform estimate

$$\phi(P)^3 \leq 64 + 10 - 3\sqrt[3]{81} = 74 - 3\sqrt[3]{81}. \tag{3.13}$$

Taking the maximum in Theorem 2.1 proves the first inequality of (1.4).

All comparisons needed for the threshold are exact. Since $81 > 64 = 4^3$, one has $3\sqrt[3]{81} > 12$, so the last expression in (3.13) is strictly less than 62. Moreover

$$27 < 62 < 64 \quad \Longrightarrow \quad 3 < \sqrt[3]{62} < 4. \tag{3.14}$$

No decimal estimate or assumption about $\widetilde{\mathrm{R}}(P)$ is used.

4 Scalar centroids and an ambient open neighborhood

Set $A = B = C = \mathbb{C}^3$ and $D = \mathbb{C}^2$. Every element of V has a unique expression

$$W = T \otimes d_0 + S \otimes d_1, \qquad T, S \in A \otimes B \otimes C. \tag{4.1}$$

The fourth-mode flattening views W as a map $D^* \to A \otimes B \otimes C$. Its rank is the dimension of $\operatorname{span}\{T, S\}$. Hence

$$X = \{W : \operatorname{rank}(W_{ABC|D}) \leq 1\} = \{T \otimes v : T \in A \otimes B \otimes C,\ v \in D\} \tag{4.2}$$

is closed, defined by the two-by-two minors of a 27-by-2 matrix.

Let $Z \subseteq V$ be the common rank-at-most-three locus of the three balanced flattenings

$$AB \mid CD, \qquad AC \mid BD, \qquad BC \mid AD. \tag{4.3}$$

These are 9-by-6 matrices, so their four-by-four minors define Z. Every point of X lies in Z: for example the $AB \mid CD$ flattening of $T \otimes (v_0 d_0 + v_1 d_1)$ is $[v_0 M \ v_1 M]$, where $M = T_{AB|C}$ has three columns.

4.1 An appended-minor matrix

Write $M_C(T)$ for the 9-by-3 flattening with row indices in $A \times B$ and column indices in C. Define $M_B(T)$ and $M_A(T)$ by the other two choices of the singled-out factor. A tensor T is *concise* if all three matrices have rank three.

For each of these three matrices, append one of the three columns of the corresponding matrix of S, and take the determinant on each choice of four rows. These expressions are linear in the 27 entries of S. Collect their coefficients into the matrix

$$L(T) : \mathbb{C}^{27} \longrightarrow \mathbb{C}^{1134}, \qquad 1134 = 3 \cdot 3\binom{9}{4}. \tag{4.4}$$

Every entry of $L(T)$ is a homogeneous cubic polynomial in the coordinates of T. Appending a column of T repeats an existing column, so

$$L(T)T = 0 \quad \text{for every } T. \tag{4.5}$$

Lemma 4.1. *If T is concise, then the tensor W in (4.1) belongs to Z if and only if $L(T)S = 0$.*

Proof. Up to column permutations, each balanced flattening of W is $[M_j(T)\ M_j(S)]$. Since $M_j(T)$ has rank three, that block matrix has rank at most three exactly when every column of $M_j(S)$ belongs to the column space of $M_j(T)$. Appending a single column gives a 9-by-4 matrix. Its four-by-four minors all vanish exactly when its rank is at most three, which, because the original three columns are independent, is precisely the required membership condition. Doing this for all columns and all three flattenings proves both directions. □

4.2 The centroid interpretation and the kernel at P

For a concise tensor T, define its centroid to be the vector space

$$\begin{aligned}\mathrm{Cent}(T) = \{(F_A, F_B, F_C) \in \mathrm{End}(A) \oplus \mathrm{End}(B) \oplus \mathrm{End}(C) : (F_A \otimes 1 \otimes 1)T \\ = (1 \otimes F_B \otimes 1)T = (1 \otimes 1 \otimes F_C)T\}.\end{aligned} \tag{4.6}$$

Taking the common value identifies this vector space with $\ker L(T)$. To see this, concision makes each map $F_j \mapsto F_j \cdot T$ injective. Its image consists precisely of those S for which the columns of $M_j(S)$ belong to the column space of $M_j(T)$: a matrix relation $M_j(S) = M_j(T)H$ corresponds to action by the transpose of H on the singled-out factor. The three simultaneous membership conditions are exactly those defining $\ker L(T)$. Under this identification, the common scalar triples $(\lambda 1_A, \lambda 1_B, \lambda 1_C)$ correspond to λT.

Lemma 4.2. *The tensor P is concise, and*

$$\ker L(P) = \mathbb{C}P, \qquad \operatorname{rank} L(P) = 26. \tag{4.7}$$

Equivalently, $\mathrm{Cent}(P)$ *consists exactly of the common scalar triples.*

Proof. The columns of $M_C(P)$ are the three independent vectors

$$e_1 \otimes e_2 + e_2 \otimes e_1, \quad e_0 \otimes e_2 + e_2 \otimes e_0, \quad e_0 \otimes e_1 + e_1 \otimes e_0.$$

Their span is the space of symmetric three-by-three matrices with zero diagonal. The same description holds for $M_A(P)$ and $M_B(P)$ by permutation symmetry, proving concision.

If $L(P)S = 0$, the column-space characterization in the preceding lemma says, for each fixed k, that the array $(S_{ijk})_{i,j}$ is symmetric with zero diagonal. The other two flattenings give the same condition for the other pairs of indices. Thus the coefficients of S are unchanged by each transposition of the three indices and vanish whenever any two indices coincide. The only possible nonzero coefficients are the six permutations of $(0, 1, 2)$; transpositions act transitively on these six, so all six coefficients agree. Therefore S is a scalar multiple of P. Conversely, (4.5) shows that every such multiple lies in the kernel. Rank-nullity gives $27 - 1 = 26$, and the centroid description follows from the preceding identification. □

4.3 Persistence in the full ambient space

Define $U \subseteq V$ by the following conditions on the first slice in (4.1):

$$U = \{T \otimes d_0 + S \otimes d_1 : \operatorname{rank} M_A(T) = \operatorname{rank} M_B(T) = \operatorname{rank} M_C(T) = 3,\ \operatorname{rank} L(T) \geq 26\}. \tag{4.8}$$

Each condition is polynomial open: maximal flattening rank is the nonvanishing of at least one three-by-three minor, and $\operatorname{rank} L(T) \geq 26$ is the nonvanishing of at least one 26-by-26 minor. The slice projection $W \mapsto T$ is linear, so these conditions indeed define an open subset of all of V. Lemma 4.2 shows that $x = P \otimes d_0 \in U$.

Proposition 4.3. *The open set U satisfies $Z \cap U \subseteq X$.*

Proof. Take $W = T \otimes d_0 + S \otimes d_1 \in Z \cap U$. Concision implies $T \neq 0$. By (4.5), the kernel of $L(T)$ contains this nonzero vector, so its rank is at most 26. The defining inequality of U forces equality. Thus $\ker L(T) = \mathbb{C}T$. Since $W \in Z$, Lemma 4.1 gives $L(T)S = 0$, and therefore $S = \lambda T$ for a scalar λ. It follows that

$$W = T \otimes (d_0 + \lambda d_1) \in X.$$

□

This is the only persistence statement needed: one nonzero minor at P, together with the universal kernel vector T, keeps the kernel one-dimensional on an open neighborhood. No assertion about arbitrary tensors with deficient single-mode ranks is used.

5 The proper relatively closed cover

We first record why all relevant matrix-rank constraints follow from asymptotic tensor rank, and why adding a fixed factor does not change the rank of a tensor.

Lemma 5.1. *For any tensor T of any order, the rank of every matrix flattening of T is at most $\widetilde{\mathrm{R}}(T)$.*

Proof. Let a flattening have rank r. The corresponding flattening of $T^{\boxtimes m}$ is, up to permutations of rows and columns, the mth matrix Kronecker power of that flattening and has rank r^m. Every rank-one tensor becomes a matrix of rank at most one under a flattening, so $r^m \leq \mathrm{R}(T^{\boxtimes m})$. Taking mth roots and the limit proves the claim, including the zero tensor. □

Lemma 5.2. *If v is a fixed nonzero vector in an additional factor, then $\mathrm{R}(T \otimes v) = \mathrm{R}(T)$ and $\widetilde{\mathrm{R}}(T \otimes v) = \widetilde{\mathrm{R}}(T)$, where the two ranks use their respective numbers of factors.*

Proof. A rank decomposition of T gives one of $T \otimes v$ with the same number of terms. Choose a linear functional ℓ with $\ell(v) = 1$. Applying it to the last factor of any decomposition of $T \otimes v$ gives a decomposition of T, with no more terms. These are the two rank inequalities. Apply the same argument to each power, with added vector $v^{\otimes m}$, and take mth roots and limits. □

Proof of Theorem 1.1. For any $W \in Y$, Lemma 5.1 bounds each balanced flattening rank by $\widetilde{\mathrm{R}}(W) \leq a < 4$. Matrix rank is integral, so all three ranks are at most three. Therefore $Y \subseteq Z$.

Equation (1.4) and Lemma 5.2 imply $\widetilde{\mathrm{R}}(x) = \widetilde{\mathrm{R}}(P) < a$, so $x \in Y$. The tensor y in (1.3) has rank at most two by its displayed decomposition. Its fourth-mode flattening has two independent columns, so Lemma 5.1 gives $\widetilde{\mathrm{R}}(y) \geq 2$. Since $\widetilde{\mathrm{R}}(y) \leq \mathrm{R}(y) \leq 2$, equality holds. In particular $y \in Y$, because $2 < a$, and $y \notin X$.

The sets

$$Y_1 = Y \cap X, \qquad Y_2 = Y \cap (V \setminus U)$$

are relatively Zariski closed in Y, since X and $V \setminus U$ are closed in V. They cover Y: if a point of Y lies in U, then it lies in $Z \cap U \subseteq X$ by Proposition 4.3; otherwise it belongs to Y_2. Both pieces are proper. The point y belongs to $Y \setminus Y_1$, and the point x belongs to $Y \setminus Y_2$ because $x \in U$. Thus Y is reducible as a topological subspace of V. Its known Zariski closedness [5, Theorem 1.2] then gives the stated interpretation as a reducible algebraic variety. □

The use of four factors occurs in the pair of slices and their balanced flattenings. The spectral and tight-support estimates are applied exclusively to three-mode tensors. The proper cover does not imply that either displayed piece is irreducible or that there are exactly two irreducible components.

6 Component growth and a boundary under grouping

For complex vector spaces $V_1, \dots, V_k$ and a real $a \geq 0$, write

$$Z_a(V_1, \dots, V_k) = \{T \in V_1 \otimes \cdots \otimes V_k : \widetilde{\mathrm{R}}(T) \leq a\}.$$

These are the full sublevels, which are Zariski closed by [5, Theorem 1.2]. Ranks and asymptotic ranks always refer to the displayed factors. Grouping a collection of factors for a matrix flattening does not change this convention.

6.1 The local rank-one chart in general dimension

Let A, B, C have the same dimension $n \geq 3$, put $W = A \otimes B \otimes C$, and let D have dimension $m \geq 2$. For $T \in W$, set

$$E_{AB}(T) = \mathrm{im}(C^* \longrightarrow A \otimes B),$$

and define $E_{AC}(T)$ and $E_{BC}(T)$ in the same way. Using the natural permutations of tensor factors, define

$$K(T) = (E_{AB}(T) \otimes C) \cap (E_{AC}(T) \otimes B) \cap (A \otimes E_{BC}(T)) \subseteq W. \tag{6.1}$$

Thus $T \in K(T)$, and T is concise precisely when the three pair-image spaces have dimension n.

The appended-minor construction of Section 4 extends without change: to each n^2-by-n flattening of T, append each column of the corresponding flattening of a variable S, and take all $(n+1)$-row determinants. Their coefficients in S form a polynomial matrix $L_n(T)$ with n^3 columns. Repeated columns give $L_n(T)T = 0$ for every T, and on the concise locus one has

$$\ker L_n(T) = K(T).$$

Indeed, a full-rank n-column matrix acquires no additional rank on appending a column exactly when that column lies in its image. Consequently

$$\Omega_n = \{T \in W : T \text{ is concise, } \operatorname{rank} L_n(T) \geq n^3 - 1\} \tag{6.2}$$

is Zariski open, and $K(T) = \mathbb{C}T$ for every $T \in \Omega_n$.

Choose a basis $d_0, \dots, d_{m-1}$ of D and write $F = \sum_{j=0}^{m-1} T_j \otimes d_j$. Define

$$\mathcal{F}_n = \{F : \operatorname{rank} F_{AB|CD},\ \operatorname{rank} F_{AC|BD},\ \operatorname{rank} F_{BC|AD} \leq n\},$$
$$X = \{F : \operatorname{rank} F_{ABC|D} \leq 1\}, \qquad U = \{F : T_0 \in \Omega_n\}.$$

Here the first line requires each of the three ranks to be at most n.

Proposition 6.1. *The full-ambient open set U satisfies*

$$\mathcal{F}_n \cap U = X \cap U. \tag{6.3}$$

If $\Omega_n \neq \varnothing$, then X is an irreducible component of $\mathcal{F}_n$ of dimension $n^3 + m - 1$. In the four-factor rank convention (A, B, C, D), for every $a \geq 0$ one has

$$Z_a(A, B, C, D) \cap X = \{T \otimes v : T \in Z_a(A, B, C),\ v \in D\}. \tag{6.4}$$

For $0 \leq a < n + 1$, its intersection with U is isomorphic to

$$\big(Z_a(A, B, C) \cap \Omega_n\big) \times \mathbb{A}^{m-1}. \tag{6.5}$$

Proof. Suppose $F \in U$. The $AB \mid CD$ flattening concatenates the columns of all the T_j, while the first slice already has an n-dimensional AB-image. This flattening has rank at most n if and only if $E_{AB}(T_j) \subseteq E_{AB}(T_0)$ for every j. The other two balanced flattenings give the corresponding conditions on AC and BC. Therefore $F \in \mathcal{F}_n$ precisely when every T_j belongs to $K(T_0) = \mathbb{C}T_0$, which is precisely $F \in X$. Conversely, every $T \otimes v$ has each balanced rank at most n, so $X \subseteq \mathcal{F}_n$. This proves (6.3).

The locus X is the rank-at-most-one variety of n^3-by-m matrices. It is closed by its two-by-two minors and irreducible as the image of $W \times D$ under $(T, v) \mapsto T \otimes v$. Its dimension is $n^3 + m - 1$: on a chart with a chosen nonzero matrix entry, a rank-one matrix is determined by its corresponding row and column, with their common entry counted once. If an irreducible closed subset C of $\mathcal{F}_n$ contains X, then $C \cap U$ is nonempty when $\Omega_n \neq \varnothing$. It is dense in C, and (6.3) places it in the closed subset X. Hence $C = X$, proving the component assertion.

For $v \neq 0$, Lemma 5.2 gives $\widetilde{\mathrm{R}}(T \otimes v) = \widetilde{\mathrm{R}}(T)$; the zero tensor is covered by taking $T = 0$. This proves (6.4). If $a < n + 1$, Lemma 5.1 and integrality give $Z_a(A, B, C, D) \subseteq \mathcal{F}_n$. The isomorphism in (6.5) is

$$(T, \lambda_1, \ldots, \lambda_{m-1}) \longmapsto T \otimes \left(d_0 + \sum_{j=1}^{m-1} \lambda_j d_j \right).$$

Its inverse reads the first slice $T \neq 0$ and the unique scalar multiples of it in the other slices. Those scalars are regular on each chart on which a selected coordinate of T is nonzero, and these expressions agree on overlaps. This also proves the claimed isomorphism globally on the indicated open subset. □

The component assertion concerns the determinantal locus $\mathcal{F}_n$; it does not assert that the whole of X lies in an asymptotic-rank sublevel. Equation (6.4) retains the actual three-mode sublevel. The exact sublevel identifications above use four factors, with D a single factor. When D abbreviates several original factors below, only the flattening statement (6.3) is used.

6.2 Scalar-kernel cores in every dimension

In this subsection $A = B = C = \mathbb{C}^n$, and $K(T)$ is the intersection of lifted two-mode images defined above. We write $\underline{\mathrm{R}}(T)$ for border rank: the least integer r for which T belongs to the Zariski closure of the rank-at-most-r locus.

For an integer $n \geq 3$ and $s \in \mathbb{C} \setminus \{0\}$, put

$$w = \sum_{i=1}^{n} e_i, \qquad T_{n,s} = \sum_{i=1}^{n} e_i \otimes e_i \otimes e_i + s\, w \otimes w \otimes w. \tag{6.6}$$

Lemma 6.2. *For every $n \geq 3$ and every nonzero complex s, the tensor $T_{n,s}$ is concise and $K(T_{n,s}) = \mathbb{C}T_{n,s}$.*

Proof. Its slices in any mode are $E_{ii} + sJ$, where J is the all-ones matrix. A linear relation among these slices has the form

$$\operatorname{diag}(x_1, \ldots, x_n) + s\Big(\sum_i x_i\Big) J = 0.$$

An off-diagonal entry gives $\sum_i x_i = 0$, and the diagonal entries then give $x_i = 0$ for every i. This proves concision in each mode.

Let $S \in K(T_{n,s})$. Every matrix in each two-mode image is symmetric, so the three lifted-image conditions make S invariant under all permutations of its indices. In the slice with third index k, write q_k for the common off-diagonal entry. Expressing this slice in the basis $E_{ii} + sJ$ gives

$$q_k = s \left(\sum_i S_{iik} - n q_k \right).$$

For distinct j, k, choose i different from both. Symmetry gives $S_{ijk} = S_{ikj}$, hence $q_k = q_j$. Let their common value be q. For $i \neq k$, symmetry also gives $S_{iik} = S_{iki} = q$. Writing $b_k = S_{kkk}$, the displayed relation becomes $q = s(b_k - q)$, so $b_k = (1 + 1/s)q$. Thus the coefficients of S are q off the full diagonal and $(1 + 1/s)q$ on it, which means $S = (q/s)T_{n,s}$. The reverse inclusion follows from $T_{n,s} \in K(T_{n,s})$. □

Lemma 6.3. *For every $n \geq 3$ and every nonzero complex s,*

$$\mathrm{R}(T_{n,s}) = \underline{\mathrm{R}}(T_{n,s}) = n + 1. \tag{6.7}$$

Proof. Equation (6.6) gives the upper bound $n+1$ for both ranks. A concise tensor of rank at most n has rank exactly n, since each single-mode flattening has rank n. In an n-term decomposition, its factors span each of the three n-dimensional modes and hence form bases. It is therefore an invertible local transform of

$$D_n = \sum_{i=1}^{n} e_i \otimes e_i \otimes e_i.$$

For D_n, the lifted-image conditions force all three indices to coincide, so $K(D_n)$ is the n-dimensional span of its diagonal summands. Invertible local maps carry $K(T)$ to the corresponding kernel for the transformed tensor. Consequently every concise tensor of rank at most n has kernel dimension n.

On the concise open subset, the appended-minor construction of Section 4 works with n^2-by-n flattenings: append each column of a variable tensor and impose every $(n+1)$-row determinant. The resulting polynomial matrix $L_n(T)$ has n^3 columns and satisfies $\ker L_n(T) = K(T)$ there. The condition $\dim K(T) \geq n$ is therefore relatively closed, being the condition $\operatorname{rank} L_n(T) \leq n^3 - n$. Rank-at-most-$n$ tensors remain dense in their border-rank closure after restriction to this open subset. Every concise tensor of border rank at most n must consequently satisfy $\dim K(T) \geq n$. Lemma 6.2 excludes this possibility for $T_{n,s}$, whose kernel has dimension one. Thus $\underline{\mathrm{R}}(T_{n,s}) > n$, proving both equalities. □

Proposition 6.4. *For every integer $n \geq 3$, define*

$$B_n = \left((n+1)^n + n^{2n/3} - \left((n+1)^n - n^n\right)^{2/3}\right)^{1/n}, \tag{6.8}$$

where the roots are positive real roots. Then $n < B_n < n+1$, and for every $s \in \mathbb{C} \setminus \{0\}$ the tensor $T_{n,s}$ is concise, with

$$K(T_{n,s}) = \mathbb{C}T_{n,s}, \qquad \mathrm{R}(T_{n,s}) = \underline{\mathrm{R}}(T_{n,s}) = n+1, \qquad \widetilde{\mathrm{R}}(T_{n,s}) \leq B_n.$$

Proof. The concision, kernel and exact-rank assertions are Lemmas 6.2 and 6.3. For the asymptotic bound we use the automatic three-direction form of the speedup theorem of Alman and Li [2, Theorem 6.1]. For a tensor Q of format $N \times N \times N$ and an integer ordinary-rank upper bound $R > 2N$, its automatic clause gives

$$\widetilde{\mathrm{R}}(Q) \leq R + N^{2/3} - (R-N)^{2/3} < R. \tag{6.9}$$

Here the auxiliary matrix-rank bound is N in each of the three directions, which is automatic for N-by-N matrices. We apply this theorem only to three-mode tensors.

Take $Q = T_{n,s}^{\boxtimes n}$, $N = n^n$ and $R = (n+1)^n$. Taking the grouped nth power of the explicit $(n+1)$-term decomposition (6.6) gives $\mathrm{R}(Q) \leq R$. The binomial expansion shows

$$(1+1/n)^n = 2 + \sum_{j=2}^{n} \binom{n}{j} n^{-j} > 2,$$

so $R > 2N$. Grouped powers satisfy $\widetilde{\mathrm{R}}(T^{\boxtimes n}) = \widetilde{\mathrm{R}}(T)^n$, directly from the defining limit (1.1). Substitution in (6.9) yields $\widetilde{\mathrm{R}}(T_{n,s})^n \leq B_n^n$.

For the endpoints put $z = R - N > N > 1$. Since $N^{2/3} < z^{2/3}$, one has $B_n^n < R$. Also

$$B_n^n - N = z - z^{2/3} + N^{2/3} > 0.$$

Taking positive nth roots gives $n < B_n < n+1$. □

Thus these cores are available throughout the interval $B_n \leq a < n+1$: the lower endpoint is included because the asymptotic-rank estimate is non-strict. The upper endpoint is excluded because the component argument uses integral flattening ranks bounded above by $a < n+1$. No exact value of $\widetilde{\mathrm{R}}(T_{n,s})$ is asserted. For dimension three, the permutation tensor P supplies the separate interval $\sqrt[3]{62} \leq a < 4$: its scalar kernel is Lemma 4.2, and its strict cost and threshold bounds are (1.4) and (3.14).

6.3 A lower bound for the number of components

Theorem 6.5. *Let $n \geq 3$ and $k \geq 4$ be integers, let $d_1, \ldots, d_k \geq 2$, and put*

$$t_n = \#\{i : d_i \geq n\}.$$

For every real a satisfying $B_n \leq a < n+1$, with B_n as in (6.8), the variety $Z_a(\mathbb{C}^{d_1}, \ldots, \mathbb{C}^{d_k})$ has at least

$$\binom{t_n}{3} + 1 \tag{6.10}$$

irreducible components. The same bound with t_3 holds for every $\sqrt[3]{62} \leq a < 4$.

Proof. First suppose that every $d_i \leq n$, and call a factor active if its dimension is n. For the first interval, choose any nonzero s and put $Q = T_{n,s}$ from Proposition 6.4. For the second interval put $n = 3$ and $Q = P$. Lemma 4.2 and (1.4) supply its scalar kernel and cost bound. Thus in both cases Q is concise, $K(Q) = \mathbb{C}Q$, and

$$\widetilde{\mathrm{R}}(Q) \leq a < n+1.$$

In particular $a \geq n \geq 3$ by Lemma 5.1.

For each triple I of active factors, let x_I be Q on those factors and a fixed basis vector e_0 on every other factor. Repeated application of Lemma 5.2 gives $\widetilde{\mathrm{R}}(x_I) = \widetilde{\mathrm{R}}(Q)$, so $x_I \in Z_a$. For linear-algebra purposes write $D = \bigotimes_{j \notin I} \mathbb{C}^{d_j}$ and choose its product basis. The complementary all-zero coordinate selects the first slice Q. Proposition 6.1 supplies a full-ambient open neighborhood U_I of x_I. Every bipartition rank on Z_a is an integer at most n, so

$$Z_a \cap U_I \subseteq X_I, \qquad X_I = \{F : \operatorname{rank} F_{I|I^c} \leq 1\}. \tag{6.11}$$

This use of D does not merge any modes in the definition of Z_a.

The closed affine variety Z_a has finitely many irreducible components. Choose a component C_I containing x_I. Its nonempty relatively open subset $C_I \cap U_I$ is dense in C_I and lies in the closed locus X_I. It follows that

$$C_I \subseteq X_I. \tag{6.12}$$

No uniqueness of a component through x_I is needed.

There is a useful stronger constraint: every single-mode rank outside I is at most one everywhere on C_I. Fix $i \in I$. The locus on C_I where the ith single-mode rank is n is nonempty, since it contains x_I, and is open, since that mode has dimension n. At any point F of this open set, (6.12) gives a nonzero factorization $F = A_I \otimes B_{I^c}$. For $j \notin I$, the Kronecker-product rank formula gives

$$\operatorname{rank} F_{\{i,j\}|\{i,j\}^c} = r_i(A_I) r_j(B_{I^c}) = n\, r_j(B_{I^c}) \leq n.$$

Here r_j denotes the appropriate single-mode flattening rank, and the last inequality follows from $F \in Z_a$. Thus $r_j(F) = 1$ on that open set. The condition $r_j(F) \leq 1$ is closed, so density extends it to all of C_I. If $I \neq J$, choose $j \in J \setminus I$. Then x_J has jth rank n, whereas every point of C_I has jth rank at most one. Therefore the chosen components C_I are pairwise distinct, whether or not the triples overlap.

Every factor has dimension at least two. The tensor

$$g = \bigotimes_{i=1}^{k} e_0 + \bigotimes_{i=1}^{k} e_1$$

has rank and asymptotic rank two: its displayed decomposition gives the upper bound, and every nontrivial bipartition has matrix rank two. Thus $g \in Z_a$. Since $k \geq 4$, each $I \mid I^c$ is nontrivial, so $g \notin X_I$ for every active triple. A component containing g is different from all the C_I, proving (6.10) in these dimensions. If there are no active triples, the bound is one and follows from the same nonempty sublevel.

For arbitrary $d_i \geq 2$, choose coordinate projections

$$p_i : \mathbb{C}^{d_i} \longrightarrow \mathbb{C}^{\min(d_i, n)}$$

with coordinate inclusions j_i as sections. The tensor-product maps $p = \bigotimes_i p_i$ and $j = \bigotimes_i j_i$ are local maps for the original k factors and satisfy $pj = 1$. Rank is nonincreasing under local maps in every power. Consequently p maps the larger sublevel onto the smaller one, and j preserves asymptotic rank on the embedded copy.

If the larger sublevel has components $C_1, \dots, C_s$, the closures of $p(C_1), \dots, p(C_s)$ in the smaller sublevel are irreducible closed sets covering it. Any irreducible component of the smaller sublevel lies in one of those finitely many closed sets, and maximality makes it equal to that set. The smaller sublevel therefore has at most s components. Its active-factor count is exactly t_n, so its established lower bound gives the same bound in the original dimensions. □

When all $d_i \geq n$, this lower bound is $\binom{k}{3} + 1$. In particular, four factors of dimensions at least three have at least five components for every $\sqrt[3]{62} \leq a < 4$. The counts are lower bounds, not complete lists of components. Their unboundedness varies the tensor format; it makes no claim of infinitely many asymptotic-rank values in a fixed format. The same section-and-retraction argument also propagates reducibility when additional factors are appended using fixed nonzero vectors.

6.4 The grouped three-mode sublevel

Proposition 6.6. *For every real* $3 \leq a < 4$*,*

$$Z_a(\mathbb{C}^9, \mathbb{C}^3, \mathbb{C}^2) = \{F : \operatorname{rank} F_{\mathbb{C}^9|\mathbb{C}^3 \otimes \mathbb{C}^2} \leq 3\}. \tag{6.13}$$

This is the irreducible rank-at-most-three determinantal variety of 9*-by-*6 *matrices.*

Proof. Lemma 5.1 gives the forward inclusion, using integrality and $a < 4$. Conversely, a tensor with this flattening rank at most three has its first-mode support in a subspace of dimension at most three. After compression and padding within that mode, it suffices to consider format $(3, 3, 2)$.

Represent such a tensor by a pair of three-by-three matrices (A, B). The set where A is invertible and $A^{-1}B$ has three distinct eigenvalues is nonempty, as shown by $(I, \operatorname{diag}(0, 1, 2))$. It is Zariski open: on $\det A \neq 0$, distinct eigenvalues are the nonvanishing of the characteristic discriminant. Thus it is dense in the affine space of all pairs. Invertible left and right changes of basis reduce every such pair to

$$(I, \operatorname{diag}(\lambda_1, \lambda_2, \lambda_3)).$$

The corresponding tensor is a sum of the three decomposable tensors $e_i \otimes e_i \otimes (d_0 + \lambda_i d_1)$. Its rank is at most three. Taking Zariski closures proves that every tensor of format $(3, 3, 2)$ has border rank at most three and hence asymptotic rank at most three. The latter implication also follows directly from closedness of Z_3: that closed set contains every rank-at-most-three tensor and therefore its closure. Local padding back to $(9, 3, 2)$ cannot increase the asymptotic rank. Since $a \geq 3$, this proves the reverse inclusion.

Every 9-by-6 matrix of rank at most three is a product of a 9-by-3 matrix and a 3-by-6 matrix. Hence the determinantal variety is the image of an irreducible affine space under matrix multiplication and is irreducible. □

Grouping any two of the three-dimensional modes of $(3, 3, 3, 2)$ gives $(9, 3, 2)$. Throughout $\sqrt[3]{62} \leq a < 4$, the full four-mode sublevel is reducible by Theorem 6.5, whereas the full grouped three-mode sublevel is irreducible by Proposition 6.6. This does not assert surjectivity from the ungrouped sublevel onto the grouped one and does not settle the general three-mode irreducibility question.

6.5 Scope and the remaining three-mode questions

The component lower bounds concern the full sublevels and allow the format to vary. They neither classify all components nor produce infinitely many asymptotic-rank values in a single fixed format. The scalar-kernel argument uses three balanced flattenings and a fourth factor containing simultaneous slices. The grouped example shows explicitly why those constraints do not transfer to three modes. A reducible linear section would also be insufficient to prove reducibility of an ambient sublevel: the irreducible quadric $xy = z^2$ has the reducible section $z = 0$.

The full three-mode asymptotic-rank conjecture asserts $\widetilde{\mathrm{R}}(T) = \max_i r_i(T)$, where $r_i(T)$ are the single-mode ranks, after compression to those mode supports. Under this conjecture, for every $a \geq 0$ the sublevel is the subspace variety defined by

$$r_i(T) \leq \min\{\dim V_i, \lfloor a \rfloor\}, \qquad i = 1, 2, 3.$$

It is irreducible: choose the three support spaces in the product of the corresponding Grassmannians and then a tensor in their tensor product. The resulting vector bundle has irreducible total space and maps onto the subspace variety. When the bound is zero, the variety is $\{0\}$. This conditional implication is already recorded in [5, Section 5]. Thus three-mode reducibility would refute the full conjecture. It would not by itself prove that the matrix-multiplication exponent ω is greater than two, or refute the narrower conjecture for concise tight cubic tensors [3, Conjectures 3 and 8].

Proposition 6.7. *In a fixed complex tensor format, finiteness of the set of asymptotic-rank values is equivalent to a uniform bound on the number of irreducible components of all its sublevels.*

Proof. Finitely many values give finitely many sublevels, each a closed affine variety with finitely many irreducible components. Conversely, suppose all sublevels in an ambient affine space of dimension N have at most C components. Associate to each one its vector of component counts in dimensions $N, N-1, \dots, 0$. For a strict inclusion of two closed subsets these vectors strictly increase lexicographically. Indeed, until the first dimension at which their components differ, all higher-dimensional components coincide. An old component in the next dimension cannot be absorbed by one of those common higher-dimensional components, and containment in a new component of equal dimension forces equality. Thus all old components of that dimension persist, and at least one new component is added. If there were no differing dimension, the two closed subsets would be equal.

There are only $\binom{C+N+1}{N+1}$ nonnegative integer vectors of length $N+1$ with coordinate sum at most C. Consequently every strict chain of sublevels has uniformly bounded finite length. Distinct attained asymptotic-rank values give distinct, totally ordered sublevels, so there are only finitely many such values. □

No uniform component bound in a fixed format is supplied by the component-growth theorem. Determining the three-mode components in the local product chart, finding a three-mode replacement for the balanced-flattening argument, and controlling components over all thresholds of a fixed format remain separate questions.

The arguments in this section are written proofs with the cited closedness and asymptotic-speedup inputs. The finite certificate described in Section 7 supports the original permutation-tensor construction. It does not formalize the general scalar cores, the component-growth theorem, the grouped determinantal description, or the fixed-format proposition above.

7 Exact verification and the finite formal layer

The finite calculations have two complementary forms: exact arithmetic certificates for the explicit coefficients and matrices, and a Lean proof of a proposition collecting their algebraic identities. All calculations use integers, rational numbers, or Gaussian integers. No floating-point tolerance, numerical rank decision, or recognition of a decimal approximation is involved.

7.1 Exact coefficient and matrix certificates

The geometry calculation reconstructs all 1134 rows of $L(P)$ from the appended determinants, using all $\binom{9}{4} = 126$ increasing four-row choices for each relevant column and flattening. Rational elimination gives rank 26, and the identity $L(P)P = 0$ identifies its one-dimensional kernel. The same calculation checks the three rank-three single-mode flattenings of P, the fourth-mode rank two of y, and the three balanced ranks two of y. A supplementary Gaussian change of basis takes P to $2\,\mathrm{cw}_2$; this identity is not needed in the geometric argument of Section 4.

There are two independent coefficient calculations for the exchange. One constructs the group, Walsh coefficients and sparse Laurent maps directly, without reading the coefficient certificate. The other contracts the specified 73-to-46-dimensional maps and evaluates all $46^3 = 97\,336$ ordered output coefficients. The source is written using 88 terms: 64 Walsh summands and 24 paired-C_8 terms. These calculations check the cancellation of every pole and the $270 = 216 + 54$ nonzero constant coefficients of $P^{\boxtimes 3} \oplus C_{18}$. They also check the ten-term cw_8 polynomial identity, the 19 injective integer weights for C_{18}, the zero weight sum on all 54 support triples, and all marginal counts. The counts are 18 at the zero index and 2 at each other index, giving the rational probabilities $1/3$ and $1/27$. The entropy evaluation and its asymptotic consequence are the written calculation (3.12) and Theorem 2.3.

The three exact checkers reconstruct the expected tensors and reject invalid input formats. Their runner binds the proof and certificate files by hashes and compares registered output in ordinary execution, optimized

Python execution, and execution with a second hash seed. One invocation includes nine positive executions and 48 external negative-control executions: each of 16 specified corruptions must be rejected in each of the three modes. The corruptions change, among other quantities, the pole-cancellation sign, Fourier normalization, a target coefficient, a degeneration scalar, the injectivity of tight weights, or a witness rank. The coefficient checker also carries out four corruption checks within each positive execution. Failure is explicit and does not depend on assertions that optimization could remove.

7.2 The precise scope of the Lean certificate

The formal proposition, with theorem name `finiteCertificate`, combines the following statements. It proves the complete enumeration of the minor rows and, for every complex pair T, S, the equality between the appended determinant and the corresponding coordinate of $L(T)S$. It proves $L(T)T = 0$ for every complex T, the rank 26 of $L(P)$, and $L(P)S = 0$ if and only if $S = cP$ for a complex scalar c. The witness flattening ranks and the supplementary Gaussian P-to-$2\,\mathrm{cw}_2$ change of basis are included.

For the exchange, the proposition proves the finite group and set identities, Walsh orthogonality, agreement of the dense and sparse contractions, and the permutation-coordinate identification of $P^{\boxtimes 3}$. It proves the Laurent coefficient identities at every output triple, including vanishing in every degree below zero or above three. The coefficient representation uses an integer normalization, with constant coefficient 64 times the specified target; division by 64 gives the normalization in (3.8). The tight weights, marginal counts and integer comparisons $27 < 62 < 64 < 81$ are included. The ten-term cw_8 identity is proved for every nonzero complex parameter and all 9^3 coordinate triples, with the required constant value at zero of its regular polynomial expression.

The formal development uses Lean `v4.34.0-rc2` with a pinned Mathlib revision. Its final theorem has precisely the logical axioms `propext`, `Classical.choice` and `Quot.sound`. Source and literal-data checks accompany a clean rebuild and include controls for altered inputs, missing proofs, altered build configuration or dependency manifests, and unexpected axiom output.

This certificate does not formalize asymptotic rank or border rank, Strassen's spectral or tight-support theorems, or the implication from entropy to asymptotic subrank. It also does not formalize the general balanced-rank-to-kernel equivalence of Lemma 4.1, the open-neighborhood argument, or the proper relatively closed cover. The paired-C_8 to diagonal cw_8 basis change in Lemma 3.1 is a separate written bridge; the formalized Gaussian change for P does not replace it. None of these omitted statements is a premise of the formal proposition. The asymptotic-rank bound and the reducibility theorem follow from the finite identities together with the stated classical inputs and the written arguments above.

Statements and Declarations

Funding. The research of Wenqi Zhu was supported by the Engineering and Physical Sciences Research Council under grant EP/Y028872/1, *Mathematical Foundations of Intelligence: An "Erlangen Programme" for AI*. Ping Nie gratefully acknowledges the support of AI4ALPHA and NetMind.AI.

Competing interests. The authors have no competing interests to declare that are relevant to the content of this article.

Data availability. The exact verification scripts, certificates and Lean 4 files supporting the finite computations are available from the corresponding author on reasonable request.

Use of generative AI and AI-assisted technologies. During the preparation of this work, the authors used AI tools (Claude, Anthropic; Codex, OpenAI; Qwen APP, Alibaba) to generate candidate constructions and code, perform computations, draft mathematical arguments, and improve the readability of the text. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication. The finite identities and the witness geometry are certified by exact arithmetic and by a Lean 4 finite certificate (Section 7); the asymptotic-rank and Zariski arguments are given in full in the text.

References

[1] Josh Alman and Baitian Li. Asymptotic rank speedup theorems, revisited. In Dana Moshkovitz, editor, *41st Computational Complexity Conference (CCC 2026)*, volume 383 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 36:1–36:42. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2026. doi: 10.4230/LIPIcs.CCC.2026.36. URL https://arxiv.org/abs/2605.21738v1. The cited arXiv numbering refers to version 1, 20 May 2026; Corollary 7.2 there is Corollary 34 in the published version.

[2] Josh Alman and Baitian Li. Asymptotic rank speedup theorems, revisited, 2026. URL https://arxiv.org/abs/2605.21738v2. Version 2, 23 September 2026; the automatic three-direction clause of Theorem 6.1.

[3] Andreas Björklund and Petteri Kaski. The asymptotic rank conjecture and the set cover conjecture are not both true, 2023. URL https://arxiv.org/abs/2310.11926v1. Version 1, 18 October 2023; Conjectures 3 and 8.

[4] Matthias Christandl, Péter Vrana, and Jeroen Zuiddam. Universal points in the asymptotic spectrum of tensors. *Journal of the American Mathematical Society*, 36(1):31–79, 2023. doi: 10.1090/jams/996. URL https://arxiv.org/abs/1709.07851v3. Numbered statements are cited from arXiv:1709.07851v3, 26 May 2018.

[5] Matthias Christandl, Koen Hoeberechts, Harold Nieuwboer, Péter Vrana, and Jeroen Zuiddam. Asymptotic tensor rank is characterized by polynomials, 2026. URL https://arxiv.org/abs/2411.15789v2. Version 2, 7 June 2026; Theorem 1.2 and Section 5. A conference version appeared in STOC 2025, pp. 750–755, doi:10.1145/3717823.3718122.

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

[7] OpenProblemsInNLA. Irreducibility of asymptotic tensor-rank sublevel varieties, 2026. URL https://github.com/ajt60gaibb/OpenProblemsInNLA/blob/d49c11a1b2091547235e4779f7e067ed4b65774d/tensor-computations/TR-23/README.md. Public catalogue entry, pinned commit d49c11a1b2091547235e4779f7e067ed4b65774d.

[8] Volker Strassen. The asymptotic spectrum of tensors. *Journal für die reine und angewandte Mathematik*, 384:102–152, 1988. doi: 10.1515/crll.1988.384.102.

[9] Volker Strassen. Degeneration and complexity of bilinear maps: Some asymptotic spectra. *Journal für die reine und angewandte Mathematik*, 413:127–180, 1991. doi: 10.1515/crll.1991.413.127.