Storage capacity

Complete current results

Detailed extension of the current status report; Attempts 1–3.

$$\boxed{1 \le \alpha_- \le \alpha_+ \le 2.}$$

The exact endpoints and their equality remain unresolved; conditional implications below retain their missing premises.

Q1. What is α_-?

Q2. What is α_+?

Q3. Which dual certificates force nonstorage?

Q4. Do α_- and α_+ coincide?

Q5. What do the finite classifications establish?

Table 1: Contributions of the three completed attempts.

Attempt	Rounds	Contribution
1	3	Storage at n; one-row upper bound; exact finite duality.
2	10	Sublinear excess; neuron support; window equivalences; 168-case recertification.
3	12	Floor transfer; graph and planting reductions; 120 new exact cases.

Proved means the stated mathematical scope; this report does not add an independent audit of every proof dependency.

Model and endpoint definitions

For integers $n \geq 2$, $1 \leq m \leq 2^n$, sample an unordered set of distinct memories uniformly:

$$D_{n,m} = \{s^1, \ldots, s^m\} \sim \mathrm{Unif}\binom{\{-1,1\}^n}{m}.$$

$$J \in \mathbb{R}^{n \times n}, \qquad J = J^\top, \qquad J_{ii} = 0, \qquad h \in \mathbb{R}^n.$$

$$F_n(m) = \{\exists (J,h) : \ s_i^\mu((Js^\mu)_i + h_i) > 0 \quad \forall i \in [n],\ \mu \in [m]\}.$$

$$\mathrm{load} = \frac{m}{n}.$$

$$p_n(m) = \Pr(F_n(m)), \qquad q_n(u) = 1 - p_n(\lfloor un \rfloor).$$

$$\begin{aligned} \alpha_- &= \sup\{u > 0 : p_n(\lfloor un \rfloor) \longrightarrow 1\}, \\ \alpha_+ &= \inf\{u > 0 : p_n(\lfloor un \rfloor) \longrightarrow 0\}. \end{aligned}$$

Every limit fixes the load u and sends n to infinity; the network may depend on the complete sample.

$$\alpha_- = a_- = \alpha_{\mathrm{sym}}, \qquad \alpha_+ = a_+.$$

The last aliases translate the filed notation; they do not assume endpoint equality.

All statements and empirical certificates below use the spin formulation.

Fields and weights are real and unrestricted; additional fixed points are allowed, and all logarithms are natural.

Constraint matrix and the symmetry experiment

Write $[k] = \{1, \dots, k\}$, $\operatorname{sym} T = (T + T^\top)/2$, and $\mathbf{1}_k = (1, \dots, 1)^\top$.

$$d = n + \binom{n}{2}, \qquad N = nm, \qquad \theta = (h, (J_{ij})_{i<j}) \in \mathbb{R}^d.$$

$$A_R = A(D_{n,m}) \in \mathbb{Z}^{N \times d}, \qquad (A_R\theta)_{\mu i} = s_i^\mu((Js^\mu)_i + h_i).$$

There is one coordinate for each field and each unordered edge; both endpoint constraints share the same edge coordinate.

$$\mathcal{C}(D) = \{\lambda \in \mathbb{R}^N_{\geq 0} : \mathbf{1}_N^\top \lambda = 1,\ A_R^\top \lambda = 0\}.$$

For a fixed sample, each neuron predicts its own sign from the other memory coordinates:

$$E_i = \left\{ \begin{array}{l} \exists\big((J_{ij})_{j\neq i}, h_i\big) \in \mathbb{R}^{n-1} \times \mathbb{R} : \\ s_i^\mu \left(\sum_{j\neq i} J_{ij} s_j^\mu + h_i\right) > 0 \quad \forall \mu \in [m] \end{array} \right\}.$$

$$E = E_n(m) = \bigcap_{i=1}^n E_i, \qquad F = F_n(m) \subseteq E.$$

E fits the rows of J without symmetry; F requires the same inequalities with $J = J^\top$, and each row stays fixed across all memories.

$$F^c = E^c \,\dot\cup\, (E \cap F^c).$$

The first event is local failure; the second is symmetry-only failure.

$$c_n(u) = \Pr(F^c \mid E) = \frac{\Pr(E \cap F^c)}{\Pr(E)} \quad \text{if } \Pr(E) > 0.$$

$$e_n(u) = \Pr(E), \qquad p_n(m) = \Pr(F) = e_n(u)[1 - c_n(u)], \qquad m = \lfloor un \rfloor.$$

The probabilities are over uniformly sampled memory sets; neuron feasibility events need not be independent.

Lemma: conditional failure bounds the endpoint

For fixed $u > 0$, assume $e_n(u) > 0$ eventually. Then

$$c_n(u) \to 1 \quad \Longrightarrow \quad \alpha_+ \leq u.$$

Proof. The identity and $0 \leq e_n(u) \leq 1$ give

$$0 \leq p_n(\lfloor un \rfloor) = e_n(u)[1 - c_n(u)] \leq 1 - c_n(u) \to 0.$$

Thus u belongs to the set whose infimum defines α_+.

Q1 · Storage constructions

Q1. What is α_-?

R2. Storage with sublinear excess

Attempt 2; Proved in the stated scope

For uniform distinct patterns and one real symmetric zero-diagonal network with free fields,

$$M_n = n + \left\lfloor \frac{n}{55296 \log n} \right\rfloor, \qquad p_n(M_n) \longrightarrow 1.$$

Since $M_n/n \to 1$, this gives no fixed storage load above one.

Proof strategy

1. Use inverse augmented covariance to write every margin as a residual.
2. Bound all residuals and singularity probabilities at the stated excess.

C1. Projection construction

Attempt 1; Proved in the stated scope

Let $1 \le m \le n$, $S = [s^1 \ \cdots \ s^m] \in \{-1, 1\}^{n \times m}$. Set $U = [S^\top \ \mathbf{1}_m]^\top \in \mathbb{R}^{(n+1)\times m}$. If deleting any of its first n rows leaves column rank m, define

$$P = U(U^\top U)^{-1} U^\top = \begin{pmatrix} B & h \\ h^\top & c \end{pmatrix}, \qquad J = B - \operatorname{diag}(B).$$

Then J is symmetric with zero diagonal, and

$$s_i^\mu((Js^\mu)_i + h_i) = 1 - P_{ii} > 0 \qquad (1 \le i \le n,\ 1 \le \mu \le m).$$

Thus one network stores every supplied pattern. For uniform distinct random patterns, the rank condition holds with probability tending to one at $m = n$; consequently

$$p_n(n) \longrightarrow 1, \qquad \alpha_- \ge 1.$$

The construction's loss of strict margins beyond n does not imply nonstorage by other networks.

Proof strategy

1. Project onto the augmented pattern span and remove the diagonal.
2. Row-deletion rank gives positive margins; random rank bounds give the limit.

C2. Symmetrization margin loss

Attempt 1; Proved in the stated scope

Let $T \in \mathbb{R}^{n\times n}$ have zero diagonal, $h \in \mathbb{R}^n$, and $s^1, \dots, s^m \in \{-1,1\}^n$, with $m \geq 1$. Define

$$J = \tfrac{1}{2}(T + T^\top), \qquad K = \tfrac{1}{2}(T - T^\top),$$

$$r_{i\mu} = s_i^\mu((Ts^\mu)_i + h_i), \qquad \ell_{i\mu} = s_i^\mu(Ks^\mu)_i.$$

The symmetric zero-diagonal network (J, h) has signed margins

$$s_i^\mu((Js^\mu)_i + h_i) = r_{i\mu} - \ell_{i\mu}.$$

It strictly stores the sample with this field exactly when $\min_{i,\mu}(r_{i\mu} - \ell_{i\mu}) > 0$. Positive directed margins alone do not suffice; finite counterexamples below remain infeasible for the averaged J even after replacing its fields. No random-sample estimate controlling this loss above load one is established here.

Proof strategy

1. Split the directed matrix into symmetric and skew parts.
2. Subtract the skew contribution from each original signed margin.

Trained primal route: setup

This route fixes $u = 9/5$, $m = \lfloor 9n/5 \rfloor$; the vector d of row weights is distinct from the scalar constraint dimension d.

$$(J_{i,-i}^{\text{free}}, h_i) = \operatorname*{argmin}_{J_{i,-i}\in\mathbb{R}^{n-1},\, t\in\mathbb{R}} (\|J_{i,-i}\|_2^2 + t^2),$$
$$\text{subject to } s_i^\mu \left(\sum_{j\neq i} J_{ij} s_j^\mu + t \right) \geq 1 \quad (\mu \in [m]).$$

The subscript $-i$ removes the diagonal entry; J^{free} denotes the fit without symmetry before constructing a symmetric J.

The strictly convex problem has a unique solution on E; put $J_{ii}^{\text{free}} = 0$ and abort outside E.

For row weights $d_i \geq 0$, $\sum_i d_i = n$, use $J = \operatorname{sym}(\operatorname{diag}(d) J^{\text{free}})$.

$$\delta_i(J) = \min_{\mu: s_i^\mu = 1} (Js^\mu)_i - \max_{\nu: s_i^\nu = -1} (Js^\nu)_i.$$

A row containing only one label imposes no opposite-label test; otherwise a free field exists exactly when $\delta_i(J) > 0$.

For the graph reduction use augmented indices $a, b \in \{0, \dots, n\}$; graph signs $z_\mu \in \{-1, 1\}$ are scalar variables, distinct from the feature vectors z_a^μ.

The graph has m pattern vertices, with all its signed edge constraints retained; consistency means every cycle has compatible signs, and isolated vertices count as components.

$$\widehat{E} = \bigcap_{a=0}^{n} \left\{ \exists A_a \in \mathbb{R}^n : \; A_a \cdot z_a^\mu \geq 1 \; \forall \mu \in [m] \right\}.$$

Here $s_0^\mu = 1$, the other signs retain their original values, and $z_a^\mu = (s_a^\mu s_b^\mu)_{b\neq a}$; augmented feasibility also imposes the row a=0 test.

Trained primal graph

Attempt 3; Proved reduction; occurrence remains open

Fix $m = \lfloor 9n/5 \rfloor$. On every-row affine feasibility, fit each unique minimum-norm unit-margin row $(J_i^{\mathrm{free}}, h_i)$, let $Q_i = \|J_i^{\mathrm{free}}\|^2 + h_i^2$, and define

$$d_i = \frac{n/Q_i}{\sum_j Q_j^{-1}}, \qquad J = \mathrm{sym}(\mathrm{diag}(d) J^{\mathrm{free}}).$$

The inverse-Q candidate succeeds when its minimum opposite-label score gap is positive, with final free fields chosen between the two score classes. It aborts on row infeasibility. For its augmented trained graph, extend each memory by $s_0^\mu = 1$; define $z_a^\mu = (s_a^\mu s_b^\mu)_{b \neq a}$, and minimize $\|A_a\|_2^2$ subject to $A_a \cdot z_a^\mu \geq 1$ for every training pattern; put $Q_a = \|A_a\|_2^2$. On augmented feasibility $\widehat{E}$, put

$$T_{ab} = A_{ab}/Q_a, \quad J^{\mathrm{aug}} = \mathrm{sym}\, T, \quad M_{a\mu} = s_a^\mu (J^{\mathrm{aug}} s^\mu)_a.$$

Each $M_{a\mu} + M_{a\nu} \leq 0$ imposes the graph equation $z_\mu z_\nu = s_a^\mu x_a^\nu$. Contradictory signs mean inconsistency; $r_n = m - \#\mathrm{components}$, including isolated vertices. For inverse-Q success C and

$$\mathcal{A}_{n,R} = \{\widehat{E}, \mathrm{consistent}, r_n \leq R\}$$

,

$$2^{-R} \Pr(\mathcal{A}_{n,R}) - o(1) \leq \Pr(C) \leq \Pr(\mathcal{A}_{n,R}) + 2^{-(R+1)} + o(1).$$

Proof strategy

1. Encode failed score pairs as signed equations on pattern vertices.
2. Use graph rank and uniform column occupancy to count successful roots.

The optimized weighted family

Attempt 3; Proved in the stated scope

Fix $m = \lfloor 9n/5 \rfloor$; on every-row feasibility E, fit minimum-norm unit-margin rows J^{free} fitted without symmetry, with fields. Let $\mathcal{P}$ consist of triples (i, μ, ν) with $s_i^\mu = 1$, $s_i^\nu = -1$. Put

$$(g_{i\mu\nu})_i = \tfrac{1}{2} \sum_{j \neq i} J_{ij}^{\mathrm{free}} (s_j^\mu - s_j^\nu), \qquad (g_{i\mu\nu})_j = \tfrac{1}{2} J_{ji}^{\mathrm{free}} (s_j^\mu - s_j^\nu) \quad (j \neq i),$$

$$\gamma_* = \max_{d \geq 0, \sum_i d_i = n} \min_{r \in \mathcal{P}} g_r \cdot d.$$

A positive optimum gives one symmetric storing matrix and final free fields.

Proof strategy

1. Express every opposite-label score gap as a linear function of the weights.
2. A positive minimum gap permits a separating field in every row.

The exact rank reduction

Attempt 3; Proved in the stated scope

The inverse-Q success probability differs by o(1) from

$$\mathbb{E}[\mathbf{1}_{\widehat{E}} \mathbf{1}_{\mathrm{consistent}} 2^{-r_n}].$$

Proof strategy

1. A consistent signed graph of rank r accepts a sign fraction 2^{-r}.
2. Uniform forest occupancy transfers this count to the trained graph.

The remaining actual-law probability

Attempt 3; Open probability premise; conditional consequence only

A positive alternative in Gap 5 or the survival estimate in Gap 12 would give

$$\alpha_+ \geq 9/5 > 17/10.$$

The negative alternative in Gap 5 retires only inverse-Q.

Proof strategy

1. Positive survival of a symmetric candidate gives positive full-storage limsup.
2. The required survival probability is still unestimated.

Q2 · Nonstorage bounds

Q2. What is α_+?

For a fixed neuron i, the input and label of pattern μ are

$$\text{input: } s^\mu_{-i} = (s^\mu_j)_{j\neq i} \in \mathbb{R}^{n-1}, \qquad \text{label: } s^\mu_i \in \{-1, 1\}.$$

$$z^\mu_i = (s^\mu_{-i}, 1) \in \mathbb{R}^n.$$

For fixed vectors z in general position and independent fair labels, Cover's homogeneous count gives

$$C(m,n) = 2\sum_{j=0}^{n-1} \binom{m-1}{j}, \qquad q^{\text{row}}_{m,n} = 2^{-m} C(m,n).$$

General position means every set of at most n vectors is linearly independent; arbitrary fixed inputs give an upper bound by the same count.

Under iid full spin patterns the labels of this one row are independent fair signs conditional on its inputs; conditioning to distinct full patterns costs at most $\binom{m}{2}2^{-n}$.

$$F \subseteq E \subseteq E_i, \qquad p_n(m) \leq q^{\text{row}}_{m,n} + \binom{m}{2} 2^{-n}.$$

Cube inputs may be degenerate, and separate feasible rows may fail reciprocal compatibility; the one-row success count below two gives no symmetric-storage lower bound.

Cover, "Geometrical and statistical properties of systems of linear inequalities" (1965), Theorem 1, p. 326, supplies the count; equality requires its general-position hypothesis.

R1. Endpoint bracket

Attempt 1; Proved in the stated scope

For all integers $n \geq 2$, $1 \leq m \leq 2^n$, the one-neuron count gives

$$p_n(m) \leq 2^{1-m} \sum_{k=0}^{n-1} \binom{m-1}{k} + \binom{m}{2} 2^{-n}.$$

Use $\binom{m-1}{k} = 0$ for $k > m-1$; the resulting bounds are

$$1 \leq \alpha_- \leq \alpha_+ \leq 2, \qquad p_n(n) \longrightarrow 1.$$

$$\forall u > 2: \quad p_n(\lfloor un \rfloor) \longrightarrow 0,$$

$$\limsup_{n\to\infty} p_n(2n) \leq \tfrac{1}{2}.$$

Proof strategy

1. Apply Cover's count to one row, then condition on distinct patterns.
2. Combine its binomial tail with the projection lower guarantee.

Positive-floor route: setup

This route fixes $1 < u < 2$, $m = \lfloor un \rfloor$, $a = n^{-4}$, with primitive columns indexed by fields i and unordered edges ij.

$$U = \begin{bmatrix} S \\ \mathbf{1}_m^\top \end{bmatrix}, \qquad S = [s^1 \ \cdots \ s^m], \qquad \kappa_{n,m} = \Pr_{\text{iid}}\{\sigma_{\min}(U) < c_u\sqrt{n}\} = o(1).$$

Here $\sigma_{\min}$ is the least row singular value and $c_u > 0$ is the conditioning constant at that fixed load.

For iid uniform spin columns define cube primitives, and independently define Gaussian primitives:

$$C_{\mu,i} = s_i^\mu, \qquad C_{\mu,ij} = s_i^\mu s_j^\mu, \qquad G \in \mathbb{R}^{m\times d}, \quad G_{\mu\alpha} \overset{\text{iid}}{\sim} N(0,1).$$

$$[A(Y)^\top \lambda]_{h_i} = \sum_{\mu=1}^m Y_{\mu,i}\lambda_{i\mu}, \qquad [A(Y)^\top \lambda]_{ij} = \sum_{\mu=1}^m Y_{\mu,ij}(\lambda_{i\mu} + \lambda_{j\mu}).$$

$$\mathcal{L}_a = \{\lambda \in \mathbb{R}^N : \lambda \geq a\mathbf{1}_N, \ \mathbf{1}_N^\top \lambda = 1\}, \qquad \rho(Y) = \min_{\lambda\in\mathcal{L}_a} \|A(Y)^\top \lambda\|_2.$$

$$\epsilon = (c_u/2)n^{-7/2}, \qquad f(Y) = \psi(\rho(Y)/\epsilon).$$

One fixed smooth $\psi : [0,\infty) \to [0,1]$ equals one on $[0,1/2]$ and zero on $[1,\infty)$; $\mathcal{L}_a$ is nonempty for all sufficiently large n.

$$\widetilde{G}_{\mu\alpha} = \frac{\sqrt{m}\, G_{\mu\alpha}}{\|G_{\cdot\alpha}\|_2}, \qquad \Delta_{n,m} = \binom{m}{2} 2^{-n} = o(1).$$

The normalized columns are independent uniform directions of length $\sqrt{m}$; cube edge columns remain products of the singleton columns.

Gaussian and cube expectations use these primitive laws, rather than independent full constraint-matrix entries.

For the Hermite statement, let $V_\alpha = \|G_{\cdot\alpha}\|_2^2$, $\Omega_\alpha = G_{\cdot\alpha}/\|G_{\cdot\alpha}\|_2$; conditioning on V or Ω means conditioning on the entire corresponding column array.

Positive-floor transfer

Attempt 3; Conditional implication

Conditional on Gap 3 and Gap 4 at one fixed load u,

$$p_n(\lfloor un \rfloor) \to 0, \qquad \alpha_+ \le u.$$

Proof strategy

1. Correct a sufficiently small positive-floor residual to an exact cube dual.
2. Transfer Gaussian occurrence through the same-load comparison estimate.

Exact reciprocal correction

Attempt 3; Proved in the stated scope

On full row rank of $U = [S; \mathbf{1}^\top]$, a seed of coordinate floor a and residual at most $ac_u\sqrt{n}/2$ becomes an exact nonnegative normalized certificate on $\sigma_{\min}(U) \ge c_u\sqrt{n}$.

Proof strategy

1. Solve the moment correction using the augmented matrix's right inverse.
2. The singular-value bound keeps the corrected multipliers nonnegative.

A finite probability bridge

Attempt 3; Proved in the stated scope

With the floor, cutoff and primitive matrix defined above,

$$p_n(m) \le 1 - \mathbb{E}_G f(G) + |\mathbb{E}_C f(C) - \mathbb{E}_G f(G)| + \kappa_{n,m} + \Delta_{n,m},$$

where both conditioning and collision errors tend to zero at fixed load.

Proof strategy

1. Bound storage by the complement of the exactified cube cutoff test.
2. Add the Gaussian comparison, conditioning and distinctness errors.

Remove Gaussian column radii

Attempt 3; Proved in the stated scope

For independent normalized Gaussian columns $\widetilde{G}$,

$$|\mathbb{E}_G f(G) - \mathbb{E} f(\widetilde{G})| = O_{u,\psi}(\sqrt{\log n/n}).$$

Positive invertible column scaling also preserves exact floor-zero occurrence samplewise.

Proof strategy

1. Normalize all Gaussian primitive columns and control their extreme radii.
2. Compare the entire cutoff test before and after this column scaling.

A restricted Gaussian construction is excluded

Attempt 3; Proved in the stated scope

For fixed $1 < u < 2$, $1 < C_0 < 1 + 1/u$, define

$$\rho_{C_0}(G) = \min_{\lambda \in \mathcal{L}_a,\ N\|\lambda\|_2^2 \le C_0} \|A(G)^\top \lambda\|_2.$$

For some fixed $k_{u,C_0} > 0$,

$$\Pr_G\{\rho_{C_0}(G) < k_{u,C_0}/\sqrt{n}\} = O_{u,C_0}(n^{-1}).$$

Since $\rho(G) \le \rho_{C_0}(G)$, this leaves full-domain occurrence open.

Proof strategy

1. Lower-bound the residual on the additional bounded-norm multiplier set.
2. The unrestricted minimum is smaller, so its occurrence remains open.

Two interaction correlations vanish

Attempt 3; Proved in the stated scope

At $u = 33/20$, write each Gaussian primitive column as $G_{\cdot\alpha} = \sqrt{V_\alpha}\,\Omega_\alpha$, and put

$$F_j = \sum_{\alpha,\mu} H_j(G_{\mu\alpha}), \qquad W_j = F_j - \mathbb{E}(F_j \mid V) - \mathbb{E}(F_j \mid \Omega), \quad j \in \{4, 6\}.$$

Here $H_4(x) = x^4 - 6x^2 + 3$, $H_6(x) = x^6 - 15x^4 + 45x^2 - 15$, and $L = \|\psi'\|_\infty$. For the same cutoff test f,

$$|\mathbb{E}(fW_4)| = O(L \log^3 n/\sqrt{n}), \qquad |\mathbb{E}(fW_6)| = O(L \log^4 n/\sqrt{n}).$$

Pure-direction terms, both paired cubic families and all retained higher degrees still require a bound.

Proof strategy

1. Subtract radial and directional conditional means from the Hermite sums.
2. Combine radial control with the cutoff bound for degrees four and six.

Supply the two probability premises

Attempt 3; Proof conditioned on Gap 3 and Gap 4

Conditional on Gap 3 and Gap 4 at $u = 33/20$,

$$\alpha_+ \le 33/20 < 17/10.$$

At $u = 9/5$, the same premises give only the compatible bound $\alpha_+ \le 9/5$.

Proof strategy

1. Insert Gaussian occurrence and the angular comparison into the bridge.
2. Both premises must hold at the same fixed load.

Q3 · Dual certificates

Q3. Which dual certificates force nonstorage?

D1. Exact finite primal–dual alternative

Attempt 1; Proved in the stated scope

For integers $n \geq 2, m \geq 1$, let $s^1, \ldots, s^m \in \{-1, 1\}^n$. These spin patterns are not jointly storable if and only if there are real multipliers

$$\lambda_i^\mu \geq 0, \qquad \sum_{i=1}^{n} \sum_{\mu=1}^{m} \lambda_i^\mu = 1,$$

such that

$$\sum_{\mu=1}^{m} \lambda_i^\mu s_i^\mu = 0 \quad (1 \leq i \leq n),$$

$$\sum_{\mu=1}^{m} (\lambda_i^\mu + \lambda_j^\mu) s_i^\mu s_j^\mu = 0 \quad (1 \leq i < j \leq n).$$

When such a witness exists, one exists with at most $n(n+1)/2+1$ positive entries. The allowed networks have symmetric real weights, zero diagonal, and free fields.

Proof strategy

1. Apply the strict linear alternative to all shared-edge margin constraints.
2. Normalize the dual; convex-hull sparsification bounds its positive entries.

D2. Coupled moment equations

Exact alternative established in Attempt 1; coupled moment forms established in Attempt 2; Proved in the stated scope

Let $S = [s^1 \ \cdots \ s^m]$, $\Lambda = (\lambda_i^\mu) \geq 0$, $B = \Lambda \odot S$, and $t = \Lambda \mathbf{1}$. Nonstorage is equivalent to existence of such multipliers satisfying

$$B\mathbf{1} = 0, \qquad BS^\top = \operatorname{diag}(t) + K, \qquad K^\top = -K, \qquad \mathbf{1}^\top t = 1.$$

This is a finite equivalence; certificate occurrence at a fixed load remains unestimated.

Proof strategy

1. Collect each field equation in the row sums of the signed multipliers.
2. Collect each shared-edge equation in the symmetric part of the moment matrix.

D3. Extensive pattern support

Attempt 1; Proved in the stated scope

Let $n \geq 8$, $1 \leq m \leq 2n$, and let D be a uniform distinct m-subset of $\{-1,1\}^n$. Set

$$r = \lfloor n/8 \rfloor, \qquad H(t) = -t\log t - (1-t)\log(1-t), \qquad c_0 = \tfrac{7}{8}\log 2 - 2H(1/16) > 0.$$

Then

$$\Pr\{\exists E \subseteq D : \ |E| \leq r, \ E \text{ nonstorable}\} \leq \frac{2ne^{-c_0 n}}{1 - \binom{r}{2}2^{-n}}.$$

For a normalized nonnegative dual witness satisfying the equations on the Exact duality page, define its active pattern set by

$$T(\lambda) = \{\mu : \ \sum_{i=1}^{n} \lambda_i^{\mu} > 0\}.$$

Every such active set is itself nonstorable. Thus, for every fixed $0 < \alpha \leq 2$ at $m = \lfloor \alpha n \rfloor$, with probability tending to one every nonstorage witness has $|T(\lambda)| > \lfloor n/8 \rfloor$. This assertion is vacuous when the full sample is storable.

Proof strategy

1. A certificate's active patterns themselves form a nonstorable subset.
2. Union-bound projection rank failures over every sufficiently small subset.

D4. Extensive neuron support

Neuron-support bound established in Attempt 2; the earlier pattern-support bound is from Attempt 1; Proved in the stated scope

For each fixed $0 < u < 2$, some $\beta(u) > 0$ satisfies, at $m = \lfloor un \rfloor$,

$$\Pr\{\exists \text{ normalized exact nonnegative dual supported on at most } \lfloor \beta(u)n \rfloor \text{ neurons}\} \longrightarrow 0.$$

This excludes small neuron supports and leaves the probability of an unrestricted certificate open.

Proof strategy

1. Reduce a fixed active-neuron set to the required random convex-hull event.
2. Apply its discrete probability bound and sum over small neuron sets.

D5. Pairwise relaxation is insufficient

Attempt 1; Proved in the stated scope

For distinct spins $s^1, \dots, s^m \in \{-1, 1\}^n$, put $d^{\mu\nu} = s^\mu - s^\nu$. Joint storage implies that some symmetric zero-diagonal real matrix J satisfies

$$(d^{\mu\nu})^\top J d^{\mu\nu} > 0 \qquad (1 \le \mu < \nu \le m).$$

Failure of the pairwise system is equivalent to the existence of real coefficients

$$c_{\mu\nu} \ge 0, \qquad \sum_{\mu<\nu} c_{\mu\nu} > 0, \qquad \sum_{\mu<\nu} c_{\mu\nu} d^{\mu\nu} (d^{\mu\nu})^\top \text{ is diagonal.}$$

Such coefficients certify nonstorage. Satisfying the pairwise inequalities does not certify storage: a six-pattern example in dimension five satisfies all pairwise inequalities but admits no storing network. For every fixed $\alpha > 2$, a uniform distinct sample of $m = \lfloor \alpha n \rfloor$ patterns is, with probability tending to one, nonstorable while satisfying the pairwise inequalities.

Proof strategy

1. Add opposite-margin inequalities to eliminate fields and obtain pair tests.
2. Use exact counterexamples and the random construction to disprove sufficiency.

Prefix route: setup

Use a uniformly ordered distinct sample and write D_r for its first r patterns, $A_r = A(D_r)$, and $Q_r = F_n(r)^c$.

$$m = \lfloor 33n/20 \rfloor, \qquad k = m - 1, \qquad \Delta_n = p_n(k) - p_n(m) = \Pr(Q_m, \neg Q_k).$$

$$\varphi_r = \max_{\lambda \ge 0,\, \mathbf{1}^\top \lambda = 1,\, A_r^\top \lambda = 0} \min_{i,\mu} \lambda_{i\mu},$$

with $\varphi_r = 0$ if the dual set is empty; Q must therefore be retained when describing dual boundaries.

$$\sigma_{\min}([S_k; \mathbf{1}_k^\top]) \ge c\sqrt{n} \quad \text{with probability tending to one,} \qquad c > 0.$$

$$a_n = \frac{n^{-4}(1 + \sqrt{n}/c)}{1 - n^{-3}} \sim c^{-1} n^{-7/2}.$$

The birth increment Δ_n is different from the defect variable $\Delta(D)$ and the collision bound $\Delta_{n,m}$.

Prefix boundary and thinness

Attempt 3; Proved reduction; occurrence remains open

For the full positive thin child event,

$$\Pr\{Q_m, 0 < \varphi_m < n^{-4}\} \le \Delta_n + \Pr(Q_k, \varphi_k = 0) + \Pr(Q_k, 0 < \varphi_k < a_n) + o(1).$$

Proof strategy

1. Extend a positive-floor prefix dual across every possible next pattern.
2. Split thin children into births, boundary prefixes and thin prefixes.

Uniform extension

Attempt 3; Proved in the stated scope

For a deterministic prefix with augmented least singular value $\sigma>0$, every exact normalized dual of positive floor a extends across every cube pattern with floor at least

$$t(a,\sigma)=\frac{a}{1+na+n/\sigma}.$$

Proof strategy

1. Add mass at the new pattern and correct the reciprocal moments.
2. Use the prefix singular value to bound the loss of the coordinate floor.

Split the full thin child event

Attempt 3; Proved in the stated scope

On conditioning, $\varphi_k \geq a_n$ excludes $0<\varphi_m<n^{-4}$. The remaining cases are birth, prefix boundary and positive prefix thinness.

Proof strategy

1. The extension bound excludes a thin child above the prefix floor threshold.
2. Partition all remaining cases and add the conditioning error.

Three unproved prefix limits

Attempt 3; Proof conditioned on Gap 7, Gap 8 and Gap 9

Conditional on Gap 7, Gap 8 and Gap 9,

$$\Pr\{Q_m, 0<\varphi_m<n^{-4}\}\to 0.$$

Exact-dual occurrence still needs its own probability theorem.

Proof strategy

1. Insert the three missing limits into the proved event decomposition.
2. This removes thin duals; it still does not prove that any dual occurs.

Q4 · Endpoint equality

Q4. Do α_- and α_+ coincide?

Uniformly order all cube vertices; D_m denotes the first m and τ_n the first nonstorable prefix length.

$$K_n(t)=\min\{m : p_n(m)\leq t\},\quad 0<t<1,\qquad b_n=K_n(1/2).$$

The filed median symbol is renamed b_n here, to distinguish it from the conditional-failure function $c_n(u)$.

For the auxiliary product law include each of $M=2^n$ cube vertices independently with probability p.

$$D_p\subseteq\{-1,1\}^n,\qquad \mu_n(p)=\Pr(D_p \text{ is nonstorable}).$$

Conditionally on its cardinality, D_p is a uniform distinct set; p is an inclusion probability, distinct from $p_n(m)$.

Use $\inf\varnothing=+\infty$ for the median-derivative assertion and a zero supremum for an empty boost family.

T1. Defect concentration and conditional sharpness

Attempt 1; Proved in the stated scope

For a uniform distinct m-subset $D \subseteq \{-1,1\}^n$, $1 \le m \le 2^n$, define

$$L(D) = \max\{|E| : E \subseteq D,\ E \text{ storable}\}, \qquad \Delta(D) = m - L(D).$$

For either $Z = L(D)$ or $Z = \Delta(D)$, and every $t \ge 0$,

$$\Pr\{|Z - \mathbb{E}Z| \ge t\} \le 2e^{-2t^2/m}.$$

Now uniformly order the entire cube, let D_m be its first m patterns, and set

$$\tau_n = \min\{m : D_m \text{ nonstorable}\}, \qquad p_n(m) = \Pr\{\tau_n > m\}.$$

A sharp proportional threshold at $c > 0$ is equivalent to

$$\forall \varepsilon > 0 : \quad \Pr\{|\tau_n/n - c| > \varepsilon\} \longrightarrow 0.$$

The displayed concentration bound for the maximum storable subset does not establish this last limit; concentration of the first failure time remains open.

Proof strategy

1. A single replacement changes the largest storable subset size by at most one.
2. Use sampling-without-replacement concentration and the prefix survival law.

T2. Window, center and endpoint equivalences

Endpoint, window and center equivalences established in Attempt 2; the network probability premises remain open; Proved in the stated scope

Uniformly order the cube, let τ_n be its first nonstorable prefix length, and set $b_n = \min\{m : p_n(m) \le 1/2\}$. For every fixed $c \in [1,2]$,

$$\alpha_- = \alpha_+ = c \quad \Longleftrightarrow \quad \tau_n/n \xrightarrow{\Pr} c,$$

$$\Longleftrightarrow \quad (\tau_n - b_n)/n \xrightarrow{\Pr} 0 \quad \text{and} \quad b_n/n \to c.$$

Neither the shrinking window nor the limiting center is established for this network.

Proof strategy

1. Express storage probabilities as the survival function of first failure.
2. Use monotonicity to separate concentration around a median from its limit.

Bounded plantings and the window

Attempt 3; Conditional implication

Include each of the $M = 2^n$ cube vertices independently with probability p, obtaining D_p, and let $\mu_n(p) = \Pr(D_p \text{ is nonstorable})$. This auxiliary product law conditional on cardinality is uniform distinct sampling. For fixed $K \geq 2$ and finite U, define

$$\mathcal{M}_{n,U} = \{p : \mu_n(p) = 1/2,\ Mp \leq Un\},$$

$$\beta_{n,K}(p) = \sup_{|H|\leq K,\ H \text{ storable}} [\Pr(D_p \cup H \text{ nonstorable}) - \mu_n(p)].$$

Use zero for an empty supremum. Conditional on Gap 6 for every fixed K,

$$\inf_{p\in\mathcal{M}_{n,U}} p\mu_n'(p) \to \infty.$$

The shrinking fixed-size window and limiting center are separate obligations.

Proof strategy

1. A bounded median derivative forces a bounded planting with positive boost.
2. Exclude all such boosts to force derivative growth at the specified medians.

Small intrinsic witnesses are rare

Attempt 3; Proved in the stated scope

For fixed K≥2, n>2K, the probability of a nonstorable subset of size at most K is at most

$$(Mp)^2 e^{-n/(8K^2)}$$

under Bernoulli inclusion, capped at one.

Proof strategy

1. Count bounded nonstorable subsets using the deterministic obstruction bound.
2. Sum their inclusion probabilities under the product sampling law.

The median-booster implication

Attempt 3; Proved in the stated scope

If $p\mu_n'(p)$ stays bounded at these medians, some bounded storable H has a positive planting boost.

Proof strategy

1. Apply the small-witness-or-booster criterion to the monotone failure event.
2. Rare intrinsic witnesses leave the bounded storable planting alternative.

Width and center still need separate estimates

Attempt 3; Proof conditioned on Gap 10 and Gap 11

Conditional on Gap 10 and Gap 11,

$$\alpha_- = \alpha_+ = c.$$

Median derivative growth alone gives neither premise.

Proof strategy

1. Combine a shrinking fixed-size window with convergence of its center.
2. The median derivative statement supplies neither of those two limits.

Q5 · Finite evidence

Q5. What do the finite classifications establish?

E1. Exact certification and finite uncertainty

Attempt 1; Proved in the stated scope

For a nonempty sample $s^1, \dots, s^m \in \{-1, 1\}^n$, let $d = n(n+1)/2$ and $\theta = ((J_{ij})_{i<j}, (h_i)_{i=1}^n) \in \mathbb{R}^d$. Define $A \in \mathbb{R}^{mn \times d}$ by

$$(A\theta)_{(\mu,i)} = s_i^\mu \left(\sum_{j \neq i} J_{ij} s_j^\mu + h_i \right), \qquad J_{ji} = J_{ij}, \quad J_{ii} = 0.$$

Then

$$\rho_* := \max\{\rho : A\theta \geq \rho \mathbf{1},\ |\theta_j| \leq 1,\ 0 \leq \rho \leq 1\} > 0 \quad \iff \quad D \text{ storable.}$$

Here $D = \{s^\mu : 1 \leq \mu \leq m\}$. A saved integer vector $z \in \mathbb{Z}^d$ with $\min(Az) > 0$, checked in exact arithmetic, certifies storage; an uncertified numerical result remains unknown.

For N independent uniform distinct-set trials, let K count valid storage certificates. For $0 < \eta < 1$, define $L_\eta(0) = 0$ and, for $1 \leq k \leq N$, the unique root in $(0, 1)$ of

$$\sum_{j=k}^{N} \binom{N}{j} L_\eta(k)^j (1 - L_\eta(k))^{N-j} = \eta.$$

Then $\Pr\{L_\eta(K) \leq p_n(m)\} \geq 1 - \eta$. Unknown trials supply no upper bound below one, and finite trials do not establish an asymptotic threshold.

Proof strategy

1. Rescale a strict feasible vector to the bounded-margin optimization.
2. Check integer certificates exactly and invert a binomial tail for uncertainty.

E2. Historical 168-case census

Attempt 2 recertified the same 168 samples from Attempt 1; no new trials were added; Proved in the stated scope

Exact primal or dual witnesses classify the 168 historical samples:

96 stored, 67 local obstructions, 5 symmetry-only obstructions, 0 unresolved.

Attempt 1 had certified 96 stored samples and retained the other 72 as unknown; these finite classifications do not locate an asymptotic endpoint.

Verification strategy

1. Reuse the original 168 cases and check saved primal or dual witnesses.
2. Separate local row obstructions from failures requiring shared-edge duals.

E3. New 120-case census

Attempt 3; Proved in the stated scope

The sample comprises $n \in \{64, 96, 128\}$, eight streams per dimension, and five nested loads

$$1.5, 1.6, 1.7, 1.75, 1.8$$

: 24 streams and 120 cases.

$$43 \text{ stored}, \qquad 77 \text{ nonstored}, \qquad 0 \text{ unknown}.$$

The 77 nonstored cases comprise 61 local obstructions and 16 symmetry-only full duals; prefixes within a stream are dependent.

Verification strategy

1. Check exact primal, local-dual or full-dual witnesses for each sampled prefix.
2. Count symmetry failures only among cases passing every row test.

Exact cohorts and conditional experiment

For each case, numerical optimization proposes a witness; exact arithmetic accepts a primal $A_R z > 0$, a normalized full dual $A_R^\top \lambda = 0$, or an analogous local-row dual.

$$\widehat{p}_n(m) = \frac{\#F}{T}, \qquad \widehat{c}_n(u) = \frac{\#(E \cap F^c)}{\#E} \quad (\#E > 0),$$

where T is the number of sampled sets at that condition; m is the number of memories in each set, and a successful trial stores all m.

Table 2: Exact classifications; historical cases are counted once.

Cohort	Stored	Local failure	Symmetry-only
Attempts 1–2: 168	96	67	5
Attempt 3: 120	43	61	16

The historical cohort used dimensions $n \in \{16, 32, 48, 64\}$; its 168 cases were recertified in Attempt 2, rather than resampled.

Attempt 3 used eight seeded streams at each $n \in \{64, 96, 128\}$, with distinct-pattern prefixes at $u \in \{1.5, 1.6, 1.7, 1.75, 1.8\}$.

$$\text{seed} = 202611080100000 + 1000n + r, \qquad r \in \{0, \dots, 7\}.$$

Each stream rejects repeated cube patterns; its five prefixes are dependent, and row tests within a case are dependent.

Table 3: Attempt 3: symmetry-only failures among row-feasible cases at load 1.7.

Neurons	Row feasible	Symmetry failures
64	1	1
96	4	4
128	7	7

In the original 120-case cohort, all three conditional denominators at load 1.8 were zero.

Partial overnight snapshot

Partial results frozen at 8:15 AM Chicago, 9 November 2026.

$$T = 32 \quad \text{target trials per } (n, u), \qquad \frac{\#F}{T} \le \widehat{p}_n(m) \le \frac{\#F + U_F}{T}.$$

Here U_F counts unresolved or unattempted storage trials; bars are computational bounds, not confidence intervals.

New certificates are solver-checked; this snapshot has no fresh independent numerical replay.

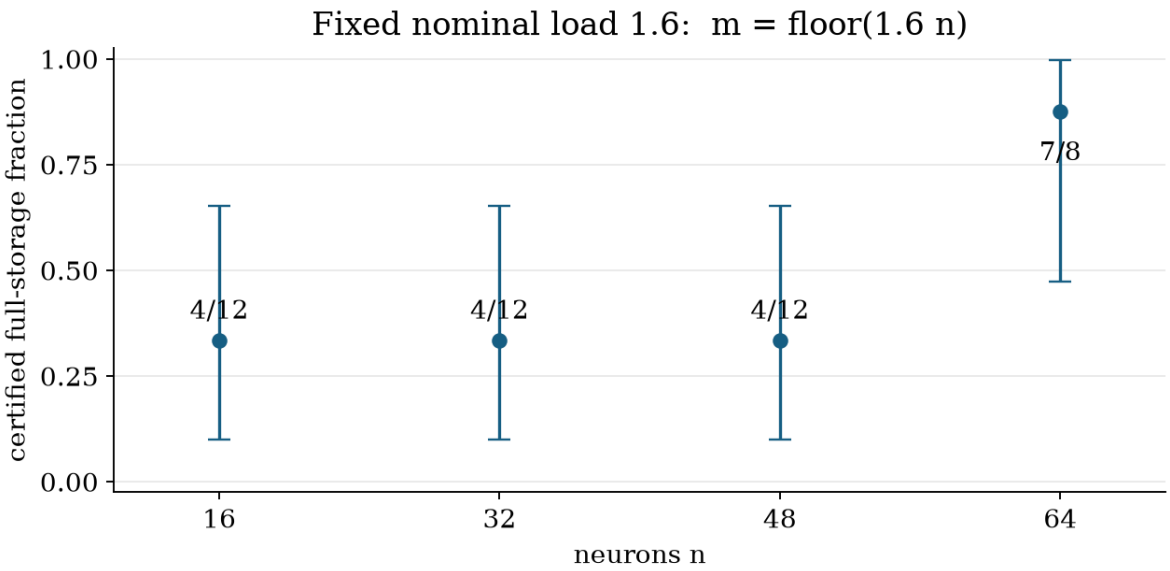


Figure 1: Historical cohort at nominal load 1.6; exact pointwise 95% binomial intervals assume independent ideal uniform-set trials.

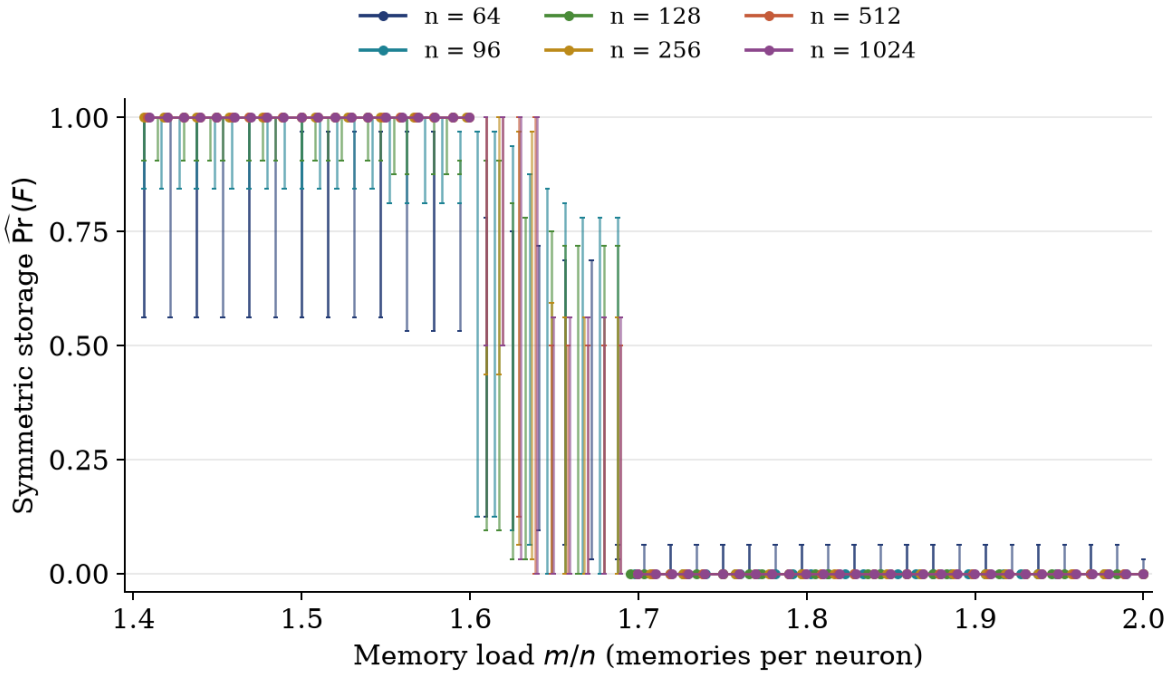


Figure 2: Current symmetric-storage fractions at the six original sizes; bars bound unresolved trials.

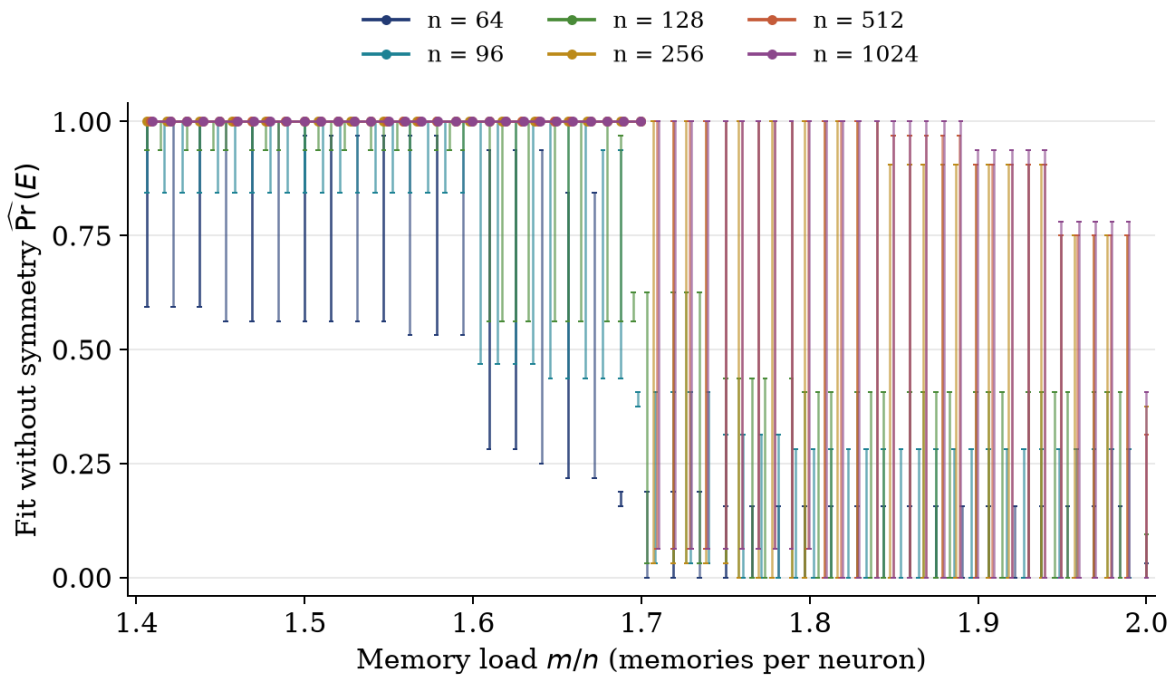


Figure 3: Current fit-without-symmetry fractions at the six original sizes; bars bound unresolved trials.

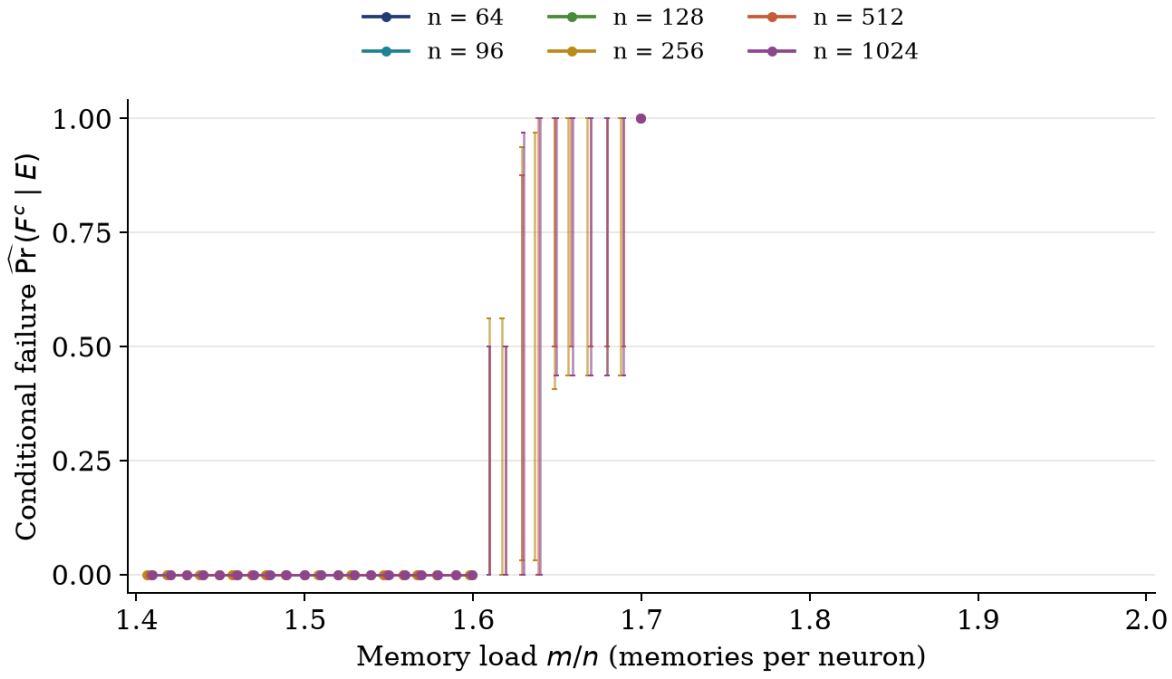


Figure 4: Current conditional symmetry failures at the six original sizes; unresolved or zero denominators are omitted.

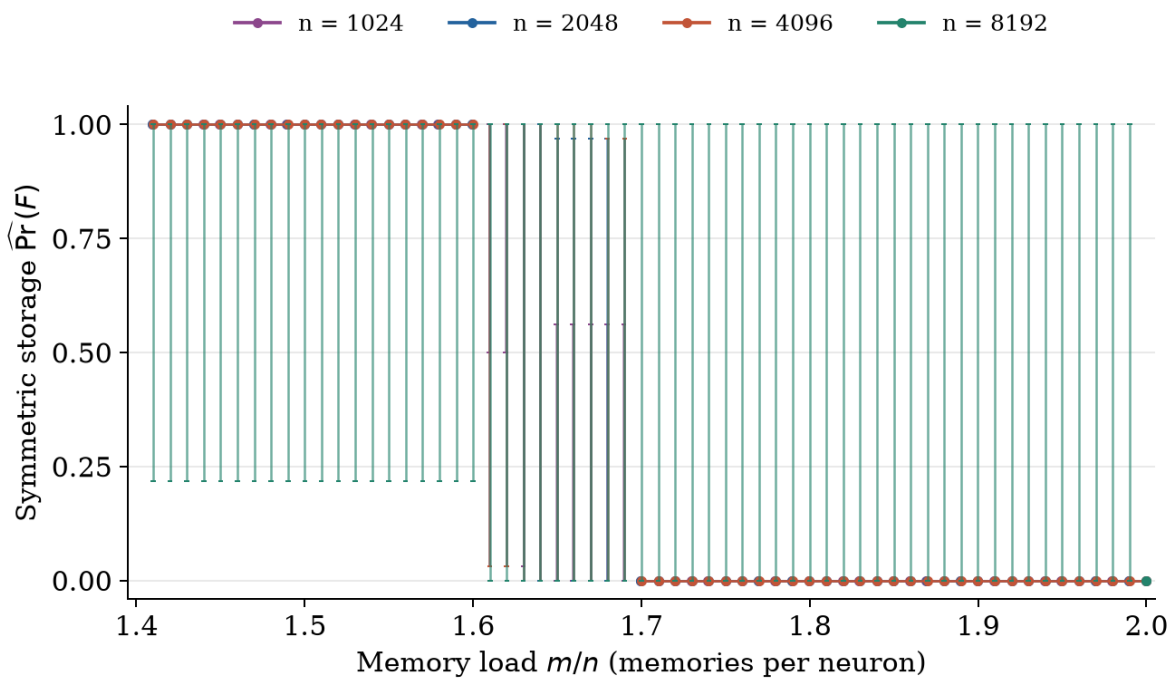


Figure 5: Partial storage fractions at n = 1024, 2048, 4096, 8192; bars retain unresolved trials.

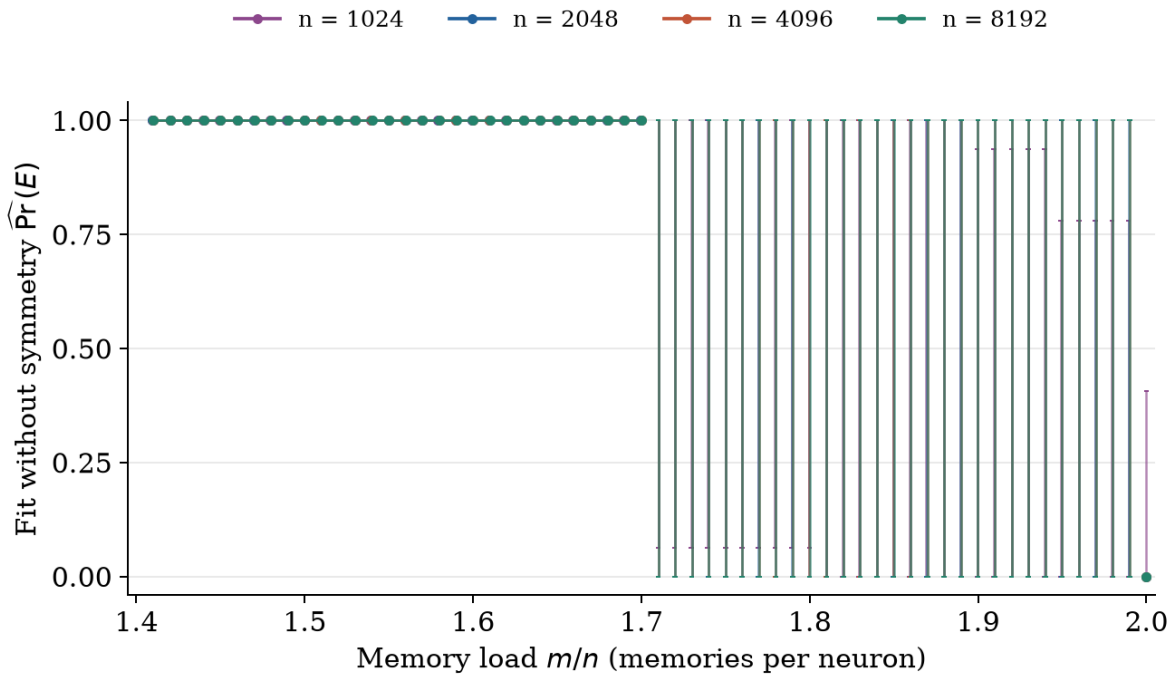


Figure 6: Partial fit-without-symmetry fractions at the four largest sizes; bars retain unresolved trials.

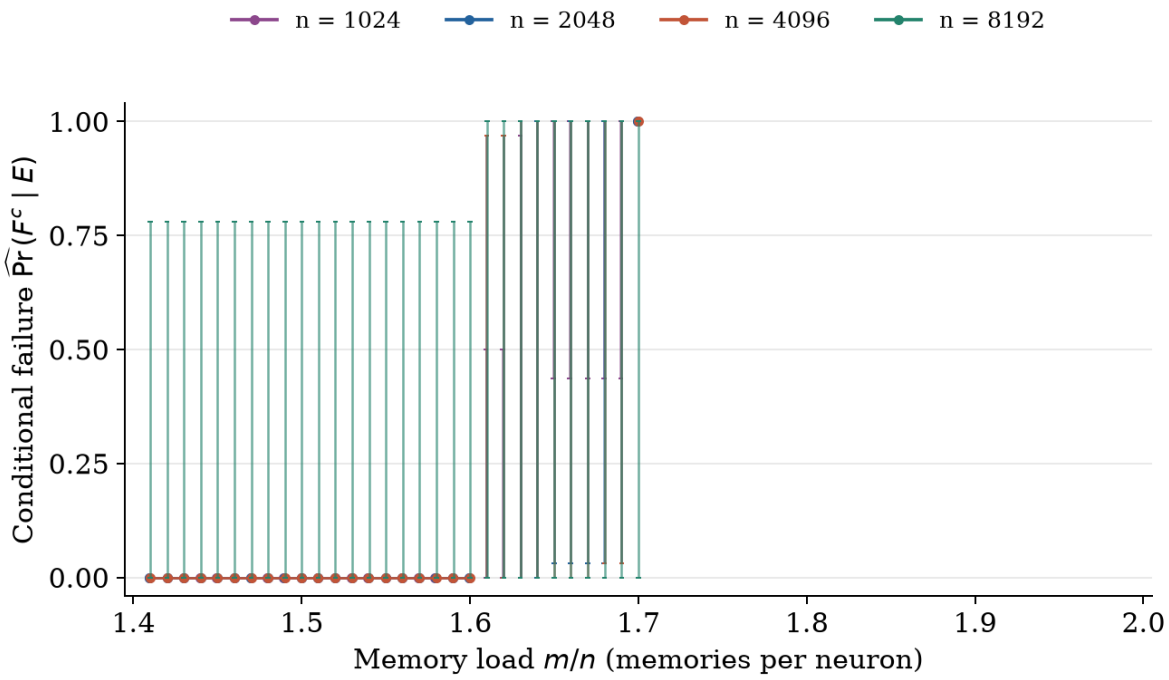


Figure 7: Partial conditional symmetry failure at the four largest sizes; unresolved or zero denominators are omitted.

The plotted intervals apply to the historical fixed condition; the Attempt 3 dots are observed fractions, and none of the plots establishes an asymptotic endpoint.

Extended sweep

The extension targets 32 seeded streams at each $n \in \{96\} \cup \{2^k : k = 6, \ldots, 13\}$, on a requested-load grid from 1.4 to 2 in steps of 0.01; prefixes within a stream remain dependent.

MLX proposes fits or full duals; original integer margins or verified residual enclosures decide acceptance, and timeouts remain unresolved.

Table 4: Partial full-storage success bounds / 32 targeted trials.

Neurons	Target u = 1.6	Target u = 1.7
64	[17,26]/32	[1,2]/32
96	[26,31]/32	0/32
128	[28,29]/32	0/32
256	32/32	0/32
512	32/32	0/32
1024	32/32	0/32
2048	32/32	0/32
4096	32/32	0/32
8192	[7,32]/32	[0,32]/32

Reading the partial plots

The run is stopped; completed results are saved and restarting requires an explicit action.

Q6. What do the bars mean?

For symmetric storage at one condition, let S count confirmed successes, D confirmed failures, and U unresolved or unattempted trials.

$$S + D + U = T.$$

$$\text{bar} = \left[\frac{S}{T}, \frac{S+U}{T}\right].$$

The bottom counts only confirmed successes; the top allows every unknown trial to succeed.
A solid dot means U = 0 for the displayed fraction; these bars are computational bounds, not confidence intervals.
At requested load 1.64,

$$n = 256, \qquad m = \lfloor 1.64 \cdot 256 \rfloor = 419.$$

$$m/n \approx 1.6367.$$

$$S = 1, \qquad D = 1, \qquad U = 30.$$

$$\left[\frac{1}{32}, \frac{31}{32}\right] = [0.03125, 0.96875].$$

Only 2 storage outcomes are resolved: 1 success and 1 failure; 30 remain unknown, including searches not yet attempted at this prefix.
All 32 trials already pass E here, so the uncertainty concerns the symmetry requirement.

Q7. How is a trial resolved?

A search proposes a network or an obstruction; only a checked certificate, or a valid implication from one, resolves the trial.
With A the signed constraint matrix for symmetric storage, the two accepted outcomes are

$$A\theta > 0 \quad \Longrightarrow \quad F \text{ and } E,$$

$$\lambda \geq 0, \quad \lambda \neq 0, \quad A^{\top}\lambda = 0 \quad \Longrightarrow \quad F^{c}.$$

Success requires strictly positive integer-checked margins at every original constraint; failure requires an exact dual or a verified enclosure proving one exists.
For E, all n separate neuron fits must pass; one checked local obstruction certifies E failure and therefore F failure.
Timeouts, rejected certificates and unattempted cases remain unknown; an optimizer failing to find a network is insufficient to certify failure.

Q8. Why are the high-load storage points resolved?

Each trial uses nested prefixes of the same sampled stream: adding memories cannot restore feasibility.

$$m_1 \le m_2 :$$
$$F(D_{m_2}) \Longrightarrow F(D_{m_1}),$$
$$F(D_{m_1})^c \Longrightarrow F(D_{m_2})^c.$$

For n = 256, all 32 streams have checked storage failures by

$$m = 435, \qquad m/n \approx 1.6992.$$

Their larger prefixes therefore have S = U = 0: a resolved dot at zero, reusing the same 32 stream certificates across loads.
Near the transition, a success at a smaller prefix and a failure at a larger one leave the intermediate prefix unknown until a certificate or implication settles it.
Resolution is specific to the displayed test: at n = 256 and u = 2, the separate E test still has 12 unresolved trials.
Conditional points are omitted when the E denominator is zero or unresolved; when the denominator is known, remaining symmetry uncertainty is shown by a bar.

Certificate checks Verified dual enclosures

Open obligations by question

All twelve obligations below are from Attempt 3; 17/10 is its historical test value, rather than the report's main goal.

Gap 2 · Q1: Lower endpoint direction

For every fixed $0 < u < 17/10$, prove

$$\limsup_{n\to\infty} p_n(\lfloor un \rfloor) > 0.$$

This is recorded obligation G2; storage through a sublinear excess above n gives no fixed load above one.

Gap 5 · Q1: Trained graph rank

Decide which alternative holds:

$$\exists R < \infty : \limsup_{n} \Pr\{\widehat{E}, \text{consistent}, r_n \le R\} > 0,$$

or

$$\forall R < \infty : \Pr\{\widehat{E}, \text{consistent}, r_n \le R\} \to 0.$$

Recorded G20. The first alternative gives positive inverse-Q survival limsup and $\alpha_+ \ge 9/5$; the second retires this learner. Neither is established. The G21 first/second-moment criterion is only a sufficient route to the second alternative.

Gap 12 · Q1: Full weighted primal survival

Prove

$$\limsup_{n} \Pr\{E, \gamma_* > 0\} > 0.$$

Recorded G5. This would yield $\alpha_+ \ge 9/5$ and strictly refute 17/10. Failure of one balanced dual seed or one chosen weighting does not decide this optimized event.

Gap 1 · Q2: Upper endpoint direction

For every fixed $v > 17/10$, prove

$$p_n(\lfloor vn \rfloor) \longrightarrow 0.$$

This is recorded obligation G1; the finite classifications and scoped construction obstructions do not prove this limit.

Gap 3 · Q2: Actual Gaussian occurrence

At a specified fixed load, prove

$$\mathbb{E}_G f(G) \longrightarrow 1.$$

Recorded G13(u) is open at both 33/20 and 9/5. A lower bound for a restricted minimum cannot lower-bound the full minimum.

Gap 4 · Q2: Cube versus spherical directions

At the same fixed load as occurrence, prove

$$\left|\mathbb{E}_C f(C) - \mathbb{E} f(\widetilde{G})\right| \longrightarrow 0.$$

Recorded G14(u). The entire-test radial normalization has vanishing error; the angular comparison is open. At 33/20, G19 is the equivalent full signed angular-plus-higher Hermite sum, with both paired cubic families retained.

Gap 7 · Q3: Exact prefix boundary

Prove

$$\Pr\{Q_k, \varphi_k = 0\} \longrightarrow 0.$$

Recorded G24 (S5.G2). Small-support exclusion leaves the dense boundary; the probability of that event remains unestimated.

Gap 8 · Q3: One-pattern birth increment

Prove

$$\Delta_n := \Pr(Q_m, \neg Q_k) = p_n(k) - p_n(m) \longrightarrow 0.$$

Recorded G23 (S5.G1). The exact expected repairing-pattern count is $m\Delta_n$. Bounds for one certificate support do not upper-bound the common support intersection over every certificate.

Gap 9 · Q3: Positive prefix thinness

Prove

$$\Pr\{Q_k, 0 < \varphi_k < a_n\} \longrightarrow 0.$$

Recorded G25 (S5.G3). The prefix threshold is larger than n^{-4}; the uniform extension theorem supplies an implication, without an occurrence estimate.

Gap 6 · Q4: All bounded storable plantings

For every fixed $K \geq 2$ and finite U, prove

$$\sup_{p \in \mathcal{M}_{n,U}} \beta_{n,K}(p) \longrightarrow 0,$$

or exhibit an actual bounded storable planting with positive boost limsup.

Recorded G22 (Open-B). The checked criterion gives median derivative growth if all boosts vanish. It does not give a probability-band estimate, a fixed-size window, a limiting center or an endpoint.

Gap 10 · Q4: A shrinking fixed-size window

For every fixed $0 < \eta < 1/2$, prove

$$K_n(\eta) - K_n(1-\eta) = o(n).$$

Recorded G9. A divergent repairing-count lower bound uniformly over the probability band (G10) would suffice. Positive liminf failure at a chosen fixed load (G11), together with this window, would bound α_+ there; positive limsup failure alone does not.

Gap 11 · Q4: Convergence of the center

Prove that some fixed c satisfies

$$K_n(1/2)/n \longrightarrow c.$$

Recorded G12. Together with Gap 10 this gives $\alpha_- = \alpha_+ = c$. Median derivative growth alone supplies neither limit nor a numerical center.

What would change an endpoint bound?

At a fixed load $u \in (1, 2)$:

$$\begin{aligned} q_n(u) \longrightarrow 1 \quad &\Longrightarrow \quad \alpha_+ \le u, \\ \limsup_n q_n(u) > 0 \quad &\Longrightarrow \quad \alpha_- \le u, \\ \limsup_n p_n(\lfloor un \rfloor) > 0 \quad &\Longrightarrow \quad \alpha_+ \ge u. \end{aligned}$$

Finite classifications, small-support exclusions and restricted-construction obstructions supply none of these asymptotic premises by themselves.