

Efficiency of the hidden fermion determinant states Ansatz in the light of different complexity measures

Björn J. Wurst,^{1,*} Dante M. Kennes,^{1,2} and Jonas B. Prof³

¹*Institut für Theorie der Statistischen Physik, RWTH Aachen University and JARA—Fundamentals of Future Information Technology, 52056 Aachen, Germany*

²*Max Planck Institute for the Structure and Dynamics of Matter, Center for Free Electron Laser Science, 22761 Hamburg, Germany*

³*Institute for Theoretical Physics, Goethe University Frankfurt, Max-von-Laue-Straße 1, D-60438 Frankfurt a.M., Germany*

(Dated: November 8, 2024)

Finding reliable approximations to the quantum many-body problem is one of the central challenges of modern physics. Elemental to this endeavor is the development of advanced numerical techniques pushing the limits of what is tractable. One such recently proposed numerical technique are neural quantum states. This new type of wavefunction based Ansätze utilizes the expressivity of neural networks to tackle fundamentally challenging problems, such as the Mott transition. In this paper we aim to gauge the universality of one representative of neural network Ansätze, the hidden-fermion Slater determinant approach. To this end, we study five different fermionic models each displaying volume law scaling of the entanglement entropy. For these, we correlate the effectiveness of the Ansatz with different complexity measures. Each measure indicates a different complexity in the absence of which a conventional Ansatz becomes efficient. We provide evidence that whenever one of the measures indicates proximity to a parameter region in which a conventional approach would work reliable, the neural network approach also works reliable and efficient. This highlights the great potential, but also challenges for neural network approaches: Finding suitable points in theory space around which to construct the Ansatz in order to be able to efficiently treat models unsuitable for their current designs.

I. INTRODUCTION

The reliable prediction of material properties is central to designing new devices with tailored functionality [1, 2]. Elemental for this is the development of numerical techniques for treating the quantum many-body problem, the solution of which would allow to fully predict how materials behave. However, the full solution to this problem is impossible to obtain even for simple materials, and therefore approximate methods have been developed. Usually one is interested in the low temperature behavior of materials. To study this regime an effective low-energy model, derived via density functional theory [3, 4] or by other means [5], is constructed. This effective model is describing the material with a minimal number of degrees of freedom. Still these simplified models are often not exactly solvable and a number of numerical techniques, each with its strengths and weaknesses, have been developed.

These numerical approaches follow different lines of thought. A broad class are based on the Luttinger-Ward functional, i.e. working on generating functionals. Examples of such approaches are dynamical mean-field theory [6–9], two particle self-consistent approach [10, 11], the Parquet approximation [12], fluctuation exchange [13], functional renormalization group [14], cluster perturbation theory [15, 16] and determinant quantum Monte-Carlo [17–19]. Another large class aims

to directly solve the stationary Schrödinger equation targeting the ground state (and a few excited ones) of the system. Examples for this class are matrix-product-state (MPS) approaches such as density matrix renormalization group (DMRG) [20–23], wavefunction based variational Monte Carlo [24–26], the Gutzwiller approximation [27–30], exact diagonalization (ED) [31] and neural quantum states (NQS) [32–35]. For the latter class the challenge is two-fold. First one has to design an algorithm allowing for an efficient ground state search: in the case of ED for example the Lanczos algorithm [36]. This of course requires a way to store the ground state, posing the second challenge as the required memory scales exponentially with the system size. This storage problem is often solved by choosing a suitable predefined basis for the states representation [21, 23] or a suitable compressing technique [37]. For example, DMRG is relying on the entanglement entropy scaling in 1d gapped Hamiltonians which ensures that the former exponential cost for diagonalizing the Hamiltonian and storing the ground state is reduced to a polynomially one, enabling numerically exact studies for specific Hamiltonian types.

Neural quantum states are a recently introduced approach of the second class [32–35]. Fueled by the fact that random neural network states chosen from a certain measure have volume law entanglement [38] a surge in numerical benchmark studies in spin [39, 40], bosonic [41–43] and fermionic systems [34, 44–55] have established this method as a versatile approach allowing to achieve high accuracy [56, 57]. However, while the numerical re-

* bjoern.wurst@rwth-aachen.de

sults are promising, a deeper understanding of what will and what will not be possible to solve in a faster than exponential scaling is lacking. Especially for fermionic systems there is currently a debate whether all physical ground states exhibiting volume law entanglement are treatable with exponential speedup or only some [48, 54]. To understand this better, it is crucial to determine what limitations are inherited by closely related approaches such as DMRG and classical VMC trial states, especially with regards to expressivity of the trial state. To phrase it differently: Currently there is no a-priori understanding whether for a given Hamiltonian the ground state can be represented efficiently or not utilizing NQS.

A hint to the conditions under which a ground state is efficiently representable might be drawn from known complexity measures and their implications for related numerical techniques, such as the bipartite entanglement and the entropy of the Born-distribution: Because NQS are usually able to represent matrix product states efficiently, they are also able to represent states with low entanglement [58–64]. On the other hand, if the Born-distribution of the ground state has a low entropy, it indicates that the relevant Hilbert space size for the ground state search is reduced. Here the proximity to configuration interaction approaches enables NQS to efficiently solve the problem at hand [65, 66]. Furthermore, since fermionic NQS often build upon Slater determinants, the limit of weak-interactions should be representable from proximity to perturbative approaches. While all these limits are understood, the question remains how far out of these limits a neural quantum state Ansatz is still reliable and efficient.

To delineate these limits, we consider a set of models with volume-law entanglement ground states, introduced in section III. For each we track the number of parameters required to reach a certain accuracy in reference to the Hilbert space dimension. Further, we correlate our findings with both the presence of correlations and entropy of these systems, see section II, highlighting that performance is indeed best when expected from proximity to efficient solvability with another method. Our results, see section IV, highlight the great versatility of neural network approaches to learn ground states of systems challenging for other numerical methods, while also outlining limitations. Further, we observe a strong correlation of the quality of our results with the quality of our Ansatz underlining the demand of model specific neural network state Ansätze in order to exploit underlying non-complexities. These non-complexities are indicating closeness to representable regions – closeness to such a region implies a reduction of the complexity of the problem, as illustrated in Fig. 1.

II. METHOD

Trial-state — For this paper we mainly want to explore NQS via the recently introduced hidden fermion




FIG. 1. Sketch of the regions in which we find that HFDS can efficiently solve the models. Green corresponds to uncorrelated states, blue to states with low entanglement and red highly fragmented states. In these areas solvability by NQSS with suitable architecture can be expected. The arrows are indicating the different models studied for which the arrow direction indicates an increase of the interaction strength. All models, except the SYK-model, are starting from a solvable region and approaching a solvable region in the large interaction limit. For all these models a reduction of the scaling of the number of parameters required compared to the dimension of the Hilbert-space will be observed. No such reduction is observed for the SYK-model.

determinant states (HFDS) [34]. This approach showed promising results for different fermionic models [34, 49], nuclear matter [67, 68] as well as neutron star matter [47]. In HFDS the expressibility of an initial variational wavefunction is systematically improved by mapping an initial state $|\Psi_0\rangle$ in an enlarged Hilbert-space \mathcal{H}' containing N visible and M hidden particles into the target space \mathcal{H} . The general form of the wave-function is given by

$$|\Psi\rangle = \sum_n |n\rangle \langle n, f(n)|\Psi_0\rangle. \quad (1)$$

By taking the initial state to be a Slater-determinant, a wavefunction of the form

$$\Psi(n) = \left\| \begin{array}{cc} \phi_v(n) & \chi_v(n) \\ \phi_h(f(n)) & \chi_h(f(n)) \end{array} \right\|. \quad (2)$$

is obtained, where $\phi_v(n)$ is a $N \times N$ matrix and $\chi_h(f(n))$ a $M \times M$ matrix. The lower part of the matrix, i.e. ϕ_h and χ_h are parameterized row wise by multi layer perceptrons as was done in Ref. [34]. We utilize

$$f(x + iy) = \text{selu}(x) + i \text{selu}(y), \quad (3)$$

as our activation function for neurons in the the hidden layers, i.e. the scaled exponential linear unit [69]

$$\text{selu}(x) = \lambda \begin{cases} x, & x \geq 0 \\ \alpha e^x - \alpha, & \text{else} \end{cases}$$

applied to the real and complex part separately. Here $\alpha \approx 1.67$ and $\lambda \approx 1.05$ are given by the default values used in the NetKet implementation [70]. An exponential activation function is used for the output layer. The

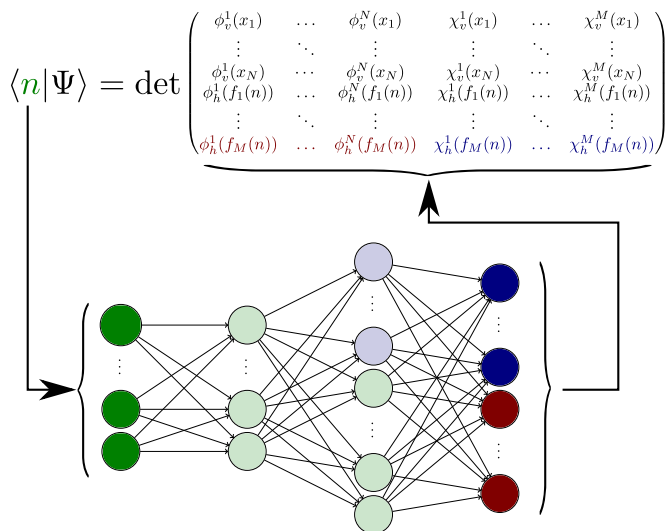


FIG. 2. Sketch of the architecture of the HFDS state. The positions x are obtained from n by choosing an order by fixing the order. The signs corresponding to an exchange of particles is absorbed into the operators. The lower part of the matrix within the Slater-determinant is parameterized row wise by multi layer perceptions, in this sketch with two hidden layers. The first half has the same dimension as the input, while the second half of the layers has the dimension of the input plus the numbers of hidden particles M .

width of the hidden layers is given by the input dimension for the first half of the hidden layers and by the input dimension plus M for the second half. The architecture is sketched in Fig. 2. Furthermore, we do also treat the visible sub sector as variational parameters opposed to using the natural orbitals corresponding to the lowest energies as was done in Ref. [49].

The necessity of building these more complex trial states in comparison to a plain neural network can be seen when considering certain non-interacting ground states, for example a fully connected random hopping model. Here no efficient representation can be found by feed forward neural networks [54], even though an efficient solution on classical computers is clearly possible.

The obtained wavefunctions are able to parameterize physically motivated trial-wavefunctions like Jastrow- or Gutzwiller-factors efficiently [34]. Furthermore, they are closely related to neural network backflow states - as suggested by Ref. [71] they parameterize a similar variational manifold. Therefore, the results we present have implications for neural backflow states as well.

Training — Our trial-states are trained on the overlap difference

$$\delta\hat{O} = 1 - \hat{P}_{\text{tar}}, \quad (4)$$

where \hat{P}_{tar} is the projector onto the target state [35]. In the following, the ground-state of the system obtained with exact diagonalization. Training on this observable allowed to obtain larger fidelities with the target state in comparison to directly training on the Hamiltonian,

especially in the case of small energy gaps to the low lying excited states. Because exact diagonalization already limits us to small system sizes and thus small Hilbert-spaces, we have calculated all observables by contractions over the full Hilbert-space instead of making use of the variational Monte Carlo method typically employed for this purpose. In doing so, we avoid an additional source for errors which would be introduced by the Monte-Carlo sampling.

For the training of the NQS we utilize the ADAM-optimizer [72] instead of the stochastic reconfiguration method more commonly employed for the optimization of NQS and variational states in general [25, 34, 35, 49, 57]. For more details on this choice we refer to appendix A, where we elaborate on the reasons for this decision.

Observables — To estimate the quality of the fit provided by the neural quantum state, we use the expectation value of our loss-function, i.e. $\langle \delta\hat{O} \rangle$. This quantity is equal to one minus the fidelity of the variational state and the target state.

In order to estimate the complexity of the target state we use two quantities: The Shannon-entropy of the Born-distribution and the von-Neumann-entropy of the one particle reduced density matrix. The Shannon-entropy of the Born-Distribution

$$S_\psi = - \sum_n |\psi(n)|^2 \log |\psi(n)|^2, \quad (5)$$

is an upper bound to the bipartition entanglement entropy with an arbitrary bipartition compatible with the basis and serves as an indicator for the ground-state becoming sparse in the computational basis, i.e. fragmentation of the Hilbert space. This quantity therefore indicates that the state can be efficiently represented by configuration interaction (CI) methods. Because of the relation between the HFDS approach and CI methods shown in Ref. [34], one expects this also for the HFDS trial states. Furthermore, we want to mention a second case here: if the fragmentation is due to an operator diagonal in the computational basis, typical examples would be density-density-interactions or Ising-type interactions, the fragmentation can be captured by particular simple methods, like Gutzwiller-correlators. These are also closely related to the HFDS-approach. A quantity which carries similar information, the inverse partition ratio, was recently also considered in Ref. [73, 74]. As pointed out there are states with a large peak in the electron density of the ground state wave functions, i.e. states having a large inverse partition ratio, hard to train. This is because it is hard to sample from the born distribution if it spans orders of magnitudes on the relevant space. Nonetheless, there exist improved sampling methods in this case making use of the closeness to CI [73–75]. Low entanglement entropy as well as sparseness, i.e. only a small fraction of the basis set of the Hilbert-space is relevant for the problem, should render the groundstate easy to represent.

The second quantity used to estimate the complexity

of the target state is given by the von Neumann entropy of the N particle reduced density matrix γ_n (n -RDM)

$$S(\gamma_n) = -\text{Tr} \gamma_n \log \gamma_n. \quad (6)$$

This quantity was analyzed in detail in Ref. [76] and is a well known indicator for complexity for fermionic systems [77–79]. The most important property for our purpose is that $S(\gamma_n)$ is bounded from below by states corresponding to Slater-determinants. This quantity therefore serves as an indicator of correlations - and consequently complexity - in the target state. For simplicity we will only consider the 1-RDM. $S(\gamma_1)$ is bounded from below by $\log(N)$, where N corresponds again to the number of fermions in the system. This lower bound corresponds to the first N natural orbitals being fully occupied and the remaining ones being empty. In the opposite limit all natural orbitals are equally occupied. Correlations are not only indicated by this quantity but also by the interaction strength. There are two main reasons for considering the 1-RDM instead: certain models have not only uncorrelated states in the non interacting limit, but also in the absence of a kinetic term: a large interaction strength does not necessarily imply strong correlations. This quantity is furthermore easier to handle when comparing different models, here correlations might be generated at different rates when raising the interaction strength. Both of these quantities, i.e. the Shannon entropy of the Born-distribution and the von-Neumann entropy of the 1-RDM, are able to capture different non-complexities. We emphasize that also other non-complexities exist but for the reasons described above we focus on these two.

III. MODELS

To gauge the predictive and expressive power of new numerical approaches, it is instructive to consider different types of prototypical models. Ultimately, the aim is to obtain an understanding of what types of models can be efficiently solved and which ones are not efficiently solvable by understanding what the underlying mechanisms for the complexity reduction are. Here, we build on Ref. [34] pointing out the relations of the used approach and known variational Jastrow based approaches as well as the known connection between MPSs and NQS [58–63]. To probe in what sense limitations of these closely related approaches are inherited, we consider SYK-like models, meaning models with fully disordered long range interactions, which are difficult for matrix product states. We introduce five different models: a pure SYK model, a random hopping Hubbard model, similar to a model already studied in Ref. [49], a density-density interaction model, which is similar to a model already studied in Ref. [54], a spin-spin interaction model and lastly a pair-pair interaction model. The five models

are given by

$$\hat{H}^{\text{SYK}} = \frac{J}{\sqrt{2N_S}} \sum_{ij} J_{ij}^{[2]} \hat{c}_i^\dagger \hat{c}_j + \frac{U}{\sqrt{2N_S}^3} \sum_{ijkl} J_{ijkl}^{[4]} \hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_l \hat{c}_k \quad (7)$$

$$\hat{H}^{\text{Hubbard}} = \frac{J}{\sqrt{N_S}} \sum_{ij\sigma} J_{ij}^{[2]} \hat{c}_{i\sigma}^\dagger \hat{c}_{j\sigma} + U \sum_i \hat{n}_i^\uparrow \hat{n}_i^\downarrow \quad (8)$$

$$\hat{H}^{\text{Density}} = \frac{J}{\sqrt{N_S}} \sum_{i \leq j, \sigma} t_{ij} \hat{c}_{i\sigma}^\dagger \hat{c}_{j\sigma} + h.c. \quad (9)$$

$$+ \frac{U}{4\sqrt{N_S}} \sum_{i \leq j} \sum_{\sigma\sigma'} v_{ij} \hat{n}_i^\sigma \hat{n}_j^{\sigma'} + h.c. \quad (10)$$

$$\hat{H}^{\text{Pair}} = \frac{J}{\sqrt{N_S}} \sum_{i \leq j, \sigma} t_{ij} \hat{c}_{i\sigma}^\dagger \hat{c}_{j\sigma} + h.c. + \frac{U}{\sqrt{N_S}} \sum_{i \leq j} v_{ij} \hat{c}_{i\uparrow}^\dagger \hat{c}_{i\downarrow}^\dagger \hat{c}_{j\downarrow} \hat{c}_{j\uparrow} + h.c. \quad (11)$$

$$\hat{H}^{\text{Spin}} = \frac{J}{\sqrt{N_S}} \sum_{i \leq j, \sigma} t_{ij} \hat{c}_{i\sigma}^\dagger \hat{c}_{j\sigma} + h.c. + \frac{U}{\sqrt{N_S}} \sum_{i \leq j} v_{ij} \hat{c}_{i\uparrow}^\dagger \hat{c}_{i\downarrow}^\dagger \hat{c}_{j\downarrow} \hat{c}_{j\uparrow} + h.c. \quad (12)$$

where the real and imaginary parts of t_{ij} and v_{ij} are drawn from normal distributions, i.e. $\mathcal{N}(0, 1) + i\mathcal{N}(0, 1)$ and $J_{ij}^{[2]}$, and $J_{ijkl}^{[4]}$ are random unitary Gaussian variables full filling $J_{ij}^{[2]} = J_{ji}^{[2]*}$ or $J_{ijkl}^{[4]} = -J_{ijlk}^{[4]} = -J_{jikl}^{[4]} = J_{klji}^{[4]*}$ and $\langle |J_{ij}^{[2]}|^2 \rangle = 1$ or $\langle |J_{ijkl}^{[4]}|^2 \rangle = 1$ respectively. When not stated otherwise, we use $J = 0$ for the SYK-model and $J = 1$ for the remaining models. N_S stands for the number of sites of the model.

From an entanglement entropy perspective, all five models follow a volume law scaling in N_S for the system sizes considered in Fig. 3. Nonetheless, they exhibit a large degree of structure; all the information necessary for their construction can be compressed down to at most $\mathcal{O}(N_S^2)$ parameters, except for the SYK-model, where only a compression into $\mathcal{O}(N_S^4)$ parameters is possible. This also corresponds to a lower bound of the complexity: the Hamiltonian can usually be reconstructed from the groundstate [80–82]. Therefore, both the variational manifold, due to the variational parameter number, as well as the set of ground states lie within a low dimensional sub manifold of the Hilbert space. Not a lot is known about the overlap of these two manifolds.

IV. RESULTS

To map out what is easy and what is hard to represent for the HFDS states, we calculate the overlap distance to the ground state as a function of the ratio of the number of parameters in the neural network and the Hilbert-space dimension. We plot these two quantities for all five




FIG. 3. Bipartition entanglement entropies for the ground states of the various models considered here: (a) SYK-model, (b) random hopping Hubbard-Model, (c) density-density-interaction, (d) pairing-interaction and (e) spin-flip interaction. $U = 3$ was used. All models seem to possess volume law entanglement. The error-bars correspond to the standard error of the mean estimated via the sample variance. Each sample contains six realizations for each system size.

models and different system sizes in Fig. 4 (a-e). We observe two qualitatively different behaviors of the considered models: for the plain SYK model the number of parameters per dimension show a collapse of the different system sizes onto a single curve, whereas in contrast, for all other models the curves for different system sizes drift towards the lower-left - indicating that larger system sizes reach the same numerical error at a smaller relative number of parameters. The speed of the movement is found to be dependent on the specific model. These two different behavior patterns indicate different complexity scaling - while for the plain SYK model we find an exponential scaling, the movement observed for all other models indicates a deviation from a scaling trend. Whether or not this scaling becomes truly polynomial, stays exponential or shows some different sub-exponential behavior is not determinable by this data alone, as emphasized in the panels (f-o) of Fig. 4. Both polynomial and exponential fits are compatible with the data. We verified this method of determining the scaling by performing the same analysis for an area law ground state, which then has a much smaller root-mean error (RME) for a polynomial scaling, see Appendix D. These results highlight an important property of NQS when compared to other simpler variational trial states, e.g. Slater-determinants with Jastrow correlators: they offer a controlled way of adding more parameters to the trial state and therefore to achieve the capacity necessary in order to achieve a given target error. The actual convergence can be prohibitively slow as indicated by the results for the SYK-model, making small errors in practice hard to achieve even for small system sizes.

While we do find a decreasing relative numbers of parameters required for reaching a certain error bound it should be noted that the absolute number of parameters still increases significantly, here the model with the density-density-interaction might be considered an exception. Reasons for this observation will be explained in more detail below. Furthermore, we find that the Hubbard-model, the model with the spin-spin-interaction and the model with the pairing interaction all require

a similar number of parameters to reach a given error as can be seen in the second and third row of Fig. 4. Here, generally a reduction in complexity had to be expected due to the reduction in entanglement entropy when compared to the SYK-model at similar dimension of the Hilbert space, as can be seen in Fig. 3 (also compare F for the Hubbard-model). Since the Hubbard-model shows a far greater entanglement entropy, while the other two models have an entanglement entropy similar to the density-density interaction model, as can be seen in Fig. 3, we conclude that the performance of the trial states can not be explained by only considering the entanglement entropy. Therefore, the pressing question arising is what determines the performance and can we assess it from inspecting the model?

Two possible quantities to consider were already outlined above, namely the von-Neumann entropy of the 1-RDM and the Shannon entropy of the Born-distribution capturing the presence of correlations and fragmentation of the Hilbert-space respectively. We aim to explain observed features in terms of these two quantities.

Before discussing our results, we briefly mention what one would naively expect of the HFDS approach. As was shown in Ref. [34], with the HFDS approach we can represent ground states which are constructed by applying Jastrow-factors onto the Slater determinant basis. Thus allowing to represent weak-coupling states. Furthermore due to the proximity to CI like methods, solving the Hamiltonian in cases in which CI works well, namely when there is a Hilbert-space fragmentation, also works well. Keeping these arguments in mind we further investigate the central question considering $\langle \delta \hat{O} \rangle$ as a function of the von-Neumann entropy of the 1-RDM and the Shannon entropy of the Born distribution for the SYK model and the fully connected Hubbard-Model.

In Fig. 5, we track the deviation for different numbers of hidden layers for different realizations of the SYK-model defined in Eq. (7): we observe a strong correlation between the increase of the deviation and the von-Neumann entropy of the 1-RDM. This indicates that the correlations are not fully captured by the Ansatz. At

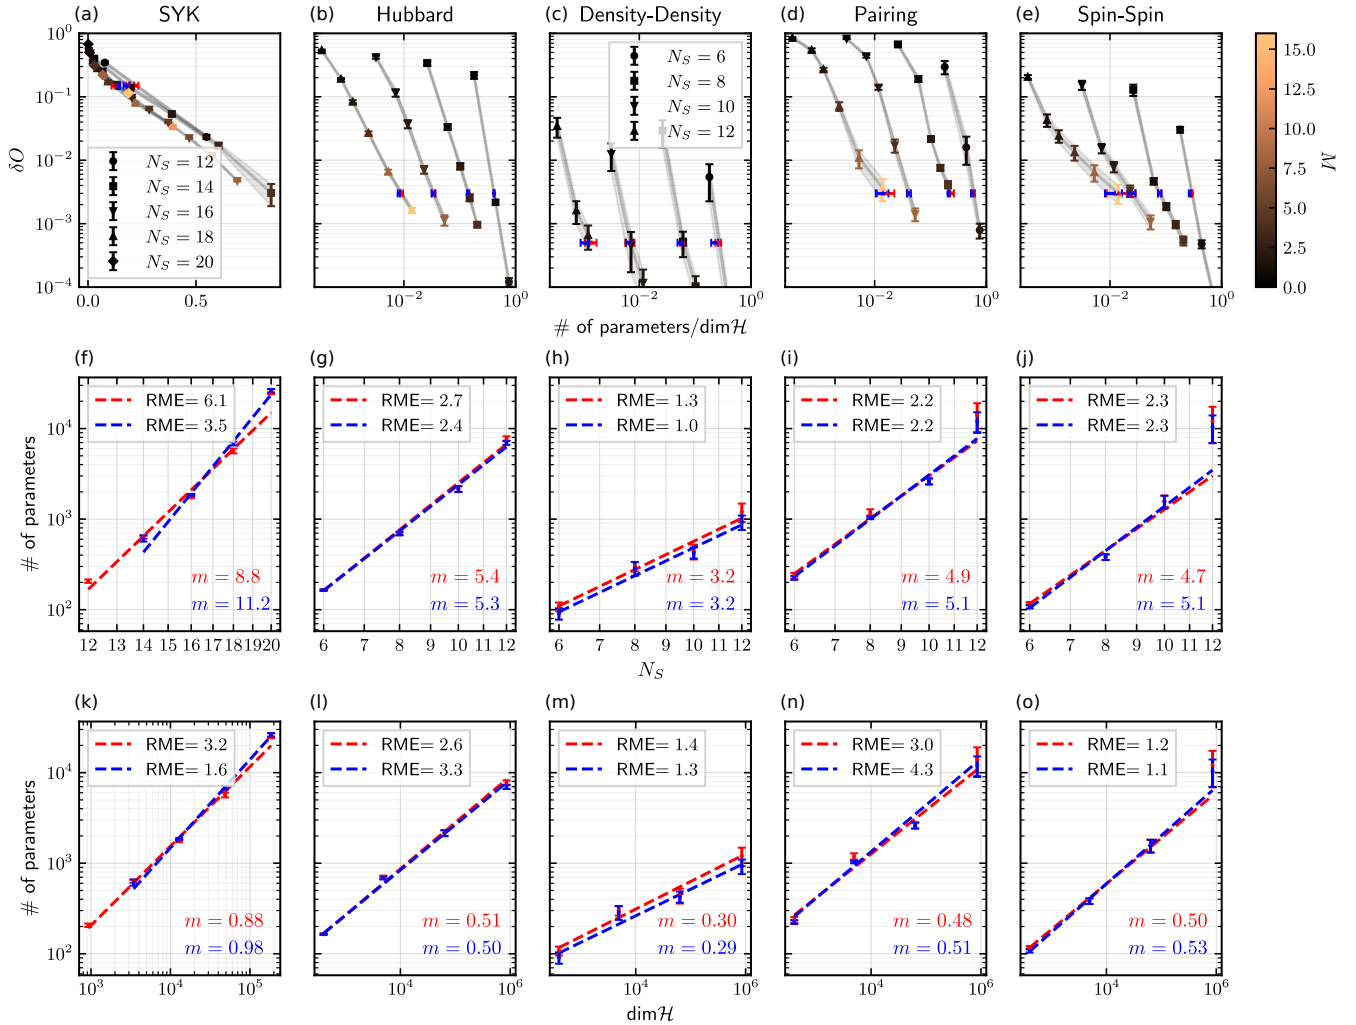


FIG. 4. (a-e) Achieved errors for the various models and System sizes at $U = 3$: (a) N_S from $N_S = 12$ to $N_S = 20$ for the SYK-model and two hidden layers and (b-e) from $N_S = 6$ to $N_S = 12$ for the remaining models without hidden layers with the total magnetization constrains enforced (see also appendix A). The expressivity is controlled via the number of hidden particles indicated by the coloring of the data points. The order of the columns is the same as in fig. 3. The blue point corresponds to the globally interpolated estimates for the number of parameters for a fixed δO and the red points correspond to the locally interpolated estimates. Details can be found in the appendix B. The color of the points encodes the number of hidden particles M . In the second (f-j) and third row (k-o) we plot the interpolated estimates of the number of parameters to reach a certain accuracy and test two different fitting hypothesis, $\propto N_S^m$ and $\propto \text{dim} \mathcal{H}^m$. In (f-j) we assume an algebraic scaling of the number of parameters with the respective *system* size, while in (k-o) we assume an algebraic scaling of the number of parameters with the respective *Hilbert*-space size (meaning exponential in system size). The quality of the fit is assessed by the root mean error (RME), for which we have rescaled the deviation by the error of the mean, which for most models is of the same quality between the two hypothesis.

the same time no strong correlation between S_ψ and the error is observed, as S_ψ stays approximately constant for all cases considered here. The observed success when representing the model in the weakly correlated case is not surprising, as it expected due to the Slater-determinant based construction of the trial state. An increase of the network's expressibility on the other hand via the depth is only leading to very minor improvements in achievable infidelity throughout the whole parameter range. Thereby indicating that the neural net-

works are struggling with the intricate correlation structure of the groundstate independent of the interaction strength. This implies that the state is not only hard to represent due to the volume law scaling of the entanglement entropy (compare Appendix D), but also due to the correlated nature of the groundstate. A slightly different view on this is outlined in Appendix E.

A somewhat different behavior can be observed in Fig. 6. While we find a similar trend, i.e. a strong correlation between entanglement of the 1-RDM and the error




FIG. 5. δO as a function of von-Neumann entropy (a) and entanglement entropy (b) for the 14 sites SYK-model for various combinations of kinetic and potential energy and seeds for the generation of the Hamiltonian. Here we fix $J+U=1$. We visualize the results for $\{0, 1, 2\}$ hidden layers symbolized by $\{o, \square, \times\}$ respectively. The trial wave function always contains one hidden particle. Independent of network depth, the error is highly correlated with the von-Neuman entropy of the 1-RDM and rather uncorrelated with the entanglement entropy.

of the trained states for small U , the behavior for large interactions is vastly different. At large interactions, we observe that the state becomes more sparse as states with double occupied sites are mostly projected out by the large interaction. This structure can be represented by the trial states considered here. The remaining part of the wavefunction only has support on a small sub-sector of the Hilbert space, leading to an easier to represent ground state. Explicitly, we observe a reduction of complexity explaining the better performance by proximity to CI-like solvability. While this implies a reduction in complexity, the exponential scaling of the problem remains, as even the spin model corresponding to the large U limit features an exponentially large Hilbert-space. An alternative interpretation of the decrease of the error at strong coupling is the following: we are now targeting a spin-liquid phase, of which certain types can be solved efficiently by Slater-determinant based states as for example done in Ref. [83]. Therefore, one can again explain the behavior by proximity to a perturbation theory around Slater determinants, exactly as in the weak coupling limit. For the chosen system size the worst performance also corresponds to the point where the approximate capacity of the model given by the number of parameters is approximately equal to the size of the relevant Hilbert space estimated via S_Ψ . Compare also Appendix E.

For the density-density-interaction model we observed a stark reduction in complexity, see Fig. 4. This can be explained from the principles introduced above: at the interaction strength tested the state does not feature strong correlations and already shows fragmentation introduced by the interaction term (see Appendix C). While the absence of correlations in the weak interaction limit is expected, this absence in the strong interaction limit is surprising at first glance. In the strong interaction limit (i.e. neglecting the kinetic part) two cases have




FIG. 6. δO as a function of von-Neumann entropy of the 1-RDM (a) and the entropy of the Born-distribution (b) for various interaction strengths for the fully connected Hubbard-Model with eight sites for $M=1$. As can be seen the achieved error is small, if either the $S(\gamma_1)$ or S_Ψ are small. Here no hidden layers were used. The t is fixed to one.

to be distinguished: either the groundstate is unique or it is two fold degenerate with the degeneracy protected by a \mathbb{Z}_2 symmetry. In the first case the groundstate is a basis state and thus uncorrelated. In the second case the degeneracy is lifted by the kinetic term for finite U and prevents the complete projection onto exclusively singly occupied states, see Appendix C. As a consequence, we observe no strong correlations in our numerical study of this model even in the large U limit, see Appendix C, which explains the reduction of the achievable error when compared to the other models.

This highlights a crucial point: not all states showing volume-law entanglement are hard to represent, but can instead be represented by simple variational wavefunctions other than MPSs. Closeness to these solvable regions can often be understood and probed by simpler observables as for instance done here via the entanglement entropy of the 1-RDM and of the Born-distribution for the density-density-interacting model. Such behavior can be understood as a strength of neural network based approaches: they are able to exploit various non complexities. Thereby, NQS can be build on already existing and well known trial states, e.g Slater-determinants with Jastrow- or Gutzwiller correlators. Moreover, they can also exploit locality and are able to find efficient compression schemes without explicit knowledge about the structure of the target state. This makes them a less biased tools when compared to other trial states. What structure is exploitable is dependent on the trial state chosen indicating a need for physically motivated Ansätze [54, 84].

In order to explore the connection between the ability of the NQS to represent the groundstate and the correlated nature of the target state we compare the different models with each other in Fig. 7. Here, we considered a range of interaction strengths. Especially for small interaction strengths a dependence between the presence of correlations and δO is observed. This strengthens our observations further – the NQS considered here are not able to fully capture the correlation induced struc-

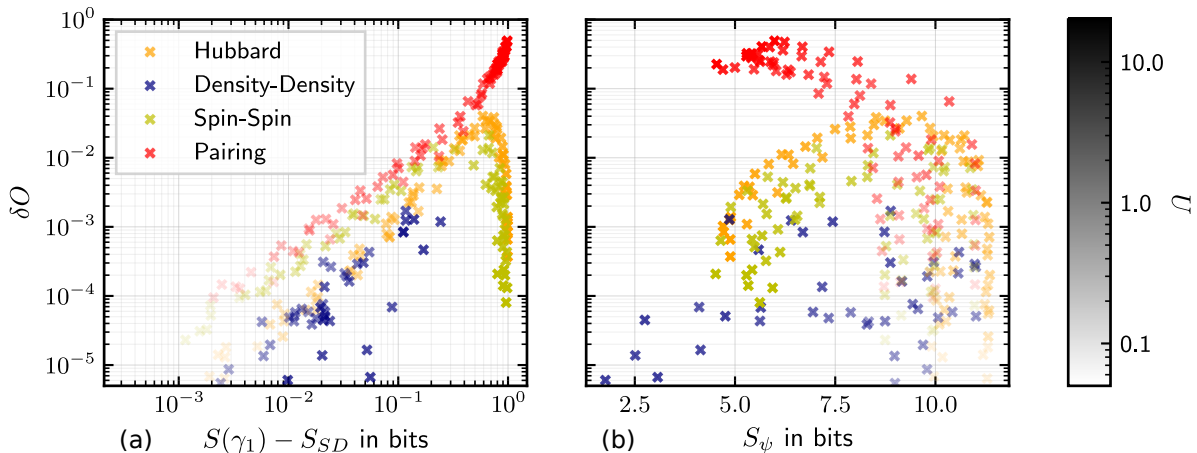


FIG. 7. Results for the different entropies: in (a) the entropy of the 1-RDM, in (b) entropies of the probability distributions. The results were obtained with $M = 1$ employing no hidden layers. The system-size was chosen as $N_S = 8$ and six random realizations were considered. U ranges from 0.1 to 20, while the kinetic coupling was kept fixed at $J = 1$. The interaction strength is indicated by the alpha values of the data-points. Small alpha values, i.e. light points, correspond to small interaction strengths, while large alpha values correspond to large interaction strengths. In subplot (a) it is clearly seen that the achievable error in the weak interaction limit is dependent on the localization of the electrons in the natural orbitals, indicating absence of strong correlation effects. Weak interaction is found on the left hand side of this plot. In subplot (b) we observe that the improvement in the strong interaction limit coincides with a reduction in complexity due to the electrons localizing in the computational basis. The data-points for strong interactions can be found on the left site of the plot.

ture of the target state. For stronger interactions, we observe a deviation from this behavior in three cases. Firstly, for the density-density-model some data-points show smaller errors than expected when only considering $S(\gamma_1)$. These points are characterized by a small entropy of the Born-distribution, compare Fig. 7 (b), and thus have a strong Hilbert space fragmentation. The second case in which a deviation is observed when comparing the spin-spin interaction model and the fully connected Hubbard model to the model featuring the pairing interaction. In the model featuring pairing interactions no increase of performance is observed for large interaction strengths, in contrast to the two other models. This deviation is understood as follows: In the large U limit the model featuring the pairing interaction is transitioning into a regime where most electrons are paired (see Appendix 11). This effect is notoriously hard to capture by Slater-determinant based approaches – here a Pfaffian based approach should achieve better results [53, 85]. Nonetheless, a rapid decrease in error can be observed if either more hidden particles are considered (compare Fig. 4) or more expressive neural networks with a single hidden layer are considered. Lastly, the density-density interaction model and the Hubbard model achieve larger fidelities for small interactions than the other two models when comparing points corresponding to similar values of $S(\gamma_1)$. This feature is explained by the structure of the trial wavefunction (see Ref. [34]): the excitations on a trial state with Jastrow-like prefactors obtained via a Laplace-expansion of the determinant strongly resemble states obtained by performing a few Lanczos-steps on the

non-interacting wavefunction. Therefore, the architecture is particularly suited to capture corrections to the noninteracting groundstates for these models.

V. CONCLUSION AND OUTLOOK

We have tested the ability of HFDS NQS to represent the groundstates of multiple prototypical models showing volume law entanglement. For these, we found two different behaviors: for the SYK-model the number of parameters necessary to achieve a given error is approximately proportional to the dimension of the Hilbert-space for all system sizes tested. This observation is in agreement with Ref. [48] and implies an exponential scaling in system size of the parameters required to achieve a given error. The problem of representing the ground state of this model with NQS is therefore not due to the inability of multilayer perceptrons to learn a complex fermionic sign structure (see also Appendix E) as speculated by Refs. [49, 54]. Our results hint at a deeper problem when trying to represent groundstates of highly complex models like the SYK-model with neural-network based approaches instead.

For all other models we observe an improvement in the performance compared to the SYK-model, reproducing the results of Ref. [54]. These models possess more structure than the SYK model, especially in the large interaction limit: the problem of representing the ground-state of the SYK-model is therefore not due to the volume law entanglement, but more generally due to the absence

of exploitable structure for the variational wavefunction Ansatz. We have not been able to determine whether this extra structure is already sufficient to achieve sub-exponential scaling in the other cases.

Furthermore, these findings highlight a fundamental strength of NQS Ansätze: They have the ability to exploit various kinds of non complexities, effectively combining the strengths of other approaches, here combining Slater-Determinants with Jastrow or Gutzwiller correlators, CI and MPSs [34]. Thereby, the NQS approaches utilize different sources of complexity reduction while other variational methods are often constructed around a single one. Whether the target state is close to such an inherited solvability region is often diagnosable by physical observables as demonstrated in this paper. For example, the absence of correlations is indicated by a small entropy of the 1-RDM, fragmentation by the entropy of the Born distribution and small entanglement by the bipartition entanglement entropy. While it is hard to understand how a neural network represents sufficiently complicated states, it is still possible to understand why an efficient representation can be found in many cases. This differentiates neural networks from interpretable methods like MPSs. When such non complexities are present, and the network is constructed such that they can be exploited, they allow for efficient compression schemes. However, when standard variational methods are not able to achieve good performance and no quantities usable for the compression are known, e.g. for the SYK-model, no efficient compression scheme is found. Furthermore, the closeness to other variational wavefunction approaches explains the good performance achieved for the density-density interaction model.

Our study suggests that while some volume law states

can be represented efficiently with neural network based approaches, as indicated by Refs. [38, 86–88], other physical volume law states seem still to be exponentially expensive and designing new approaches will be required to resolve these issues. For example, to make progress on treating SYK like models, the most natural step is to extend known wavefunctions proposed in the literature [89, 90] with a neural network based Ansatz. For finding such architectures and furthermore understand inherited limitations a better understanding of when and why current approaches fail will be necessary. Here our study has shed some light onto this boundary by studying the scaling of the parameters required for finding suitable solutions for various systems. Efficient solutions exist in case of weak correlations, strong fragmentation or local entanglement and all cases sufficiently close to such a limit.

ACKNOWLEDGEMENTS

We thank Giacomo Passetti for fruitful discussions in the early stages of the project, further we thank Roser Valentí and Giuseppe Carleo for comments our manuscript. Results were obtained using NetKet [70, 91], a Jax [92] based library for NQS. Computations were performed with computing resources granted by RWTH Aachen University under project thes1588. JBP, DMK acknowledge funding by the DFG under RTG 1995, within the Priority Program SPP 2244 “2DMP” — 443274199. JBP gratefully acknowledge support from the DFG through FOR 5249 (QUAST, Project No. 449872909, TP5). DMK acknowledges support by the Max Planck-New York City Center for Nonequilibrium Quantum Phenomena.

-
- [1] A. Kitaev, Fault-tolerant quantum computation by anyons, *Annals of Physics* **303**, 2 (2003).
 - [2] D. Castelvecchi, How would room-temperature superconductors change science?, *Nature* **621**, 18–19 (2023).
 - [3] W. Kohn and L. J. Sham, Self-consistent equations including exchange and correlation effects, *Phys. Rev.* **140**, A1133 (1965).
 - [4] P. Hohenberg and W. Kohn, Inhomogeneous electron gas, *Phys. Rev.* **136**, B864 (1964).
 - [5] L. Hedin, New method for calculating the one-particle green’s function with application to the electron-gas problem, *Phys. Rev.* **139**, A796 (1965).
 - [6] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions, *Rev. Mod. Phys.* **68**, 13 (1996).
 - [7] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, Electronic structure calculations with dynamical mean-field theory, *Rev. Mod. Phys.* **78**, 865 (2006).
 - [8] G. Rohringer, H. Hafermann, A. Toschi, A. A. Katanin, A. E. Antipov, M. I. Katsnelson, A. I. Lichtenstein, A. N. Rubtsov, and K. Held, Diagrammatic routes to nonlocal correlations beyond dynamical mean field theory, *Rev. Mod. Phys.* **90**, 025003 (2018).
 - [9] M. Vandelli, J. Kaufmann, M. El-Nabulsi, V. Harkov, A. Lichtenstein, and E. Stepanov, Multi-band d-trilex approach to materials with strong electronic correlations, *SciPost Physics* **13**, 10.21468/scipostphys.13.2.036 (2022).
 - [10] Y. M. Vilchik and A.-M. Tremblay, Non-perturbative many-body approach to the hubbard model and single-particle pseudogap, *Journal de Physique I* **7**, 1309–1368 (1997).
 - [11] K. Zantout, S. Backes, and R. Valentí, Effect of nonlocal correlations on the electronic structure of lifeas, *Physical Review Letters* **123**, 10.1103/physrevlett.123.256401 (2019).
 - [12] N. E. Bickers and S. R. White, Conserving approximations for strongly fluctuating electron systems. ii. numerical results and parquet extension, *Phys. Rev. B* **43**,

- 8044 (1991).
- [13] G. Esirgen and N. E. Bickers, Fluctuation-exchange theory for general lattice hamiltonians, *Phys. Rev. B* **55**, 2122 (1997).
- [14] W. Metzner, M. Salmhofer, C. Honerkamp, V. Meden, and K. Schönhammer, Functional renormalization group approach to correlated fermion systems, *Rev. Mod. Phys.* **84**, 299 (2012).
- [15] C. Gros and R. Valentí, Cluster expansion for the self-energy: A simple many-body method for interpreting the photoemission spectra of correlated fermi systems, *Phys. Rev. B* **48**, 418 (1993).
- [16] D. Sénéchal, D. Perez, and M. Pioro-Ladrière, Spectral weight of the hubbard model through cluster perturbation theory, *Phys. Rev. Lett.* **84**, 522 (2000).
- [17] S. Zhang and H. Krakauer, Quantum monte carlo method using phase-free random walks with slater determinants, *Phys. Rev. Lett.* **90**, 136401 (2003).
- [18] W. M. C. Foulkes, L. Mitras, R. J. Needs, and G. Rajagopal, Quantum monte carlo simulations of solids, *Rev. Mod. Phys.* **73**, 33 (2001).
- [19] K. Van Houcke, E. Kozik, N. Prokof'ev, and B. Svistunov, Diagrammatic monte carlo, *Physics Procedia* **6**, 95–105 (2010).
- [20] S. R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [21] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Annals of Physics* **326**, 96–192 (2011).
- [22] J. I. Cirac, D. Pérez-García, N. Schuch, and F. Verstraete, Matrix product states and projected entangled pair states: Concepts, symmetries, theorems, *Rev. Mod. Phys.* **93**, 045003 (2021).
- [23] B. Bruognolo, J.-W. Li, J. von Delft, and A. Weichselbaum, A beginner's guide to non-abelian iPEPS for correlated fermions, *SciPost Phys. Lect. Notes* , 25 (2021).
- [24] C. Gros, Physics of projected wavefunctions, *Annals of Physics* **189**, 53 (1989).
- [25] S. Sorella, Wave function optimization in the variational monte carlo method, *Phys. Rev. B* **71**, 241103 (2005).
- [26] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, 2017).
- [27] M. C. Gutzwiller, Correlation of electrons in a narrow s band, *Phys. Rev.* **137**, A1726 (1965).
- [28] N. Lanatà, H. U. R. Strand, X. Dai, and B. Hellsing, Efficient implementation of the gutzwiller variational method, *Phys. Rev. B* **85**, 035133 (2012).
- [29] N. Lanatà, T.-H. Lee, Y.-X. Yao, and V. Dobrosavljević, Emergent bloch excitations in mott matter, *Phys. Rev. B* **96**, 195126 (2017).
- [30] D. Guerci, M. Capone, and M. Fabrizio, Exciton mott transition revisited, *Phys. Rev. Mater.* **3**, 054605 (2019).
- [31] M. Caffarel and W. Krauth, Exact diagonalization approach to correlated fermions in infinite dimensions: Mott transition and superconductivity, *Phys. Rev. Lett.* **72**, 1545 (1994).
- [32] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017), <https://www.science.org/doi/pdf/10.1126/science.aag2302>.
- [33] H. Lange, A. V. de Walle, A. Abedinnia, and A. Bohrdt, From architectures to applications: A review of neural quantum states (2024), arXiv:2402.09402 [cond-mat.dis-nn].
- [34] J. Robledo Moreno, G. Carleo, A. Georges, and J. Stokes, Fermionic wave functions from neural-network constrained hidden states, *Proceedings of the National Academy of Sciences* **119**, 10.1073/pnas.2122059119 (2022).
- [35] M. Medvidović and J. R. Moreno, Neural-network quantum states for many-body physics (2024).
- [36] C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators, *J. Res. Natl. Bur. Stand. B* **45**, 255 (1950).
- [37] M. K. Ritter, Y. Núñez Fernández, M. Wallerberger, J. von Delft, H. Shinaoka, and X. Waintal, Quantics tensor cross interpolation for high-resolution parsimonious representations of multivariate functions, *Phys. Rev. Lett.* **132**, 056501 (2024).
- [38] D.-L. Deng, X. Li, and S. Das Sarma, Quantum entanglement in neural network states, *Phys. Rev. X* **7**, 021021 (2017).
- [39] K. Choo, T. Neupert, and G. Carleo, Two-dimensional frustrated J_1-J_2 model studied with neural network quantum states, *Phys. Rev. B* **100**, 125124 (2019).
- [40] F. Ferrari, F. Becca, and J. Carrasquilla, Neural gutzwiller-projected variational wave functions, *Physical Review B* **100**, 10.1103/physrevb.100.125131 (2019).
- [41] K. McBrien, G. Carleo, and E. Khatami, Ground state phase diagram of the one-dimensional bose-hubbard model from restricted boltzmann machines, *Journal of Physics: Conference Series* **1290**, 012005 (2019).
- [42] J. Stokes, S. De, S. Veerapaneni, and G. Carleo, Continuous-variable neural network quantum states and the quantum rotor model, *Quantum Machine Intelligence* **5**, 12 (2023).
- [43] Z. Denis and G. Carleo, Accurate neural quantum states for interacting lattice bosons (2024), arXiv:2404.07869 [quant-ph].
- [44] K. Choo, A. Mezzacapo, and G. Carleo, Fermionic neural-network states for ab-initio electronic structure, *Nature Communications* **11**, 2368 (2020).
- [45] J. Stokes, J. R. Moreno, E. A. Pnevmatikakis, and G. Carleo, Phases of two-dimensional spinless lattice fermions with first-quantized deep neural-network quantum states, *Phys. Rev. B* **102**, 205122 (2020).
- [46] G. Pescia, J. Han, A. Lovato, J. Lu, and G. Carleo, Neural-network quantum states for periodic systems in continuous space, *Phys. Rev. Res.* **4**, 023138 (2022).
- [47] B. Fore, J. M. Kim, G. Carleo, M. Hjorth-Jensen, A. Lovato, and M. Piarulli, Dilute neutron star matter from neural-network quantum states, *Phys. Rev. Res.* **5**, 033062 (2023).
- [48] G. Passetti, D. Hofmann, P. Neiteimeier, L. Grunwald, M. A. Sentef, and D. M. Kennes, Can neural quantum states learn volume-law ground states?, *Phys. Rev. Lett.* **131**, 036502 (2023).
- [49] C. Gauvin-Ndiaye, J. Tindall, J. R. Moreno, and A. Georges, Mott transition and volume law entanglement with neural quantum states (2023), arXiv:2311.05749 [cond-mat.str-el].
- [50] H. Lange, F. Döschl, J. Carrasquilla, and A. Bohrdt, Neural network approach to quasiparticle dispersions in doped antiferromagnets, *Communications Physics* **7**, 187 (2024).

- [51] F. Döschl, F. A. Palm, H. Lange, F. Grusdt, and A. Bohrdt, Neural network quantum states for the interacting hofstadter model with higher local occupations and long-range interactions (2024), arXiv:2405.04472 [cond-mat.dis-nn].
- [52] J. Nys, G. Pescia, A. Sinibaldi, and G. Carleo, Ab-initio variational wave functions for the time-dependent many-electron schrödinger equation (2024), arXiv:2403.07447 [cond-mat.str-el].
- [53] J. Kim, G. Pescia, B. Fore, J. Nys, G. Carleo, S. Gandolfi, M. Hjorth-Jensen, and A. Lovato, Neural-network quantum states for ultra-cold fermi gases, *Communications Physics* **7**, 148 (2024).
- [54] Z. Denis, A. Sinibaldi, and G. Carleo, Comment on "can neural quantum states learn volume-law ground states?" (2024), arXiv:2309.11534 [quant-ph].
- [55] S. Dash, F. Vicentini, M. Ferrero, and A. Georges, Efficiency of neural quantum states in light of the quantum geometric tensor (2024), arXiv:2402.01565 [quant-ph].
- [56] D. Wu, R. Rossi, F. Vicentini, N. Astrakhantsev, F. Becca, X. Cao, J. Carrasquilla, F. Ferrari, A. Georges, M. Hibat-Allah, M. Imada, A. M. Läuchli, G. Mazzola, A. Mezzacapo, A. Millis, J. R. Moreno, T. Neupert, Y. Nomura, J. Nys, O. Parcollet, R. Pohle, I. Romero, M. Schmid, J. M. Silvester, S. Sorella, L. F. Tocchio, L. Wang, S. R. White, A. Wietek, Q. Yang, Y. Yang, S. Zhang, and G. Carleo, Variational benchmarks for quantum many-body problems (2023), arXiv:2302.04919 [quant-ph].
- [57] A. Chen and M. Heyl, Empowering deep neural quantum states through efficient optimization, *Nature Physics* (2024).
- [58] J. Chen, Equivalence of restricted Boltzmann machines and tensor network states, *Physical Review B* **97**, 10.1103/PhysRevB.97.085104 (2018).
- [59] S. R. Clark, Unifying neural-network quantum states and correlator product states via tensor networks, *Journal of Physics A: Mathematical and Theoretical* **51**, 135301 (2018).
- [60] D.-L. Deng, Quantum Entanglement in Neural Network States, *Physical Review X* **7**, 10.1103/PhysRevX.7.021021 (2017).
- [61] Y. Huang and J. E. Moore, Neural Network Representation of Tensor Network and Chiral States, *Physical Review Letters* **127**, 170601 (2021).
- [62] O. Sharir, Neural tensor contractions and the expressive power of deep neural quantum states, *Physical Review B* **106**, 10.1103/PhysRevB.106.w.205136 (2022).
- [63] D. Wu, From tensor-network quantum states to tensorial recurrent neural networks, *Physical Review Research* **5**, 10.1103/PhysRevResearch.5.L032001 (2023).
- [64] T.-H. Yang, M. Soleimanifar, T. Bergamaschi, and J. Preskill, When can classical neural networks represent quantum states? (2024), arXiv:2410.23152 [quant-ph].
- [65] J. M. Foster and S. F. Boys, Canonical configurational interaction procedure, *Rev. Mod. Phys.* **32**, 300 (1960).
- [66] C. David Sherrill and H. F. Schaefer, The configuration interaction method: Advances in highly correlated approaches, in *Advances in Quantum Chemistry* (Elsevier, 1999) p. 143–269.
- [67] A. Gnech, B. Fore, and A. Lovato, Distilling the essential elements of nuclear binding via neural-network quantum states (2023), arXiv:2308.16266 [nucl-th].
- [68] A. Lovato, C. Adams, G. Carleo, and N. Rocco, Hidden-nucleons neural-network quantum states for the nuclear many-body problem, *Phys. Rev. Res.* **4**, 043178 (2022).
- [69] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, Self-normalizing neural networks, in *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17* (Curran Associates Inc., Red Hook, NY, USA, 2017) p. 972–981.
- [70] F. Vicentini, D. Hofmann, A. Szabó, D. Wu, C. Roth, C. Giuliani, G. Pescia, J. Nys, V. Vargas-Calderón, N. Astrakhantsev, and G. Carleo, NetKet 3: Machine Learning Toolbox for Many-Body Quantum Systems, *SciPost Phys. Codebases*, **7** (2022).
- [71] Z. Liu and B. K. Clark, Unifying view of fermionic neural network quantum states: From neural network backflow to hidden fermion determinant states, *Phys. Rev. B* **110**, 115124 (2024).
- [72] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, *CoRR* **abs/1412.6980** (2014).
- [73] A. Malyshev, M. Schmitt, and A. I. Lvovsky, Neural quantum states and peaked molecular wave functions: Curse or blessing? (2024), arXiv:2408.07625 [quant-ph].
- [74] X. Li, J.-C. Huang, G.-Z. Zhang, H.-E. Li, C.-S. Cao, D. Lv, and H.-S. Hu, A Nonstochastic Optimization Algorithm for Neural-Network Quantum States, *Journal of Chemical Theory and Computation* **19**, 8156 (2023).
- [75] X. Cao, Z. Zhong, and Y. Lu, Vision transformer neural quantum states for impurity models (2024), arXiv:2408.13050 [cond-mat.str-el].
- [76] E. A. Carlen, E. H. Lieb, and R. Reuvers, Entropy and entanglement bounds for reduced density matrices of fermionic states, *Communications in Mathematical Physics* **344**, 655 (2016).
- [77] A. Bulgac, M. Kafker, and I. Abdurrahman, Measures of complexity and entanglement in many-fermion systems, *Phys. Rev. C* **107**, 044318 (2023).
- [78] C. Robin, M. J. Savage, and N. Pillet, Entanglement rearrangement in self-consistent nuclear structure calculations, *Phys. Rev. C* **103**, 034325 (2021).
- [79] N. Gigena and R. Rossignoli, Entanglement in fermion systems, *Phys. Rev. A* **92**, 042326 (2015).
- [80] E. Chertkov and B. K. Clark, Computational inverse method for constructing spaces of quantum models from wave functions, *Phys. Rev. X* **8**, 031029 (2018).
- [81] X.-L. Qi and D. Ranard, Determining a local hamiltonian from a single eigenstate, *Quantum* **3**, 159 (2019).
- [82] S.-Y. Hou, N. Cao, S. Lu, Y. Shen, Y.-T. Poon, and B. Zeng, Determining system hamiltonian from eigenstate measurements without correlation functions, *New Journal of Physics* **22**, 083088 (2020).
- [83] W.-J. Hu, F. Becca, A. Parola, and S. Sorella, Direct evidence for a gapless Z_2 spin liquid by frustrating néel antiferromagnetism, *Phys. Rev. B* **88**, 060402 (2013).
- [84] D. Luo and B. K. Clark, Backflow transformations via neural networks for quantum many-body wave functions, *Phys. Rev. Lett.* **122**, 226401 (2019).
- [85] M. Bajdich, L. Mitás, L. K. Wagner, and K. E. Schmidt, Pfaffian pairing and backflow wavefunctions for electronic structure quantum monte carlo methods, *Phys. Rev. B* **77**, 115112 (2008).
- [86] Y. Levine, O. Sharir, N. Cohen, and A. Shashua, Quantum entanglement in deep learning architectures, *Phys. Rev. Lett.* **122**, 065301 (2019).

- [87] X.-Q. Sun, T. Nebabu, X. Han, M. O. Flynn, and X.-L. Qi, Entanglement features of random neural network quantum states, *Phys. Rev. B* **106**, 115138 (2022).
- [88] G. Passetti and D. M. Kennes, Entanglement transition in deep neural quantum states (2023), arXiv:2312.11941 [quant-ph].
- [89] A. Haldar, O. Tavakol, and T. Scaffidi, Variational wave functions for sachdev-ye-kitaev models, *Phys. Rev. Res.* **3**, 023020 (2021).
- [90] V. Bettaque and B. Swingle, Nora: A tensor network ansatz for volume-law entangled equilibrium states of highly connected hamiltonians, *Quantum* **8**, 1362 (2024).
- [91] G. Carleo, K. Choo, D. Hofmann, J. E. T. Smith, T. Westerhout, F. Alet, E. J. Davis, S. Efthymiou, I. Glasser, S.-H. Lin, M. Mauri, G. Mazzola, C. B. Mendl, E. van Nieuwenburg, O. O'Reilly, H. Théveniaut, G. Torlai, F. Vicentini, and A. Wietek, Netket: A machine learning toolkit for many-body quantum systems, *SoftwareX*, 100311 (2019).
- [92] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: composable transformations of Python+NumPy programs (2018).
- [93] R. Eldan and O. Shamir, The power of depth for feed-forward neural networks, in *29th Annual Conference on Learning Theory*, Proceedings of Machine Learning Research, Vol. 49, edited by V. Feldman, A. Rakhlin, and O. Shamir (PMLR, Columbia University, New York, New York, USA, 2016) pp. 907–940.
- [94] H. W. Lin, M. Tegmark, and D. Rolnick, Why does deep and cheap learning work so well?, *Journal of Statistical Physics* **168**, 1223 (2017).
- [95] X. Gao and L.-M. Duan, Efficient representation of quantum many-body states with deep neural networks, *Nature Communications* **8**, 662 (2017).
- [96] G. Carleo, Y. Nomura, and M. Imada, Constructing exact representations of quantum many-body systems with deep neural networks, *Nature Communications* **9**, 5322 (2018).
- [97] Here deep architecture means deep Boltzmann machines. Deep Boltzmann machines might lead to exponential increase of cost, even for only polynomial parameter numbers. The situation considered there is therefore different to feed forward neural networks, where a polynomial increase in parameters implies a polynomial increase in computation time for an evaluation of the network. Deep Boltzmann machines are in this regard similar to higher dimensional tensor networks.
- [98] R. Rende, L. L. Viteritti, L. Bardone, F. Becca, and S. Goldt, A simple linear algebra identity to optimize large-scale neural network quantum states, *Communications Physics* **7**, 260 (2024).
- [99] M. B. Hastings, An area law for one-dimensional quantum systems, *Journal of Statistical Mechanics: Theory and Experiment* **2007**, P08024 (2007).
- [100] T. Westerhout, N. Astrakhantsev, K. S. Tikhonov, M. I. Katsnelson, and A. A. Bagrov, Generalization properties of neural network approximations to frustrated magnet ground states, *Nature Communications* **11**, 1593 (2020).




FIG. 8. Comparison of the performance of deep and shallow networks at the same parameter number. On the left we show the results for the complex hopping Hubbard model with $N_S = 10$ and on the right for the SYK model with $N_S = 18$. For the former, we consider shallow networks with no hidden layers and various numbers of hidden particles compared to deep networks with only one hidden particles. For the SYK model, we consider two hidden particles for the deep network. In general, shallow networks are outperforming deep ones at the same number of parameters.

Appendix A: Variational Ansatz and Optimization

Deep and Shallow Architecture — For many of the calculations shown in this paper shallow neural networks were employed, even though it is known that deep neural networks are able to outperform shallow networks in their ability to approximate functions efficiently in some cases. I.e. for some functions an exponential increase in efficiency can be observed [93, 94]. Furthermore, various works imply an advantage of deep architectures over shallow architectures specifically for the problem of targeting physical wavefunctions [95, 96][97]. Therefore we want to address possible concerns here that our observations are due to considering shallow networks instead of deep ones.

When choosing our architectures we mainly had two things in mind: Firstly for the non SYK-model defined in eq. 7 with conserved total magnetization the number of parameters compared to the dimension of the Hilbert space would be large for small systems and deep networks. This leads to the ability to learn the wavefunction up to very high accuracy even for a single hidden particle. Nonetheless, the ability to learn the target state can not be interpreted as the neural network finding efficient representations. The number of parameters of the trial state is already approximately equal to the dimension of the Hilbert space. In other words: it is not necessary to learn the structure of the target state. The neural networks already have enough capacity achieve high accuracies on noise. Secondly, we did not observe a significant influence of the depth in the ability to learn the target states when keeping the parameter number fixed, see Fig. 8. This observation is in line with Ref. [34] where the authors also found that adding more hidden particles generally leads to lower expectation values for the energy than growing out the neural networks.

Optimization Algorithm — For our optimization we have used the Adam algorithm [72] instead of the typi-




FIG. 9. Comparison of the loss functions during the training of the random hopping Hubbard model with 8 sites and various seeds for the construction of the Hamiltonian for networks with no hidden layer and one/three hidden particle on the left/right. The learning rates are 0.05 and 0.005 respectively.




FIG. 10. Training curves for the data points used for the Hubbard model in fig. 4 with $N_S = 12$ and $M = 16$. A flattening out of the training curves can be observed at the end of the training.

cally employed stochastic reconfiguration algorithm (SR) for the training of NQS. This does not seem to have an influence on the achievable model performance as can be seen in Fig. 9.

Because we are not able to employ the kernel trick proposed in Refs. [57, 98], due to contracting over the whole Hilbert space instead of making use of variational Monte-Carlo for the calculation of the quantities necessary, SR becomes very expensive per step in comparison to the ADAM-algorithm for large models. Even when needing fewer steps to converge, ADAM is still cheaper in terms of computational resources. Furthermore, the optimization via SR did sometimes get stuck at $\delta O \sim 0$. Which is most likely an effect of the spectral properties of the Projector.

Lastly we have observed that the quantum geometric tensor is ill behaved during some steps of the optimization, especially for the SYK-models. Thus we need to apply large regularization corrections to the quantum geometric tensor leading to sub optimal performance of this method. The low rank behavior of the quantum geometric tensor in certain cases was also studied in Ref. [55].

Because of these reasons we concluded that the ADAM-optimizer is a better fit for our use-case.

Convergence — For larger systems with many parameters, fully converging the training curves is becoming increasingly expensive. For this reason the training was only performed until the training process was flattening

out, see Fig. 10. Here the training is transitioning into a regime where a small improvement in training accuracy would require an exponentially fast increasing number of training steps.

Appendix B: Interpolation

The data in Fig. 4 was interpolated (and extrapolated) in order to estimate the number of parameters needed to achieve a given accuracy for various system sizes. Here, the target error was chosen as small as possible without relying on extrapolation of our data to predict the number of parameters necessary for the error chosen. For the purpose of estimating the number of parameters necessary to reach a certain accuracy bound two methods were employed: a piece-wise and a global interpolation.

Piecewise interpolation — For the piecewise interpolation we estimated the mean δO as well as the error of the mean for each data point first and obtained the estimates for the number of parameters as well as the errors via a simple interpolation process between neighboring points. Here we used an algebraic fit, i.e. Ax^{-m} for all models except the SYK-model 7, where we have observed a better agreement for the function Ae^{-mx} , where x corresponds to the number of parameters.

A problem of this method is that errors due to convergence problems for large models are not considered as these are hard to train. The introduced bias towards larger parameter numbers for trial wave functions of large systems results in a bias towards the exponential scaling hypothesis. This problem should at least be partially cured by the global interpolation.

Global interpolation — For the models discussed in the main text we have observed a good agreement of our data for a given system size to our fitting functions if one excludes the $M = 0$ point, i.e. the pure Slater-determinants. Therefore, we additionally interpolated the data points globally for a given system size and a given seed used for the generation of the matrix elements appearing in the Hamiltonian. This allows us again to estimate the number of parameters given δO . The reported estimate now corresponds to the mean of these predictions with the error bars estimated by the sample variance.

This approach is not as responsive to small errors due to convergence problems in the largest models. Still a good agreement between both methods can be observed for our data points. Therefore, we conclude that the increase of the computational cost observed above is not due to not fully converging out the largest systems.

Because $M = 0$ and $M = 1$ are the only two data-points relevant for the smallest system with pure density-density-interaction, we kept the $M = 0$ point in this case. A similar problem occurred for the smallest system size for the SYK-model. Here we decided to skip this system size for this model for the global interpolation as we were able to achieve results for the same number of system sizes as for the other models considered.

The difference in the fitting function for the global fit is necessary due to data-constraints. For the piece-wise interpolation corrections are in general still small (see also Fig. 15). Our observed scaling for the SYK-model opposed to all of the other models observed in Fig. 4 is therefore not an artifact of differences in the interpolation method.

Appendix C: Large U -Behaviour




FIG. 11. Expectation value of \hat{x} for the groundstate of the model with the density-density-interaction (a), spin-spin interaction (b), pairing interaction (c) averaged over various realizations. The scale of the interaction was fixed at $U = 1$. For the spin-spin model double occupations are completely suppressed, while for the pairing based model we see that single occupations are suppressed. For the density-density model the behavior is more complicated due to the possible \mathbb{Z}_2 symmetry.

In order to gain a better understanding of the behavior of the models featuring a density-density interaction, a spin-spin interaction and a pairing interaction in the limit of large coupling strengths, we calculate

$$\hat{x} = \frac{1}{N_S} \sum_i 1 + 2\hat{n}_{i\uparrow}\hat{n}_{i\downarrow} - \hat{n}_{i\downarrow} - \hat{n}_{i\uparrow} \equiv \frac{1}{N_S} \sum_i \hat{x}_i \quad (\text{C1})$$

for various system sizes and ratios J/U . The results are visualized in Fig. 11. The \hat{x}_i is defined in such a way that it returns zero for single occupied sites and one for doubly occupied ones. We observe in the model with the spin-spin interaction double occupations are suppressed for large interactions. The model becomes effectively a spin model.

A similar behavior can be observed for the model with the pairing interaction; here single occupations are suppressed. After performing a particle-hole transformation of the spin-down species, the model can be mapped back to the spin-spin model (in the absence of a kinetic term). Therefore, the model is transitioning into a regime with where all electrons are paired and single occupied sites are projected out.

For the density-density interaction model the expectation value of \hat{x} is not converging towards 1 for $J \rightarrow 0$ for all system sizes considered here. This can be traced to the interacting Hamiltonian being already diagonal in the computational basis and we find that at least for some of the realizations the ground state is degenerate. The corresponding symmetry-group here is \mathbb{Z}_2 : This symmetry




FIG. 12. Correlations measured by $S(\gamma_1)$ (a) and fragmentation measured by S_ψ divided by the maximal possible value (b) for the groundstate of the density-density-interaction model for various system sizes averaged over multiple realizations. Correlation stay small for all system sizes considered here, even though they are not tending to zero. This indicates a degenerate groundstate in the $J \rightarrow 0$ limit. For the strongly interacting case strong fragmentation can be observed.

is originating from a computational basis state mapping either onto itself or a different computational basis state under a global spin flip. In the latter case we observe $\langle \hat{x} \rangle \neq 1$ in the former $\langle \hat{x} \rangle = 1$. The degeneracy is lifted by the kinetic term. In Fig.C (a) it can be seen that the degenerate groundstate appears for all system sizes considered. Subplot (b) shows that the target space is increasingly peaked in the strong interaction limit as would be expected. While the support of the target space seems to be growing with the dimension of the Hilbert space even in the strongly interacting case in the limit $J = 0$ this behavior brakes down. While seemingly still leading to an exponential grows of the relevant sub-sector the prefactor becomes arbitrarily small in this limit, indicating the (maximally) two dimensional space spanned by the ground states in this limit.

Furthermore, as indicated by $S(\gamma_1)$ in subplot (a) of Fig.C, the groundstate of this model is never strongly correlated – $S(\gamma_1)$ stays far below the maximal possible value. This renders these states simpler to learn for Slater-determinant based approaches than expected from considering the fragmentation.

These effects are explaining the ability to represent the groundstate of the density-density interaction model with fewer parameters compared to all other models. The state has exploitable structure, i.e. it is close to a gaussian state and fragmented.

Appendix D: Application to Area-Law States

This appendix serves the purpose of testing our method on one dimensional area-law states [99]. If we were not able to observe a reduction in complexity in this case, it would indicate a problem with our analysis.




FIG. 13. Entanglement entropy for various system sizes for the model defined in eq. D1 with $J = 0$. Even for these small system sizes the entanglement entropy does not seem to be dependent on the system size, indicating area-law scaling of the entanglement entropy for this model.

To this end we consider

$$\hat{H} = \frac{J}{2} \sum_{(i,j) \in S_2} J_{ij}^{[2]} \hat{c}_i^\dagger \hat{c}_j + \frac{U}{\sqrt{8^3}} \sum_{(i,j,k,l) \in S_4} J_{ijkl}^{[4]} \hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_l \hat{c}_k, \quad (\text{D1})$$

where $J^{[2]}$ and $J^{[4]}$ are sampled from the same distribution as they were for Eq. 7. S_2 is the set of tuples with indices at most one site apart, i.e. on-site and neighboring sites, and S_4 is the set of quadruples where the elements are at most four sites apart. Results for the entanglement entropy can be found in Fig. 13 and the results for the degree of correlations and sparsity of the wavefunction are shown in Fig. 14. In what follows we will be concerned with the case $J = 0$. The target state is supported by a large subset of our computational basis, contains correlations and shows a small entanglement entropy that stays constant in system size, i.e. the groundstate of the system is a one-dimensional area-law state.

Due to these properties of the model we are able to use it to test the influence of low entanglement on the representability of ground states within the HFDS approach – neither of the other two complexities implies that the state is particular easy to learn.




FIG. 14. Correlations measured by the von-Neumann entropy of the 1-RDM (a) and the entropy of the Born-distribution (b) for various system sizes and various kinetic energy strengths given by J . The interaction scaled was kept fixed at $U = 1$. In what follows we will be concerned with the case $J = 0$. In this case S_ψ indicates a broad born distribution as can be seen in subplot (b). Also $S(\gamma_1)$ is indicating a correlated groundstate.

The scaling of the number of parameters with the system size is shown in Fig. 15. Here we have chosen to show the results for the exponential interpolation and the algebraic interpolation. While the fit for the algebraic methods seems to be working better, the final result for the scaling of the number of parameters with the system size is largely independent of this choice.

Furthermore, beyond the reduction of the scaling in comparison to all models (except the density-density interaction model) a strong tendency towards polynomial scaling can be observed. This is indicated by the root mean error of the polynomial fit. From the reduction of computational cost, we conclude that the increase of the computational cost is not due to the method chosen for the estimation of parameters, but due to the complexity of the target states. In other words, Area-law scaling is an exploitable non complexity in one dimension. In order to make other states learnable other non complexities are necessary. In this light the trend towards polynomial scaling can be expected and small entanglement entropy can be added to the set of exploitable quantities.

While a scaling with exponent $m \sim 3.6$ seems very expensive when compared to tensor network methods, we still want to highlight that in order to make this wavefunction extensive/size-consistent one has to choose $M \propto N_S$, already implying $m = 3$ asymptotically. Still, this does lead to an increase of error in system size, as the error on different subsystems also needs to be reduced in order to keep the total error constant, implying $m > 3$.

Appendix E: Phase structure of the Target State

It was already observed that the ability of NQS to find suitable approximations for the target state is in practice often hindered by their ability to obtain sufficiently good approximations for the sign- or phase-structure of the target state [100]. To link this observation back to our observed behavior, we try to learn the sign structure of the ground state in a Slater-determinant basis. For this purpose we will train the trial state

$$\langle n | \Psi_{var} \rangle = \frac{\langle n | \Phi \rangle}{|\langle n | \Phi \rangle|} |\langle n | \Psi \rangle|, \quad (\text{E1})$$

where $|\Psi\rangle$ corresponds to the target state and $|\Phi\rangle$ is chosen to be the trainable Slater-determinant. The trial is now non-continuous when elements of $\langle n | \Phi \rangle$ are passing through zero, which complicates the training procedure.

If the state is close to a Slater-determinant, as indicated by a small value for $S(\gamma_1)$, one already trivially expects a high fidelity for this approach, as the whole wavefunction including the amplitude can be recovered approximately by only considering Slater-determinants. A Slater-determinant can only be insufficient to determine the phase-structure if $S(\gamma_1)$ is large enough. In other words, with this approach we probe how strongly the sign structure of the ground-state wave-function de-

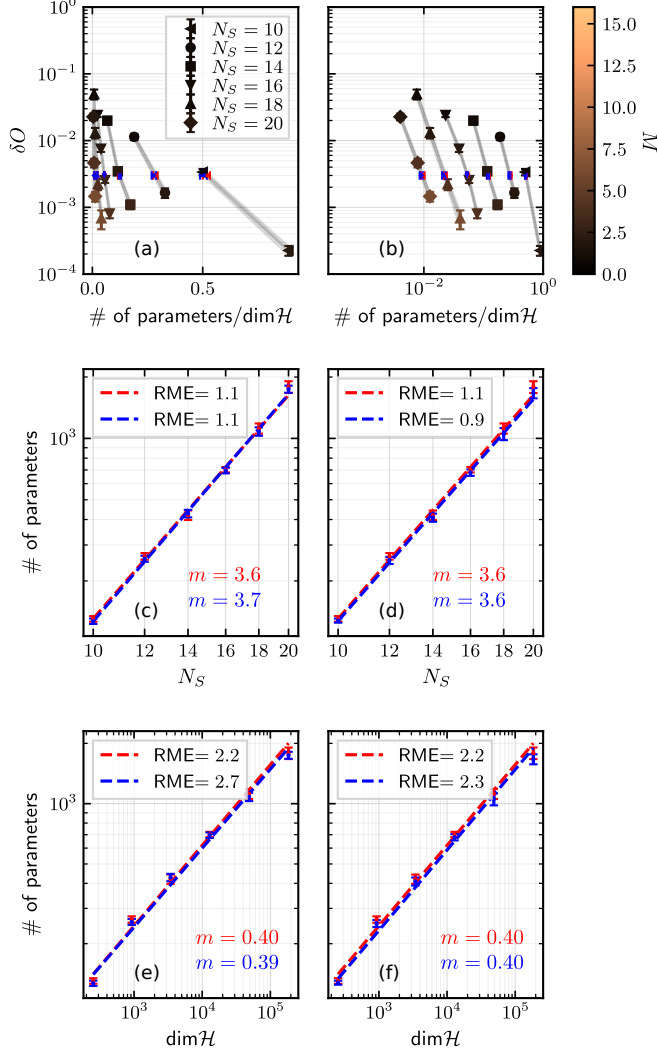


FIG. 15. Estimates for the scaling of the number of parameters with the system size for model Eq. D1. For each system size 30 realizations of the model were used. Here no hidden layers were used. This choice was taken due to the necessity of considering models with fewer parameters in comparison to the fully connected SYK-model. For the left column the exponential hypothesis was used for the interpolation and for the right column the algebraic one was used (see Appendix B). While the algebraic hypothesis provided a better fit for the data, the final results (c-f) show now strong dependence on the method chosen. Here the estimates for the scaling of the parameters can be found in subplots (c-f), first row with the polynomial scaling hypothesis and in the second row with the exponential scaling hypothesis. Here the polynomial is clearly favored over the exponential one. The RME is expected to take the value of one assuming a Gaussian distribution of the estimated means around their actual values. The deviation of the data point for the largest system in (c-f) can again be explained by the difficulty of training these larger models.

viates from the one of the Slater-determinant of the single particle orbitals.

Results for the SYK-model can be found in Fig. 16 (a). For this model, we observe that, while able to achieve a high fidelity in the weak interaction limit, a Slater-determinant is insufficient to capture the complex phase structure of the SYK-model in the strongly interacting limit. Yet, even when the proper sign-structure is restored, are still not able to represent the groundstate of the SYK-model in the absence of a kinetic term (see Fig. 17).




FIG. 16. Achieved infidelity when using the phase of a Slater-determinant and the amplitude of the exact wavefunction for the groundstate of the SYK-model (a), the Hubbard model (b) and the model featuring a spin interaction (c) for various system sizes and interaction strengths U and J fixed at one. The results were averaged over ten realizations. We observe that a Slater determinant not able to capture the sign structure of the SYK-model in the large interaction limit, while the performance on the other models is far better throughout the whole parameter range.

Qualitatively different behavior can be observed in Fig. 16 (b-c). These plots show the results for the trial states defined in eq. E1 trained on the groundstate of the Hubbard-model and the model with the spin interaction respectively. For these Slater-determinants are not only able to capture the sign structure in the weak interaction limit but also in the strong interaction limit. This indicates a deviation from the behavior observed for the SYK-model. Nonetheless, it is not clear whether this is due to the ground state again obtaining a more Slater-determinant like sign-structure or due to the ability of the Slater-determinant to more easily over fit the phase structure of the most important subspace due to an effective reduction of dimension in this limit, see Appendix G.

If the phase-structure of the groundstate is approximately given by a Slater-determinant, a Slater-determinant with a Gutzwiller style correlator is a reasonable variational ansatz. This ansatz is able to make use of both these properties when constructing approximations for the groundstate. In the following we compare the HFSD ansatz to the an inhomogeneous Gutzwiller correlator acting on a Slater-determinant:

$$\langle n | \Psi_{iGW} \rangle \propto \langle n | e^{-\sum_i \lambda_i \hat{n}_i^\uparrow \hat{n}_i^\downarrow} | \Phi \rangle. \quad (\text{E2})$$

Results for this wavefunction on the Hubbard model can

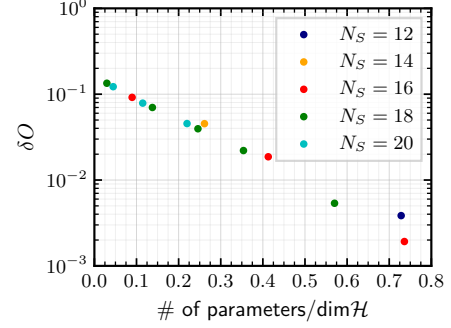


FIG. 17. Here we have used a real multilayer perceptron with selu-hidden activation functions and $4N_S$ neurons per hidden layer and various depths for the logarithm of the amplitude of the wavefunction, while the phase was fixed by the exact phase of the groundstate obtained via exact diagonalization. No error-bars are shown, because the errors are smaller than the markers used in all cases. The relative number of hidden parameters required to achieve a given accuracy is again approximately proportional to the dimension of the Hilbert space indicating exponential scaling of the parameters required in the system size. The problems of NQSs when applied to the groundstates of the SYK-model can not be explained by the inability of neural quantum states to learn the correct sign structure.

be found in Fig. 18 and on the model featuring the spin-spin interaction in Fig. 19. The ability to achieve low infidelities on the groundstate indicates that the actual target state can be approximated reasonably well by a simple Gutzwiller-style wavefunction. This indicates that for both models the groundstate is close to a Gutzwiller-style wavefunction, at least for the small system sizes tested here, leading again to a reduction in complexity.

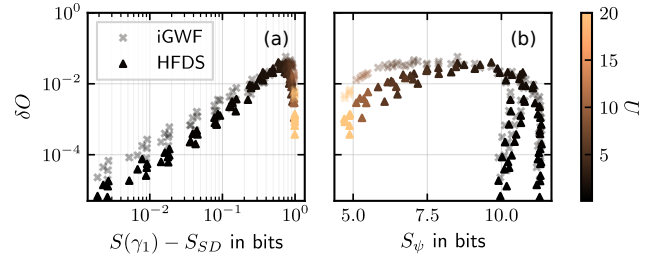


FIG. 18. δO as a function of von-Neumann entropy of the 1-RDM (a) and the entropy of the Born-distribution (b) for various interaction strengths for the fully connected Hubbard-Model with eight sites for $M = 1$ as well as for the inhomogeneous Gutzwiller wavefunction. Here no hidden layers were used. The J controlling the kinetic term is fixed to one.

Remarkably, in case of the Hubbard model the achieved error with the inhomogeneous Gutzwiller wavefunction is already similar to the HFDS's error, even though it requires far fewer parameters. Furthermore, we observed that the overlap only improves in the large

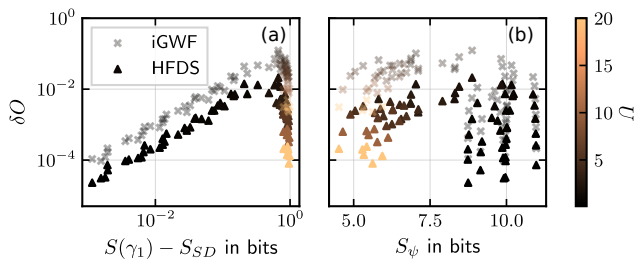


FIG. 19. δO as a function of von-Neumann entropy of the 1-RDM (a) and the entropy of the Born-distribution (b) for various interaction strengths for the model with the spin interaction channel with eight sites for $M = 1$ as well as for the inhomogenous Gutzwiller wavefunction. Here no hidden layers were used. The J controlling the kinetic term is fixed to one.




FIG. 20. Entanglement for the Hubbard model when choosing the bipartition between the two spin sectors instead of splitting based on lattice sites for $U = 3$. This reduces the entanglement compared to what was shown in Fig. 3, indicating that its ground-state is expected to be far more compressible than the SYK-model.

interaction limit, once the the problems becomes effectively fragmented pointing again to this effect being due to the fragmentation and not due to the phase becoming slater-like.

Appendix F: Entanglement structure of the Hubbard-model

For the Hubbard model the entanglement within the spin subsectors is generated by the kinetic term. This entanglement can be captured due to the Slater-determinant based construction of the trial state used. In contrast the interaction term, the part for which makes the usage of the neural networks necessary, generates, to lowest order, local entanglement between the two spin sectors. This distinguishes the Hubbard-model from the remaining models, because for the remaining models also the interaction can generate non-local entanglement. For this reason have we also calculated the entanglement of the Hubbard model for a bipartition between the spin sectors. Considering this bipartition instead of a local




FIG. 21. Mutual information for the fully connected Hubbard-model with $N_S = 8$ at half filling for various U . The kinetic energy scale was fixed by setting $J = 1$. The upper row contains the results for $\sigma = \sigma'$ and the lower row the results for $\sigma \neq \sigma'$. The results are averaged over six realizations. Here the localization of the entanglement can clearly be seen.

bipartition leads to a reduction in entanglement as can be seen in Fig. 20 when compared to Fig. 3.

Yet, parts of the entanglement shown in Fig. 20 are still not exponentially hard to capture, as for instance the suppression of double occupation in the large U -limit, an effect that could be described by a simple Gutzwiller-correlator as done in Fig. 18, is generating entanglement. To further understand this, one can considering the mutual information

$$I_{ij}^{\sigma\sigma'} = S(\rho_i^\sigma) + S(\rho_j^\sigma) - S(\rho_{ij}^{\sigma\sigma'}), \quad (\text{F1})$$

where ρ_i^σ is the density matrix corresponding to the orbital i with spin σ and $\rho_{ij}^{\sigma\sigma'}$ is the density matrix corresponding to the orbitals (i, σ) and (j, σ') respectively. Here we have defined $\rho_{ii}^{\sigma\sigma} := \rho_i^\sigma$, see Fig. 21 for the Hubbard model case. In subplots (b,e) we observe that the local Hubbard-style interaction are leading to local entanglement on the sites between the spin-sectors.

This general structure of the entanglement is further highlighted in Fig. 22. The spectrum corresponding to a local bipartition of the system can be captured by a slater-determinant, while a clear deviation from the exact results can be seen for the bipartition with respect to the spin sectors. However, the entanglement between the spin sectors is captured by a inhomogenous Gutzwiller wavefunction. Furthermore, HFDS are already able to improve significantly once a single hidden particle is considered.

Appendix G: Away from half filling

In this appendix, we consider the effect of changed filling on the ability to represent the wavefunction. For this purpose we consider the Hubbard model with one electron per spin species removed and compare these results




FIG. 22. Entropy spectrum of the singular values ξ_i of the reduced density matrix calculated for a realization of the fully connected Hubbard model with $N_S = 10$ and $U/J = 3$. Here the bipartition by spin sector is compared to a local bipartition. We compare a HFDS without hidden layers for $M = 1$ as well as $M = 1$ to exact results as well as the inhomogeneous Gutzwiller wavefunction. $\xi_i < 10^{-14}$ were discarded.




FIG. 24. Achieved infidelity when using the phase of a Slater-determinant and the amplitude of the exact wavefunction trained on the groundstate of the Hubbard-model with 3 particles per spin species for various interaction strengths U . J is fixed at one. The results were averaged over ten realizations.

to Fig. 18. Naively one would expect worse performance in the large U limit as the ground-state sector is now larger. This is confirmed in our calculations, see Fig. 23. Furthermore, this model features a phase-structure for which the Slater-determinant achieves worse performance in the large interaction limit as shown in Fig. 24 when compared to the case of half filling shown in Fig. 16 (b) also leading to worse performance of the inhomogeneous Gutzwiller wavefunction in the large interaction limit as shown in Fig. 16 (b) when compared to the case of half filling. However, it is not clear whether this is an actual feature of the sign structure or simply due to the dimension of the ground state block being strictly separated from the Slater-determinants in terms of parameter number.

Appendix H: Complexities and Compression




FIG. 23. Same as Fig. 18, but with 3 instead of 4 electrons per spin species. δO as a function of von-Neumann entropy of the 1-RDM (a) and the entropy of the Born-distribution (b) for various interaction strengths for the fully connected Hubbard-Model with eight sites for HFDS with $M = 1$ as well as the inhomogeneous Gutzwiller wavefunction. Here again no hidden layers were used. The J controlling the kinetic term is fixed to one.

The three complexity measures, entanglement-entropy, entropy of the Born-distribution and the von-Neumann-entropy of the 1-RDM, considered above all allow an efficient compression of the ground-state. However, as shown in Fig. 25, the ground-state is approximated by less correlated, less entangled states with lower entropy of the Born distribution. Furthermore, in the second row of Fig. 25 we demonstrate that we are able to converge these differences to zero by increasing the number of hidden particles for the largest system size considered here. In general, we find the rather simple statement that complex states are approximated by simpler ones. Furthermore, while the exact compression performed by neural networks usually can not be understood, these measures give insights how to improve the networks. Especially in case of the entanglement entropy it would be interesting to investigate whether it is more localized than for the exact ground state. Fig. 25 also shows the growths of the entanglement entropy of the trained trial state.




FIG. 25. Comparison of the complexity measures for the SYK-model for the HFDS with two hidden layers and the exact state. In the first row (a-c) are the results for $M = 1$ and various system sizes. In the second row (d-f) are the results for $N_S = 20$ and various M . For the estimation of the means and errors we average over six realizations. The results for the entanglement-entropy of the 1-RDM are found in the first column (a,d), for the Shannon-entropy of the born distribution in the second column (b,e) and for the bipartition entanglement entropy in the last column (c,f).