NEURAL COMPUTATION: OPTIMAL PAIRWISE-MAX RELU REPRESENTATIONS

MICHAEL SCHROEDER

ABSTRACT. The maximum function offers a concrete benchmark for the cost of exact neural computation. For five and six inputs in the pairwise-max model of side size two (PM2), we prove the exact directional widths

$$W^*(5,2) = 6, \qquad W^*(6,2) = 33.$$

The features are all pairwise maxima, including the input coordinates; each hinge applies a ReLU to the difference between two sums of two features, with repetitions allowed. We count distinct nonparallel hinge directions and allow a free feature-linear term. Explicit rational identities attain both bounds on all real inputs. Optimality holds for arbitrary coefficient signs and magnitudes, without symmetry assumptions. The lower bounds exploit cancellation: every kink introduced where the maximum is linear must disappear from the sum. This condition yields exact linear equations. Small dependencies among binary kink profiles settle the five-input case; a canonical reduction that never increases support and a complete analysis of local support costs settle the six-input case. Both rational optimality theorems are fully formalized in Lean. A separate rationalization argument gives the same minima for real coefficients in the fixed dictionary.

1. INTRODUCTION

The maximum function gives a concrete test of exact neural computation. The identity $\max(a,b) = b + \rho(a-b)$, where $\rho(t) = \max(0,t)$, implements one comparison with a ReLU hinge. Repeating it gives a comparison tree. Signed sums offer a different mechanism: individual hinges may introduce bends that disappear only when their outputs are added. How many directions does this cancellation require?

We answer that question in a fixed pairwise-max dictionary. The features are $m_{ij}(x) = \max(x_i, x_j)$, including $m_{ii} = x_i$. Each hinge compares two sums of two features, with repetitions allowed. Width counts distinct nonparallel directions after parallel terms have been combined; a feature-linear term is free. This is the *pairwise-max model of side size two*, abbreviated PM2.

Date: 9 November 2026; Version 1.0.

Key words and phrases. ReLU neural networks, Neural network expressivity, Maximum function, Sparse representations, Lower bounds, Lean formalization.

Zenodo DOI: 10.5281/zenodo.24051030.

Theorem 1.1 (Exact directional widths)**.** *In the rational PM2 model of Definition 2.1,*

$$\boxed{W^*(5,2) = 6 \quad \textit{and} \quad W^*(6,2) = 33.}$$

The identities hold on $\mathbb{R}^5$ and $\mathbb{R}^6$, respectively. There is no sign, symmetry, or magnitude restriction on the coefficients. The same minima hold when the free-term and outer hinge coefficients are allowed to be real.

These are the last two nontrivial finite cases in PM2. The dimension obstruction of Rueß et al. [5, Proposition B.1] rules out every maximum with seven or more inputs in this fixed model, even with signed real coefficients and a free feature-linear term. Together with the elementary smaller cases, our results give

n	1	2	3	4	5	6	≥ 7
$W^*(n,2)$	0	0	1	1	6	33	$+\infty$.

Corollary 3.5 proves this consequence and specifies its source. The two new rational optimality statements have complete Lean proofs; the real-coefficient extension is the separate rationalization argument of Lemma 3.2.

The identities also yield ordinary two-hidden-layer ReLU networks of widths $(10, 8)$ for MAX5 and $(16, 35)$ for MAX6: respectively 18 and 51 hidden neurons. Their weights are rational, their biases are zero, and they have no skip connections (Corollaries 4.2 and 4.3). These are construction bounds. The exact directional minima do not establish neuron optimality when network weights are freely chosen.

1.1. **Why six terms become thirty-three.** In the five-input construction, three hinges touch all five coordinates and three more correct their lower-order content. For six inputs, the canonical equations prescribe the sum of the coefficients of a specific star-shaped family at every coordinate. Supplying this *star mass* creates further cancellation demands. Private terms supply the mass; alternating terms meet the resulting local demands.

The optimal six-input construction makes the accounting concrete:

$$\underbrace{25}_{\text{private core}} + \underbrace{5}_{\text{alternating terms}} + \underbrace{3}_{\text{four-set corrections}} = 33.$$

Section 9 explains why this family cannot save a term: width 32 leaves room for only four alternatings, and every such completion fails. The global argument also excludes all other support families. A classification of every optimal identity remains open.

1.2. **The proof in three steps.** First, cancellation of every interior kink turns an identity on real space into exact linear equations on the coordinate-order chambers. Boundary values determine the remaining linear function. Second, a reduction that never increases support replaces the six-input dictionary by 1,530 retained directions and 15 four-set corrections. Third, local

support costs and a complete case analysis exclude every smaller support. Coefficients remain arbitrary signed rationals throughout.

For five inputs, small dependencies among binary kink profiles suffice; exact row separators exclude their possible extensions. For six inputs, a correction graph organizes the cases. Successive lower bounds are proved in order, so each support reduction invokes an already established bound. Figure 3, at the global assembly, shows these dependencies.

1.3. **Position among neural representation results.** Exact depth and exact size are different questions. Arora et al. [1] gave a logarithmic-depth upper bound for continuous piecewise-linear functions. Bakaev et al. [4] constructed MAX5 with two hidden layers; Rueß et al. [5] extended this through MAX12. Brunck et al. [7] now prove that two hidden layers suffice for every continuous piecewise-linear function with finitely many pieces whose affine coefficients lie in a subfield $K \subseteq \mathbb{R}$ containing $\mathbb{Q}$; the network coefficients can be taken in the same field. Zimmermann's explicit caterpillar construction [9] gives maximum networks in arbitrary dimension and has a Lean formalization. Our contribution is exact directional size in the specified finite dictionary, including exhaustive lower bounds.

Restrictions remain decisive. Safran [8] proves quadratic width lower bounds as the input dimension grows for two hidden layers. Braid-fan-compatible networks obey different depth bounds [2, 10]; PM2 permits interior kinks that cancel at the output. Integer-weight and fixed-denominator-ring restrictions also give depth obstructions [11, 12]. Our outer coefficients range over all rationals; integer dictionary directions do not make the resulting network weights integral.

Wang–Basu's side-sum formulas [6] and Myrisiotis's hexagon identity [13] give further explicit PM2 comparisons. Table 1 counts directions after expanding labeled terms, combining parallel coefficients and absorbing orientation changes into the free term. Supplement S12.9 gives the full comparison, its additional reductions and exact verification. Myrisiotis's tailored MAX6 network has widths $(17, 41)$, totaling 58 hidden neurons; our $(16, 35)$ construction uses seven fewer. Neither count is an unrestricted minimum.

Selected explicit PM2 representation	MAX5	MAX6
Rueß et al.: expanded certificate	105	495
Wang–Basu: first ordered-pattern identity	240	810
Myrisiotis: normalized hexagon identity	—	39
Optimal identities in this paper	**6**	**33**

TABLE 1. Directional widths under the same projective convention. These are counts for selected certificates, not neuron minima. Full reconstruction and further reductions appear in Supplement S12.9.

1.4. **Reading routes.** Sections 2–4 give the model, all-real identity test and constructions. Sections 5–8 prove the lower bounds. Readers seeking the mechanism can read the sharp example in Section 9 before the full case analysis. The integral supplement uses the prefix S: S3–S5 establishes transfer and local laws; S6–S10 organizes the exclusions; S13–S15 gives their detailed proofs. S11–S12 states the finite predicates and execution records, and S12.7 maps the written argument to Lean declarations. The formal proof includes transfer and exhaustive assembly; it does not subsume the separately stated real-coefficient extension.

2. The model and its normalization

Let $[n] = \{1, \dots, n\}$ and

$$E_n = \{(i,j) : 1 \le i \le j \le n\}, \qquad m_{ij}(x) = \max(x_i, x_j).$$

The features are ordered lexicographically; $m_{ii} = x_i$. Directions are vectors in feature-coefficient space $\mathbb{Q}^{E_n}$. Their nonparallelness is not a condition on one fixed hyperplane in input space: a hinge argument has a different linear form on different chambers. For a multiset A of features, χ_A is its multiplicity vector. Put

$$\mathcal{D}(n,2) = \{\chi_B - \chi_A \neq 0 : |A| = |B| = 2\}.$$

Definition 2.1 (PM2 directional width)**.** An admissible representation of width W is an identity

$$\mathrm{MAX}_n(x) = \ell \cdot m(x) + \sum_{t=1}^{W} c_t \rho(d_t \cdot m(x)) \quad \text{for every } x \in \mathbb{R}^n, \tag{2.1}$$

where $\ell \in \mathbb{Q}^{E_n}$, $c_t \in \mathbb{Q} \setminus \{0\}$, and $d_t \in \mathcal{D}(n,2)$ are pairwise nonparallel. The least possible W is $W^*(n,2)$, with value $+\infty$ if no representation exists. The second argument records the size of each feature multiset. More generally, $W^*(n,k)$ uses multisets of size k with the same coefficient field, support convention and free term.

Figure 1 shows the allowed representation. The free term is feature-linear, and need not be linear in x. There is no independent bias. Since every direction is balanced, every hinge is unchanged by $x \mapsto x + a\mathbf{1}$. Thus a valid free term has $\sum_e \ell_e = 1$.

Lemma 2.2 (Projective normalization)**.** *The dictionary can be represented by primitive balanced integer vectors of positive mass at most two, modulo sign. Consolidating a finite PM2 sum in these coordinates never increases width.*

Proof. Cancel common features and divide the remaining integer vector by its positive greatest common divisor. Absorb positive scale into the outer coefficient. Reverse orientation using

$$\rho(-g) = \rho(g) - g. \tag{2.2}$$

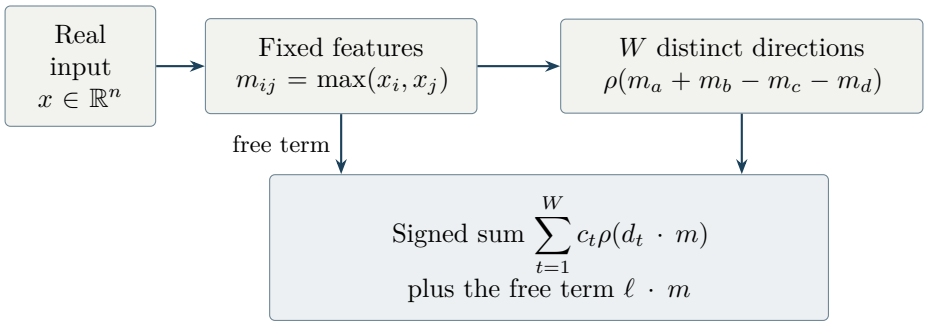


FIGURE 1. The PM2 representation. The labels a, b, c, d denote features and may repeat. Width counts the nonparallel hinge directions in the top right box. A pairwise maximum is itself implemented by ReLU units using Proposition 2.3; the diagram displays the mathematical features, not a neuron count for that implementation.

The last term is feature-linear; the outer hinge coefficient keeps its sign. Combine parallel terms and remove zero coefficients. Conversely, the positive and negative parts of a balanced vector of mass at most two give the two multisets; a mass-one vector is padded by a common feature. □

For $M = |E_n| = \binom{n+1}{2}$, the primitive types are

$$e_a - e_b, \qquad 2e_a - e_b - e_c, \qquad e_a + e_b - e_c - e_d,$$

with the displayed labels distinct. Hence the dictionary size is

$$\binom{M}{2} + M\binom{M-1}{2} + 3\binom{M}{4}. \tag{2.3}$$

It equals 5,565 at $n = 5$ and 22,155 at $n = 6$. These are counts of all labeled projective directions, including loops.

For size-two multisets A, B, the conversion is exact:

$$\max\left(\sum_{e\in A} m_e, \sum_{e\in B} m_e\right) = \sum_{e\in A} m_e + \rho((\chi_B - \chi_A)\cdot m). \tag{2.4}$$

Proposition 2.3 (Conversion to an ordinary network)**.** *A PM2 representation of width W yields a ReLU network with two hidden layers of widths $\binom{n}{2} + 2n$ and at most $W + 2$.*

Proof. The first layer computes $\rho(x_i - x_j)$ for $i < j$ and $\rho(x_i), \rho(-x_i)$ for each i. Then $x_i = \rho(x_i) - \rho(-x_i)$ and $m_{ij} = x_j + \rho(x_i - x_j)$. Each prescribed hinge is therefore one second-layer unit. Two more units realize the free term by $z = \rho(z) - \rho(-z)$ with $z = \ell \cdot m$. □

3. An exact test on the whole real space

On the braid chamber $x_{\pi(1)} < \cdots < x_{\pi(n)}$, use the common translation coordinate $a = x_{\pi(1)}$ and positive gaps $t_k = x_{\pi(k+1)} - x_{\pi(k)}$. A balanced direction becomes

$$d \cdot m(x) = \sum_{k=1}^{n-1} g_{\pi,k}(d) t_k, \qquad g_{\pi,k}(d) = d \cdot m\big(\mathbf{1}_{\{\pi(k+1),\dots,\pi(n)\}}\big). \tag{3.1}$$

The gap normal crosses the positive orthant precisely when it has entries of both signs. Only such a crossing can create an interior kink.

Proposition 3.1 (Chamber criterion)**.** *Let $F = \ell \cdot m + \sum_t c_t \rho(d_t \cdot m)$ with balanced integer directions. In every chamber, group crossing gap normals by primitive normal modulo sign. If $g_\pi(d_t) = s_t h$ in a group, require*

$$\sum_t c_t |s_t| = 0. \tag{3.2}$$

Then $F = \mathrm{MAX}_n$ on $\mathbb{R}^n$ if and only if every such requirement holds, $\sum_e \ell_e = 1$, and $F(\mathbf{1}_T) = 1$ for every nonempty $T \subseteq [n]$.

Proof. At a point on one interior hyperplane away from the others, the gradient jump of its group is the scalar in (3.2) times a fixed normal. Every jump must vanish for a chamber-linear function. Conversely, vanishing jumps make the gradients on adjacent cells of the finite hyperplane arrangement agree. Thus F is linear throughout the chamber and, by continuity, on its closure.

The closed chamber is generated by the common-translation line and its Boolean suffix rays. The translation slope and the stated ray values therefore determine its linear function, giving exactly MAX_n. All chambers together cover the real space, including their boundaries. □

The criterion separates shape from values. Kink cancellation makes the function linear on each chamber. The translation slope and Boolean values then identify that linear function. Checking values without first checking shape can miss an interior bend, as the following example shows.

Why Boolean values alone cannot certify an identity. For $n \geq 3$, consider the admissible signed hinge sum

$$h(x) = \rho(x_1 + x_2 - 2x_3) - \rho(x_1 - x_3) - \rho(x_2 - x_3).$$

It vanishes at every Boolean input: when $x_3 = 0$ all three arguments are nonnegative, and when $x_3 = 1$ they are all nonpositive. Yet $h(1, -1, 0, \dots, 0) = -1$. Adding this expression to an identity preserves every Boolean value and the translation slope but changes the function. The chamber kink equations detect precisely the missing condition.

Lemma 3.2 (Rationalization on a fixed support)**.** *For the fixed dictionary $\mathcal{D}(n, 2)$, allowing real free-term and outer hinge coefficients leaves the rational minimum $W^*(n, 2)$ unchanged.*

Proof. Fix the directions of a real-coefficient identity. All their gap normals and Boolean values are integers. Proposition 3.1 therefore expresses the identity as a finite system $Az = b$ with rational data, linear in the free-term and outer coefficients. Row reduction over $\mathbb{Q}$ shows that a rational system consistent over $\mathbb{R}$ has a rational solution: choose rational values for its free variables and back-substitute. Some hinge coefficients may become zero; deleting those terms cannot increase support. Thus the rational minimum is at most the real minimum. The reverse inequality follows by inclusion. The directions themselves remain in the specified integer dictionary throughout. □

We may consequently prove the lower bounds over $\mathbb{Q}$. This reduction is specific to fixed directions; it does not rationalize freely variable first-layer weights in an unrestricted network.

Lemma 3.3 (Binary kink columns)**.** *For a primitive PM2 direction, each nonzero interior kink has primitive integer strength one. Its full kink profile is therefore binary.*

Proof. A chamber coordinate normal has positive and negative masses at most two. If its content is at least two, it is $2(e_i - e_j)$ up to sign. That form has fixed sign on a braid chamber, so it cannot produce an interior kink. Successive partial sums and differences preserve content when passing between balanced coordinate normals and gap normals. □

For a continuous positively homogeneous function define, for nonempty S,

$$\alpha_S(f) = -\sum_{U \subseteq S} (-1)^{|S|-|U|} f(\mathbf{1}_{[n]\setminus U}). \tag{3.3}$$

Call f *braid-linear* if it is linear on each braid chamber.

The following standard reconstruction is Boolean Möbius inversion combined with the chamberwise linear extension of a set function (the Lovász extension; see [3]). We include the argument to fix our complement and sign conventions.

Lemma 3.4 (Reconstruction by maxima)**.** *A braid-linear function has the unique expansion*

$$f = \sum_{\varnothing \neq S \subseteq [n]} \alpha_S(f)\, \mathrm{MAX}_S, \qquad \mathrm{MAX}_S(x) = \max_{i \in S} x_i. \tag{3.4}$$

If f does not depend on x_v, all coefficients with $v \in S$ vanish. Consequently, a braid-linear function with no coefficients of degree at least three is feature-linear.

Proof. Boolean Möbius inversion gives the unique coefficients because $\mathrm{MAX}_S(\mathbf{1}_{[n]\setminus U}) = 1 - \mathbf{1}_{S \subseteq U}$. The resulting sum has the same Boolean values as f. These determine each chamber-linear function on its translation line

and suffix rays, as in Proposition 3.1. Independence of x_v pairs opposite terms in (3.3), proving the vanishing assertion. □

3.1. **The arity ceiling of the fixed model.**

Corollary 3.5 (Complete fixed-PM2 width sequence)**.** *For rational free and outer coefficients,*

$$W^*(n,2) = \begin{cases} 0, & n = 1, 2, \\ 1, & n = 3, 4, \\ 6, & n = 5, \\ 33, & n = 6, \\ +\infty, & n \geq 7. \end{cases}$$

The same sequence holds for real free and outer coefficients in this dictionary. The impossibility for $n \geq 7$ is a consequence of the signed support-function obstruction in [5, Proposition B.1]*.*

Proof. The first two maxima are features. For $n = 3, 4$, respectively,

$$\mathrm{MAX}_3 = m_{12} + \rho(x_3 - m_{12}), \qquad \mathrm{MAX}_4 = m_{12} + \rho(m_{34} - m_{12}).$$

These mass-one directions belong to PM2 by common padding. Width zero is impossible: a feature-linear function has zero full-set Boolean Möbius coefficient for $n \geq 3$, whereas MAX_n has coefficient one in Lemma 3.4. Theorem 1.1 supplies $n = 5, 6$.

For a polytope $P \subset \mathbb{R}^n$, write $h_P(x) = \max_{p \in P} p \cdot x$. A feature m_{ij} is the support function of $[e_i, e_j]$. For a size-two multiset A, put $Z_A = \sum_{(i,j) \in A} [e_i, e_j]$, a Minkowski sum of dimension at most two. Every allowed hinge satisfies

$$\rho(h_{Z_B} - h_{Z_A}) = h_{\mathrm{conv}(Z_A \cup Z_B)} - h_{Z_A}.$$

The convex hull has dimension at most $2 + 2 + 1 = 5$. Thus any PM2 representation, including its free feature-linear term, expresses MAX_n as a signed real sum of support functions of polytopes of dimension at most five. But MAX_n supports the $(n-1)$-simplex $\mathrm{conv}\{e_1, \ldots, e_n\}$. The cited obstruction forbids expressing a simplex support function as such a sum of lower-dimensional polytopes. This proves impossibility for $n \geq 7$, without symmetry or sign assumptions. □

This corollary uses a cited geometric theorem in addition to the two Lean-verified rational optima. No formal verification of the corollary as a whole is asserted.

4. The two constructions

The upper bounds are explicit. Tables 2 and 3 give every hinge argument and its rational coefficient; the free terms appear in the identities below. The tables are generated directly from the rational data used by the exact checker, so the displayed formulas and the verified constructions agree term for term.

For five inputs,

$$\mathrm{MAX}_5 = -\tfrac{1}{2}m_{15} + \tfrac{1}{2}m_{35} + x_5 + \sum_{t=1}^{6} c_t\rho(g_t), \tag{4.1}$$

with the six terms in Table 2.

t	c_t	g_t
1	$\frac{1}{2}$	$m_{34} - x_5$
2	$\frac{1}{2}$	$m_{15} - m_{25} - m_{35} + m_{45}$
3	$-\frac{1}{2}$	$m_{12} - m_{15} - m_{34} + m_{35}$
4	$-\frac{1}{2}$	$m_{12} - m_{25} - m_{34} + m_{45}$
5	1	$m_{12} - m_{34}$
6	$\frac{1}{2}$	$m_{12} - x_5$

TABLE 2. The six terms for MAX5.

For six inputs,

$$\mathrm{MAX}_6 = -\tfrac{1}{3}m_{14} + \tfrac{4}{3}m_{56} + \sum_{t=1}^{33} c_t\rho(g_t). \tag{4.2}$$

Table 3 gives all thirty-three terms. In each identity the directions are primitive, have positive mass at most two, and are pairwise nonparallel. The six-input identity has three mass-one directions and thirty mass-two directions.

Proposition 4.1 (Upper bounds)**.** *Equations* (4.1) *and* (4.2) *hold for all real inputs and are admissible PM2 representations. Thus* $W^*(5,2) \le 6$ *and* $W^*(6,2) \le 33$.

Exact finite verification. The supplied independent checker verifies dictionary membership and projective distinctness, then applies Proposition 3.1 using integer gap normals and rational arithmetic. It checks all 120 and 720 chambers, respectively, and all 32 and 64 Boolean points. The six-input check groups 1,472 interior normals; all group sums vanish. The translation slope is one in both cases. These checks were executed for this edition; the input hashes and output are retained. Both upper bounds are also included in the formal optimality theorems. □

t	c_t	g_t	t	c_t	g_t
1	$\frac{2}{3}$	$m_{34} - m_{56}$	18	$-\frac{1}{3}$	$m_{12} - m_{34} - m_{36} + m_{56}$
2	$-\frac{1}{6}$	$m_{15} + m_{25} - m_{34} - m_{36}$	19	$-\frac{1}{3}$	$m_{12} - m_{34} - m_{46} + m_{56}$
3	$-\frac{1}{6}$	$m_{15} + m_{25} - m_{34} - m_{46}$	20	$\frac{2}{3}$	$m_{12} - m_{34}$
4	$\frac{1}{3}$	$m_{15} + m_{25} - m_{34} - m_{56}$	21	$\frac{1}{3}$	$m_{12} - m_{36} - m_{46} + m_{56}$
5	$\frac{1}{6}$	$m_{15} + m_{25} - m_{36} - m_{46}$	22	$\frac{2}{3}$	$m_{12} - m_{56}$
6	$\frac{1}{3}$	$m_{14} - m_{24} + m_{34} - m_{56}$	23	$\frac{1}{6}$	$m_{12} + m_{25} - m_{34} - m_{36}$
7	$-\frac{1}{3}$	$m_{14} + m_{23} - m_{24} - m_{56}$	24	$\frac{1}{6}$	$m_{12} + m_{25} - m_{34} - m_{46}$
8	$\frac{1}{3}$	$m_{14} + m_{23} - m_{34} - m_{56}$	25	$-\frac{1}{3}$	$m_{12} + m_{25} - m_{34} - m_{56}$
9	$-\frac{1}{3}$	$m_{13} - m_{14} - m_{23} + m_{56}$	26	$-\frac{1}{6}$	$m_{12} + m_{25} - m_{36} - m_{46}$
10	$\frac{1}{3}$	$m_{13} - m_{23} + m_{34} - m_{56}$	27	$-\frac{1}{3}$	$m_{12} + m_{24} - m_{34} - m_{56}$
11	$-\frac{1}{3}$	$m_{12} - m_{13} - m_{34} + m_{56}$	28	$\frac{1}{6}$	$m_{12} + m_{15} - m_{34} - m_{36}$
12	$\frac{1}{3}$	$m_{12} - m_{13} + m_{14} - m_{56}$	29	$\frac{1}{6}$	$m_{12} + m_{15} - m_{34} - m_{46}$
13	$\frac{1}{3}$	$m_{12} - m_{14} - m_{23} + m_{56}$	30	$-\frac{1}{3}$	$m_{12} + m_{15} - m_{34} - m_{56}$
14	$-\frac{1}{3}$	$m_{12} - m_{14} - m_{34} + m_{56}$	31	$-\frac{1}{6}$	$m_{12} + m_{15} - m_{36} - m_{46}$
15	$-\frac{2}{3}$	$m_{12} - m_{23} + m_{34} - m_{56}$	32	$-\frac{1}{3}$	$m_{12} + m_{14} - m_{34} - m_{56}$
16	$\frac{1}{3}$	$m_{12} - m_{23} + m_{24} - m_{56}$	33	$\frac{1}{3}$	$2m_{12} - m_{34} - m_{56}$
17	$\frac{1}{3}$	$m_{12} - 2m_{34} + m_{56}$			

TABLE 3. The complete MAX6 construction. The free term is given in (4.2).

Corollary 4.2 (A MAX5 network with eighteen hidden neurons)**.** *There is an ordinary ReLU network with architecture* $5 \to 10 \to 8 \to 1$ *that computes* MAX_5 *on all of* $\mathbb{R}^5$*. Its weights are rational, its biases are zero, and it has no skip connections.*

Proof. Write $r_{ab} = \rho(x_a - x_b)$. The first layer computes

$$r_{12}, r_{15}, r_{25}, r_{34}, r_{35}, r_{45}, r_{52}, r_{54}, \rho(x_5), \rho(-x_5).$$

Linear readouts recover $x_5 = \rho(x_5) - \rho(-x_5)$, $x_2 = x_5 + r_{25} - r_{52}$ and $x_4 = x_5 + r_{45} - r_{54}$. Hence they recover $m_{12} = x_2 + r_{12}$, $m_{34} = x_4 + r_{34}$, and $m_{i5} = x_5 + r_{i5}$ for $i = 1, 2, 3, 4$. These are all features used in (4.1) and Table 2. Six second-layer units compute the six prescribed hinges; two compute $\rho(L), \rho(-L)$ for $L = -m_{15}/2 + m_{35}/2 + x_5$. The output uses the six displayed coefficients and $L = \rho(L) - \rho(-L)$. Supplement S12.1 specifies the matrices and their independent exact check. □

Corollary 4.3 (A MAX6 network with fifty-one hidden neurons)**.** *There is an ordinary ReLU network with architecture* $6 \to 16 \to 35 \to 1$ *that computes* MAX_6 *on all of* $\mathbb{R}^6$*. It has rational weights, zero biases, and no skip connections.*

Proof. Write $r_{ab} = \rho(x_a - x_b)$. The eleven pairwise maxima appearing in (4.2) and Table 3 are recovered using the orientations

$$\mathcal{E} = \{(1,2),(1,3),(1,4),(1,5),(2,3),(2,4),\\ (2,5),(3,4),(6,3),(6,4),(6,5)\}.$$

The first layer computes r_{ab} for $(a,b) \in \mathcal{E}$, together with r_{32}, r_{42}, r_{52} and $q_+ = \rho(x_2), q_- = \rho(-x_2)$. These are sixteen ReLUs. Their outputs linearly recover

$$x_2 = q_+ - q_-, \qquad x_j = x_2 - r_{2j} + r_{j2} \quad (j = 3,4,5).$$

Every second endpoint in $\mathcal{E}$ belongs to $\{2,3,4,5\}$, so $m_{ab} = r_{ab} + x_b$ recovers every required feature. Hence each of the thirty-three arguments g_t is a linear form in the first-layer outputs. The second layer computes their ReLUs and the two additional units $\rho(L), \rho(-L)$, where $L = -m_{14}/3 + 4m_{56}/3$. The output is $\sum_t c_t \rho(g_t) + \rho(L) - \rho(-L)$, which equals MAX_6 by (4.2). Every operation uses rational coefficients. □

The ordered gates, feature readouts, complete weight matrices, and an independent exact check of the exported network are given in Supplement S12.1.

5. Why five directions cannot represent MAX5

For each normalized direction d, let $k(d)$ be its binary kink column and $b(d) = (\alpha_S(\rho(d \cdot m)))_{|S|\geq 3}$. Any identity must satisfy

$$\sum_d c_d k(d) = 0, \qquad \sum_d c_d b(d) = e_{[5]}. \tag{5.1}$$

The free term disappears from these equations. The checked quotient maps the 5,565 dictionary directions to 3,245 stored columns with a pool of 2,626 kink profiles, including the zero profile. The map preserves both kink and Boolean data and cannot increase support. Keeping both kinds of data is necessary: equality of kink profiles alone does not identify the represented functions.

5.1. **A small solution must contain a small circuit.** A circuit is a nonzero linear dependence with minimal support. Every kernel vector is a sum of circuits supported within its support: subtract a multiple of one circuit to kill a coordinate and argue by induction. If a solution of (5.1) had at most five nonzero entries, some circuit in this decomposition would have nonzero top Boolean content. Call it an anchor. For this complete, normalized five-input PM2 dictionary, the finite profile lemma in Supplement S2 proves that a zero kink profile has zero top content and that equal profiles have equal top content. Its proof computes the top content from the histogram of a profile over the 120 chambers; these are dictionary-specific facts. Consequently a one-column circuit, or a two-column difference of equal binary profiles, has zero top content. An anchor therefore has size three, four,

or five. Its coefficients need not equal the original solution coefficients; only containment of its support is used.

Distinct nonzero binary columns admit one three-column circuit pattern, $(1, 1, -1)$, and four four-column patterns, up to permutation and common scale:

$$(-1,1,1,1), \quad (-2,1,1,1), \quad (-1,-1,1,1), \quad (-2,-1,1,1). \tag{5.2}$$

For a five-column circuit, cofactors of a four-row basis give integer coefficients of absolute value at most three. Indeed, a 4×4 binary determinant is at most three in absolute value: adjoining a border of ones and converting to a 5×5 sign matrix gives $16|\det A| \le 5^{5/2} < 64$. A finite binary-row rank test leaves seventeen patterns. Supplement S2 gives the complete list and the small-circuit argument.

5.2. **Exact exclusion of every extension.** Stack the two matrices in (5.1) as A and write $y = (0, e_{[5]})$. A proposed support J is impossible if a rational row vector u satisfies

$$u^T A_J = 0, \qquad u^T y \neq 0. \tag{5.3}$$

Such a vector excludes all rational coefficients on that support at once. The finite proof covers every anchor and every extension up to five columns: no extension for a five-anchor, one for a four-anchor, and two for a three-anchor. Zero-profile extensions and coincident catalogue indices are included. Coordinate permutations act on the whole support, including the extensions.

Lemma 5.1 (Finite MAX5 exclusion)**.** *For every support of at most five columns in the checked quotient, $A_J c = y$ is inconsistent over $\mathbb{Q}$.*

Computer-assisted proof. The anchor argument reduces the claim to the three exhaustive reference families just described. Their coverage and exact separators, including quotient-projection versions of (5.3), are composed in the formal declaration `stored_solution_support_gt_five`. Its coverage premise is discharged by `canonical_saved_reference_coverage`. The final roots are `PM2.WStar_eq_six` and `PM2.width_six_optimal`. Supplement S2 specifies the finite families; Supplement S12 describes the recorded kernel check and its precise limits. □

An identity of width at most five would transfer to a solution forbidden by Lemma 5.1. Together with Proposition 4.1, this proves $W^*(5, 2) = 6$ over $\mathbb{Q}$. Lemma 3.2 extends the minimum to real coefficients, completing the five-input assertion of Theorem 1.1.

6. From the full dictionary to canonical form

The six-input proof begins by replacing the full dictionary with a smaller one. The essential requirement is a comparison of supports: every original identity must yield a reduced identity with at most as many directions. We

establish that transfer here, before deriving lower bounds in the reduced system. A hinge sum is *closed* when all its interior kinks cancel.

6.1. **Separation by the coordinates touched.** For a normalized direction d, let $V(d)$ be the union of the vertices in its nonzero features, after common multiset entries have been canceled. This is its raw touched set. A direction is covering if $V(d) = [6]$.

Lemma 6.1 (Touched-set separation)**.** *In a closed PM2 sum, the terms with any fixed raw touched set form a closed sum themselves.*

Proof. Fix v and put $x_v = T > \max_{i \neq v} x_i$. A touching argument is $s_v(d)T + g_d(y)$, where $s_v(d)$ is the sum of coefficients on incident features, counting a loop once. For $s_v < 0$ the hinge eventually vanishes; for $s_v > 0$ its intercept is feature-linear. If $s_v = 0$, its positive and negative incident masses agree and are at least one. The remaining argument has mass at most one. Its hinge is zero or has the form

$$\rho(m_A - m_B) = \mathrm{MAX}_{A \cup B} - m_B, \tag{6.1}$$

involving at most four remaining vertices.

The intercept of the whole closed sum is continuous and braid-linear in y. Locally a common sufficiently large T works for all its finitely many terms. Subtracting the touching intercepts shows that the sum avoiding v is closed. Repeat for a set of avoided coordinates and use inclusion-exclusion to isolate an exact touched set. □

This is where side size two matters. With side size three, the remaining argument can still have mass two, and the intercept need not be braid-linear. For a covering closed sum on six inputs, the same argument shows that its five-set coefficients vanish. It follows that each raw-touch-five block in a MAX6 identity has zero top coefficient on its five-set: only that block and the covering block could contribute to it.

6.2. **Compress the covering equations.** There are 1,980 covering directions. In their full kink matrix a row with one surviving entry kills that component; a row with two entries fixes their coefficient ratio. Iterating these operations on the *global* matrix removes 180 directions, groups 270 into 90 disjoint triples, and leaves 1,530 singleton directions. Branch-specific sparsity is not used in this dictionary-wide compression.

A *star* is a retained direction represented by $m_{ua} + m_{ub} - m_{uv} - m_{cd}$, with six distinct coordinates. Its center is u; its star mass is its stored coefficient. The mass at a center is the sum over all such stars centered there. Glues and the other retained types are defined in Section 7 and Supplement S4.

The exact trace and Boolean-content checks establish the following facts (Supplement S3). Each surviving component has a fixed rational lifting vector with no zero entry, so its support cost is exactly its number of raw members. Each triple is closed, has zero total star coefficient, and contributes to at most one four-set. For a closed retained sum H, its total star

coefficient equals $2\alpha_{[6]}(H)$, and its three-set content is zero. When the total star coefficient is two, every center has star mass $1/3$, and

$$b_{uv} = -\alpha_{[6]\setminus\{u,v\}}(H), \qquad \sum_{v\neq u} b_{uv} = \tfrac{2}{3}. \tag{6.2}$$

The other required finite fact is a hosting bound. A closed raw-touch-five block of h terms with zero top content has at most h nonzero four-set coefficients. Here is why only small relations require checking. Equal kink profiles have equal degree-four and degree-five content and may be merged. There is no zero profile in this family. For three distinct profiles the only dependence is the three-column binary pattern; for four, use (5.2). Exact content checks give at most one and three four-sets, respectively. With five or more profiles the bound is automatic because a five-set has only five four-subsets. The enumeration uses exact profile comparisons and covers arbitrary common rational scale.

Theorem 6.2 (Canonical support transfer)**.** *Every width-W rational PM2 identity for MAX6 has a canonical identity of width at most W, consisting of retained covering terms and one chosen mass-one correction direction for each nonzero four-set coefficient. Equivalently,*

$$W^*(6,2) = \min_H \left(|\operatorname{supp} H| + \#\{S : |S| = 4,\ \alpha_S(H) \neq 0\}\right), \tag{6.3}$$

where H ranges over rational closed retained sums with total star coefficient two.

Proof. Normalize an arbitrary identity. Its covering block is closed by Lemma 6.1. Write its compressed support cost as $r + 3q_0 + N_0 \leq W$, where r is retained support, q_0 counts nonzero triples, and N_0 counts proper touched-set terms. Delete the triples to obtain H. Proper blocks cannot contribute to degree six. The content identities therefore give top coefficient one, star sum two, and zero degree-three and degree-five coefficients for H.

Its degree-four coefficients were canceled in the original identity by the deleted triples and proper blocks. A triple supplies at most one four-set; a block on at most four variables supplies at most one; a block on five variables supplies at most its number of terms. Thus, writing $E(H)$ for the number of nonzero four-set coefficients,

$$E(H) \leq q_0 + N_0, \qquad r + E(H) \leq W - 2q_0 \leq W. \tag{6.4}$$

To construct the new identity, choose a split $S = A_S \sqcup B_S$ into pairs for each four-set and put $b_S = -\alpha_S(H)$. By Lemma 3.4, the remaining discrepancy after adding $\sum b_S \operatorname{MAX}_S$ is feature-linear. Replace each MAX_S using (6.1) and absorb $b_S m_{B_S}$ into that free term. This gives an actual rational identity with exactly $r + E(H)$ directions. The same construction works for every H in (6.3), proving the reverse implication as well. □

Proper blocks are removed jointly. The argument requires no separate low-cost replacement for the degree-three content of each block. Equation (6.4)

supplies the decisive support comparison: a lower bound for canonical identities is a lower bound in the original dictionary.

7. The local geometry of a canonical identity

Multiply the canonical hinge and correction coefficients by three. This preserves every support count and makes the star mass at each center one. The complementary-edge correction weights $w_{uv} = 3b_{uv}$ have signed weighted degree two. Let $e = |\operatorname{supp} w|$ and $R = |\operatorname{supp} H|$, so $W = R + e$.

The local equations now serve two purposes. They constrain how the fixed star mass can be supplied, and they identify the additional terms needed to cancel the resulting block demands. The notation below separates these coefficient equations from the support counts they force.

Symbol	Meaning
R, e, W	Retained directions, correction edges, and total width $W = R + e$.
Π, A, C, X	Private terms including helpers; alternating, loop, and mixed terms.
u, σ, r	A center, a partition into triples, and a private-line marginal.
K, m	Centers carrying nonalternating terms and their number.
G, U, B	A perfect matching, one of its pairs, and the block $B = (G, U)$.
t_U^G, Ω_G	Alternating sum in a source array and its common residual.
β_B, δ_B	Helper sum and signed flow of a block.
H^0, w	Effective private edge weights and correction edge weights.

We reserve G here for a perfect matching. Write g_{tot}, s_{tot} and n_{tot} for the total numbers of glue, star and nonalternating columns (the counts called G, S, N in the supplement).

7.1. **Private lines, helpers, and the support budget.** The retained directions have nine incidence types; their representatives and exact registry appear in Supplement S4. The notation records the roles used in the proof:

$$\underbrace{G, S, N}_{\text{glues, stars, nonalternating terms}}, \quad \underbrace{D, Z}_{\text{helpers}},$$

$$\underbrace{A}_{\text{alternating}}, \quad \underbrace{L, E}_{\text{loop terms}}, \quad X \text{ (mixed terms)}.$$

Counts always mean actual nonzero columns. Write

$$R = \Pi + A + C + X, \qquad \Pi = g_{\mathrm{tot}} + s_{\mathrm{tot}} + n_{\mathrm{tot}} + d + z, \qquad C = L + E. \tag{7.1}$$

The letters $d, z, h = d + z$ count D, Z, and all helpers. Uppercase family labels name column types; these counts name their supports.

Two examples explain the geometry. For the partition $123 \mid 456$, the glue indexed by centers $1, 4$ has argument $m_{12} + m_{13} - m_{45} - m_{46}$. The star in the same position has argument $m_{12} + m_{13} - m_{14} - m_{56}$. Glues couple the two triples; stars connect a center to its two halfmates and compare that

fork with two disjoint edges. These are roles of literal dictionary columns, not assumptions about coefficient signs.

Each of the ten partitions of [6] into two triples carries a 3×3 matrix of glue coefficients. Its rows and columns are called private lines; each has a physical center and a containing triple, called its half. On a line, let l_i be its three glue entries, call $r = \sum_i l_i$ its *marginal*, and let ν_i be the sum of its actual N entries in position i. The three star coefficients obey

$$S_i = l_i - r - \nu_i. \tag{7.2}$$

Write g, s, n for the numbers of nonzero glue, star, and individual nonalternating entries on this line. Its doubled support cost is $F = g+2s+2n$. For example, with no nonalternating entry, $l = (-1/2, 0, 0)$ gives $S = (0, 1/2, 1/2)$ and $F = 1 + 2 \cdot 2 = 5$. If $f(u)$ is the sum of line costs at center u, double-counting glues gives

$$2(\Pi - h) = \sum_{u=1}^{6} f(u). \tag{7.3}$$

Let K be the centers with any selected N and $m = |K|$. Balanced canceling N still count. The local equations give $f(u) \geq 5$ off K and $f(u) \geq 8$ on K; a nonzero marginal sum at an active center raises the latter bound to nine. The small-cost equality forms are proved and classified in Supplement S4–S5.

7.2. **Two different ways of grouping alternating terms.** A block is a perfect matching $G = \{U, V, W\}$ together with a distinguished source pair U. There are $15 \cdot 3 = 45$ blocks. For $U = \{u, v\}$ define coefficient sums

$$\zeta_{B,u} = D_B + Z_{B,u}, \qquad \beta_B = 2D_B + Z_{B,u} + Z_{B,v}.$$

The actual diagonal vertex of a Z argument is its attachment. Coefficients retain their stored sign when a hinge argument is reversed.

An alternating direction consists of an alternating four-vertex path and a disjoint edge U. Join U to the two end edges of the path; these three pairs form its matching G. The four directions with these same labels form the *source array* (G, U), and t_U^G is their coefficient sum. Its block ownership is determined separately by the literal local equations. These two partitions have different members, as the literal example in Figure 2 shows. The equations imply that $t_U^G - \beta_{(G,U)}$ is independent of $U \in G$; define this common value to be Ω_G. Thus

$$t_U^G = \Omega_G + \beta_{(G,U)}, \qquad w = \sum_{G \in \mathcal{M}} \alpha_G \mathbf{1}_G, \qquad \alpha_G = -\Omega_G, \quad \sum_G \alpha_G = 2. \tag{7.4}$$

Here is the cancellation behind the common value. Fix $G = \{12, 34, 56\}$. The transport and private rows cancel all columns except the two source arrays and their helpers, leaving

$$\begin{aligned} t_{34}^G - t_{12}^G = {} & 2(D_{(G,34)} - D_{(G,12)}) \\ & + Z_{(G,34),3} + Z_{(G,34),4} - Z_{(G,12),1} - Z_{(G,12),2}. \end{aligned}$$

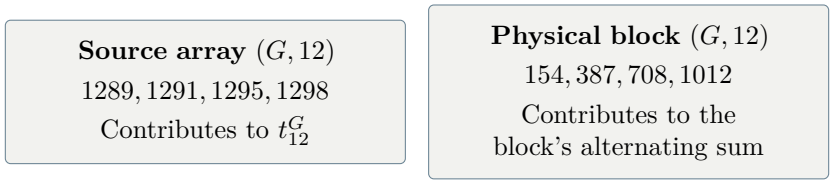


FIGURE 2. Literal alternating-column memberships (Supplement S10.4). Coordinates are numbered $1, \dots, 6$; column IDs retain the zero-based registry. Each column has one array and one block owner, but these partitions differ. Their support bounds may be added only after the relevant supports have been proved disjoint.

The right side is exactly $\beta_{(G,34)} - \beta_{(G,12)}$. Exchanging the pairs gives the same identity for the third source, so all three corrected array sums equal Ω_G. Supplement S4.2 prints the full-row certificate and its expansion. The Boolean-content rows give the displayed decomposition of w. Summing its edge weights gives $3\sum_G \alpha_G = 6$, because the six vertex degrees are all two. This proves the normalization $\sum_G \alpha_G = 2$ without a sign assumption.

Independently, each block satisfies

$$\sum_{S \text{ in } B} S + \sum_{A \text{ in } B} A = \beta_B. \tag{7.5}$$

The two groupings give different counting arguments. A nonzero array sum requires a term in that array; a nonzero block demand requires a term in that block. Since a column belongs to both groupings, their lower bounds can be added only after the relevant sets of columns have been proved disjoint.

Let $q = \#\{G : \Omega_G \neq 0\}$. A helper can cancel at most one of the $3q$ initially nonzero array sums, so $A \geq 3q - h$. Combining this with (7.3) gives the basic estimate

$$R \geq 15 + \left\lceil \frac{3m}{2} \right\rceil + 3q + C + X. \tag{7.6}$$

The argument uses support counts, not positivity of coefficients.

7.3. **Flow and the physical projection.** Define the directed private mass and its symmetrization by

$$B_{uv} = -2 \sum_{\substack{\text{lines at } u \\ v \text{ in their half}}} r, \qquad H^0_{uv} = \tfrac{1}{2}(B_{uv} + B_{vu}).$$

Orient the side pairs of a block from V to W and denote its loop flow by δ_B. Its edge divergence has entries $-\delta_B$ at V and $+\delta_B$ at W. The literal

canonical equations imply

$$\operatorname{div}(\delta)_a = H^0_a + \sum_B \beta_B(\mathbf{1}_{a=U} + \tfrac{1}{2}\mathbf{1}_{a=V} + \tfrac{1}{2}\mathbf{1}_{a=W}) - w_a, \tag{7.7}$$

$$2\sum_B \delta_B(\mathbf{1}_W - \mathbf{1}_V) = \sum_{B=(G,\{u,v\})} (Z_{B,v} - Z_{B,u})(e_u - e_v). \tag{7.8}$$

The first equation is indexed by edges; the second uses vertex incidence vectors. Full-column row identities establish both, including the cancellation of mixed entries (Supplement S4). In the matching normal form, a sole Z requires at least six loop columns. With $Z = 0$, nonzero flow also costs at least six. Zero net flow may still contain nonzero loop columns.

At fixed correction data, the pure kernel move

$$(D, Z_u, Z_v, E) \longmapsto (D + a, Z_u - a, Z_v - a, E + a) \tag{7.9}$$

allows a representative minimizing (R, z, d) lexicographically. Integer support counts ensure existence, without bounding rational coefficients. Its consequences are applied only within their proved scope.

8. The global lower bound for MAX6

It remains to exclude canonical identities of total width at most 32. The proof proceeds in two stages: first rule out widths at most 31, then resolve width 32. Figure 3 records the dependencies of the assembly. This order is essential because some normalizations at width 32 decrease support. Their contradictions use the earlier lower bound, which must already be established.

8.1. **The correction graph organizes the cases.** No vertex of the correction graph is isolated, because its weighted degree is two. Three edges therefore form a perfect matching, with all weights two. Four edges are impossible. Indeed, a simple four-edge graph on six nonisolated vertices is a forest with component sizes $2 + 4$ or $3 + 3$. Summing the signed degree equations over the two classes of a bipartite component forces equal class sizes. The remaining four-vertex path then has two end weights two and a zero middle weight, contradicting its support. Also, $q \leq 2$ in (7.4) gives at most six correction edges.

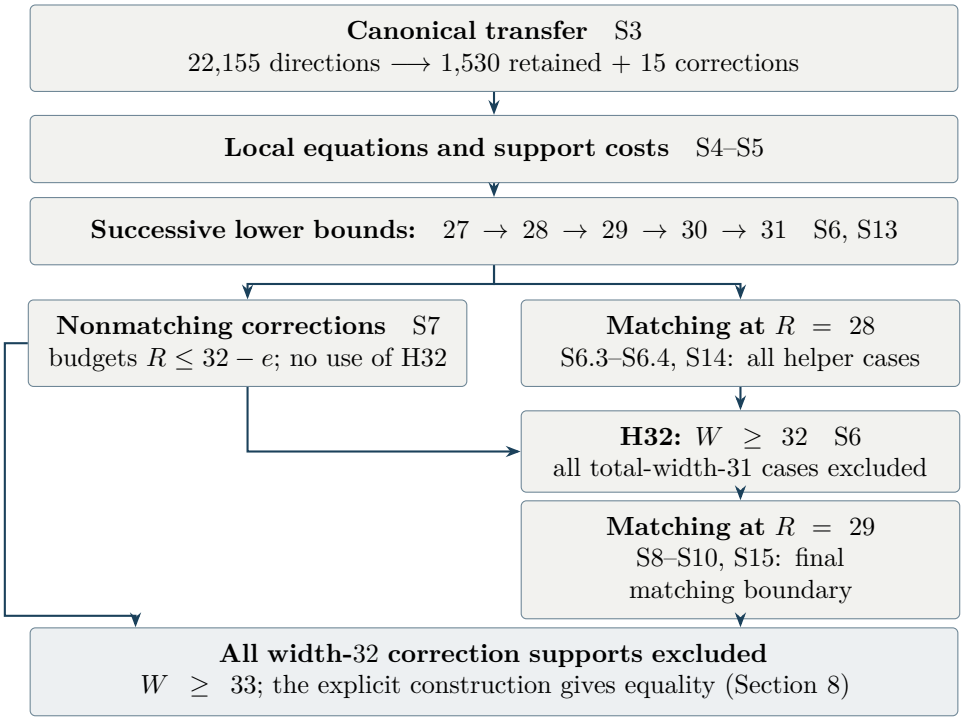


FIGURE 3. The MAX6 proof dependencies. Here $W = R+e$, with R retained directions and e corrections. The nonmatching branch includes S7 and the degree and cost cases of Section 8; it is proved before H32 and reused in the final join. The matching $R = 29$ argument uses H32 only after it has been established. Labels beginning with S refer to the technical supplement.

Lemma 8.1 (Width-31 exclusion)**.** *The canonical system has no rational identity of width at most* 31*.*

Computer-assisted proof; Supplement S6. The preliminary lower bound 31 is established by the complete chain in Supplement S13. Its finite ingredients include the retained-cap-21 search, whose complete branching and necessary pruning predicates are specified in Supplement S6.1. At total width 31, the degree census and the established branch bounds leave precisely

$$(e, R) = (3, 28), (5, 26), (6, 25), (7, 24).$$

For five and six edges, the independent budget lemmas of Supplement S7 exclude every $R \leq 32 - e$, including the smaller values here; their proofs assume no lower bound of 32. The seven-edge case is excluded by the local costs. For $R = 28$ with matching corrections, the proof first separates zero helpers at every attachment count and both flow states. With nonzero

helpers, it separates D-only from Z-present cases, then all possible attachment counts and source geometries. The local equations, support costs, and exact residual systems exclude each branch. Supplement S6 gives the routing and the physical-Z arguments, including arbitrary star positions and the shared line partner; Supplement S14 supplies the detailed component proofs. Strict support decreases contradict only the preliminary bound 31; the later matching-29 exclusion is not a premise. □

Lemma 8.2 (Matching boundary)**.** *There is no normalized canonical identity with three correction edges and retained support at most* 29*.*

Computer-assisted proof; Supplement S8–S10.

Let M be the correction matching. A nonzero signed relation among perfect matching incidence vectors has support at least six (Supplement S4.5). If $\Omega \neq -2e_M$, then $\Omega + 2e_M$ is such a relation and $q \geq 5$. Equation (7.6) would force $R \geq 30$. Thus $\Omega = -2e_M$ and $w = 2\mathbf{1}_M$.

The finite integer envelope in Supplement S8 bounds the possible support costs. It routes hypotheses to proofs; it does not enumerate representations. Zero helpers are treated first. At zero flow the effective private matching lemma requires private support at least 27, and three nonzero arrays give $R \geq 30$. At nonzero flow the loop cost confines the private layer to support at most 20; its complete equality forms fail block demands or literal row equations.

For D only, source degrees and local line costs reduce the proof to single-source, repeated-source, path, and disjoint-source cases, with all balanced attachment centers retained. The sharp private-25 case is proved in Section 9. For Z present, the one- and three-attachment cases have forced local support classifications. The remaining two-attachment cases split into same attachment, both blocks in M, one in and one out, both out, and the residual three-helper cases. They use the physical projection (7.8), exact containing-column spaces, and separate cost arguments where those spaces are consistent. Supplement S8–S10 records each subdivision, its hypotheses, and its terminal contradiction, with the further derivations in Supplement S15. This exhausts the normalized support envelope. □

8.2. **Nonmatching corrections at width thirty-two.** Assume a canonical identity of width at most 32. By Lemma 8.1 its width is 32. The complete correction-support split is in Table 4; we explain the smaller analytic branches here.

For $e = 5$, the signed reduction derives $q = 2$; it does not assume that every five-edge signed graph is a four-cycle plus an isolated edge. The signed double-star alternative is included in the reduction. In the surviving two-matching geometry, the local arguments eliminate helpers. Zero flow requires effective private support at least 24, whereas $A \geq 6$ leaves at most 21. Nonzero flow forces $(\Pi, A, C, X) = (15, 6, 6, 0)$. The minimum private-line forms give 60 dense and 120 sparse cores. Each would need one alternating

e	$R = 32 - e$	Contradiction
0–2	—	A vertex would be isolated.
3	29	Matching boundary, Lemma 8.2.
4	28	Signed forest census.
5	27	Signed reduction and five-edge equality cores.
6	26	Small private layers for $q = 3$; six-edge exclusion for $q = 2$.
7	25	Three supported matchings yield eight or nine edges.
8	24	Twelve distinct star blocks require $A \geq 12$, but $A = 9$.
≥ 9	≤ 23	$q \geq 3$ and (7.6) require $R \geq 24$.

TABLE 4. Every possible correction support at width thirty-two. The five- and six-edge cases are proved in Supplement S7.

in each of six source arrays and six demanded blocks. Exact Hall deficits rule out all such assignments (Supplement S7.1).

For $e = 6$, (7.6) excludes $q \geq 4$. With $q = 3$, (7.6) gives $m \leq 1$. If $m = 1$, it also gives $C = X = 0$: a D needs two active endpoints, and a sole Z needs at least six loops (S13.5, Part G, item 1). Thus helpers are absent. If $m = 0$, there are no helpers and $C + X \leq 2$; nonzero helper-free flow would need at least six loops (S4.3). In either case the flow is zero. The $q = 3$ groups require $A \geq 9$, so $\Pi \leq 17$. Eliminating balanced N preserves center mass and H^0 and does not increase private support (S4.1; S13.4, item 2, "Balanced transport"). The effective layer therefore has support at most 17. When $N = 0$, the line spectrum is

$$r \neq 0 : \ F \in \{5, 7, 8, 9\}, \qquad r = 0 : \ F \in \{0, 6, 9\}.$$

By the private-support gap and line spectrum (S5.1), each center needs a nonzero marginal, while seven such lines would cost at least $35 > 2 \cdot 17$. Hence each center has one marginal $-1/2$; B has two ones per row and $H^0 = w \geq 0$. Each of the three distinct matchings has an edge unique to it, so its weight is positive. Their union has at least seven edges, contradicting $e = 6$. For $q = 2$, the matchings form a six-cycle. The HE6 reduction leaves 33 effective private forms and 47 actual forms after balanced-N erasures are restored. The two surviving alternating completions each fail eight literal T rows (Supplement S7.2).

For $e = 7$, the budget forces $q = 3, m = 0, C + X \leq 1$. Effective private support is at most 16. The same S5.1 six-line argument, with center mass one and zero flow, gives $w \in \{0, 1/2, 1\}$. Each of three distinct matchings has an edge in neither of the other two: each other matching shares at most one of its three edges. These unique edges force weights $(1, 1/2, 1/2)$. The weight-one matching must be disjoint from both others, so the union has eight or nine edges. For $e = 8$, equality gives $\Pi = 15, A = 9, C = X = 0$. The eight-edge equality census in S5.1 ("Eight-edge equality census") classifies

every such private layer. Its halves, in center order and up to relabeling, are

$$125, \quad 234, \quad 346, \quad 124, \quad 156, \quad 356.$$

Each line has a single complementary selected half, hence one glue $-1/2$ and two stars $+1/2$. Their source pairs are respectively $25, 34, 46, 12, 16, 35$, all distinct. The two stars at a center have different matching owners. Thus twelve distinct blocks need alternating terms by (7.5), contradicting $A = 9$.

Completion of Theorem 1.1. The five-input result was proved in Section 5. For six inputs, Theorem 6.2 transfers any identity of width at most 32 to a canonical one. Lemma 8.1 removes all widths below 32. Table 4 and the preceding arguments remove every correction graph at width 32, with Lemma 8.2 treating the matching. Therefore $W^*(6, 2) \geq 33$. Proposition 4.1 gives equality over $\mathbb{Q}$, and Lemma 3.2 gives the asserted real-coefficient minimum. □

9. At the sharp boundary: a twenty-five-term private core

The final example shows how the lower-bound equations lead directly to the width-33 construction. It isolates one family at the boundary; the preceding sections supply the global exclusion. Assume $R \leq 29$, four active centers, $Z = 0$, and two disjoint nonzero D sources. The local cost and source equations at $\Pi \leq 25$ force both coefficients to be one, their common matching to be $M = 12 \mid 34 \mid 56$, and the sources to be 12 and 34. There are four active centers of cost nine and two inactive centers of cost five. Supplement S15.3 and Supplement S14.5 derive these equalities without a sign assumption.

At each active center, exactly two N coefficients, both -1, occupy distinct lines. An additional balanced pair cannot fit. One line has cost four, no glue, and an actual star $+1$. The other has cost five. Equation (7.2) and marginal $1/2$ force its glue vector to be $(-1/2, 1/2, 1/2)$, with no stars and the negative entry at its N position.

The dense half contains its source pair and an inactive vertex. If it contained a vertex of the other source pair, a crossing glue would reach that vertex's mate on a separating line, impossible in the cost-nine form. The crossings therefore force the common split $125 \mid 346$. With rows $(1, 2, 5)$ and columns $(3, 4, 6)$ the glue matrix is

$$\Lambda = \frac{1}{2} \begin{pmatrix} 1 & 1 & -1 \\ 1 & 1 & -1 \\ -1 & -1 & 1 \end{pmatrix}. \tag{9.1}$$

The dense N is fixed at each active center. Its other line has two choices, so the entire family consists of exactly $2^4 = 16$ private cores, each with $9G + 6S + 8N + 2D = 25$ terms. These choices follow from the equations before the finite check starts.

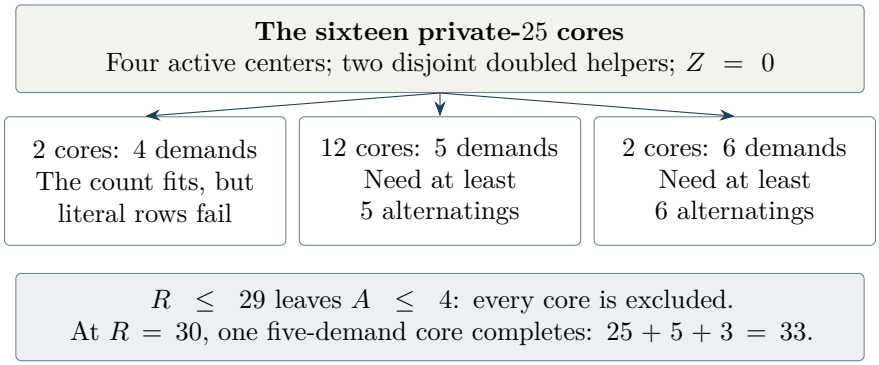


FIGURE 4. How private support creates alternating demand in the family of Section 9. The last three terms are four-set corrections. This is a scoped boundary census, not a classification of all optimal identities.

Among these sixteen cores, two have four nonzero block demands, twelve have five, and two have six (Figure 4). Every demand needs an alternating term in its own block. Thus $A \geq 4$, and the total budget forces

$$29 \geq 25 + A + C + X \quad \Longrightarrow \quad A = 4,\ C = X = 0. \tag{9.2}$$

Fourteen cores are already impossible. In either remaining core, all allowed alternating columns lie in the four demanded blocks. Their sixteen columns are retained in the exact equations; they are not narrowed to four guessed columns. Seven literal row substitutions have the form

$$\begin{gathered} u_0 = 0, \qquad u_0 + u_1 = 1, \qquad u_2 + u_1 = 0, \qquad u_2 + u_3 = 0, \\ u_3 + u_4 = 0, \qquad u_5 = -1, \qquad u_6 = -1. \end{gathered} \tag{9.3}$$

The six distinct coefficients $u_1, \ldots, u_6$ are nonzero, contradicting $A = 4$. The actual column identifiers and row substitutions are given in Supplement S9.4; the renaming hides no identification of columns.

The hypothesis $R \leq 29$ is essential. Only (9.2) removes loops and mixed terms from the literal equations. At $R = 30$, one of the five-demand cores accepts five alternatings. Together with three corrections it gives the explicit width-33 construction in Table 3.

10. Verification and further questions

10.1. **The formal theorem and its scope.** Both rational optimality results have Lean proofs, using Lean and Mathlib [14, 15]. The MAX6 theorem is the closed declaration

```
PM2.ContinuationMax6Optimality20261029.
  rational_PM2_MAX6_optimality :  PM2.IsOptimal6 33.
```

Unfolding this statement gives both an admissible 33-direction identity and the lower bound for every admissible identity on all of $\mathbb{R}^6$. The coefficient

field is $\mathbb{Q}$; the directions are integer differences of two size-two feature multisets; the free term is feature-linear. There are no remaining coverage or certificate hypotheses in the final statement. The passage to real coefficients is the written proof of Lemma 3.2.

On 31 October 2026, all 16,057 PM2 modules in the MAX6 theorem's actual dependency closure completed a fresh source rebuild. The closed theorem and the two explicit all-real upper- and lower-bound projections then passed independent kernel replay with checking enabled and trust level zero. The only axioms were `propext`, `Classical.choice`, and `Quot.sound`. Scope and completion audits matched the statement to the original model and confirmed the complete case assembly. The MAX5 result has a separate successful replay of its preserved closed proofs. Supplement S12 gives the records and distinguishes these verification modes.

10.2. **Exact arithmetic and the role of the supplement.** An independent rational checker verifies both displayed identities on every chamber. The finite lower-bound computations regenerate the MAX6 dictionary and matrix, check the local row identities and containing spaces, and exhaust the retained-cap-21 search, which visits 16,008,595 nodes with no admissible support. Their predicates and recorded executions are specified in S11–S12. The written local arguments and case coverage are developed in S3–S10 and S13–S15.

The formal proof establishes the rational theorem end to end; the supplement explains its mathematics. In particular, the two roles of a finite certificate remain visible: the proof must first show that its column set contains the case under consideration, and the certificate must then exclude that set. A numerical separation without containment would not establish the lower bound.

10.3. **Questions beyond the two exact values.** The fixed side-size-two sequence is complete by Corollary 3.5. Three questions remain beyond that classification. Which width-33 identities exist up to coordinate permutation and projective normalization? What is $W^*(7,3)$, the directional optimum when each side may contain three features? And which reductions survive when more than two features are allowed on each side of a hinge? For the last question, touched-set separation pinpoints the difficulty: after one variable tends to infinity, a side-size-three hinge can retain an interior kink. Extending the method therefore requires a new reduction. At fixed side size two there is no finite asymptotic width sequence: representability ends at six inputs. An asymptotic study must let the side size grow or enlarge the dictionary. Unrestricted network optimality remains a separate problem.
Data and code availability. The compact Zenodo package contains both papers, editable manuscript sources, the direct exact-arithmetic checkers and their inputs, and the complete Lean source compendium with its finite proof data and accepted verification records. No Lean proof case is omitted. Lean

4.32.2 and the pinned external libraries must be installed separately for compilation. The large historical research archives and the macOS offline verification environment are retained locally and are not part of this deposit. Supplement S12.2 separates checks supported by the deposit from those requiring the retained archives. The record `evidence/max6-formalization.json` identifies the formal theorem, its scope, and the dated completion audits. The DOI for this paper is 10.5281/zenodo.24051030. The papers are licensed under CC BY 4.0; the author's Lean proofs and code are licensed under Apache 2.0. Existing third-party licenses are preserved.

Acknowledgments

The preparation of this manuscript and its computational companion benefited from assistance by AI systems developed by OpenAI and Anthropic. Responsibility for the mathematical claims, implementation, and presentation rests with the author.

References

[1] R. Arora, A. Basu, P. Mianjy, and A. Mukherjee. *Understanding Deep Neural Networks with Rectified Linear Units*. International Conference on Learning Representations, 2018. arXiv:1611.01491.

[2] C. Hertrich, A. Basu, M. Di Summa, and M. Skutella. *Towards Lower Bounds on the Depth of ReLU Neural Networks*. SIAM Journal on Discrete Mathematics 37(2), 997–1029, 2023. doi:10.1137/22M1489332.

[3] F. Bach. *Learning with Submodular Functions: A Convex Optimization Perspective*. Foundations and Trends in Machine Learning 6(2–3), 145–373, 2013. doi:10.1561/2200000039.

[4] E. Bakaev, F. Brunck, C. Hertrich, J. Stade, and A. Yehudayoff. *Better Neural Network Expressivity: Subdividing the Simplex*. Proceedings of the 58th Annual ACM Symposium on Theory of Computing, 2026. arXiv:2505.14338v3, 19 February 2026.

[5] K. Rueß, G. Averkov, F. Brunck, M. Grillo, C. Hertrich, G. Loho, J. Stade, M. Stargalla, M. Sun, and M. Winter. *Shallower ReLU Network Representations via Exact Linear Algebra*. arXiv:2607.21651v2, 1 September 2026.

[6] Z. Wang and A. Basu. *Representing MAX Functions Using Two-Hidden-Layer ReLU Networks*. arXiv:2608.25221v1, 25 August 2026.

[7] F. Brunck, C. Hertrich, K. Rueß, J. Stade, and A. Yehudayoff. *Two ReLU Layers is All You Need*. Preprint, arXiv:2610.28587v1, 20 October 2026.
[8] I. Safran. *A Depth Hierarchy for Computing the Maximum in ReLU Networks via Extremal Graph Theory*. Proceedings of the 39th Conference on Learning Theory, Proceedings of Machine Learning Research 336, 5721–5742, 2026. PMLR 336: Safran.
[9] J. P. Zimmermann. *Two Hidden Layers Suffice: The Caterpillar Net*. Preprint, arXiv:2610.26795v1, 19 October 2026.
[10] M. Grillo, C. Hertrich, and G. Loho. *Depth-Bounds for Neural Networks via the Braid Arrangement*. Advances in Neural Information Processing Systems 38, 2025. doi:10.52202/085713-4257.
[11] C. Haase, C. Hertrich, and G. Loho. *Lower Bounds on the Depth of Integral ReLU Neural Networks via Lattice Polytopes*. International Conference on Learning Representations, 2023. arXiv:2302.12553.
[12] G. Averkov, C. Hojny, and M. Merkert. *On the Expressiveness of Rational ReLU Neural Networks With Bounded Depth*. International Conference on Learning Representations, 2025. ICLR 2025 proceedings.
[13] D. Myrisiotis. *A Hand-Checkable Proof That Two Hidden ReLU Layers Compute the Maximum of Six Numbers*. Preprint, arXiv:2610.04256v1, 3 October 2026.
[14] L. de Moura and S. Ullrich. *The Lean 4 Theorem Prover and Programming Language*. In *Automated Deduction—CADE 28*, pp. 625–635, 2021. doi:10.1007/978-3-030-79876-5_37.
[15] The mathlib Community. *The Lean Mathematical Library*. In *Certified Programs and Proofs*, 2020. doi:10.1145/3372885.3373824.

Independent Researcher