



Entanglement and expressivity in neural quantum states

Nisarga Paul
California Institute of Technology

Interplay between Tensor Networks and Neural Networks
July 8th, 2026



Caltech

- I. Introduction
- II. Entanglement in NQS
- III. Expressivity and choice of basis in NQS

I. Introduction

II. Entanglement in NQS

III. Expressivity and choice of basis in NQS

Variational approach for the quantum many body problem

Quantum many-body problem:


- Intractable in general due to exponential dimension of many-body Hilbert space

Variational approach:

- Restrict to a polynomially small variational subspace $|\Psi_\theta\rangle$
- Apply the variational principle, minimize E_θ :


$$E_\theta = \frac{\langle \Psi_\theta | H | \Psi_\theta \rangle}{\langle \Psi_\theta | \Psi_\theta \rangle} \geq E_0$$

- Physically motivated choices include: Slater determinants (Hartree-Fock), matrix product states / tensor networks (DMRG)



Neural quantum state variational approach

- Success of ML has inspired a new variational wavefunction which parameterizes quantum states using neural networks (Carleo & Troyer, 2017)
- Simplest incarnation:




RESEARCH ARTICLE

MANY-BODY PHYSICS

Solving the quantum many-body problem with artificial neural networks

Giuseppe Carleo^{1*} and Matthias Troyer^{1,2}

- Estimating and minimizing energy (gradients): (Markov chain) Variational Monte Carlo
- **Great empirical success**




Neural quantum states: working definition

Example: multilayer perceptron (MLP) NQS

Definition: We consider pure states on n spin degrees of freedom $\mathbf{s} = (s_1, s_2, \dots, s_n)$

$$|\Psi\rangle = \sum_{\mathbf{s}} \Psi_{\theta}(\mathbf{s}) |\mathbf{s}\rangle$$

where $\Psi_{\theta}(\mathbf{s})$ is specified by a directed, acyclic computational graph (V, E) where $V \sim \text{nodes}$, $E \sim \text{edges}$



Value of each node $v \in V$ is defined recursively as

$$v(\mathbf{s}) = \sigma_v \left[\underset{\text{bias}}{b_v} + \sum_{(u,v) \in E} \underset{\text{previous nodes}}{w_{u,v}} u(\mathbf{s}) + \sum_i \underset{\text{inputs}}{w_{i,v}} s_i \right]$$

bias

previous nodes

inputs

Example of computational graph of NQS

where $\sigma_v = \text{scalar nonlinearity}$, $\{b_v, w_{u,v}, w_{i,v}\} = \theta$ (parameters).

What part of many-body Hilbert space do NQS capture?

[Yang,...,Preskill 2024],

[Sharir et al. 2021],




...

our focus: entanglement in NQS

Motivation: Area law constraint for tensor network states

- Entanglement entropy area law of a region A for MPS/TNS
- Are there analogous entanglement constraints on NQS?

Yes



Based on:

NP, *Bound on entanglement in neural quantum states*, Phys. Rev. Lett. 136, 120403 (2026)

I. Introduction


II. Entanglement in NQS

III. Expressivity and choice of basis in NQS


Entanglement in NQS: a simple numerical experiment

Numerical examples: SN-NQS, MLP-NQS, T-NQS


Single-nonlinearity NQS



Multilayer perceptron NQS




Transformer NQS




Entanglement in NQS: a simple numerical experiment

Numerical examples: SN-NQS, MLP-NQS, T-NQS


Single-nonlinearity NQS



Multilayer perceptron NQS




Transformer NQS




1. Randomly initialize n -qubit state
2. Form reduced density matrix of a subregion
3. Plot entanglement entropy


Entanglement in NQS: a simple numerical experiment



Subregion entanglement entropy, varying $|A|$, $n = 22$, averaged over random initializations



Fixed $|A| = n/2$, varying n , averaged over random initializations



Lesson from numerical experiment:

Log entanglement scaling seems generic for small NQS

Numerics details:

- SN-NQS: log-amplitudes, up to normalization, given by $\sigma(\mathbf{w} \cdot \mathbf{s} + b)$ where $\sigma = \sigma_1 + i\sigma_2$ and σ_1, σ_2 are real tanh nonlinearities
- MLP-NQS: vanilla neural network of width 3, depth 2, linear projections to the real and imaginary parts of the log-amplitude
- T-NQS: splits $\mathbf{s}$ into $O(1)$ patches, processes them using sparse self-attention, concatenates vectors, linearly projects onto real and imaginary parts of log-amplitude

Bound on entanglement in NQS: statement

Largely architecture-agnostic *logarithm law* bound on entanglement in NQS; scaling in n is tight.
This provides an NQS analog of the area-law constraint for MPS.

- **The bound.** Under certain well-behavedness conditions on the nonlinearities but otherwise general feed-forward architectures, we prove that the von Neumann entanglement entropy of any subregion A is bounded by

$$S_A = O(k \log n)$$

where $k = \#$ of elementary nonlinearities, $n = \#$ spins.

Bound on entanglement in NQS: statement

Largely architecture-agnostic *logarithm law* bound on entanglement in NQS; scaling in n is tight.
This provides an NQS analog of the area-law constraint for MPS.

- **The bound.** Under certain well-behavedness conditions on the nonlinearities but otherwise general feed-forward architectures, we prove that the von Neumann entanglement entropy of any subregion A is bounded by

$$S_A = O(k \log n)$$

where $k = \#$ of elementary nonlinearities, $n = \#$ spins.

- **Scaling is tight.** We show, analytically and numerically, the scaling in n is saturated by various NQS.
- **Corollary.** No NQS with $k \sim O(1)$ can reach volume law entanglement (which requires $S_A \sim O(n)$ for $|A| \sim n$).

This places a rare restriction on NQS, but also affirms their substantial expressive power.

Bound on entanglement in NQS: statement

Largely architecture-agnostic *logarithm law* bound on entanglement in NQS; scaling in n is tight.
This provides an NQS analog of the area-law constraint for MPS.

- **The bound.** Under certain well-behavedness conditions on the nonlinearities but otherwise general feed-forward architectures, we prove that the von Neumann entanglement entropy of any subregion A is bounded by

$$S_A = O(k \log n)$$

where $k = \#$ of elementary nonlinearities, $n = \#$ spins.

This gives (at least partial) explanation of the simple numerical experiment

- **Scaling is tight.** We show, analytically and numerically, the scaling in n is saturated by various NQS.
- **Corollary.** No NQS with $k \sim O(1)$ can reach volume law entanglement (which requires $S_A \sim O(n)$ for $|A| \sim n$).

This places a rare restriction on NQS, but also affirms their substantial expressive power.

Bound on entanglement in NQS: outline

Two steps:


- 1) **Feature reduction** from graph-theoretic arguments on the computational graph of the NQS
- 2) **Polynomial approximation** of the NQS and derivation of the bound from a close approximant

Step 1: Feature reduction

Structure of the NQS (again). We consider pure states on n spin degrees of freedom $\mathbf{s} = (s_1, s_2, \dots, s_n)$

$$|\Psi\rangle = \sum_{\mathbf{s}} \Psi(\mathbf{s}) |\mathbf{s}\rangle$$

where $\Psi(\mathbf{s})$ is specified by a directed, acyclic computational graph (V, E) where $V \sim$ nodes, $E \sim$ edges



Value of each node $v \in V$ is defined recursively as

$$v(\mathbf{s}) = \sigma_v \left[\underset{\text{bias}}{b_v} + \sum_{(u,v) \in E} \underset{\text{previous nodes}}{w_{u,v}} u(\mathbf{s}) + \sum_i \underset{\text{inputs}}{w_{i,v}} s_i \right]$$

bias


previous nodes

inputs

Example of computational graph of NQS

where $\sigma_v =$ scalar nonlinearity, $b_v, w_{u,v}, w_{i,v} =$ parameters.

Step 1: Feature reduction




Affine features.

- Define an *affine feature* to be a scalar $t(\mathbf{s})$ of the form $t(\mathbf{s}) = \mathbf{w} \cdot \mathbf{s} + b$
- We wish to write the NQS as a function of affine features: $\Psi(\mathbf{s}) = \mathcal{G}(t_1(\mathbf{s}), \dots, t_\mu(\mathbf{s}))$

Claim: we can achieve this with # of affine features $\mu \leq k + 1$, where $k = |V|$ is the # of nonlinear nodes.

Step 1: Feature reduction




Affine features.

- Define an *affine feature* to be a scalar $t(\mathbf{s})$ of the form $t(\mathbf{s}) = \mathbf{w} \cdot \mathbf{s} + b$
- We wish to write the NQS as a function of affine features: $\Psi(\mathbf{s}) = \mathcal{G}(t_1(\mathbf{s}), \dots, t_\mu(\mathbf{s}))$

Claim: we can achieve this with # of affine features $\mu \leq k + 1$, where $k = |V|$ is the # of nonlinear nodes.

Proof idea. A directed acyclic graph (V, E) has an ordering of nodes $O = (v_1, \dots, v_k)$ such that if $(v_i, v_j) \in E$ then v_i comes before v_j in O . We prove by induction on O .

v_1 must be of the form $v_1(t_1(\mathbf{s}))$. v_2 must be of the form $v_2(v_1, t_2(\mathbf{s})) = \widetilde{v}_2(t_1(\mathbf{s}), t_2(\mathbf{s}))$. Carrying on, $v_k(\mathbf{s})$ depends on k affine features, and $\Psi(\mathbf{s}) = \mathcal{G}(t_1(\mathbf{s}), \dots, t_{k+1}(\mathbf{s}))$ depends on at most one more.



Step 2: Polynomial approximation

We have that $\Psi(\mathbf{s}) = \mathcal{G}(t_1(\mathbf{s}), \dots, t_\mu(\mathbf{s}))$ for $\mu \leq k + 1$ affine features $t(\mathbf{s}) = \mathbf{w} \cdot \mathbf{s} + b$.

Well-behavedness. Some further assumption on $\Psi(\mathbf{s})$ is necessary. Different choices can be made, but we choose:

Assumption. $\mathcal{G}(t_1, t_2, \dots, t_\mu): \mathbf{R}^\mu \rightarrow \mathbf{C}$ is analytic in a complex neighborhood of each of its arguments, and bounded on this neighborhood, in a manner independent of n .

Colloquially, the nonlinearities must be analytic and bounded on a neighborhood of their domains, in a manner independent of n . This prohibits any single nonlinearity from hiding a complexity which grows with n .

Corollary: (Classical approx. theory) There exists a polynomial $P_d: \mathbf{R}^\mu \rightarrow \mathbf{C}$ of degree $\leq d$ in each variable which approximates $\mathcal{G}$ on its domain $\mathcal{D} \subset \mathbf{R}^\mu$ to $O(e^{-\gamma d})$:

$$\max_{\mathbf{x} \in \mathcal{D}} |\mathcal{G}(\mathbf{x}) - P_d(\mathbf{x})| < \beta e^{-\gamma d}$$

where $\beta, \gamma > 0$ are $O(1)$ constants.

Step 2: Polynomial approximation


Theorem 1. von Neumann entropy $S_A = -\text{Tr } \rho_A \log \rho_A$ of any subregion A satisfies $S_A = O(\mu \log n) = O(k \log n)$.

Proof ideas.

1. Bound the Schmidt rank (across any bipartition) of an auxiliary state $|\tilde{\Psi}\rangle = \mathcal{N}^{-1/2} \sum_{\mathbf{s}} P_d(t_1(\mathbf{s}), \dots, t_\mu(\mathbf{s})) |\mathbf{s}\rangle$.
2. Transfer the bound to $|\Psi\rangle$ using standard tricks

Analytical example: Dicke state

We show the scaling of the bound in n is tight by explicit example.





Dicke state. The equal-weight superposition of all zero-magnetization states on n qubits:

$$|\Psi_D\rangle = \sum_{\mathbf{s}} \delta_{\sum_i s_i, 0} |\mathbf{s}\rangle$$



Realizable as NQS with a *single nonlinearity*: e.g. $\sigma(x) = \text{ReLU}(x - 1) - 2\text{ReLU}(x) + \text{ReLU}(x + 1)$, $x = \sum_i s_i$.

Entanglement entropy of Dicke state: $S_A = \frac{1}{2} \log |A|$. Hence for $|A| = cn$, $S_A = \frac{c}{2} \log n$.



Hence the bound
cannot be better than
logarithmic in n

Approaching volume law entanglement



Subregion entanglement entropy for CosNet states with $n=22$ qubits, averaged over 20 random initializations. Dashed black line $\sim$ Page value, i.e. the Haar-random state prediction.

CosNet. Defined by $\Psi(s) = \Psi_1(s) + i \Psi_2(s)$ with each of $\Psi_{1,2}$ of the form

$$f(s) = \sum_{i=1}^k a_i \cos(\mathbf{w}_i \cdot \mathbf{s} + b_i)$$

with $a_i \sim N\left(0, \frac{\sigma_a^2}{k}\right)$, $\mathbf{w}_i \sim N\left(0, \frac{\sigma_w^2}{n}\right)$, $b_i \sim U[-\pi, \pi]$.

(Sub)logarithmic entanglement scaling for $k \ll n$ and volume law only when $k \gg n$, **consistent with our bound** $S \leq O(k \log n)$

Bound on entanglement in NQS: recap

Under quite general conditions, an NQS acting on n spins with k applications of scalar nonlinearities satisfies a bound on entanglement over any subregion A :

$$S_A = O(\mu \log n) = O(k \log n).$$

- Here μ is the dimension of an intermediate feature space (i.e. # of affine features)
- This bound prohibits volume-law entanglement for $k \sim O(1)$. Consistent with known examples of volume-law entangled NQS.
- However, this bound is **generically saturated**. Hence NQS *generically* have more entanglement capacity than tensor networks (log law vs area law)

Forthcoming work with:

Di Luo, *Tsinghua U.*

Haimeng Zhao, *Caltech*,
others

I. Introduction

II. Entanglement in NQS

III. Expressivity and choice of basis in NQS

Expressivity and choice of basis in neural quantum states

- Entanglement is local basis independent. But NQS are highly basis dependent.

For example, the ease of learning a ferromagnetic product state depends on the basis:

$$|+X\rangle^{\otimes n} = |\rightarrow\rangle^{\otimes n} \propto \sum_{\mathbf{s} \in \{\uparrow, \downarrow\}^n} |\mathbf{s}\rangle \quad \text{Constant function}$$


$$|+Z\rangle^{\otimes n} = |\uparrow\rangle^{\otimes n} \propto \sum_{\mathbf{s} \in \{\uparrow, \downarrow\}^n} \delta_{\mathbf{s}, \uparrow \dots \uparrow} |\mathbf{s}\rangle \quad \text{Delta function}$$

$$|-X\rangle^{\otimes n} = |\leftarrow\rangle^{\otimes n} \propto \sum_{\mathbf{s} \in \{\uparrow, \downarrow\}^n} (-1)^{|\mathbf{s}|_1} |\mathbf{s}\rangle \quad \text{Parity function}$$

Experiment: Hardness of learning a ferromagnet

- $\bigotimes_i | +X \rangle_i = +X$, $\bigotimes_i | \uparrow \rangle_i = +Z$, $\bigotimes_i | -X \rangle_i = -X$ We try to learn $+X$, $+Z$, $-X$ using MLP-NQS.

epochs to 0.99 fidelity per qubit vs N



Proposal of a measure of hardness to learn: spikiness

$$\psi(s) : \{0, 1\}^n \rightarrow \mathbb{C} \quad \hat{\psi}(r) = \frac{1}{2^n} \sum_{s \in \{0, 1\}^n} (-1)^{r \cdot s} \psi(s) \quad p(r) = \frac{|\hat{\psi}(r)|^2}{\sum_r |\hat{\psi}(r)|^2}$$

$$\mathcal{S}_\alpha^{\text{fix}} = \frac{1}{\alpha} \log \left(\sum_r e^{\alpha |r|_1} p(r) \right) = \frac{1}{\alpha} \log \mathbb{E}_p[e^{\alpha |r|_1}] \quad \textit{Fixed-basis spikiness}$$

$$\mathcal{U}_{1,1} = \{U_1 \otimes \cdots \otimes U_n \mid U_j \in \text{SU}(2)\}$$

$$\mathcal{S}_\alpha(\psi) = \min_{U \in \mathcal{U}_{1,1}} \mathcal{S}_\alpha^{\text{fix}}(U\psi) \quad \textit{(intrinsic) spikiness}$$

Definition 2 (equivalent)

Factors nicely over qubits:

$$U_j^\dagger X_j U_j = \hat{\mathbf{n}}_j \cdot \boldsymbol{\sigma}_j$$

Minimization is that of a classical Heisenberg spin model with many-body exchange couplings coming from correlations in $|\psi\rangle$


$$\mathcal{S}_\alpha(\psi) = \min_{\hat{\mathbf{n}}_j \in S^2} \frac{1}{\alpha} \log \left\langle \psi \left| \prod_j (c_{\alpha,+} I - c_{\alpha,-} \hat{\mathbf{n}}_j \cdot \boldsymbol{\sigma}_j) \right| \psi \right\rangle, \quad c_{\alpha,\pm} = \frac{e^\alpha \pm 1}{2} \quad \text{spikiness}$$

Theoretical properties of spikiness:

***Fixed-basis spikiness
is small***



***Exists a small MLP-NQS
representation***



$$\frac{1}{\alpha} \log \frac{1 + e^{\alpha}}{2}$$


Entanglement vs magic vs spikiness

	Depends on	Invariant under	Probes
Entanglement	• Bipartition $A B$ • Reduced density matrix	• $U_A \otimes U_B$ (unitaries within A or B)	• Bipartite correlation across cut
Magic	• Pauli spectrum (e.g. stabilizer Renyi entropy)	• Clifford group	• Nonstabilizerness • Non-Clifford resources
Spikiness	• Walsh degree distribution	• Single-qubit product unitaries $\otimes_j U_j$	• Irreducible sign structure in amplitudes, cost to a low-capacity learner

Transverse-field Ising model spikiness

Consider ground state of L-site TFIM (pbc). Fixed-basis spikiness vs (intrinsic) spikiness:


$$H = - \sum_{j=1}^L Z_j Z_{j+1} - h \sum_{j=1}^L X_j$$




Experiment: Transverse-field Ising model

Rotating to spikiness minimizing basis helps to learn the wavefunction!

RBM perspective: Cortes et al., *Basis dependence of neural quantum states for the transverse field Ising model*, SciPost Physics Core 9 (2), 027



TFIM Learnability ($L = 12$, 48 seeds, geometric mean). Epochs to per-qubit fidelity 0.99 in computational vs $\alpha = 1$ minimum-spikiness basis (width128depth3tanhMLP). Bands = standard error of mean



Thank you!

Entanglement:

NP, *Bound on entanglement in neural quantum states*, Phys. Rev. Lett. 136, 120403 (2026)

Other aspects of expressivity:

Ongoing work with:

- Di Luo, *Tsinghua U.*



- Haimeng Zhao, *Caltech*



Full proof of Theorem 1

Theorem 1. von Neumann entropy $S_A = -\text{Tr } \rho_A \log \rho_A$ of any subregion A satisfies $S_A = O(\mu \log n) = O(k \log n)$.

- 1) Introduce a short-hand notation $z_{j,s} = t_j(\mathbf{s}) = \mathbf{w}_j \cdot \mathbf{s} + b_j$. The NQS is

$$|\Psi\rangle \equiv \sum_{\mathbf{s}} \mathcal{G}(z_{1,\mathbf{s}}, \dots, z_{\mu,\mathbf{s}}) |\mathbf{s}\rangle$$

whose reduced density matrix we denote $\rho_A = \text{Tr}_{A^c} |\Psi\rangle\langle\Psi|$. We introduce a normalized auxiliary state

$$|\tilde{\Psi}\rangle \equiv Q^{-1/2} \sum_{\mathbf{s}} P_d(z_{1,\mathbf{s}}, \dots, z_{\mu,\mathbf{s}}) |\mathbf{s}\rangle$$

whose reduced density matrix is $\tilde{\rho}_A$ and its von Neumann entropy by $\tilde{S}_A$.

$$\begin{aligned} 2) \quad |\tilde{\Psi}\rangle &= Q^{-1/2} \sum_{\mathbf{s}} P_d(z_{1,\mathbf{s}}, \dots, z_{\mu,\mathbf{s}}) |\mathbf{s}\rangle = Q^{-1/2} \sum_{\mathbf{u}, \mathbf{v}} \tilde{M}_{\mathbf{u}, \mathbf{v}} |\mathbf{u}\rangle_A |\mathbf{v}\rangle_{\bar{A}} \\ P_d(\mathbf{t}) &= \sum_{0 \leq n_1, \dots, n_\mu \leq d} \alpha_{\mathbf{n}} t_1^{n_1} \dots t_\mu^{n_\mu} \quad \tilde{M}_{\mathbf{u}, \mathbf{v}} = P_d(x_{1,u} + y_{1,v}, \dots, x_{\mu,u} + y_{\mu,v}) \\ \tilde{M} &= \sum_{\mathbf{u}, \mathbf{v}} \tilde{M}_{\mathbf{u}, \mathbf{v}} |\mathbf{u}\rangle\langle\mathbf{v}| \\ &= \sum_{\mathbf{u}, \mathbf{v}} \sum_{0 \leq n_1, \dots, n_\mu \leq d} \alpha_{\mathbf{n}} (x_{1,u} + y_{1,v})^{n_1} \dots (x_{\mu,u} + y_{\mu,v})^{n_\mu} |\mathbf{u}\rangle\langle\mathbf{v}| \\ &= \sum_{\mathbf{u}, \mathbf{v}} \sum_{0 \leq n_1, \dots, n_\mu \leq d} \alpha_{\mathbf{n}} \prod_{\ell=1}^{\mu} \sum_{k_\ell=0}^{n_\ell} \binom{n_\ell}{k_\ell} x_{\ell,u}^{n_\ell - k_\ell} y_{\ell,v}^{k_\ell} |\mathbf{u}\rangle\langle\mathbf{v}| \\ &= \sum_{0 \leq n_1, \dots, n_\mu \leq d} \sum_{k_1, \dots, k_\mu=0}^{n_1, \dots, n_\mu} \alpha_{\mathbf{n}} \left(\sum_{\mathbf{u}} \prod_{\ell=1}^{\mu} \binom{n_\ell}{k_\ell} x_{\ell,u}^{n_\ell - k_\ell} |\mathbf{u}\rangle \right) \left(\sum_{\mathbf{v}} \prod_{\ell=1}^{\mu} y_{\ell,v}^{k_\ell} \langle\mathbf{v}| \right) \\ &\equiv \sum_{0 \leq n_1, \dots, n_\mu \leq d} \sum_{k_1, \dots, k_\mu=0}^{n_1, \dots, n_\mu} \tilde{m}_{\mathbf{n}, k_1, \dots, k_\mu}^{\text{rank-1}} \end{aligned}$$

$$\longrightarrow \text{rank } \tilde{M} \leq \sum_{0 \leq n_1, \dots, n_\mu \leq d} \sum_{k_1, \dots, k_\mu=0}^{n_1, \dots, n_\mu} 1 = ((d+1)(d+2)/2)^\mu \longrightarrow \tilde{S}_A = O(\mu \log d)$$

Full proof of Theorem 1, cont.'d

We transfer the bound from the auxiliary state to the original state $|\Psi\rangle$, taking $d = \Theta(n)$. We write $\max |\mathcal{G} - P_d| \leq \varepsilon_{\text{poly}}$

1. Define $|\delta\rangle \equiv |\Psi\rangle - |\tilde{\Psi}\rangle$; it satisfies

$$\begin{aligned} \|\delta\|^2 &\leq 4 \left(\sum_{\mathbf{s}} |\mathcal{G}(z_{1,\mathbf{s}}, \dots, z_{\mu,\mathbf{s}}) - P_d(z_{1,\mathbf{s}}, \dots, z_{\mu,\mathbf{s}})|^2 \right)^{1/2} \\ &\leq 4\varepsilon_{\text{poly}} 2^{n/2} \end{aligned}$$

2. The trace distance is bounded by

$$\frac{1}{2} \|\rho_A - \tilde{\rho}_A\|_1 \leq \frac{1}{2} \|\rho - \tilde{\rho}\|_1 \leq \|\delta\|_2 \leq 2\sqrt{\varepsilon_{\text{poly}}} 2^{n/4}$$

3. The Fannes-Audenaert inequality gives us

$$\begin{aligned} |S_A - \tilde{S}_A| &\leq \frac{1}{2} \|\rho_A - \tilde{\rho}_A\|_1 \log(2^{|A|} - 1) + 1 \\ |S_A - \tilde{S}_A| &\leq 2\sqrt{\varepsilon_{\text{poly}}} 2^{n/4} \log(2^{|A|}) + 1 \\ &\leq 2\sqrt{\varepsilon_{\text{poly}}} 2^{n/4} |A| + 1 \\ &\leq 2\sqrt{\varepsilon_{\text{poly}}} 2^{n/4} n + 1 \\ &= O(1) \end{aligned}$$

Conclude: von Neumann entropy $S_A = -\text{Tr } \rho_A \log \rho_A$ of any subregion A satisfies $S_A = O(\mu \log n) = O(k \log n)$.



taking $d = \Theta(n)$ in the last step, where $\varepsilon_{\text{poly}} \sim e^{-\beta d}$.