

Generating Generalized Ground-State *Ansätze* from Few-Body Examples

Matt Lourens^{1,*}, Ilya Sinayskiy^{2,3}, Johannes N. Kriel¹, and Francesco Petruccione^{1,3,4}

¹Department of Physics, Stellenbosch University, Stellenbosch, South Africa

²School of Agriculture and Science, University of KwaZulu-Natal, Durban, South Africa

³National Institute for Theoretical and Computational Sciences (NITheCS), Stellenbosch, South Africa

⁴School of Data Science and Computational Thinking, Stellenbosch University, Stellenbosch, South Africa



(Received 3 April 2025; revised 13 August 2025; accepted 16 October 2025; published 20 November 2025)

We introduce a method that generates ground-state *Ansätze* for quantum many-body systems that are both analytically tractable and accurate over wide parameter regimes. Our approach leverages a custom symbolic language to construct tensor network states via an evolutionary algorithm. This language provides operations that allow the generated tensor network states to automatically scale with system size. Consequently, we can evaluate *Ansatz* fitness for small systems, which is computationally efficient, while favoring structures that continue to perform well with increasing system size. This ensures that the *Ansatz* captures robust features of the ground-state structure. Remarkably, we find analytically tractable *Ansätze* with a degree of universality, which encode correlations, capture finite-size effects, accurately predict ground-state energies, and offer a good description of critical phenomena. We demonstrate this method on the Lipkin-Meshkov-Glick model and the quantum transverse-field Ising model, where the same *Ansatz* was independently generated for both. The simple structure of the *Ansatz* allows us to obtain exact expressions for the expectation values of local observables as well as for correlation functions. In addition, it permits symmetries that are broken in the *Ansatz* to be restored, which provides a systematic means of improving the accuracy of the *Ansatz*.

DOI: [10.1103/5vc5-9f4d](https://doi.org/10.1103/PhysRevLett.135.210401)

Obtaining an exact ground-state solution for an interacting quantum many-body system is generally a very difficult, if not completely intractable, task. As a result, approaches to this problem are often based on a variational *Ansatz*, i.e., a simplified functional form of the ground state intended to capture the latter's essential physical features. A structurally simple *Ansatz* with few parameters allows for analytic calculations and provides qualitative insights often at the expense of quantitative accuracy. Conversely, a structurally and variationally complex *Ansatz* requires a fully numeric approach but offers improved quantitative accuracy at the expense of qualitative insight. Balancing these qualitative and quantitative extremes is challenging. Consequently, constructing *Ansätze* that permit an analytic treatment while yielding accurate results over a wide range of system parameters is highly desirable.

For qualitative insights, a simple *Ansatz* is to construct a product state with minimal variational parameters, as is typical in mean-field theory (MFT) [1–5]. This approach

and its extensions [6–12] function by neglecting fluctuations, and offer a low-cost procedure to obtain analytic insight into a system. However, in the vicinity of critical points these fluctuations become large. Here, a different method is required, such as the renormalization group [13–15].

For quantitative accuracy, a powerful class of variational *Ansätze* is tensor network states (TNSs). A special case, matrix product states (MPSs) [16–18], naturally represent low-energy states of systems in 1D with local interactions [19,20]. Other classes of network states include projected entangled pair states [21–24], the multi-scale-entanglement-renormalization *Ansatz* [25], and tree tensor networks [26–29]. While TNSs are broadly applicable, their parameter count typically scales linearly with system size and polynomially with bond dimension, limiting analytical tractability.

For both qualitative insights and quantitative accuracy, we introduce a method to generate tensor network states with minimal structural and variational complexity while preserving high accuracy. We leverage a domain-specific language (DSL)—specific syntax and rules for compactly expressing TNSs via modular building blocks—implemented as an open source Python package [30,31]. These blocks encode the system's size scaling, spatial homogeneity and correlations. An evolutionary algorithm exploits this DSL to construct low-energy tensor network

*Contact author: lourensmattj@gmail.com

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. An overview of our method. A domain-specific language enables *Ansatz* generation via an evolutionary algorithm. (a) A primitive (cycle, mask, or pivot) is an edge generation pattern associated with a tensor. (b) Composition: sequences of primitives form motifs; sequences of motifs form higher-level motifs. (c) A primitive applied to an initialized state (specified with the “init” command) forms a tensor network. (d) A specified network, being itself a tensor, can again be associated with an edge generation pattern to form a new primitive. (e) The evolutionary algorithm mutates and crosses over motifs each generation. (f) Once the *Ansatz* is found, broken symmetries can be restored.

states for a given Hamiltonian. Our approach shares several complementary threads with recent works. In variational quantum algorithms, searching over circuit structures yields *Ansätze* with fewer gates yet comparable performance [32–34]. Using few-site information to build scalable wave functions has proved effective in many-body methods such as density matrix embedding theory [35] and variational Monte Carlo [36], likewise in scalable variational quantum eigensolver constructions [37–39]. The Tequila framework [40] demonstrates the power of a domain-specific language and abstract data structures in the context of quantum algorithm development.

Method—The elements of our method are outlined in Fig. 1. Within the DSL the elementary building blocks are called “primitives,” of which we show three: a “cycle,” “pivot,” and “mask.” Examples of DSL usage are provided in Sec. V of Supplemental Material (SM) [41]. See also Refs. [30,31] for a more detailed account. In brief, the cycle and pivot primitives are defined by size-independent properties, the main two being an edge generation pattern with an associated tensor (a). Other such properties include edge order, weight sharing, and boundary conditions. The mask is a special primitive that hides or unhides physical indices based on some pattern. These primitives can be composed sequentially (b) to create a sequence of primitives called a “motif.” In turn, multiple motifs can be composed to form higher-level motifs, and so on. In this way motifs are built without specifying a system size, and size scaling is captured within the properties of primitives. System size is specified through the init command, which initializes the state with a tensor (c). When a motif is applied to an initialized state, a tensor network is formed. Since a tensor network is itself a tensor, it can again be associated with an edge generation pattern, thereby forming a new primitive, and allowing larger networks to be built hierarchically from subnetworks (d).

Our evolutionary algorithm (e) follows the design of those in neural architecture search [42,43] which is based on tournament selection [44]. The optimal *Ansatz*

produced by this algorithm will generally not exhibit the same symmetries as the system Hamiltonian. However, with sufficient complexity penalties the resulting structure will be simple enough for broken symmetries to be restored (f). The algorithm starts with a randomly initialized pool of primitives. The tensors associated with these primitives are chosen from a fixed set and contain variational parameters. A tournament is then repeatedly held between a set fraction $\rho \in [0, 1]$ of the total population. Here, ρ denotes the selection pressure, with large values facilitating exploitation. During each tournament, the winners (defined below) are used to generate offspring through mutation, crossover, and other genetic operations using the DSL. The winners are typically the two fittest individuals, but there is a probability $\epsilon \in [0, 1]$ that two random individuals are declared winners instead. Large values of ϵ therefore facilitate a wider exploration of possible network structures. We evaluate the fitness $f(\psi_a)$ of an *Ansatz* ψ_a over system sizes $n \in \mathcal{N}$ with penalties on variational and structural complexity according to

$$f(\psi_a) \equiv \sum_{n \in \mathcal{N}} (E(\psi_a^{(n)}) + l_1 S(\psi_a^{(n)}) + l_2 V(\psi_a^{(n)})) w_n. \quad (1)$$

Here, E , S , and V are respectively the per-site energy expectation value, the structural complexity, and variational complexity of an individual. The parameters l_1 and l_2 control the penalties associated with these two complexities, while $w_n \in [0, 1]$ are normalized weights assigned to the various system sizes in $\mathcal{N}$. The variational complexity $V(\psi_a^{(n)})$ of an individual is the number of variational parameters it contains. The structural complexity $S(\psi_a^{(n)})$ is the sum of the ranks of the tensors that make up the *Ansatz*, divided by $2n$. During the search, a main worker process maintains a queue of unevaluated *Ansätze*, of which the fitnesses have not yet been calculated, together with a pool of evaluated ones. Each tournament enqueues unique *Ansätze* while duplicates only increase the multiplicities




FIG. 2. The *Ansatz* generated by our method for the LMG and TFIM models.

of network structures already in the pool. We define one evolutionary step as the addition of ten unique evaluated *ansätze* to the pool.

Ansatz structure and expectation values—We demonstrate our method on the Lipkin-Meshkov-Glick (LMG) model [45–47] and the quantum transverse-field Ising model (TFIM) [48]. For both models, the search yields the *Ansatz* shown in Fig. 2. For N spins it produces the state

$$|\psi_o\rangle = \left(\prod_{k=0}^{N-1} C_{k,k+1}^\theta R_{k+1}^\theta \right) \left(\prod_{j=0}^{N-1} R_j^\phi \right) |z, +\rangle^{\otimes N}, \quad (2)$$

where ϕ and θ are variational parameters and $|z, \pm\rangle$ are the eigenstates of the Pauli-Z matrix with eigenvalues ± 1 . The operator products in Eq. (2) are ordered with the k and j indices decreasing from left to right. The two unitary operators appearing in $|\psi_o\rangle$ are

$$C_{ij}^\theta = e^{i\frac{\theta}{2}Z_i Y_j} \quad \text{and} \quad R_j^\theta = e^{-i\frac{\theta}{2}Y_j}. \quad (3)$$

Here, (X_i, Y_i, Z_i) are the Pauli spin matrices associated with the spin at site $i \in \{0, \dots, N-1\}$, obeying periodic boundary conditions: $i + N \equiv i$. As shown in Sec. I of SM [41], it is possible to obtain an exact representation of $|\psi_o\rangle$ as an MPS. In this form the *Ansatz* reads

$$|\psi_o\rangle = 2^{-N/2} \sum_{\vec{\sigma}} \text{Tr}(B^{\sigma_0} A^{\sigma_1} \dots A^{\sigma_{N-1}}) |y, \vec{\sigma}\rangle, \quad (4)$$

where $|y, \sigma_i\rangle$ are the eigenstates of the Pauli- Y_i matrix with eigenvalues $\sigma_i = \pm 1$, and

$$A^+ = \begin{bmatrix} \cos \frac{\theta}{2} & \cos \frac{\theta}{2} \\ i \sin \frac{\theta}{2} & -i \sin \frac{\theta}{2} \end{bmatrix} \begin{bmatrix} e^{-\frac{i(\theta+\phi)}{2}} & 0 \\ 0 & e^{\frac{i(\theta+\phi)}{2}} \end{bmatrix}, \quad (5)$$

$$B^+ = \begin{bmatrix} \cos \frac{\theta}{2} & \cos \frac{\theta}{2} \\ i \sin \frac{\theta}{2} & i \sin \frac{\theta}{2} \end{bmatrix} \begin{bmatrix} e^{-\frac{i(\theta+\phi)}{2}} & 0 \\ 0 & e^{-\frac{i(\theta-\phi)}{2}} \end{bmatrix}, \quad (6)$$

$$A^- = (A^+)^*, \quad B^- = (B^+)^*. \quad (7)$$

Both the TFIM and LMG models exhibit translational invariance. However, this property is not shared by $|\psi_o\rangle$

due to the lone $B^\pm$ matrix appearing in Eq. (4). We will restore this symmetry through a minimal modification of the original *Ansatz* $|\psi_o\rangle$ by replacing the B^{σ_0} with A^{σ_0} in Eq. (4). This modification marginally improves results for small systems, while still yielding the same ground-state energy as the original *Ansatz* for both models in the thermodynamic limit. This results in the translationally invariant *Ansatz*,

$$|\psi_t\rangle = \frac{1}{M} \sum_{\vec{\sigma}} \text{Tr}(A^{\sigma_0} \dots A^{\sigma_{N-1}}) |y, \vec{\sigma}\rangle, \quad (8)$$

where M is a normalization factor. The structure of both $|\psi_o\rangle$ and $|\psi_t\rangle$ in Eqs. (4) and (8) allows us to obtain closed-form expressions for expectation values using a transfer matrix approach. See Sec. II.A of SM [41]. For $|\psi_t\rangle$ we find

$$\langle X_i \rangle = \frac{1}{M^2} \left[\frac{c^2(s-t)}{st-1} + \frac{d^2(s-t)}{t^2(st-1)} (st)^N \right], \quad (9)$$

$$\langle Z_i \rangle = \frac{1}{M^2} \left[\frac{cd((st)^N - 1)}{st-1} \right], \quad (10)$$

$$\langle Z_i Z_{i+r} \rangle = \frac{1}{M^2} [f(r) + (st)^N f(-r)], \quad (11)$$

where

$$f(r) = \frac{c^2 d^2 + (s-t)^2 (st)^r}{(st-1)^2}, \quad (12)$$

$$M^2 = 1 + (st)^N, \quad |st| \neq 1 \quad (13)$$

and

$$c = \cos(\theta), \quad s = \sin(\theta), \quad (14)$$

$$d = \cos(\theta + \phi), \quad t = \sin(\theta + \phi). \quad (15)$$

Results for the LMG model—The LMG Hamiltonian for N spin- $\frac{1}{2}$ particles reads

$$H = -\frac{J}{4N} \sum_{i<j} Z_i Z_j - \frac{h}{2} \sum_{i=0}^{N-1} X_i, \quad (16)$$

where J and h set the strengths of the spin-spin interaction and external field respectively. The all-to-all nature of the spin interaction results in the system's mean-field description becoming exact for certain predictions in the thermodynamic limit. We first show that our *Ansatz* shares this property. Figure 2 shows that our approach contains the mean-field result as a special case. Specifically, the first layer of R^ϕ rotations generates a product state amounting to a mean-field *Ansatz*. When $\theta \neq 0$, the second layer of $C^\theta R^\theta$




FIG. 3. The rms magnetization of Eq. (21) vs h/J for the LMG model. Exact results compared to the symmetrized *Ansatz*.

rotations then introduces correlations beyond the mean-field level. To proceed, we calculate the energy per spin in the thermodynamic limit with respect to $|\psi_t\rangle$ using Eqs. (9) and (11). See Sec. III.A of SM for details. This yields

$$\lim_{N \rightarrow \infty} \frac{\langle H \rangle}{N} = -\frac{c^2 d^2}{8(st-1)^2} - \frac{h(s-t)c^2}{2(st-1)}, \quad (17)$$

which is a function of θ and ϕ via Eqs. (14) and (15). Minimizing this expression with respect to these angles produces

$$\sin(\phi) = \begin{cases} 2h & |2h| \leq 1 \\ \text{sgn}(h) & \text{otherwise} \end{cases} \quad \text{and} \quad \theta = 0, \quad (18)$$

as shown in Sec. III.B of SM. The vanishing of θ implies that our *Ansatz* reduces to a product state generated by the first layer of R^ϕ rotations. Inserting this into Eq. (10) for $\langle Z_i \rangle$ yields the spontaneous magnetization,

$$\lim_{N \rightarrow \infty} \frac{1}{2N} \sum_i \langle Z_i \rangle = \begin{cases} \pm \frac{1}{2} \sqrt{1-4h^2} & |2h| \leq 1 \\ 0 & \text{otherwise} \end{cases}, \quad (19)$$

from which we identify the critical value of h as $h_c = 1/2$. This field strength marks the transition between the paramagnetic ($|h| > h_c$) and ferromagnetic ($|h| < h_c$) phases. For the energy per spin we find

$$\lim_{N \rightarrow \infty} \frac{\langle H \rangle}{N} = \begin{cases} -\frac{1}{2}(h^2 + \frac{1}{4}) & |h| \leq h_c \\ -\frac{|h|}{2} & |h| > h_c \end{cases}. \quad (20)$$

Both Eqs. (19) and (20) are exact results for the thermodynamic limit.




FIG. 4. Relative ground-state energy error vs N for the LMG model at different field strengths h . Compares symmetrized *Ansatz* (solid) and MFT (dashed).

For finite systems, the optimal value of θ is nonzero, and the layer of $C^\theta R^\theta$ rotations in Eq. (2) will introduce correlations between the spins. This brings about a major improvement in accuracy compared to the product state mean-field *Ansatz*. We further this improvement by restoring in our *Ansatz* the symmetries present in the LMG Hamiltonian (16). Specifically, H exhibits permutation symmetry under the exchange of any two spins, and also parity symmetry under a π rotation about the x axis, which sends (X_i, Y_i, Z_i) to $(X_i, -Y_i, -Z_i)$. We enforce these symmetries on the *Ansatz* $|\psi_t\rangle$ by projecting it onto the relevant symmetry subspaces. As shown in Sec. III.C of SM [41], this yields a state $|\psi_s\rangle$ within the $(2S+1)$ -dimensional subspace corresponding to the maximum magnitude $S = N/2$ of the total spin. The analytic expression for $|\psi_s\rangle$, parametrized by θ and ϕ , now serves as a refined version of the original *Ansatz*. We use this symmetrized *Ansatz* to estimate the ground-state energy as well as the rms magnetization,

$$M_{\text{rms}} = \frac{1}{2N} \sqrt{\left\langle \left(\sum_i Z_i \right)^2 \right\rangle}. \quad (21)$$

Figure 3 shows the result of this calculation of M_{rms} for different numbers of spins N . Remarkably, there is no visible difference between our *Ansatz*-based result and the exact value of the magnetization. This suggests that the symmetrized *Ansatz* $|\psi_s\rangle$ captures finite-size effects very accurately. Figure 4 shows the relative error in the ground-state energy, $\epsilon_{\text{rel}} = |E_{\text{pred}} - E_{\text{exact}}|/|E_{\text{exact}}|$, for different field strengths h , plotted on a logarithmic scale as N increases. For our symmetrized *Ansatz*, this error is at most of order 10^{-6} for $h = h_c = 1/2$ and about $N = 25$, and tends to zero as N increases. The result of using the mean-field product state *Ansatz* (with $\theta = 0$) is also shown. While this too becomes exact in the thermodynamic limit, it fares




FIG. 5. The density of particular structural classes of discovered *Ansätze* at different evolutionary steps. Each subplot shows results for a different complexity penalty pair (l_1, l_2) (increasing to the right). At each step (vertical slice) the lengths of intervals of a particular color indicate the fraction of runs that yielded a fittest individual within a specific class. The classes are determined by the structural complexity $S(\psi)$ of the fittest individual ψ . White represents the original *Ansatz* ψ_o in Eq. (2). Green and Red denote *Ansätze* with structural complexity less than the original, with red being the mean-field solution. Blue indicates *Ansätze* with structural complexity equal to the original. Pink and yellow are more complex: yellow contains the original *Ansatz* as a building block, while pink does not.

much worse than the symmetrized *Ansatz* for finite system sizes.

Robustness of the search—Figure 5 shows how varying the complexity penalties l_1 and l_2 in the fitness function (1) affects the search outcome. Each subplot contains data from $79 \pm (3)$ runs where the selection pressure and exploration parameters are fixed at $\rho = 0.01$ and $\epsilon = 0.33$, and where the fitness is evaluated for the set of system sizes $\mathcal{N} = \{5, 6, 7\}$. The fittest individual at each evolutionary step is classified into one of six classes indicated by color. The classes are based on the structural complexity of that individual $S(\psi)$ relative to that of the original *Ansatz* $S(\psi_o)$: green and red denote lower complexity, blue and white equal complexity, and pink and yellow higher complexity. Between the mean-field (red) and original *Ansatz* (white) there are competing local minima (blue) with exactly the same structural complexity as the original, namely $S(\psi_o^{(n)}) = 3$ for all $n \in \mathcal{N}$ but with lower fitness. We see that for the largest penalties l_1 and l_2 , the fittest individual after 200 steps was the original *Ansatz* in 83% and 55% of runs for the Ising model and LMG model, respectively. As we decrease the complexity penalties more complex structures begin to emerge, shown in yellow and pink. These structures have $S(\psi) > 3$, and are distinguished by whether they contain the original *Ansatz* as a building block (yellow) or not (pink). The prevalence of the former yellow class for lower complexity penalties shows that the original *Ansatz* is a robust structure that functions as a building block for more complex *Ansätze*. Figure 5 also shows that the effect of the complexity penalties is model dependent. With decreasing complexity penalties the LMG model is dominated by complex structures more readily than the Ising model. This results from the smaller energy difference found between different network structures for the LMG model. Still, the original *Ansatz* (white) is a prominent beyond-mean-field structure in all cases.

Conclusion—We have introduced a general method for constructing ground-state *Ansätze* that are both analytically tractable and quantitatively accurate across a wide range of

system parameters. Our approach can be applied to any physical system that is amenable to a variational treatment in terms of tensor network states. For example, finding time evolution approximations would entail a train-by-example strategy. The subsequent work [49] uses our framework with this strategy and shows how to automatically generate the N -qubit realizations of popular quantum algorithms (Deutsch-Jozsa, QFT, Grover) by considering only examples up to five qubits. The core of our approach lies in the interplay between the domain-specific language and the fitness criteria. The former enables fitness evaluation on small system sizes, which is computationally efficient and allows capturing system-size scaling. The latter favors *Ansätze* with low variational and structural complexity while preserving accuracy. This results in expressive *Ansätze* that tend to break the underlying model's symmetries, but due to their simple structure, these symmetries can be restored analytically. This provides a systematic way to improve the *Ansatz*.

Remarkably, our method yields the same *Ansatz* for both the LMG and TFIM models, showcasing its ability to identify network structures with some degree of universality. Only small system sizes ($\mathcal{N} = \{3, 4, 5\}$) are required and the search takes about 6 CPU hours to find the *Ansatz*. The resulting *Ansatz* is interpretable—our method autonomously constructs a mean-field treatment and extends it to incorporate correlations. For the LMG model we obtain highly accurate results for finite systems, far surpassing that of a mean-field treatment, and which become exact in the thermodynamic limit. For the TFIM we obtain accurate results across all system sizes and greatly improve upon the mean-field treatment. The TFIM case and the computational cost of the method are discussed in the End Matter.

Acknowledgments—The authors thank Amy Rouillard for valuable discussions. Support from the Oppenheimer Memorial Trust, the Department of Science, Technology and Innovation, and the National Research Foundation of South Africa is kindly acknowledged. Support for the Python

package development by the Unitary Foundation is also kindly acknowledged. The authors acknowledge the Centre for High Performance Computing (CHPC), South Africa, for providing computational resources to this research project.

Data availability—The data that support the findings of this Letter are openly available [50].

- [1] P. Weiss, L'hypothèse du champ moléculaire et la propriété ferromagnétique, *J. Phys. Theor. Appl.* **6**, 661 (1907).
- [2] H. A. Bethe and W. L. Bragg, Statistical theory of superlattices, *Proc. R. Soc. A* **150**, 552 (1997).
- [3] L. Onsager, Electric moments of molecules in liquids, *J. Am. Chem. Soc.* **58**, 1486 (1936).
- [4] R. Kikuchi, A theory of cooperative phenomena, *Phys. Rev.* **81**, 988 (1951).
- [5] T. Oguchi, A theory of antiferromagnetism, II, *Prog. Theor. Phys.* **13**, 148 (1955).
- [6] D. Yamamoto, Correlated cluster mean-field theory for spin systems, *Phys. Rev. B* **79**, 144427 (2009).
- [7] G. M. Wysin and J. Kaplan, Correlated molecular-field theory for Ising models, *Phys. Rev. E* **61**, 6399 (2000).
- [8] G. M. Wysin, Onsager reaction-field theory for magnetic models on diamond and HCP lattices, *Phys. Rev. B* **62**, 3251 (2000).
- [9] A. Du, Y. Q. Yü, and H. J. Liu, Expanded Bethe–Peierls approximation for the Blume–Capel model, *Physica A (Amsterdam)* **320**, 387 (2003).
- [10] K. K. Zhuravlev, Molecular-field theory method for evaluating critical points of the Ising model, *Phys. Rev. E* **72**, 056104 (2005).
- [11] I. Etxebarria, L. Elcoro, and J. M. Perez-Mato, Generalized boundary conditions for periodic lattice systems: Application to the two-dimensional Ising model on a square lattice, *Phys. Rev. E* **70**, 066133 (2004).
- [12] D. C. Mattis, Molecular-field theory with correlations, *Phys. Rev. B* **19**, 4737 (1979).
- [13] K. G. Wilson, Renormalization group and critical phenomena. I. Renormalization group and the kadanoff scaling picture, *Phys. Rev. B* **4**, 3174 (1971).
- [14] M. E. Fisher, Renormalization group theory: Its basis and formulation in statistical physics, *Rev. Mod. Phys.* **70**, 653 (1998).
- [15] L. P. Kadanoff, W. Götzke, D. Hamblen, R. Hecht, E. A. S. Lewis, V. V. Palciauskas, M. Rayl, J. Swift, D. Aspnes, and J. Kane, Static phenomena near critical points: Theory and experiment, *Rev. Mod. Phys.* **39**, 395 (1967).
- [16] M. Fannes, B. Nachtergaele, and R. F. Werner, Finitely correlated states on quantum spin chains, *Commun. Math. Phys.* **144**, 443 (1992).
- [17] D. Perez-Garcia, F. Verstraete, M. Wolf, and J. Cirac, Matrix product state representations, *Quantum Inf. Comput.* **7**, 401 (2007).
- [18] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys. (Amsterdam)* **326**, 96 (2011), January 2011 Special Issue.
- [19] M. B. Hastings, An area law for one-dimensional quantum systems, *J. Stat. Mech.* (2007) P08024.
- [20] F. Verstraete and J. I. Cirac, Matrix product states represent ground states faithfully, *Phys. Rev. B* **73**, 094423 (2006).
- [21] F. Verstraete, V. Murg, and J. Cirac, Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems, *Adv. Phys.* **57**, 143 (2008).
- [22] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions, *arXiv:cond-mat/0407066*.
- [23] V. Murg, F. Verstraete, and J. I. Cirac, Variational study of hard-core bosons in a two-dimensional optical lattice using projected entangled pair states, *Phys. Rev. A* **75**, 033605 (2007).
- [24] V. Murg, F. Verstraete, and J. I. Cirac, Exploring frustrated spin systems using projected entangled pair states, *Phys. Rev. B* **79**, 195119 (2009).
- [25] G. Vidal, Class of quantum many-body states that can be efficiently simulated, *Phys. Rev. Lett.* **101**, 110501 (2008).
- [26] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, *Phys. Rev. A* **74**, 022320 (2006).
- [27] P. Silvi, V. Giovannetti, S. Montangero, M. Rizzi, J. I. Cirac, and R. Fazio, Homogeneous binary trees as ground states of quantum critical Hamiltonians, *Phys. Rev. A* **81**, 062335 (2010).
- [28] M. Gerster, P. Silvi, M. Rizzi, R. Fazio, T. Calarco, and S. Montangero, Unconstrained tree tensor network: An adaptive gauge picture for enhanced performance, *Phys. Rev. B* **90**, 125154 (2014).
- [29] M. Gerster, M. Rizzi, P. Silvi, M. Dalmonte, and S. Montangero, Fractional quantum Hall effect in the interacting Hofstadter model via tensor networks, *Phys. Rev. B* **96**, 195123 (2017).
- [30] M. Lourens, Hierarqcal, <https://github.com/matt-lourens/hierarqcal> (2024).
- [31] M. Lourens, I. Sinayskiy, D. K. Park, C. Blank, and F. Petruccione, Hierarchical quantum circuit representations for neural architecture search, *npj Quantum Inf.* **9**, 1 (2023).
- [32] L. Li, M. Fan, M. Coram, P. Riley, and S. Leichenauer, Quantum optimization with a novel Gibbs objective function and ansatz architecture search, *Phys. Rev. Res.* **2**, 023074 (2020).
- [33] Y. Du, T. Huang, S. You, M.-H. Hsieh, and D. Tao, Quantum circuit architecture search for variational quantum algorithms, *npj Quantum Inf.* **8**, 62 (2022).
- [34] M. Bilkis, M. Cerezo, G. Verdon, P. J. Coles, and L. Cincio, A semi-agnostic ansatz with variable structure for variational quantum algorithms, *Quantum Mach. Intell.* **5**, 43 (2023).
- [35] G. Knizia and G. K.-L. Chan, Density matrix embedding: A simple alternative to dynamical mean-field theory, *Phys. Rev. Lett.* **109**, 186404 (2012).
- [36] L. Yang, W. Hu, and L. Li, Scalable variational Monte Carlo with graph neural ansatz, *arXiv:2011.12453*.
- [37] R. C. Farrell, M. Illa, A. N. Ciavarella, and M. J. Savage, Scalable circuits for preparing ground states on digital quantum computers: The Schwinger model vacuum on 100 qubits, *PRX Quantum* **5**, 020315 (2024).

- [38] R. C. Farrell, M. Illa, A. N. Ciavarella, and M. J. Savage, Quantum simulations of hadron dynamics in the Schwinger model using 112 qubits, *Phys. Rev. D* **109**, 114510 (2024).
- [39] E. Gustafson, K. Sherbert, A. Florio, K. Shirali, Y. Chen, H. Lamm, S. Valgushev, A. Weichselbaum, S. E. Economou, R. D. Pisarski, and N. M. Tubman, Surrogate-constructed scalable-circuits adaptive variational quantum eigensolver in the Schwinger model, *Phys. Rev. Appl.* **23**, 064002 (2025).
- [40] J. S. Kottmann, S. Alperin-Lea, T. Tamayo-Mendoza, A. Cervera-Lierta, C. Lavigne, T.-C. Yen, V. Verteletskyi, P. Schleich, A. Anand, M. Degroote, S. Chaney, M. Kesibi, N. G. Curnow, B. Solo, G. Tsilimigkounakis, C. Zendejas-Morales, A. F. Izmaylov, and A. Aspuru-Guzik, TEQUILA: A platform for rapid development of quantum algorithms, *Quantum Sci. Technol.* **6**, 024009 (2021).
- [41] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/5vc5-9f4d> for detailed analytical derivations and domain-specific language examples.
- [42] H. Liu, K. Simonyan, O. Vinyals, C. Fernando, and K. Kavukcuoglu, Hierarchical representations for efficient architecture search, in *International Conference on Learning Representations* (OpenReview.net, Vancouver, BC, Canada, 2018).
- [43] E. Real, S. Moore, A. Selle, S. Saxena, Y. L. Suematsu, J. Tan, Q. V. Le, and A. Kurakin, Large-scale evolution of image classifiers, in *Proceedings of the 34th International Conference on Machine Learning, Proceedings of Machine Learning Research Vol. 70*, edited by D. Precup and Y. W. Teh (PMLR, Cambridge, MA, 2017), pp. 2902–2911.
- [44] D. E. Goldberg and K. Deb, A comparative analysis of selection schemes used in genetic algorithms, in *Foundations of Genetic Algorithms*, edited by G. J. E. Rawlins (Elsevier, New York, 1991), Vol. 1, pp. 69–93.
- [45] H. J. Lipkin, N. Meshkov, and A. J. Glick, Validity of many-body approximation methods for a solvable model: (I). Exact solutions and perturbation theory, *Nucl. Phys.* **62**, 188 (1965).
- [46] N. Meshkov, A. J. Glick, and H. J. Lipkin, Validity of many-body approximation methods for a solvable model: (II). Linearization procedures, *Nucl. Phys.* **62**, 199 (1965).
- [47] A. J. Glick, H. J. Lipkin, and N. Meshkov, Validity of many-body approximation methods for a solvable model: (III). Diagram summations, *Nucl. Phys.* **62**, 211 (1965).
- [48] S. Sachdev, *Quantum Phase Transitions*, (Cambridge University Press, Cambridge, 1999), Vol. 12.
- [49] A. Rouillard, M. Lourens, and F. Petruccione, Automated quantum algorithm synthesis, [arXiv:2503.08449](https://arxiv.org/abs/2503.08449).
- [50] M. Lourens, Github repository for a code example, https://github.com/matt-lourens/ansatz_search (2025).
- [51] S. Dusuel and J. Vidal, Finite-size scaling exponents of the Lipkin-Meshkov-Glick model, *Phys. Rev. Lett.* **93**, 237204 (2004).
- [52] T. Barthel, S. Dusuel, and J. Vidal, Entanglement entropy beyond the free case, *Phys. Rev. Lett.* **97**, 220402 (2006).
- [53] R. Orús, S. Dusuel, and J. Vidal, Equivalence of critical scaling laws for many-body entanglement in the Lipkin-Meshkov-Glick model, *Phys. Rev. Lett.* **101**, 025701 (2008).
- [54] J. Gray, QUIMB: A Python library for quantum information and many-body calculations, *J. Open Source Software* **3**, 819 (2018).

End Matter

LMG discussion—Both the translationally invariant *Ansatz* $|\psi_t\rangle$ and its symmetrized counterpart $|\psi_s\rangle$ produce exact results for M_{rms} and the ground-state energy within the thermodynamic limit. However, these quantities probe limited features of the two states, and it turns out that $|\psi_t\rangle$ and $|\psi_s\rangle$ have fundamentally different characters, even in this limit. Specifically, we found that $|\psi_t\rangle$ reduces to a mean-field product state as $N \rightarrow \infty$ due to the optimal value of θ vanishing. In contrast, optimizing $|\psi_s\rangle$ yields a nonzero θ , even in the $N \rightarrow \infty$ limit, thereby retaining the entanglement from the $C^\theta R^\theta$ rotations. This aligns with Refs. [51–53], where it is shown that the *exact* ground state in the paramagnetic phase always contains nontrivial entanglement and does not reduce to a product state as $N \rightarrow \infty$. These observations underscore that the symmetrization step can fundamentally alter the correlations present in the *Ansatz*, and that optimizing the variational parameters before versus after this step can yield very different results.

Results for the TFIM—The Hamiltonian for the TFIM with N spin- $\frac{1}{2}$ particles on a periodic chain is

$$H = -\frac{J}{4} \sum_{i=0}^{N-1} Z_i Z_{i+1} - \frac{h}{2} \sum_{i=0}^{N-1} X_i, \quad (\text{B1})$$

with J and h again the interaction and external field strengths. We set $J = 1$ as before. For the LMG model it was seen that the mean-field *Ansatz* with $\theta = 0$ in Eq. (2) was sufficient for correctly predicting the critical value of the external field strength and for calculating the order parameter M_{rms} in the thermodynamic limit. For the TFIM the situation is quite different. Here, even in the thermodynamic limit, the $C^\theta R^\theta$ rotations play a crucial role in introducing correlations between spins, and are essential for shifting the estimate for the critical field strength closer to its true value. Using the *Ansatz* in Eq. (8) together with Eqs. (9)–(11) we find the energy per spin in the thermodynamic limit to be

$$\lim_{N \rightarrow \infty} \frac{\langle H \rangle}{N} = -\frac{s-t}{4(st-1)} ((s-t) + 2hc^2) - \frac{1}{4}, \quad (\text{B2})$$

which is a function of θ and ϕ via Eqs. (14) and (15). By minimizing this expression with respect to θ and ϕ we identify a critical field strength of




FIG. 6. The long-range correlation function $\rho_{N/2}^z$ versus h/J for the TFIM. Results are shown for the symmetrized *Ansatz* (finite N , $N \rightarrow \infty$), together with exact values and the mean-field prediction.

$$h_c = \frac{1 + \sqrt{2}}{4} \approx 0.604, \quad (\text{B3})$$

above which the magnetization $\langle Z_i \rangle$ in Eq. (10) vanishes. See Sec. IV of SM for the details of this calculation. This estimate for h_c is indeed closer to the exact value $h_c^{\text{ex}} = 0.5$ when compared to the mean-field result of $h_c^{\text{mf}} = 1$, which would follow from setting $\theta = 0$ and only varying ϕ . While the TFIM Hamiltonian lacks the permutation symmetry of the LMG model, it retains the parity symmetry. We again restore this symmetry by projecting the *Ansatz* Eq. (8) onto the positive symmetry subspace to produce a modified *Ansatz* $|\psi_p\rangle$. See Sec. II.D of SM [41]. Using $|\psi_p\rangle$ we calculate the long-range correlation function $\rho_{N/2}^z = \frac{1}{4} \langle \psi_p | Z_0 Z_{N/2} | \psi_p \rangle$, which serves as an order parameter for characterizing the model's two phases. Figure 6 shows the result this calculation for various system sizes. While our *Ansatz*-based result matches the exact one closely for small N , it begins to deviate from it as N increases. This is to be expected due to the error in the *Ansatz*'s prediction of the critical field strength. The mean-field result, with its prediction of $h_c^{\text{mf}} = 1$, is also shown.

Other choices of operator basis—The original *Ansatz* from Eq. (2) is represented as a product of unitary operators, meaning it is equivalent to a quantum circuit. This is because the initial operator pool was generated from the Pauli matrices $\sigma \in \{I, X, Y, Z\}$ in the form $e^{i\phi\sigma_i}$, $e^{i\theta\sigma_i\otimes\sigma_j}$, where $\theta, \phi \in \mathbb{R}$. We can change the operator pool to include nonunitary tensors to find more general tensor networks. For example, using an operator basis that contains ladder operators, i.e., $\{I, \sigma^+, \sigma^-, Z\}$, our method finds another efficient *Ansatz* shown in




FIG. 7. *Ansatz* discovered with the ladder-operator basis $\{I, \sigma^+, \sigma^-, Z\}$. The search returns a translationally invariant, two-parameter periodic MPS Eq. (C1), and it attains the same energy expectation values as the original *Ansatz* in Fig. 2.

Fig. 7. With this operator basis the search was able to obtain a translationally invariant *Ansatz*. Numerically, this *Ansatz* yields the same energy expectation values as the original *Ansatz* in Fig. 2 for all system sizes. We can express this *Ansatz* as the following two-parameter periodic MPS:

$$|\psi_l\rangle = \frac{1}{M_l} \sum_{\vec{\sigma}} \text{Tr}(C^{\sigma_0} C^{\sigma_1} \dots C^{\sigma_{N-1}}) |z, \vec{\sigma}\rangle, \quad (\text{C1})$$

where M_l is a normalization factor and the matrices $C^\pm$ are given by

$$C^\pm = a^{\frac{1\pm 1}{2}} \begin{bmatrix} \cosh(b) & \pm \sinh(b) \\ \pm \sinh(b) & \sinh(b) \end{bmatrix}, \quad (\text{C2})$$

with $a, b \in \mathbb{R}$. Expectation values may be obtained in the same fashion as detailed for the original *Ansatz* in Sec. II of SM [41]. The choice of operator basis will depend on the context of the problem. For example, if the goal is to obtain short-depth quantum circuits for state preparation on a specific quantum device, then using native gates for that hardware will be beneficial.

Computational cost—The main bottleneck in the evaluation of individuals during the evolutionary search is memory, as the bond dimension between two sites grows exponentially with the number of tensors that connect them. Increasing the complexity penalty limits the exploration of overly large network structures and, in turn, allows larger system sizes to be considered. Because the DSL automatically scales over system sizes, it can explore structures that, in principle, extend well beyond what we can evaluate explicitly for large systems. Figures 8 and 9 report the computational cost of our method over 75 searches for an Ising-model *Ansatz*, organized into five size groups and averaged over 15 runs each. We used $l_1 = l_2 = 7 \times 10^{-4}$, $\rho = 0.01$, and $\epsilon = 0.33$ for all runs. In 65 of the 75 cases (86%), the fittest individual at step 200 was the original *Ansatz*, close to the 83% ratio observed at step 200 for system sizes $n \in \{5, 6, 7\}$ in Fig. 5 with the same penalty. This suggests that the converged *Ansatz* is largely insensitive to the system sizes considered. We evaluated system




FIG. 8. Runtime of the evolutionary search for the Ising model. We consider five system-size groups, each averaged over 15 runs (75 runs total), using complexity penalties $l_1 = l_2 = 7 \times 10^{-4}$ and selection and exploration parameters $\rho = 0.01$ and $\epsilon = 0.33$.

sizes up to $n = 21$ spins, enabled by the complexity penalty. From Fig. 5, beyond-mean-field structures begin to dominate around 100 steps; in Fig. 8 this corresponds to 6.27 CPU hours on average for the smallest group, $n \in \{3, 4, 5\}$, and about 13.8 CPU hours for the largest, $n \in \{19, 20, 21\}$. Figure 9 shows the per-worker memory footprint. For 100 evolutionary steps and $n \in \{3, 4, 5\}$, we required less than 0.57 GiB per worker on average, with each worker occupying one CPU core. For 100 steps and $n \in \{19, 20, 21\}$, each worker required less than 1.06 GiB on average.

After finding an *Ansatz* with the search algorithm, we optimize its variational parameters for each system size. This optimisation may not always be required—for example, with the Ising model the optimal angles essentially




FIG. 9. Per-CPU memory usage during evolutionary search. Fitness evaluations are run in parallel, so each CPU corresponds to one worker. Memory therefore depends on both the number of workers and the complexity penalty.

become constant once $N > 20$. For the LMG model in the ferromagnetic phase, numerical results suggest that the leading-order finite-size corrections scale like $1/N$. Using only system sizes $N \in \{50, 100, 150\}$, we can extrapolate to arbitrary system sizes without losing accuracy. For the paramagnetic phase, extrapolation to the thermodynamic limit is no longer reliable. The symmetrized *Ansatz* produces essentially exact results in this phase, which might cause the optimal parameter values to depend more sensitively on the system size. The Python package *hierarqcal* [30] was used to implement the DSL (genetic operations, system-size scaling, and *ansatz* representation). We also used QUIMB [54] to determine tensor-network contraction paths and to evaluate energies.