Matrix Product States for Efficient Representation of Quantum Many-Body Systems

How entanglement structure determines when MPS compresses well

Jasper Sands
Columbia University

April 30, 2026

Research Question and Thesis

Core question

How does entanglement structure determine the efficiency of Matrix Product State representations?

- A generic N-qubit pure state needs 2^N complex amplitudes.
- MPS replaces that dense tensor by a chain of low-rank tensors linked by bond dimensions χ_k.
- The real issue is not just *whether* an MPS exists; it is whether the required bond dimensions remain small.
- Thesis: **MPS is efficient when bipartite entanglement is structured and low-rank, and inefficient when many Schmidt modes are significant across most cuts.**

From Schmidt Decomposition to MPS

For a cut after site k,

$$|\psi\rangle = \sum_{\alpha=1}^{\chi_k} \lambda_\alpha^{(k)} |\phi_\alpha^{[1\ldots k]}\rangle \otimes |\phi_\alpha^{[k+1\ldots N]}\rangle.$$

Key consequences:

- χ_k is the Schmidt rank across that bipartition.
- The entanglement entropy is

$$S_k = -\sum_\alpha p_\alpha^{(k)} \log_2 p_\alpha^{(k)}, \qquad p_\alpha^{(k)} = |\lambda_\alpha^{(k)}|^2.$$

- More spread-out entanglement requires larger bond dimensions.

An MPS stores local tensors and the correlation data carried by the internal bond dimensions.

Implementation and Metrics

Sequential SVD algorithm

1. Reshape the wavefunction at cut k into a $(2^k) \times (2^{N-k})$ matrix.
2. Compute $\Psi = USV^\dagger$.
3. Store U as the next site tensor.
4. Propagate $SV^\dagger$ to the right.
5. Repeat until all N site tensors are constructed.

Metrics reported

- Entanglement entropy vs cut position
- Bond dimension profile and max bond dimension
- Parameter count relative to the dense 2^N vector
- Truncated-MPS fidelity
- L_2 reconstruction error

Experimental Design

State families

- Product baseline: $|00\dots0\rangle$
- GHZ: $\left(|00\dots0\rangle + |11\dots1\rangle\right)/\sqrt{2}$
- Haar-random states
- Open-boundary transverse-field Ising ground states

System sizes

- All four families: $N \in \{4, 6, 8, 10, 12, 14, 16\}$
- Haar averages use 5 random seeds per N

Representative profile plots

Entropy-vs-cut plots are shown at $N = 10$ for Product, GHZ, one Haar-random sample, and Ising ground states with $h \in \{0.5, 1.0, 1.5\}$.

Experimental Goals and Compression Study

Hypotheses

- Product states should be trivial to compress.
- GHZ should show nonzero entropy but low bond dimension.
- Haar-random states should approach volume-law behavior and resist compression.
- Ising ground states should remain more compressible because the entanglement is generated by a local Hamiltonian.

Truncation sweep

Sweep $\chi_{\max} \in \{1, 2, 4, 8, 16, 32\}$ and record:

- fidelity,
- L_2 reconstruction error,
- parameter-count ratio.

Entropy Profiles Reveal the Entanglement Pattern

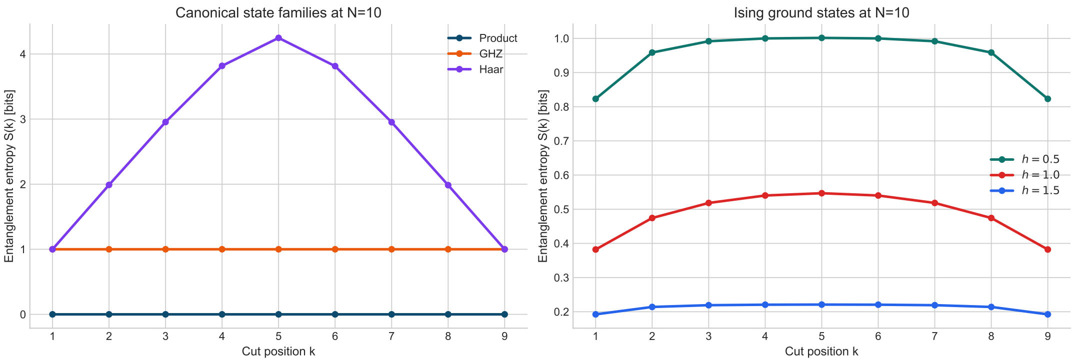

- Product states stay at zero entropy across every cut.
- GHZ is flat at 1 bit across internal cuts: global correlations, but only two Schmidt modes matter.
- Haar-random states show high entanglement nearly everywhere.
- The Ising chain stays much lower than Haar-random, and in this finite-size example the $h = 0.5$ case shows the largest peak entropy of the three fields shown.

Scaling Results: Exact MPS Gets Expensive When Entanglement Spreads

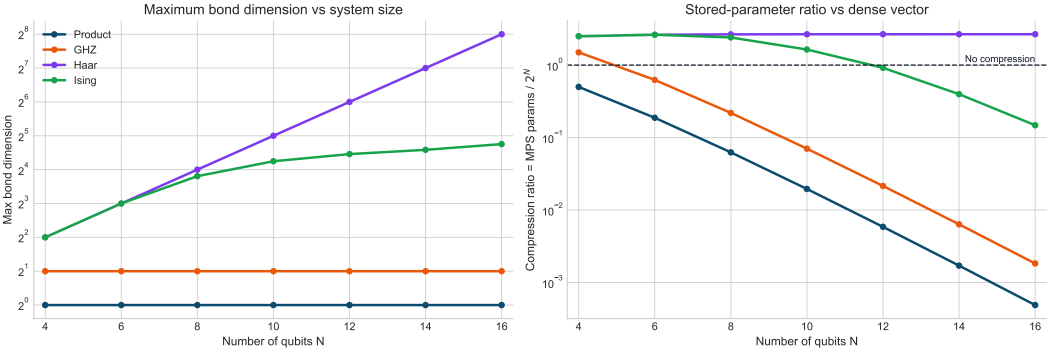

- Product states remain at $\chi = 1$, independent of system size.
- GHZ stays at $\chi = 2$, so its exact MPS grows only linearly in N.
- Haar-random states drive the bond dimension to the full internal rank across the whole sweep, and the exact MPS storage stays larger than the dense vector.
- Ising ground states still grow far more gently than Haar-random states; by $N = 16$ their exact MPS storage is below the dense-vector cost again.

Truncation Study: Compression Helps Structured States First

- GHZ becomes exact as soon as $\chi_{\max} = 2$.
- The Ising ground state at $N = 10, h = 1.0$ already reaches fidelity ≈ 0.997 at $\chi_{\max} = 2$.
- The Haar-random state still has fidelity only ≈ 0.483 at $\chi_{\max} = 8$, showing that low-rank truncation is much less effective there.

Representative Numbers at $N = 10$

State family	Max entropy (bits)	Max bond dim	Compression ratio
Product	0.0000	1	0.0195
GHZ	1.0000	2	0.0703
Haar (5-seed mean)	4.2895	32.0	2.6641
Ising $h = 0.5$	1.0015	22	2.0391
Ising $h = 1.0$	0.5469	22	2.0391
Ising $h = 1.5$	0.2209	20	1.8047

- The Haar-random state is the clearest failure mode: exact MPS is not a compression at all here.
- The Ising family sits in between the easy structured cases and the hard random case.

Answer to the Research Question

Main result

Entanglement structure controls MPS efficiency because the Schmidt spectrum at each cut determines the bond dimension required for an accurate representation.

- **Low-rank entanglement** leads to efficient MPS: product and GHZ states are the clearest examples.
- **Local Hamiltonians** can still generate nontrivial entanglement, but the structure is often compressible: that is what we saw for the Ising chain.
- **Broad, volume-law-like entanglement** makes MPS inefficient: that is what the Haar-random states demonstrate.

Limitations and Next Steps

Limitations

- Small-system exact diagonalization only
- One-dimensional open-boundary setting
- Exact MPS construction from a full state vector, not variational optimization

Extensions

- DMRG-style variational ground-state methods
- Time evolution with TEBD
- Stronger finite-size scaling near criticality
- Comparison with other tensor-network architectures

References

G. Vidal,
Efficient Classical Simulation of Slightly Entangled Quantum Computations,
Physical Review Letters, 91, 147902 (2003).

G. Vidal,
Efficient Simulation of One-Dimensional Quantum Many-Body Systems,
Physical Review Letters, 93, 040502 (2004).

R. Orús,
A Practical Introduction to Tensor Networks: Matrix Product States and Projected Entangled Pair States,
Annals of Physics, 349, 117–158 (2014).

Validation checks

- Product states reconstruct exactly with $\chi = 1$
- GHZ reconstructs exactly with $\chi = 2$
- Full-rank decompositions reproduce the state to numerical precision
- Truncated fidelity is monotone in $\chi_{\max}$
- Ising ground states remain normalized with finite entropies

Code interface

- `product_state`, `ghz_state`, `haar_random_state`
- `ising_ground_state`
- `mps_decompose`
- `reconstruct_from_mps`
- `entanglement_entropies`
- `bond_dimensions`
- `compression_stats`