NQOF: NEURAL OPERATOR FLOW FOR QUANTUM OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Ground energy estimation is central to quantum chemistry and many-body physics, yet its computational cost grows rapidly with system size. State-centered methods depend strongly on the wavefunction ansatz, while operator-flow methods identify energy-lowering directions within an exponentially large operator space. To address this challenge, we reformulate ground energy optimization as learning a Brockett-Wegner flow on the Hamiltonian's operator manifold. Building on this formulation, we introduce Neural Quantum Operator Flow (NQOF), a learn-and-propagate framework that uses an autoregressive neural network to discover a compact set of Pauli operators without enumerating the full operator basis. A quadratic projection then determines their coefficients to approximate the energy-lowering direction of the Brockett-Wegner flow. The resulting unitary updates preserve the Hamiltonian spectrum while progressively driving the variational energy toward the ground energy. We conduct systematic experiments on molecular and quantum many-body systems with up to 30 qubits. The results show that NQOF achieves competitive accuracy while requiring fewer learnable parameters, and remains robust across different systems. Our approach establishes operator-flow learning as a new paradigm for quantum optimization and opens a promising avenue toward solving complex quantum problems.

1 INTRODUCTION

Quantum optimization like ground energy estimation underpins fundamental problems ranging from predicting molecular ground states and correlated phases of matter to solving combinatorial problems encoded in quantum Hamiltonians (Cao et al., 2019; Carleo & Troyer, 2017; Lucas, 2014). However, the exponential growth of the Hilbert space, compounded by strong correlations and entanglement, makes the efficient exploration and optimization of quantum systems a central computational challenge (Eisert et al., 2010; Carleo & Troyer, 2017). The existing approaches for ground energy estimation can be broadly categorized into state- and operator-centered. State-centered methods optimize or approximate the wavefunction, ranging from exact diagonalization and quantum Monte Carlo to tensor networks, neural quantum states, and variational quantum algorithms, but typically face a trade-off among accuracy, trainability, and computational cost(Wietek & Läuchli, 2018; Carleo & Troyer, 2017; Cirac et al., 2021). Operator-centered methods instead transform the Hamiltonian to lower the variational energy (Kowalski & Bauman, 2025; Commeau et al., 2020; Robbiati et al., 2026a). However, the previous approaches may suffer from rapid operator proliferation and errors introduced by heuristic or predefined generator selection, limiting their accuracy and scalability. Besides, the quantum algorithms further require deep circuit and noise control. These limitations motivate a central question: *Can we develop a scalable and expressive framework for ground energy optimization?*

Therefore, we are motivated by the observation that ground energy optimization can be equivalently reformulated as an isospectral evolution of the Hamiltonian while keeping a reference state fixed. This operator-centered formulation provides a geometrically structured descent dynamics on the unitary orbit. It shifts the computational task from directly representing and optimizing a many-body state to approximating a local operator-flow direction towards the low-energy eigenspace. Building on this, we develop Neural Quantum Operator Flow (NQOF), a framework that exploits the structure by learning a compact Pauli generator via an autoregressive neural network (Barrett et al., 2022; Vaswani et al., 2017) that captures the relevant tangent direction at each flow step, rather than

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

explicitly constructing the direction in the full operator space. The Hamiltonian is then propagated along the flow by the learned unitary update, while reducing the variational energy. Repeating this learn-and-propagate procedure drives the Hamiltonian toward the lowest-energy state, thereby converting ground energy search into a sequence of tractable local operator optimizations. NQOF attempts to alleviate the problems of greedy operator selection or rapid operator proliferation that may exist in previous studies, while ensuring high accuracy.

We then conduct systematic numerical experiments on ground energy estimation across diverse molecular and many-body quantum systems to evaluate the performance of NQOF. The results show that NQOF consistently and reliably approaches the ground energies of systems containing up to 30 qubits while requiring substantially $11.8\times$ fewer parameters than Transformer-based neural quantum state. Moreover, by learning a compact subset of relevant Pauli operators, NQOF avoids greedily searching the full Pauli space and mitigates operator proliferation without compromising accuracy. Together, these results demonstrate the scalability and effectiveness of NQOF as a learning-based, operator-centered alternative framework for ground energy optimization.

Contributions. In this work, our main contributions are threefold: **1**) Building on the operator-flow perspective, we develop NQOF, a learning-based framework for ground optimization that approximates an isospectral Hamiltonian flow rather than directly optimizing a many-body wavefunction. **2**) By exploiting the geometric structure of the operator manifold and restricting optimization to a compact tangent subspace, NQOF constructs effective descent directions with substantially fewer trainable parameters than state-based approaches such as NQS. **3**) We conduct extensive numerical experiments over molecular and quantum many-body systems to demonstrate the accuracy, parameter efficiency, and scalability of NQOF, showing that competitive ground energies can be achieved without greedily exploring the full operator space.

2 PRELIMINARY

In this section, we briefly introduce the basic concept of quantum optimization of ground energy estimation and flow over operator manifold.

Optimization in studying quantum systems. Many computational problems in quantum science can be formulated as optimization problems of the general form $\min_{\boldsymbol{x}\in\mathcal{X}} \mathcal{L}(\boldsymbol{x})$, where $\boldsymbol{x}$ denotes the variables to be optimized. Depending on the task, $\boldsymbol{x}$ may represent the parameters of a quantum circuit, a time-dependent control signal, or a discrete configuration. A typical problem that has been widely studied is ground energy estimation. Consider an n-qubit quantum system governed by a Hamiltonian represented by a Hermitian operator $\boldsymbol{H} \in \mathbb{C}^{N\times N}, N = 2^n$, which encodes the energy spectrum and interactions of the system. Generally, the Hamiltonian is defined by the linear combination of Pauli strings, i.e., $\boldsymbol{H} = \sum_i \alpha_i \boldsymbol{P}_i$, where $\boldsymbol{P}_i \in \mathcal{P}_n = \{\boldsymbol{I}, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}\}^{\otimes n}$, $\boldsymbol{X} = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right), \boldsymbol{Y} = \left(\begin{smallmatrix} 0 & -i \\ i & 0 \end{smallmatrix}\right), \boldsymbol{Z} = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$. By the variational principle, the ground energy of the system, mathematically the lowest eigenvalue of $\boldsymbol{H}$, is given by minimizing the variational energy E,

$$\min_{\boldsymbol{u}} E(\boldsymbol{u}) = \langle \boldsymbol{u}|\boldsymbol{H}|\boldsymbol{u}\rangle, \tag{1}$$

where $|\boldsymbol{u}\rangle$ denotes a tunable quantum state expressed by a unit complex vector in Hilbert space $\mathcal{H}$, while $\langle \boldsymbol{u}|$ refers to its Hermitian conjugate.

Operator flow on a manifold. An operator flow on a differentiable manifold $\mathcal{M}$ is a smooth one-parameter evolution $t \mapsto \boldsymbol{O}_t$ such that $\boldsymbol{O}_t \in \mathcal{M}$ for all t, which satisfies

$$\frac{d\boldsymbol{O}_t}{dt} = \Phi_t(\boldsymbol{O}_t), \tag{2}$$

where $\Phi_t : \mathcal{M} \to \mathcal{M}$ is a time-dependent super-operator. The flow is said to be generated by Φ_t with a few important special cases, such as Hamiltonian flow: $\Phi_t(\rho) = -\frac{i}{\hbar}[\boldsymbol{H}, \rho]$ and dissipative flow: $\Phi_t(\rho) = \mathcal{L}(\rho)$, where ρ and $\mathcal{L}$ separately refer to the density matrix and the Lindblad operator, describing the dynamics of a quantum system. Operator flows extend the viewpoint from trajectories of states to trajectories in operator space. Within this framework, geometric constraints, such as spectral invariance under unitary flows, can also be incorporated.

Autoregressive model. An autoregressive neural network (ANN) $f_{\boldsymbol{W}}$ is a type of generative model that can be used to represent the joint probability distribution over multiple random variables

$\mathbf{x} = (x_1, \cdots, x_n)$ by decomposing it into a product of conditional distributions, $P_{\boldsymbol{W}}(\mathbf{x}) = f_{\boldsymbol{W}}(\mathbf{x}) = \prod_{i=1}^{n} P_{\boldsymbol{W}}(x_i \mid x_1, \ldots, x_{i-1})$, where each variable x_i is predicted from the variables preceding it under a chosen ordering (Larochelle & Murray, 2011; Germain et al., 2015). These conditional distributions can be parameterized using recurrent neural networks, masked neural networks, convolutional networks, or Transformers (Vaswani et al., 2017). This factorization provides a normalized probability distribution, enables explicit likelihood evaluation, and permits direct sampling. Consequently, ANNs have been widely applied to language modeling, time-series forecasting, image/audio generation, and the representation and sampling of quantum states(van den Oord et al., 2016; Sharir et al., 2020).

2.1 Related Works

Most existing approaches to ground energy estimation can be broadly categorized into state-centered and operator-centered methods. State-centered approaches optimize or approximate the many-body wavefunction directly, ranging from exact diagonalization and FCI to more scalable approximations such as mean-field, coupled-cluster with single-double excitation (CCSD), tensor-network, quantum Monte Carlo, and neural quantum state (NQS) methods (Perdew et al., 2021; Tubman et al., 2020; Neufeld & Berkelbach, 2023; Cirac et al., 2021; Needs et al., 2020; Carleo & Troyer, 2017). While these methods can substantially reduce the cost of exact diagonalization, their performance may depend on the expressivity of the chosen wavefunction representation and the difficulty of the associated optimization problem (Sharir et al., 2022; Wu et al., 2024). In contrast, operator-centered approaches transform the Hamiltonian itself to expose its low-energy structure. Quantum implementations include double-bracket quantum algorithms (DBQA), feedback-based control, and spectral-filtering methods (Gluza et al., 2026; Robbiati et al., 2026b; Larsen et al., 2024; Dong et al., 2022). These approaches, however, can suffer from operator growth, truncation errors, or demanding circuit-depth and noise requirements. Refer to Appendix. B for more discussion.

3 Neural Operator Flow for Quantum Optimization

In this section, we first recast quantum optimization from the conventional state-centered formulation into a flow-learning problem on an operator manifold, and then introduce the Neural Quantum Operator Flow (NQOF) framework to construct a tractable approximation of the target operator flow and leverage the resulting dynamics to determine the ground state energy.

3.1 Reformulating Quantum Optimization in Operator Space

To determine the ground energy of a quantum system with a given Hamiltonian, we depart from the conventional state-centered formulation in Eq. (1), where the optimization is performed over a parameterized trial state $|\boldsymbol{u}\rangle$. Instead, we adopt an operator perspective and reformulate the task as learning the Brockett-Wegner flow (Brockett, 1991) on the manifold $\mathcal{M}$.

Theorem 1 (Brockett-Wegner flow)**.** *Let $\boldsymbol{H}, \boldsymbol{H}_t \in \mathbb{C}^{d\times d}$ be Hermitian, and $|\boldsymbol{v}\rangle$ be a fixed complex vector. Define the Brockett-Wegner flow*

$$\frac{d\boldsymbol{H}_t}{dt} = -[[\rho, \boldsymbol{H}_t], \boldsymbol{H}_t], \quad \boldsymbol{H}_0 = \boldsymbol{H}, \tag{3}$$

where $\rho = |\boldsymbol{v}\rangle\langle\boldsymbol{v}|$, $|\boldsymbol{v}\rangle$ is a unit complex vector and $[\cdot,\cdot]$ denotes the commutation between operators, i.e. $[\boldsymbol{A}, \boldsymbol{B}] = \boldsymbol{AB} - \boldsymbol{BA}$. Then the following hold.

(i)(Isospectrality.) *The solution remains on the unitary orbit $\boldsymbol{H}_t \in \mathcal{M} = \{\boldsymbol{U}_t^\dagger \boldsymbol{H} \boldsymbol{U}_t, \boldsymbol{U}_t \in \mathcal{U}(d)\}$. Consequently, the eigenvalue spectrum of $\boldsymbol{H}_t$ is invariant with $\boldsymbol{H}$ along the flow, i.e. $\mathrm{spec}(\boldsymbol{H}_t) = \mathrm{spec}(\boldsymbol{H})$ for all t.*

(ii)(Monotonic decrease.) *Define $E(\boldsymbol{H}_t) = \langle\boldsymbol{v}|\boldsymbol{H}_t|\boldsymbol{v}\rangle$, which is non-increasing along the flow, i.e. $\frac{d}{dt}E(\boldsymbol{H}_t) = -\left\|[\boldsymbol{H}_t, \rho]\right\|_F^2 \leq 0$.*

Assumption 1. *Let $|\psi_0\rangle$ be the eigenvector with lowest eigenvalue of the $\boldsymbol{H}$, we assume that the reference state $|\boldsymbol{v}\rangle$ has nonzero overlap with $|\psi_0\rangle$, i.e. $\|\langle\boldsymbol{v}|\psi_0\rangle\|^2 > 0$.*

Assumption 2. *Let E_0 and E_1 be the ground and excited energy of $\boldsymbol{H}$ respectively (corresponding to the two smallest eigenvalues). We assume the spectral gap $\Delta = \|E_1 - E_0\| > 0$.*

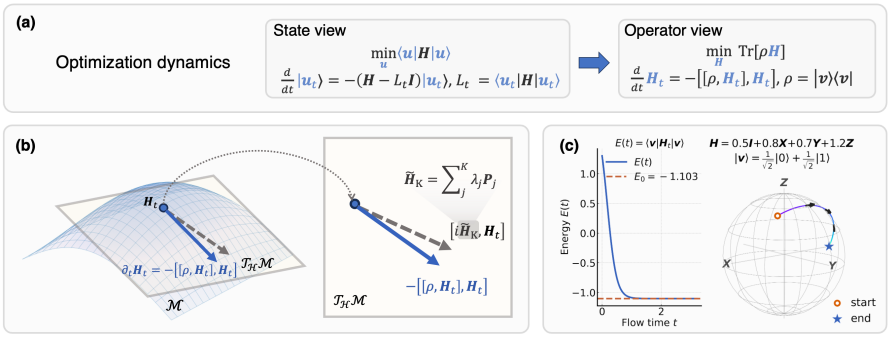


Figure 1: The illustration of reformulating ground energy optimization from a state-centered to an operator-centered perspective. (a) For a given Hamiltonian $\boldsymbol{H}$, conventional state-centered optimization evolves the state $|\boldsymbol{u}_t\rangle$ while keeping $\boldsymbol{H}$ fixed. Equivalently, in the operator view, a reference state $|\boldsymbol{v}\rangle$ is fixed while the Hamiltonian evolves according to the Brockett–Wegner flow $\partial_t \boldsymbol{H}_t = -[[\rho, \boldsymbol{H}_t], \boldsymbol{H}_t]$, where $\rho = |\boldsymbol{v}\rangle\langle\boldsymbol{v}|$. (b) The ideal Brockett-Wegner direction lies in the tangent space $\mathcal{T}_{\boldsymbol{H}_t}\mathcal{M}$. NQOF approximates this direction with the induced tangent vector $[i\tilde{\boldsymbol{H}}_K, \boldsymbol{H}_t]$, where $\tilde{\boldsymbol{H}}_K$ is a compact Pauli generator. (c) A one-qubit example of operator-centered ground energy optimization. The left panel shows that the variational energy decreases monotonically and converges exponentially to the ground energy, while the right panel depicts the corresponding isospectral trajectory on the operator sphere.

The above assumptions are mild for finding ground energy. The nonzero ground state overlap is a standard assumption in ground state preparation, eigenstate filtering, and related double-bracket methods (Lin & Tong, 2020; Gluza et al., 2026), otherwise the task would become QMA-hard. And the gap assumption is standard in adiabatic evolution (Jansen et al., 2007).

Proposition 1 (Exponential convergence of the Brockett-Wegner flow)**.** *Under Assumption 1 and 2, the energy error satisfies* $\epsilon_t = \mathcal{O}(e^{-2t\Delta})$*. Hence,* $E(H_t)$ *converges exponentially to* E_0*. If the initial state additionally has nonzero projection onto the first excited eigenspace, the asymptotic decay exponent is* 2Δ*.*

In a state-centered view, given a Hamiltonian $\boldsymbol{H}$, the ground state energy can be estimated via minimizing Eq. (1). As an arbitrary state $|\boldsymbol{u}\rangle$ can be written as an unitary $\boldsymbol{U}$ applied on a fixed state $|\boldsymbol{v}\rangle$, i.e. $|\boldsymbol{u}\rangle = \boldsymbol{U}|\boldsymbol{v}\rangle$. Consequently, the ground energy optimization in Eq. (1) can be equivalently written as

$$\min_{|\boldsymbol{u}|=1} \underbrace{\langle \boldsymbol{u}|\boldsymbol{H}|\boldsymbol{u}\rangle}_{E(\boldsymbol{u})} = \min_{\boldsymbol{U}\in \mathrm{U}(d)} \langle \boldsymbol{v}|\boldsymbol{U}^\dagger \boldsymbol{H}\boldsymbol{U}|\boldsymbol{v}\rangle = \min_{\boldsymbol{H}'\in\mathcal{M}} \underbrace{\mathrm{Tr}(\rho \boldsymbol{H}')}_{E(\boldsymbol{H}')}, \tag{4}$$

where $\rho = |\boldsymbol{v}\rangle\langle\boldsymbol{v}|$ and $\mathcal{M} = \{\boldsymbol{U}^\dagger \boldsymbol{H}\boldsymbol{U} : \boldsymbol{U} \in \mathrm{U}(d)\}$. This reformulation transfers the optimization from the state space to the operator space. By Theorem 1 and Proposition 1, the variational energy $E(\boldsymbol{H}'_t) = \mathrm{Tr}(\rho \boldsymbol{H}'_t)$ decreases monotonically along the Brockett-Wegner flow, while $\boldsymbol{H}'_t$ remains isospectral to the problem Hamiltonian $\boldsymbol{H}$. Consequently, the flow in Eq. (3) defines a continuous descent dynamics on $\mathcal{M}$ that drives the variational energy in Eq. (4) toward the its minimum, yielding $\lim_{t\to\infty} E(\boldsymbol{H}'_t) = E_0$. Therefore, the minimization problem in Eq. (4) can be reformulated as solving the Brockett-Wegner flow dynamics as shown in Fig. 1(a).

Nevertheless, direct integration of the Brockett-Wegner flow requires constructing nested commutators in the full operator space, which become exponentially expensive with system size. Therefore, instead of explicitly evaluating the Brockett-Wegner flow, we approximate its local direction within a compact tangent subspace illustrated in Fig. 1(b). Specifically, we aim to construct $\tilde{\boldsymbol{H}}_K = \sum_{j=1}^{K} \lambda_j \boldsymbol{P}_j$, where $K \ll |\mathcal{P}_n|$, $\boldsymbol{P}_j$ and λ_j denote a selected Pauli string and corresponding coefficient, respectively, such that the induced tangent vector approximates the target flow direction, i.e., minimizing the residual R

$$R(\mathbb{S}_{\boldsymbol{P}}, \boldsymbol{\lambda}) = \|[i\tilde{\boldsymbol{H}}_K, \boldsymbol{H}_t] + [[\rho, \boldsymbol{H}_t], \boldsymbol{H}_t]\|_F^2, \tag{5}$$

where $\mathbb{S}_{\boldsymbol{P}} = \{\boldsymbol{P}_1, \cdots, \boldsymbol{P}_K\}$ and $\boldsymbol{\lambda} = (\lambda_1, \cdots, \lambda_K)$. It measures the squared Frobenius distance between the local flow direction induced by $\tilde{\boldsymbol{H}}_K$ and the ideal direction $\boldsymbol{D}_t = -[[\rho, \boldsymbol{H}_t], \boldsymbol{H}_t]$. Rather than exploring all possible tangent directions on $\mathcal{T}_{\boldsymbol{H}_t}\mathcal{M}$, we search for a compact approximate $[i\tilde{\boldsymbol{H}}_K, \boldsymbol{H}_t]$ within a smaller subspace, making the local flow construction substantially more tractable.

3.2 Quantum Neural Operator Flow

Recalling the ground energy optimization in Eq. (4), and motivated by its equivalence to the Brockett-Wegner flow, we propose NQOF to efficiently approximate the flow dynamics by learning a compact tangent direction at each step t, thereby enabling ground energy estimation.

Specifically, it contains three main stages as depicted in Fig. 2: (i) first, an ANN defines a probability distribution P over Pauli strings and samples K candidate operators according to P; (ii) then, through separately optimizing the coefficients $\boldsymbol{\lambda}$ and the parameters $\boldsymbol{W}$ of the neural network, it forms a compact generator so that the induced tangent direction closely approximates the local target flow direction; (iii) the resulting generator is used to propagate the current Hamiltonian along the learned flow. Repeating these steps yields a discrete trajectory that approximates the target evolution without explicitly constructing its full operator-valued direction.

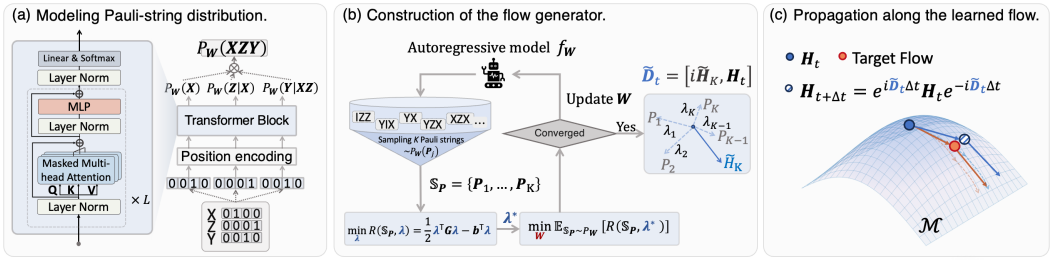


Figure 2: An overview of NQOF. (a) An ANN is used to parameterize a probability distribution over Pauli strings. (b) An alternating optimization procedure to progressively construct the approximate direction $\tilde{\boldsymbol{D}}_t = [i\tilde{\boldsymbol{H}}_K, \boldsymbol{H}_t]$. (c) Propagating the current Hamiltonian $\boldsymbol{H}_t$ via $\boldsymbol{U}_t = \exp(-i\tilde{\boldsymbol{H}}_K \Delta t)$ along the learned flow.

Autoregressive modeling of Pauli-string distribution. Since an arbitrary n-qubit operator can be expanded in the complete Pauli basis, which contains 4^n possible Pauli strings, exhaustively searching this exponentially large space to identify operators relevant to the local target-flow direction is computationally impractical. Instead, we employ an ANN to define a learnable probability distribution P over the space of Pauli strings, enabling relevant operators to be sampled without enumerating the full Pauli basis.

Concretely, we employ an ANN $f_{\boldsymbol{W}}$ to parameterize a probability distribution over n-site Pauli strings $\mathcal{P}_n$. Here, instead of directly feeding the matrix of Pauli strings $\boldsymbol{P}$ into $f_{\boldsymbol{W}}$, we encode each Pauli string $\boldsymbol{P}$ as a matrix $\boldsymbol{M} \in \mathbb{R}^{n\times d}$, whose k-th row is a one-hot vector representing the Pauli operator acting on the k-th site. The probability $P_{\boldsymbol{W}}(\boldsymbol{P})$ assigned to $\boldsymbol{P}$ is expressed by a product of site-wise conditional probabilities, $P_{\boldsymbol{W}}(\boldsymbol{P}) = \prod_{i=1}^{n} P_{\boldsymbol{W}}\left(\boldsymbol{P}^{(i)} \middle| \boldsymbol{P}^{(1)}, \cdots, \boldsymbol{P}^{(i-1)}\right) = f_{\boldsymbol{W}}(\boldsymbol{M})$, where $\boldsymbol{P}^{(i)}$ denotes the Pauli operator acting on the i-th site, and $P_{\boldsymbol{W}}\left(\boldsymbol{P}^{(i)} \mid \boldsymbol{P}^{(1)}, \ldots, \boldsymbol{P}^{(i-1)}\right)$ is the corresponding conditional probability predicted by the network for $\boldsymbol{P}^{(i)}$ given the operators on $i-1$ preceding sites. The role of the autoregressive model is therefore not to represent the entire operator space, but we can progressively concentrate sampling probability on Pauli strings whose induced tangent vectors are useful for approximating the local target flow by tuning the parameters $\boldsymbol{W}$. Besides, such autoregressive structure allows us to efficiently sample Pauli strings $\{\boldsymbol{P}_j\}_{j=1}^{K}$ used for the compact tangent subspace. Here, we utilize a Transformer to build $f_{\boldsymbol{W}}$; refer to Appendix. C for the details about the neural architecture of $f_{\boldsymbol{W}}$ and the sampling method.

Construction of the flow generator. To find an effective approximation to the local target flow direction $\boldsymbol{D}_t$ at flow step t, we aim to minimize the residual R in Eq. (5) with respect to $\boldsymbol{\lambda}$ and $\boldsymbol{W}$, which reduces the construction of the flow generator to a continuous optimization problem on an adaptive subspace by a neural operator $\tilde{\boldsymbol{H}}_K$ expressed by $\tilde{\boldsymbol{H}}_K = \sum_{j=1}^{K} \lambda_j \boldsymbol{P}_j$,

$$\mathcal{L} = \min_{\boldsymbol{\lambda}, \boldsymbol{W}} \mathbb{E}_{\boldsymbol{P}_j \in \mathbb{S}_{\boldsymbol{P}} \sim P_{\boldsymbol{W}}} [R(\mathbb{S}_{\boldsymbol{P}}, \boldsymbol{\lambda})], \tag{6}$$

where $\mathbb{S}_{\boldsymbol{P}} = \{\boldsymbol{P}_j\}_{j=1}^{K}$ denotes an active support set containing the K Pauli strings sampled from $P_{\boldsymbol{W}}$ used to construct the neural operator neural operator $\tilde{\boldsymbol{H}}_K$. The parameters $\boldsymbol{W}$ of the autoregressive model determine a probability distribution $P_{\boldsymbol{W}}$ over the discrete space of Pauli strings and therefore

control which operators are selected, whereas the coefficients $\boldsymbol{\lambda}$ determine how the selected operators are combined to approximate the target tangent direction. Since $\boldsymbol{W}$ and $\boldsymbol{\lambda}$ serve distinct roles in constructing the approximate flow direction, we optimize them separately.

Therefore, at each flow step t, NQOF adopts an *alternating optimization* procedure. Specifically, starting from the current autoregressive distribution $P_{\boldsymbol{W}}$, we first sample a Pauli string set $\mathbb{S}_{\boldsymbol{P}}$, which defines a local tangent subspace through the directions $[i\boldsymbol{P}_j, \boldsymbol{H}_t]$ where $\boldsymbol{P}_j \in \mathbb{S}_{\boldsymbol{P}}$. Within this sampled fixed support, we determine the coefficients $\boldsymbol{\lambda}$ of the compact generator $\tilde{\boldsymbol{H}}_K$ by minimizing the residual R. Since the induced direction depends linearly on $\boldsymbol{\lambda}$, we can reduce it to a convex quadratic programming, i.e. $\min_{\boldsymbol{\lambda}} R(\mathbb{S}_{\boldsymbol{P}}, \boldsymbol{\lambda}) = \boldsymbol{\lambda}^{\mathsf{T}}\boldsymbol{G}\boldsymbol{\lambda} - 2\boldsymbol{b}^{\mathsf{T}}\boldsymbol{\lambda} + \|\boldsymbol{D}_t\|^2$, where $\boldsymbol{G}_{i,j} = \langle [i\boldsymbol{P}_i, \boldsymbol{H}_t], [i\boldsymbol{P}_j, \boldsymbol{H}_t] \rangle_F$ which encodes the geometry of the selected tangent subspace, while $\boldsymbol{b}_i = \langle [i\boldsymbol{P}_i, \boldsymbol{H}_t], \boldsymbol{D}_t] \rangle_F$ measures the relevance of $\boldsymbol{P}_i$ to the desired local flow. Then, the optimal coefficients $\boldsymbol{\lambda}^*$ can be obtained by solving the linear system $(\boldsymbol{G} + \mu \boldsymbol{I})\boldsymbol{\lambda}^* = \boldsymbol{b}$. After obtaining optimal $\boldsymbol{\lambda}^*$, we can update the autoregressive model by minimizing the residual $R(\mathbb{S}_{\boldsymbol{P}}, \boldsymbol{\lambda}^*)$, i.e. $\min_{\boldsymbol{W}} \mathbb{E}_{\mathbb{S}_{\boldsymbol{P}} \sim P_{\boldsymbol{W}}}[R(\mathbb{S}_{\boldsymbol{P}}, \boldsymbol{\lambda}^*)]$, which provides feedback about which sampled Pauli strings are useful, through the gradient $\nabla_{\boldsymbol{W}} \mathcal{L} = \mathbb{E}_{\boldsymbol{P} \sim P_{\boldsymbol{W}}}[R(\mathbb{S}_{\boldsymbol{P}}, \boldsymbol{\lambda}^*) \sum_{j=1}^{N} \nabla_{\boldsymbol{W}} \log P_{\boldsymbol{W}}(\boldsymbol{P}_j)]$. By repeating the above procedure, the distribution $P_{\boldsymbol{W}}$ progressively concentrates on operators to build $\tilde{\boldsymbol{H}}_K$ that are effective for approximating the current local flow. For efficient training, we replace the operator-space mismatch with a state-projected loss and surrogate gradient in our implementation, refer to Appendix C.3 for details.

Hamiltonian propagation along the learned flow. Once the compact generator $\tilde{\boldsymbol{H}}_K$ has been obtained, we use it to advance the current Hamiltonian, $\boldsymbol{H}_t = \sum_j \alpha_j \boldsymbol{Q}_j$, along the learned tangent direction. For a finite step size Δt, the compact Hermitian generator $\tilde{\boldsymbol{H}}_K$ induces the approximate flow direction. We then define the unitary transformation $\boldsymbol{U}_t = \exp(-i\Delta t\, \tilde{\boldsymbol{H}}_K)$ and update the Hamiltonian as

$$\boldsymbol{H}_{t+1} \leftarrow \boldsymbol{U}_t^{\dagger} \boldsymbol{H}_t \boldsymbol{U}_t. \tag{7}$$

It exactly preserves the spectrum and keeps the trajectory on the operator manifold. Here, we utilize the Trotterized Pauli propagation to implement such unitary conjugation on $\boldsymbol{H}_t$. Concretely, we first decompose the unitary $\boldsymbol{U}_t$ into a product of exponentials of individual terms in $\tilde{\boldsymbol{H}}_K$ via first-order Trotterization, i.e. $\boldsymbol{U}_t = (\prod_{k=1}^{K} \boldsymbol{V}_k)^N$, where $\boldsymbol{V}_k = \exp(-\frac{i\Delta t}{N} \lambda_k \boldsymbol{P}_k), N \in \mathbb{N}^+$, and then we perform Pauli propagation(Rudolph et al., 2026) to sequentially apply $\boldsymbol{V}_k$ on $\boldsymbol{H}_t$ to analytically evaluate $\boldsymbol{U}_t^{\dagger} \boldsymbol{H}_t \boldsymbol{U}_t$. Besides, to avoid operator proliferation, we choose a reasonable threshold τ to drop the terms with very small coefficients during the Pauli propagation. Repeating this rule, it yields the propagated Hamiltonian as another weighted Pauli expansion, avoiding the explicit construction and multiplication of exponentially large matrices. We defer more details to Appendix. C.4.

Theorem 2 (Informal)**.** *Let τ be the coefficient truncation threshold in the above procedure, and let $\boldsymbol{H}_M^{(\tau)}$ and $\boldsymbol{H}_M$ denote the Hamiltonians obtained after M iterations with and without truncation, respectively, using the same sequence of Pauli rotations. Under the assumptions in Appendix F.3, there exists a constant threshold $t^\star > 0$ such that, for $T = M\Delta t < t^\star$,*

$$\|\boldsymbol{H}_M^{(\tau)} - \boldsymbol{H}_M\|_{\mathrm{op}} \le C_n (M\tau)^{\nu_n},$$

where $C_n > 0$ and $\nu_n \in (0, 1)$ are specified in the appendix, and $\|\cdot\|_{\mathrm{op}}$ denotes the operator norm.

The above theorem indicates that coefficient truncation introduces a controllable approximation error within the stated short-time regime. In particular, the error bound scales as a sublinear power of $M\tau$, with its system-size dependence specified in the Appendix.F.3. Therefore, sufficiently small Pauli coefficients can be discarded to limit operator growth and reduce propagation cost while keeping the truncation error below a prescribed tolerance by appropriately choosing τ. Refer to Appendix.F.3 for rigorous proof.

4 Experiments

To evaluate the performance of NQOF, we consider the task of ground energy estimation for various molecular and quantum many-body systems. We defer the details about the settings of experiments, such as the architecture of the ANN, learning rate, and hyperparameters to Appendix. D. To further

validate the outperformance and wide applicability of NQOF, we also present additional experiments in Appendix. G, including convergence of NQOF, timing comparison between NQOF and NQS, and application to combinatorial optimization.

4.1 EXPERIMENTAL SETTINGS

Molecular systems. For a molecular system, the electronic Hamiltonian is commonly expressed in second-quantized form using fermionic creation and annihilation operators to describe electron motion and interactions. To represent the Hamiltonian on qubits, it can be expressed as a weighted sum of Pauli strings $\boldsymbol{H} = \sum_j h_j \boldsymbol{P}_j$. In our experiments, we focus on molecules, including LiH, H_2O, NH_3, CH_4, N_2 and Li_2O with system size varying from 12 to 30 qubits.

Quantum many-body systems. A quantum many-body system consists of many interacting particles whose collective behavior gives rise to complex phenomena. Here, we mainly consider 1) **The transverse-field Ising model (TFIM)**. It describes interacting spins subject to an external field transverse to the interaction direction, the Hamiltonian of which is typically written as $\boldsymbol{H}_{\text{TFIM}} = -J\sum_{i,j} \boldsymbol{Z}_i\boldsymbol{Z}_j - h\sum_i \boldsymbol{X}_i$, where $\boldsymbol{Z}_i$ and $\boldsymbol{X}_i$ refer to the Pauli operators $\boldsymbol{Z}$ and $\boldsymbol{X}$ acting on the i-th qubit, respectively; J controls the spin interaction and h sets the transverse-field strength. 2) **Quantum Heisenberg model (QHM).** It is a model to study the interacting spins through exchange couplings represented by the Hamiltonian $\boldsymbol{H}_{\text{QH}} = \frac{1}{2}\sum_j (J_{\boldsymbol{X}}\boldsymbol{X}_j\boldsymbol{X}_{j+1} + J_{\boldsymbol{Y}}\boldsymbol{Y}_j\boldsymbol{Y}_{j+1} + J_{\boldsymbol{Z}}\boldsymbol{Z}_j\boldsymbol{Z}_{j+1})$, where $J_{\boldsymbol{X}}, J_{\boldsymbol{Y}}, J_{\boldsymbol{Z}}$ are the strength of coupling and external magnetic field, respectively.

Baseline methods. To evaluate NQOF against representative and closely related approaches, we select benchmarks spanning exact, approximate, state-centered learning, and operator-centered quantum methods. First, we employ *(FCI)*(Knowles & Handy, 1984) as the exact reference, and select *(CCSD) and (HF)*(Purvis & Bartlett, 1982; Roothaan, 1951) as standard approximate quantum-chemistry method that offers a practical balance between accuracy and computational cost. We further consider *(NQS)*(Carleo & Troyer, 2017) as a learning-based, state-centered baseline, enabling a direct comparison between state and operator-centered neural optimization. Finally, we select the *(DBQA)* (Robbiati et al., 2026a) as the most closely related operator-centered quantum baseline because it implements double-bracket transformations through quantum circuits to find the ground energy. The implementation details of these reference methods are deferred to Appendix. E.

Evaluation metrics. To quantify the performance of NQOF in terms of energy accuracy and its ability to approximate the target Brockett-Wegner flow, we consider the following metrics. 1) **Accuracy.** We define the accuracy as the absolute error between the energy $\hat{E}$ estimated by NQOF and the ground energy E_0, i.e., $\Delta E = |\hat{E} - E_0|$. For molecular systems, we adopt chemical accuracy (**CA**) as the threshold, a conventional benchmark in quantum chemistry indicating that a computed energy differs from the ground energy by no more than 1.6×10^{-3} (Ha). 2) **Flow fidelity.** We measure the agreement between the NQOF-generated flow and the target Brockett-Wegner flow using the state fidelity $F = |\langle \boldsymbol{v}^* | \hat{\boldsymbol{v}} \rangle|^2$, where $|\boldsymbol{v}^*\rangle$ and $|\hat{\boldsymbol{v}}\rangle$ denote the states evolved under the NQOF and Brockett-Wegner flows, respectively.

4.2 RESULTS

We then evaluate the performance of NQOF against the reference methods over various systems with different settings.

NQOF outperforms the benchmark models in accuracy. First, we evaluated the accuracy of NQOF on different molecular systems and many-body systems, respectively. The molecular systems are selected from H_2O (14-qubit) to Li_2O (30-qubit), and the many-body systems include the 20-qubit QHM and 20-qubit TFIM with $J/h = 1$. The absolute ground energy errors of NQOF with those of the baseline methods are shown in Table. 1. In all molecular cases, NQOF remains well below the **CA** threshold, demonstrating consistently competitive performance for molecular system. Compared with DBQA, the accuracy of the NQOF is better since NQOF approximates the target Brockett-Wegner direction through a locally optimized projection at each step which provides more flexibility under a limited resource budget. These results demonstrate that NQOF provides robust and highly accurate ground energy estimates across diverse quantum systems.

Table 1: The accuracy $\Delta E = |\hat{E} - E_0|$ comparision between NQOF and benchmark methods. The **bold number** refers to the best results.

Type	System	E_0 (Ref.)	CCSD	NQS	DBQA	NQOF
Molecule	H_2O	-75.0132	1.18×10^{-4}	7.79×10^{-6}	1.5929×10^{-2}	$\mathbf{6.14 \times 10^{-6}}$
	NH_3	-55.4997	1.89×10^{-4}	6.30×10^{-5}	1.6252×10^{-2}	$\mathbf{2.15 \times 10^{-6}}$
	CH_4	-39.8060	2.34×10^{-4}	1.39×10^{-4}	2.5021×10^{-2}	$\mathbf{1.13 \times 10^{-5}}$
	N_2	-107.6541	3.925×10^{-3}	3.35×10^{-4}	3.3430×10^{-2}	$\mathbf{1.66 \times 10^{-4}}$
Many-body system	QHM	-34.7299	–	$\mathbf{2.30 \times 10^{-3}}$	7.22×10^{-2}	1.83×10^{-2}
	TFIM	-25.1078	–	1.05×10^{-2}	4.920×10^{-1}	$\mathbf{6.50 \times 10^{-3}}$

Scalability of NQOF. To evaluate the scalability of NQOF, we consider molecular systems ranging from 12-qubit LiH to 30-qubit Li_2O, together with $10 \sim 30$ qubits TFIMs at $J/h \in \{0.5, 1.0, 2.0\}$. As shown in Fig. 3(a), the mean molecular energy errors remain below the **CA** threshold across all tested sizes, while the TFIM errors remain on the order of 10^{-3} or lower. The larger error at $J/h = 1$ reflects the increased difficulty near the critical regime. In Fig. 3(b), it shows that, for Li_2O, NQOF rapidly lowers the energy into the **CA** window, with decreasing variation across independent runs as the optimization converges. Furthermore, Fig. 3(c) shows that NQOF accurately reconstructs the Li_2O potential-energy curve, closely following the FCI reference over the entire range of internuclear separations, even around the equilibrium region where the energy is hard to estimate. These results demonstrate that NQOF maintains reliable convergence and high accuracy as system size increases.

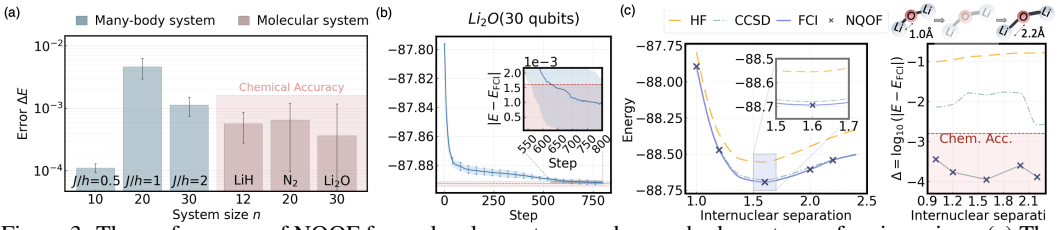


Figure 3: The performance of NQOF for molecular systems and many-body systems of various sizes. (a) The accuracy ΔE obtained by NQOF for various systems. (b) Optimization curves for 30-qubit Li_2O. The solid curves show the average energy obtained by NQOF, while the shaded regions represent multiple runs. The horizontal dashed lines denote the FCI, and the shaded red regions indicate the **CA** window. (c) Potential-energy curve of the 30-qubit Li_2O molecule as a function of the internuclear separation. The left panel compares the energies obtained by NQOF with HF and CCSD, with the inset enlarging the region around the equilibrium geometry. The right panel reports the logarithmic energy error, where the red dashed line marks the **CA** threshold.

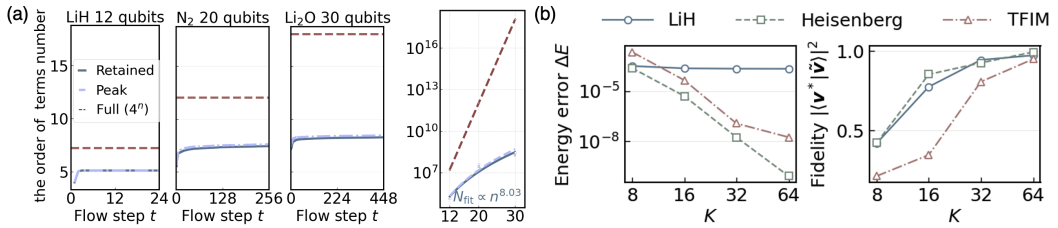


Figure 4: Operator complexity of $\boldsymbol{H}_t$ and sensitivity of K. (a) The left panel is the number of Pauli terms of $\boldsymbol{H}_t$ during the flow of NQOF for LiH, N_2, and Li_2O. The solid and dotted blue curves denote the retained and maximum numbers of Pauli terms, respectively, while the dark red dashed lines refer to the size of the complete n-qubit Pauli space. The right panel shows an empirical scaling of the terms of $\boldsymbol{H}_t$. (b) The impact of K on the energy error ΔE and state fidelity along the flow over different systems.

The impact of the K on performance. Furthermore, we investigate how the size K of the sampled Pauli basis affects the performance of NQOF for LiH, QHM, and TFIM with 8-qubit shown in Fig. 4(b). Increasing K enlarges the accessible generator subspace and therefore allows for a more accurate approximation of the target flow direction. Consequently, the energy errors for the Heisenberg model and TFIM decrease by orders of magnitude as K increases from 8 to 64, while the state fidelity approaches 1. For LiH, NQOF attains **CA** with only $K = 8$, despite the relatively

low flow fidelity. This behavior arises because $K = 8$ is sufficient to capture the dominant energy-lowering directions and attain **CA**, but not to reproduce the complete Brockett-Wegner trajectory. Increasing K recovers the additional directions and improves the flow fidelity, even though the energy error has already saturated. Overall, these results reveal that we can make a trade-off between K and approximation quality to obtain an accurate estimation.

The term number of $\boldsymbol{H}_t$ maintained. We first examine the number of terms of the transformed Hamiltonian $\boldsymbol{H}_t$ required by NQOF to achieve **CA** as the system size increases. To this end, we consider three molecular systems of different sizes. As shown in Fig. 4, the number of Pauli terms in $\boldsymbol{H}_t$ initially increases rapidly but subsequently saturates for all cases, indicating that NQOF manipulates only a limited subset of the full operator space throughout the flow. Moreover, the gap between the number of retained terms and the size of the complete Pauli basis becomes increasingly pronounced as the system size grows. In particular, for the 30-qubit Li_2O system, the full operator space contains $4^{30} \approx 1.15 \times 10^{18}$ Pauli strings, whereas NQOF retains only approximately 10^9 terms. An empirical fit over the considered system sizes further suggests that the number of retained operators scales approximately as $N_{\text{op}} \propto n^{8.03}$ rather than exponential growth over the tested range.

Exact-flow validation of NQOF. We investigate whether NQOF generates a trajectory that matches the target Brockett-Wegner flow. As shown in Fig. 5, the NQOF can closely reproduce the state evolved along the exact flow with the same initial state. The two trajectories exhibit nearly identical changes in their matrix-element structures across all selected flow steps. At $t = 80$, the element-wise error remains below approximately 10^{-5}, indicating that the approximation error stays well controlled over the evolution, which demonstrates that NQOF accurately captures the target flow dynamics.

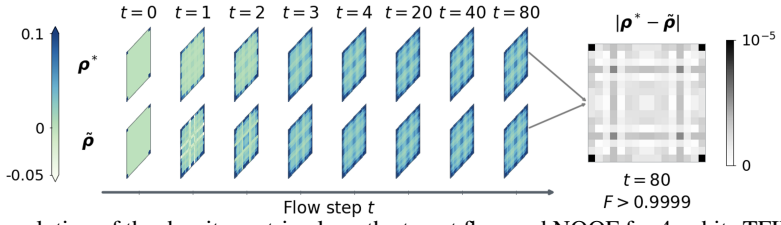


Figure 5: The evolution of the density matrix along the target flow and NQOF for 4 qubits TFIM model. The density matrix of initial state is $\boldsymbol{\rho}_0 = (|+\rangle\langle+|)^{\otimes 4}$, where $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. The upper and lower rows show the density matrices $\boldsymbol{\rho}_t^\star$ and $\tilde{\boldsymbol{\rho}}_t$ respectively obtained from the exact flow and NQOF at selected flow steps $t \in \{0, 1, 2, 3, 4, 20, 40, 80\}$. The right panel shows the element-wise error matrix $|\boldsymbol{\rho}^\star - \tilde{\boldsymbol{\rho}}|$ at the final step.

NQOF achieves competitive performance with fewer parameters. To investigate the efficiency of NQOF, we compare the parameter efficiency of NQOF with that of NQS (Shang et al., 2025) under the same accuracy requirement (10^{-3}) over various molecular systems, as shown in Fig. 6. Transformer-based NQS requires approximately 0.85M parameters across all six molecules, whereas NQOF attains **CA** using only around 70k parameters. Specifically, NQOF uses at least $11.8\times$ fewer parameters for these molecules, corresponding to reductions of around 91%. This substantial reduction suggests that directly learning a compact set of relevant operators and their coefficients can be considerably more parameter-efficient than representing the full many-body wavefunction.

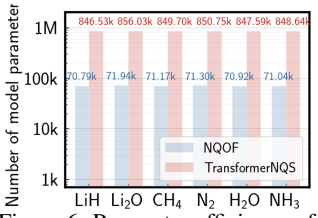


Figure 6: Parameter efficiency of NQOF compared with NQS.

5 Conclusion

In this work, we introduced NQOF, a learning-based framework for ground state energy optimization that approximates the Brockett-Wegner flow using a compact, adaptively learned Pauli subspace. By combining autoregressive operator selection, quadratic coefficient optimization, and Hamiltonian propagation, NQOF avoids direct optimization of a many-body wavefunction and reduces the need to explore the full operator space. The results on molecular and many-body systems demonstrate accurate ground energy estimation with high parameter efficiency, highlighting learned operator flows as a promising and complementary paradigm for scalable quantum optimization.

REFERENCES

Thomas D Barrett, Aleksei Malyshev, and Alexander I Lvovsky. Autoregressive neural-network wavefunctions for ab initio quantum chemistry. *Nature Machine Intelligence*, 4(4):351–358, 2022.

Nicholas P. Bauman and Karol Kowalski. Coupled cluster downfolding methods: The effect of double commutator terms on the accuracy of ground-state energies. *The Journal of Chemical Physics*, 156(9):094106, 2022. doi: 10.1063/5.0076260. URL https://doi.org/10.1063/5.0076260.

Roger W. Brockett. Dynamical systems that sort lists, diagonalize matrices, and solve linear programming problems. *Linear Algebra and its Applications*, 146:79–91, 1991. doi: 10.1016/0024-3795(91)90021-N.

Marin Bukov, Markus Schmitt, and Maxime Dupont. Learning the ground state of a non-stoquastic quantum hamiltonian in a rugged neural network landscape. *SciPost Physics*, 10(6):147, 2021. doi: 10.21468/SciPostPhys.10.6.147. URL https://doi.org/10.21468/SciPostPhys.10.6.147.

Yudong Cao, Jonathan Romero, Jonathan P. Olson, Matthias Degroote, Peter D. Johnson, Mária Kieferová, Ian D. Kivlichan, Tim Menke, Borja Peropadre, Nicolas P. D. Sawaya, Sukin Sim, Libor Veis, and Alán Aspuru-Guzik. Quantum chemistry in the age of quantum computing. *Chemical Reviews*, 119(19):10856–10915, 2019. doi: 10.1021/acs.chemrev.8b00803.

Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017. doi: 10.1126/science.aag2302. URL https://doi.org/10.1126/science.aag2302.

Andrew M. Childs, Yuan Su, Minh C. Tran, Nathan Wiebe, and Shuchen Zhu. Theory of trotter error with commutator scaling. *Physical Review X*, 11(1):011020, 2021. doi: 10.1103/PhysRevX.11.011020.

Kenny Choo, Antonio Mezzacapo, and Giuseppe Carleo. Fermionic neural-network states for ab-initio electronic structure. *Nature Communications*, 11(1):2368, 2020. doi: 10.1038/s41467-020-15724-9.

J. Ignacio Cirac, David Pérez-García, Norbert Schuch, and Frank Verstraete. Matrix product states and projected entangled pair states: Concepts, symmetries, and theorems. *Reviews of Modern Physics*, 93(4):045003, 2021. doi: 10.1103/RevModPhys.93.045003. URL https://doi.org/10.1103/RevModPhys.93.045003.

Benjamin Commeau, M. Cerezo, Zoë Holmes, Lukasz Cincio, Patrick J. Coles, and Andrew Sornborger. Variational hamiltonian diagonalization for dynamical quantum simulation. *arXiv preprint arXiv:2009.02559*, 2020. doi: 10.48550/arXiv.2009.02559. URL https://arxiv.org/abs/2009.02559.

Duy-Khoi Dang, Joshua A. Kammeraad, and Paul M. Zimmerman. Advances in parallel heat bath configuration interaction. *The Journal of Physical Chemistry A*, 127(1):400–411, 2023. doi: 10.1021/acs.jpca.2c07949. URL https://doi.org/10.1021/acs.jpca.2c07949.

Yulong Dong, Lin Lin, and Yu Tong. Ground-state preparation and energy estimation on early fault-tolerant quantum computers via quantum eigenvalue transformation of unitary matrices. *PRX Quantum*, 3(4):040305, 2022. doi: 10.1103/PRXQuantum.3.040305. URL https://doi.org/10.1103/PRXQuantum.3.040305.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.

Yuxuan Du, Min-Hsiu Hsieh, and Dacheng Tao. Efficient learning for linear properties of bounded-gate quantum circuits. *Nature Communications*, 16(1):3790, 2025.

Jens Eisert, Marcus Cramer, and Martin B. Plenio. Colloquium: Area laws for the entanglement entropy. *Reviews of Modern Physics*, 82(1):277–306, 2010. doi: 10.1103/RevModPhys.82.277.

Francesco A. Evangelista. A driven similarity renormalization group approach to quantum many-body problems. *The Journal of Chemical Physics*, 141(5):054109, 2014. doi: 10.1063/1.4890660. URL https://doi.org/10.1063/1.4890660.

Hong Gao, Satoshi Imamura, Akihiko Kasagi, and Eiji Yoshida. Distributed implementation of full configuration interaction for one trillion determinants. *Journal of Chemical Theory and Computation*, 20(3):1185–1192, 2024. doi: 10.1021/acs.jctc.3c01190. URL https://doi.org/10.1021/acs.jctc.3c01190.

Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. MADE: Masked autoencoder for distribution estimation. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 881–889. PMLR, 2015.

Marek Gluza. Double-bracket quantum algorithms for diagonalization. *Quantum*, 8: 1316, 2024. doi: 10.22331/q-2024-04-09-1316. URL https://doi.org/10.22331/q-2024-04-09-1316.

Marek Gluza, Jeongrak Son, Bi Hong Tiang, René Zander, Raphael Seidel, Yudai Suzuki, Zoë Holmes, and Nelly H. Y. Ng. Double-bracket quantum algorithms for quantum imaginary-time evolution. *Physical Review Letters*, 136(2):020601, 2026. doi: 10.1103/rw81-k8vk. URL https://doi.org/10.1103/rw81-k8vk.

Gene H Golub, Per Christian Hansen, and Dianne P O'Leary. Tikhonov regularization and total least squares. *SIAM journal on matrix analysis and applications*, 21(1):185–194, 1999.

H. Hergert, S. K. Bogner, T. D. Morris, A. Schwenk, and K. Tsukiyama. The in-medium similarity renormalization group: A novel ab initio method for nuclei. *Physics Reports*, 621:165–222, 2016. doi: 10.1016/j.physrep.2015.12.007. URL https://doi.org/10.1016/j.physrep.2015.12.007.

Sabine Jansen, Mary-Beth Ruskai, and Ruedi Seiler. Bounds for the adiabatic approximation with applications to quantum computation. *Journal of Mathematical Physics*, 48(10):102111, October 2007. doi: 10.1063/1.2798382. URL https://doi.org/10.1063/1.2798382.

Pascual Jordan and Eugene Wigner. Über das paulische äquivalenzverbot. *Zeitschrift für Physik*, 47 (9):631–651, 1928.

Peter J. Knowles and Nicholas C. Handy. A new determinant-based full configuration interaction method. *Chemical Physics Letters*, 111(4-5):315–321, 1984. doi: 10.1016/0009-2614(84)85513-X.

Karol Kowalski and Nicholas P. Bauman. Quantum flow algorithms for simulating many-body systems on quantum computers. *Physical Review Letters*, 131(20):200601, 2023. doi: 10.1103/PhysRevLett.131.200601. URL https://doi.org/10.1103/PhysRevLett.131.200601.

Karol Kowalski and Nicholas P. Bauman. Resource-adaptive quantum flow algorithms for quantum simulations of many-body systems: Sub-flow embedding procedures. *Molecular Physics*, 123 (19–20):e2489774, 2025. doi: 10.1080/00268976.2025.2489774. URL https://doi.org/10.1080/00268976.2025.2489774.

Ashutosh Kumar, Ayush Asthana, Conner Masteran, Edward F. Valeev, Yu Zhang, Lukasz Cincio, Sergei Tretiak, and Pavel A. Dub. Quantum simulation of molecular electronic states with a transcorrelated hamiltonian: Higher accuracy with fewer qubits. *Journal of Chemical Theory and Computation*, 18(9):5312–5324, 2022. doi: 10.1021/acs.jctc.2c00520. URL https://doi.org/10.1021/acs.jctc.2c00520.

Hannah Lange, Anka Van de Walle, Atiye Abedinnia, and Annabelle Bohrdt. From architectures to applications: A review of neural quantum states. *Quantum Science and Technology*, 9(4):040501, 2024. doi: 10.1088/2058-9565/ad7168. URL https://doi.org/10.1088/2058-9565/ad7168.

Hugo Larochelle and Iain Murray. The neural autoregressive distribution estimator. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pp. 29–37. PMLR, 2011.

James B. Larsen, Matthew D. Grace, Andrew D. Baczewski, and Alicia B. Magann. Feedback-based quantum algorithms for ground state preparation. *Physical Review Research*, 6(3):033336, 2024. doi: 10.1103/PhysRevResearch.6.033336. URL https://doi.org/10.1103/PhysRevResearch.6.033336.

Lin Lin and Yu Tong. Near-optimal ground state preparation. *Quantum*, 4:372, December 2020. doi: 10.22331/q-2020-12-14-372. URL https://doi.org/10.22331/q-2020-12-14-372.

Andrew Lucas. Ising formulations of many NP problems. *Frontiers in Physics*, 2:5, 2014. doi: 10.3389/fphy.2014.00005.

Alicia B. Magann, Kenneth M. Rudinger, Matthew D. Grace, and Mohan Sarovar. Feedback-based quantum optimization. *Physical Review Letters*, 129(25):250502, 2022. doi: 10.1103/PhysRevLett.129.250502. URL https://doi.org/10.1103/PhysRevLett.129.250502.

Matija Medvidović and Javier Robledo Moreno. Neural-network quantum states for many-body physics. *The European Physical Journal Plus*, 139:631, 2024. doi: 10.1140/epjp/s13360-024-05311-y. URL https://doi.org/10.1140/epjp/s13360-024-05311-y.

R. J. Needs, M. D. Towler, N. D. Drummond, P. López Ríos, and J. R. Trail. Variational and diffusion quantum monte carlo calculations with the CASINO code. *The Journal of Chemical Physics*, 152(15):154106, 2020. doi: 10.1063/1.5144288. URL https://doi.org/10.1063/1.5144288.

Verena A. Neufeld and Timothy C. Berkelbach. Highly accurate electronic structure of metallic solids from coupled-cluster theory with nonperturbative triple excitations. *Physical Review Letters*, 131(18):186402, 2023. doi: 10.1103/PhysRevLett.131.186402. URL https://doi.org/10.1103/PhysRevLett.131.186402.

John P. Perdew, Adrienn Ruzsinszky, Jianwei Sun, Niraj K. Nepal, and Aaron D. Kaplan. Interpretations of ground-state symmetry breaking and strong correlation in wavefunction and density functional theories. *Proceedings of the National Academy of Sciences*, 118(4):e2017850118, 2021. doi: 10.1073/pnas.2017850118. URL https://doi.org/10.1073/pnas.2017850118.

George D. Purvis and Rodney J. Bartlett. A full coupled-cluster singles and doubles model: The inclusion of disconnected triples. *The Journal of Chemical Physics*, 76(4):1910–1918, 1982. doi: 10.1063/1.443164.

Weiluo Ren, Weizhong Fu, Xiaojie Wu, and Ji Chen. Towards the ground state of molecules via diffusion monte carlo on neural networks. *Nature Communications*, 14:1860, 2023. doi: 10.1038/s41467-023-37609-3. URL https://doi.org/10.1038/s41467-023-37609-3.

Matteo Robbiati, Edoardo Pedicillo, Andrea Pasquale, Xiaoyue Li, Oriel Kiss, Andrew Wright, Renato M. S. Farias, Khanh Uyen Giang, Jeongrak Son, Johannes Knörzer, Siong Thye Goh, Jun Yong Khoo, Nelly H. Y. Ng, Zoë Holmes, Stefano Carrazza, and Marek Gluza. Double-bracket quantum algorithms for high-fidelity ground-state preparation. *Physical Review Research*, 8(3):033078, 2026a. doi: 10.1103/jz88-32rb. URL https://doi.org/10.1103/jz88-32rb.

Matteo Robbiati, Edoardo Pedicillo, Andrea Pasquale, Xiaoyue Li, Oriel Kiss, Andrew Wright, Renato MS Farias, Khanh Uyen Giang, Jeongrak Son, Johannes Knörzer, et al. Double-bracket quantum algorithms for high-fidelity ground-state preparation. *Physical Review Research*, 8(3):033078, 2026b.

Clemens C. J. Roothaan. New developments in molecular orbital theory. *Reviews of Modern Physics*, 23(2):69–89, 1951. doi: 10.1103/RevModPhys.23.69.

Manuel S Rudolph, Tyson Jones, Yanting Teng, Armando Angrisani, and Zoë Holmes. Pauli propagation: A computational framework for simulating quantum systems. *PRX Quantum*, 7(3): 032001, 2026.

Tobias Schäfer. Ground states for metals from converged coupled cluster calculations. *The Journal of Physical Chemistry Letters*, 16(1):17–23, 2025. doi: 10.1021/acs.jpclett.4c03134. URL https://doi.org/10.1021/acs.jpclett.4c03134.

Honghui Shang, Chu Guo, Yangjun Wu, Zhenyu Li, and Jinlong Yang. Solving the many-electron schrödinger equation with a transformer-based framework. *Nature Communications*, 16(1):8464, 2025.

Or Sharir, Yoav Levine, Noam Wies, Giuseppe Carleo, and Amnon Shashua. Deep autoregressive models for the efficient variational simulation of many-body quantum systems. *Physical Review Letters*, 124(2):020503, 2020. doi: 10.1103/PhysRevLett.124.020503.

Or Sharir, Amnon Shashua, and Giuseppe Carleo. Neural tensor contractions and the expressive power of deep neural quantum states. *Physical Review B*, 106(20):205136, 2022. doi: 10.1103/PhysRevB.106.205136. URL https://doi.org/10.1103/PhysRevB.106.205136.

James Shee, Matthias Loipersberger, Adam Rettig, Joonho Lee, and Martin Head-Gordon. Regularized second-order Møller–Plesset theory: A more accurate alternative to conventional MP2 for noncovalent interactions and transition metal thermochemistry for the same computational cost. *The Journal of Physical Chemistry Letters*, 12(50):12084–12097, 2021. doi: 10.1021/acs.jpclett.1c03468. URL https://doi.org/10.1021/acs.jpclett.1c03468.

Chinmay Shrikhande, Arnab Bachhar, Aarón Rodríguez Jiménez, and Nicholas J. Mayhall. Rapid ground state energy estimation with a sparse pauli dynamics-enabled variational double bracket flow. *arXiv preprint arXiv:2511.21651*, 2025. doi: 10.48550/arXiv.2511.21651. URL https://arxiv.org/abs/2511.21651.

S. R. Stroberg, T. D. Morris, and B. C. He. In-medium similarity renormalization group with flowing 3-body operators, and approximations thereof. *Physical Review C*, 110(4):044316, 2024. doi: 10.1103/PhysRevC.110.044316. URL https://doi.org/10.1103/PhysRevC.110.044316.

Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations. *Physical Review Letters*, 94(17):170201, 2005. doi: 10.1103/PhysRevLett.94.170201. URL https://doi.org/10.1103/PhysRevLett.94.170201.

Norm M. Tubman, C. Daniel Freeman, Daniel S. Levine, Diptarka Hait, Martin Head-Gordon, and K. Birgitta Whaley. Modern approaches to exact diagonalization and selected configuration interaction with the adaptive sampling CI method. *Journal of Chemical Theory and Computation*, 16(4):2139–2159, 2020. doi: 10.1021/acs.jctc.8b00536. URL https://doi.org/10.1021/acs.jctc.8b00536.

Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1747–1756. PMLR, 2016.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, pp. 5998–6008, 2017.

Steven R. White. Density-matrix algorithms for quantum renormalization groups. *Physical Review B*, 48(14):10345–10356, 1993. doi: 10.1103/PhysRevB.48.10345.

Alexander Wietek and Andreas M. Läuchli. Sublattice coding algorithm and distributed memory parallelization for large-scale exact diagonalizations of quantum many-body systems. *Physical Review E*, 98(3):033309, 2018. doi: 10.1103/PhysRevE.98.033309. URL https://doi.org/10.1103/PhysRevE.98.033309.

Dian Wu, Riccardo Rossi, Filippo Vicentini, Nikita Astrakhantsev, Federico Becca, Xiaodong Cao, Juan Carrasquilla, Francesco Ferrari, Antoine Georges, Mohamed Hibat-Allah, Masatoshi Imada, Andreas M. Läuchli, Guglielmo Mazzola, Antonio Mezzacapo, Andrew Millis, Javier Robledo Moreno, Titus Neupert, Yusuke Nomura, Jannes Nys, Olivier Parcollet, Rico Pohle, Imelda Romero, Michael Schmid, J. Maxwell Silvester, Sandro Sorella, Luca F. Tocchio, Lei Wang, Steven R. White, Alexander Wietek, Qi Yang, Yiqi Yang, Shiwei Zhang, and Giuseppe Carleo. Variational benchmarks for quantum many-body problems. *Science*, 386(6719):296–301, 2024. doi: 10.1126/science.adg9774. URL https://doi.org/10.1126/science.adg9774.

Li Xiaoyue, Matteo Robbiati, Andrea Pasquale, Edoardo Pedicillo, Andrew Wright, Stefano Carrazza, and Marek Gluza. Strategies for optimizing double-bracket quantum algorithms. *arXiv preprint arXiv:2408.07431*, 2024.

Ning Zhang, Wenjian Liu, and Mark R. Hoffmann. Further development of iCIPT2 for strongly correlated electrons. *Journal of Chemical Theory and Computation*, 17(2):949–964, 2021. doi: 10.1021/acs.jctc.0c01187. URL https://doi.org/10.1021/acs.jctc.0c01187.

AI USE STATEMENT

In this work, we used generative AI tools for minor language polishing and expression refinement. We have not used generative AI tools for any scientific ideas, experiments, or results do not apply to this work. We have reviewed all AI-assisted work. We take responsibility for the final content of this work, including those influenced by AI-assisted text.

ETHICS STATEMENT

The research presented in this paper conform with the ICLR Code of Ethics. There is no societal impact of the work performed. Our work does not describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse. The paper does not involve crowdsourcing nor research with human subjects.

REPRODUCIBILITY STATEMENT

To support reproducibility, we release our source code at https://anonymous.4open.science/r/NQO_flow-EF78. The main text describes the NQOF framework and optimization procedure, while the appendix provides implementation details and experimental settings, including model architectures, hyperparameters, benchmark Hamiltonians, and evaluation protocols. The assumptions underlying our theoretical results and their complete proofs are also provided in the appendix.

APPENDIX

A NOTATION AND MATHEMATICAL PRELIMINARIES

States and operators. An n-qubit system has Hilbert space $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ of dimension $d = 2^n$. A pure state is represented by a normalized vector $|\boldsymbol{u}\rangle \in \mathcal{H}$, with $\langle \boldsymbol{u}| = |\boldsymbol{u}\rangle^\dagger$ and $\langle \boldsymbol{u}|\boldsymbol{u}\rangle = 1$. The symbols $(\cdot)^\dagger$, $(\cdot)^T$, Tr, and Re denote the conjugate transpose, transpose, trace, and real part, respectively. An operator $\mathbf{A}$ is Hermitian if $\mathbf{A}^\dagger = \mathbf{A}$, and anti-Hermitian if $\mathbf{A}^\dagger = -\mathbf{A}$. A unitary operator satisfies $\mathbf{U}^\dagger\mathbf{U} = \mathbf{U}\mathbf{U}^\dagger = \mathbf{I}$. The unitary group is denoted by $\mathcal{U}(d)$. We write $[\mathbf{A}, \mathbf{B}] = \mathbf{AB} - \mathbf{BA}, \{\mathbf{A}, \mathbf{B}\} = \mathbf{AB} + \mathbf{BA}$ for the commutator and anticommutator. For Hermitian $\mathbf{A}$ and $\mathbf{B}$, their commutator is anti-Hermitian, whereas $[i\mathbf{A}, \mathbf{B}]$ is Hermitian.

Norms and the reference state. The Euclidean norm of a state vector is $\| |\boldsymbol{u}\rangle \|_2 = \sqrt{\langle \boldsymbol{u}|\boldsymbol{u}\rangle}$. For an operator, the Frobenius and operator norms are $\|\mathbf{A}\|_F = \sqrt{\mathrm{Tr}(\mathbf{A}^\dagger\mathbf{A})}$ and $\|\mathbf{A}\|_{\mathrm{op}} = \sup_{\langle \boldsymbol{u}|\boldsymbol{u}\rangle=1} \|\mathbf{A}|u\rangle\|_2$. Throughout, $|\boldsymbol{v}\rangle$ is a fixed normalized reference state and $\rho = |\boldsymbol{v}\rangle\langle \boldsymbol{v}|$ is its rank-one projector. The state-dependent seminorm used in the fitting objective is $\|\mathbf{A}\|_{\boldsymbol{v}} := \|\mathbf{A}|\boldsymbol{v}\rangle\|_2 = \sqrt{\mathrm{Tr}(\rho\mathbf{A}^\dagger\mathbf{A})}$. It is a seminorm because a nonzero operator can annihilate $|v\rangle$. The following inequalities relate these quantities: $|\langle \boldsymbol{v}|\mathbf{A}|\boldsymbol{v}\rangle| \le \|\mathbf{A}\|_{\boldsymbol{v}} \le \|\mathbf{A}\|_{\mathrm{op}} \le \|\mathbf{A}\|_F$.

Pauli representation. The single-qubit Pauli matrices are

$$\mathbf{I} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{X} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \mathbf{Y} = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \mathbf{Z} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

An n-qubit Pauli string is $\mathbf{P} = \bigotimes_{j=1}^n \mathbf{P}^{(j)}$, where $\mathbf{P}^{(j)} \in \{\mathbf{I}, \mathbf{X}, \mathbf{Y}, \mathbf{Z}\}$. The set $\mathcal{P}_n$ of all 4^n such strings forms an orthogonal operator basis: $\mathbf{P}^\dagger = \mathbf{P}, \mathbf{P}^2 = \mathbf{I}$. Consequently, every Hermitian operator admits the expansion $\mathbf{A} = \sum_{\mathbf{P}\in\mathcal{P}_n} a_{\mathbf{P}}\mathbf{P}, a_{\mathbf{P}} = \frac{1}{d}\mathrm{Tr}(\mathbf{PA}) \in \mathbb{R}, \|\mathbf{A}\|_F^2 = d\sum_{\mathbf{P}\in\mathcal{P}_n} a_{\mathbf{P}}^2$. Any two Pauli strings either commute or anticommute.

Unitary orbits and operator flows. The unitary orbit of $\boldsymbol{H}$ is $\mathcal{M}(\boldsymbol{H}) = \{\boldsymbol{U}^\dagger\boldsymbol{H}\mathbf{U} : \boldsymbol{U} \in \mathcal{U}(d)\}$. Every operator on this orbit has the same spectrum. At $\boldsymbol{A} \in \mathcal{M}(\boldsymbol{H})$, the tangent space is $T_{\boldsymbol{A}}\mathcal{M}(\boldsymbol{H}) = \{[\boldsymbol{K}, \boldsymbol{A}] : \boldsymbol{K}^\dagger = -\boldsymbol{K}\} = \{[i\boldsymbol{B}, \boldsymbol{A}] : \boldsymbol{B}^\dagger = \boldsymbol{B}\}$. Thus, an operator flow satisfies $\dot{\mathbf{H}}_t = \Phi_t(\mathbf{H}_t)$ with $\Phi_t(\mathbf{H}_t) \in T_{\mathbf{H}_t}\mathcal{M}(\mathbf{H})$. Here, $t \ge 0$ denotes continuous flow time.

B RELATED WORKS

State-centered methods. Most existing approaches adopt a state-centered perspective, seeking the ground state through exact, approximate, or compressed representations of the wavefunction (Wu et al., 2024; Medvidović & Robledo Moreno, 2024). Exact diagonalization (full configuration interaction, FCI) and Krylov-subspace methods directly solve the eigenvalue problem and provide reliable reference results, but their computational and memory costs grow exponentially with system size (Wietek & Läuchli, 2018; Tubman et al., 2020; Gao et al., 2024). More scalable approaches reduce this cost through structured approximations, compressed representations, or stochastic sampling, including mean-field (Perdew et al., 2021), perturbative approaches (Shee et al., 2021), configuration-interaction (Zhang et al., 2021; Dang et al., 2023), coupled-cluster (Neufeld & Berkelbach, 2023; Schäfer, 2025)), tensor networks (Cirac et al., 2021), quantum Monte Carlo (Needs et al., 2020; Ren et al., 2023), and neural quantum states (Carleo & Troyer, 2017; Lange et al., 2024; Medvidović & Robledo Moreno, 2024). Nevertheless, their accuracy and efficiency can remain sensitive to the expressivity of the chosen wavefunction representation and optimization challenges (Sharir et al., 2022; Bukov et al., 2021; Wu et al., 2024), fermionic quantum Monte Carlo methods additionally face the sign or phase problem (Troyer & Wiese, 2005; Ren et al., 2023).

Operator-centered methods. Operator-centered methods aim to transform the Hamiltonian such that its ground energy can be extracted more efficiently, and can be mainly divided into classical and quantum approaches. Classical methods follow two main strategies: flow-based methods, such as SRG, IMSRG, DSRG, vDBF, progressively diagonalize or decouple the Hamiltonian (Hergert

et al., 2016; Evangelista, 2014; Stroberg et al., 2024; Shrikhande et al., 2025; Robbiati et al., 2026b), whereas downfolding and transcorrelated methods construct compact effective Hamiltonians that capture the relevant low-energy (Bauman & Kowalski, 2022; Kumar et al., 2022). Their accuracy can be limited by operator proliferation, active-space selection, truncation of higher-body interactions, and greedy selection of the Pauli basis. Quantum approaches realize related Hamiltonian transformations or operator-guided energy reduction through quantum circuits, including QFlow, variational Hamiltonian diagonalization, double-bracket algorithms, feedback-based control, and spectral-filtering methods (Kowalski & Bauman, 2023; 2025; Commeau et al., 2020; Gluza, 2024; Xiaoyue et al., 2024; Gluza et al., 2026; Robbiati et al., 2026b; Magann et al., 2022; Larsen et al., 2024; Dong et al., 2022). However, large-scale implementation of these methods typically requires deep quantum circuits and levels of noise suppression beyond the capabilities of current noisy devices.

C Implementation Details

C.1 Neural Architecture

To efficiently explore the exponentially large space of 4^n Pauli strings, we parameterize the probability distribution over Pauli strings using an autoregressive Transformer $f_{\boldsymbol{W}}$. Given a Pauli string $\boldsymbol{P} = \boldsymbol{P}^{(1)}, \ldots, \boldsymbol{P}^{(n)}$, each local operator $\boldsymbol{P}^{(i)} \in \{\boldsymbol{I}, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}\}$ is represented by a four-dimensional one-hot vector and augmented with positional information. Concretely, we only utilize the decoder part of the Transformer, where the masked self-attention captures correlations among Pauli operators at different sites while preserving causality, such that the conditional distribution at site i depends only on the prefix $\boldsymbol{P}^{(1)} \cdots \boldsymbol{P}^{(i-1)}$. As illustrated in Fig. 7, the autoregressive model $f_{\boldsymbol{W}}$ consists of $L = 2$ Transformer layers, each with four masked attention heads, a hidden dimension of $d = 64$, and a feed-forward dimension of 128. The output layer has dimension 4 and produces the categorical probabilities of $\{\boldsymbol{I}, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}\}$ at each site. This architecture defines a normalized and expressive probability distribution over Pauli strings, allowing exact likelihood evaluation and direct ancestral sampling without enumerating the full operator space or relying on Markov-chain sampling.

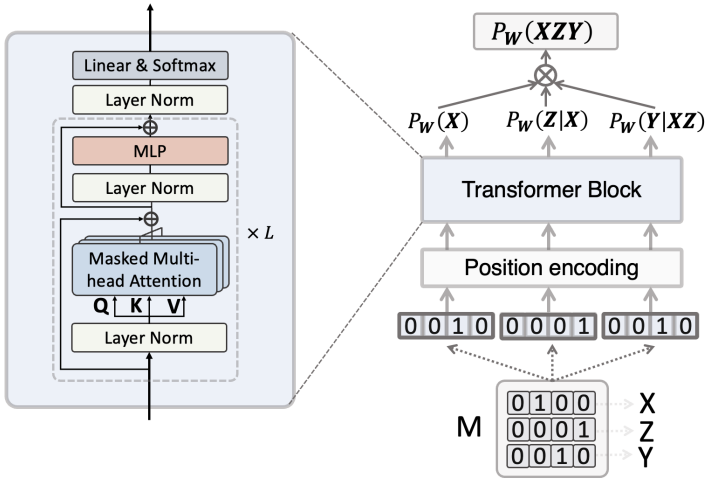


Figure 7: Transformer-based autoregressive parameterization of probabilities over Pauli strings. Each local Pauli operator is represented by a four-dimensional one-hot vector, forming the input matrix $\boldsymbol{M}$. After positional information is added, the encoded sequence is processed by L Transformer layers, each comprising masked multi-head self-attention, a multilayer perceptron, layer normalization, and residual connections. The causal attention mask ensures that the probability at each site depends only on the preceding Pauli operators. A linear layer followed by a softmax produces the conditional distributions whose product gives the probability of the complete string, such as $P_{\boldsymbol{W}}(\boldsymbol{XZY}) = P_{\boldsymbol{W}}(\boldsymbol{X})P_{\boldsymbol{W}}(\boldsymbol{Z} \mid \boldsymbol{X})P_{\boldsymbol{W}}(\boldsymbol{Y} \mid \boldsymbol{XZ})$.

C.2 Sampling Method

For a n-qubit Pauli string $\boldsymbol{P} = \boldsymbol{P}^{(1)}, \ldots, \boldsymbol{P}^{(n)}$, the parametrized probability distribution $P_{\boldsymbol{W}}$ can be autoregressively expressed as $P_{\boldsymbol{W}}(\boldsymbol{P}) = \prod_{i=1}^{n} P_{\boldsymbol{W}}\big(\boldsymbol{P}^{(i)} \mid \boldsymbol{P}^{(1:i-1)}\big)$, where $\boldsymbol{P}^{(1:i-1)} = \boldsymbol{P}^{(1)}, \ldots, \boldsymbol{P}^{(i-1)}$.

Conditional Pauli string generation. Here, we utilize a Transformer $f_{\boldsymbol{W}}$ to parametrized the conditional probability $P_{\boldsymbol{W}}(\boldsymbol{P}^{(i)} = \boldsymbol{\alpha} \mid \boldsymbol{P}^{(1:i-1)})$ via the above chain-rule factorization,

$$P_{\boldsymbol{W}}(\boldsymbol{P}^{(i)} = \boldsymbol{\alpha} \mid \boldsymbol{P}^{(1:i-1)}) = \mathrm{softmax}_{\boldsymbol{\alpha}}\left(\frac{\ell_{\boldsymbol{W}}(\boldsymbol{P}^{(1:i-1)})}{\tau}\right), \tag{8}$$

where $\boldsymbol{\alpha} \in \{\boldsymbol{I}, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}\}$, $\ell_{\boldsymbol{W}}(\boldsymbol{P}^{(1:i-1)})$ is the output before softmax with $\boldsymbol{P}^{(1:i-1)}$ as input, $\mathrm{softmax}_{\boldsymbol{\alpha}}$ assigns the probability for $\boldsymbol{P}^{(i)} = \boldsymbol{\alpha}$, τ is the softmax temperature factor used to control the randomness of predictions; here we fix $\tau = 0.8$. Sampling begins with an empty prefix. At site i, the Transformer $f_{\boldsymbol{W}}$ takes the previously sampled prefix $\boldsymbol{P}^{(1:i-1)}$ as input and outputs the logits for the next Pauli symbol. After applying temperature scaling and the validity mask, a symbol $\boldsymbol{\alpha} \in \{\boldsymbol{I}, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}\}$ is drawn from the resulting categorical distribution $P_{\boldsymbol{W}}\left(\boldsymbol{P}^{(i)} = \boldsymbol{\alpha} \mid \boldsymbol{P}^{(1:i-1)}\right)$, appended to the prefix, and used to generate the subsequent symbol. Repeating this sequential procedure with independent randomness yields independent samples of Pauli strings. This autoregressive sampling process requires neither Markov-chain burn-in nor enumeration of the full 4^n Pauli space.

C.3 Implementation details of NQOF

At each flow step t, we optimize $\boldsymbol{\lambda}$ and $\boldsymbol{W}$ using a state-projected loss. The original residual in Eq. (5) compares the two tangent directions globally and therefore penalizes discrepancies across all matrix elements of the operator space. However, the energy $E_t = \langle \boldsymbol{v}|\boldsymbol{H}_t|\boldsymbol{v}\rangle$ depends only on how the evolving Hamiltonian acts on the fixed reference state. To reduce the cost of evaluating the global operator mismatch while retaining direct control over the instantaneous energy change, we instead employ

$$\mathcal{L}(\boldsymbol{\lambda}, \mathbb{S}_{\boldsymbol{P}}) = \frac{1}{2}\left\|\left(-[[\rho, \boldsymbol{H}_t], \boldsymbol{H}_t] - [i\tilde{\boldsymbol{H}}_K, \boldsymbol{H}_t]\right)|\boldsymbol{v}\rangle\right\|_2^2, \qquad \tilde{\boldsymbol{H}}_K = \sum_{j=1}^{K} \lambda_j \boldsymbol{P}_j, \tag{9}$$

where $\mathbb{S}_{\boldsymbol{P}} = (\boldsymbol{P}_1, \cdots, \boldsymbol{P}_K)$, $\rho = |\boldsymbol{v}\rangle\langle\boldsymbol{v}|$. Defining the tangent mismatch as $\delta\boldsymbol{D}_t = \boldsymbol{D}_t - [i\tilde{\boldsymbol{H}}_K, \boldsymbol{H}_t]$, we can rewrite the the loss as

$$\mathcal{L} = \frac{1}{2}\|\delta\boldsymbol{D}_t|\boldsymbol{v}\rangle\|_2^2 = \frac{1}{2}\,\mathrm{Tr}\left(\rho\,\delta\boldsymbol{D}_t^{\dagger}\delta\boldsymbol{D}_t\right). \tag{10}$$

Thus, $\mathcal{L}$ measures the tangent mismatch using the state-dependent seminorm $\|\boldsymbol{A}\|_{\boldsymbol{v}} = \|\boldsymbol{A}|\boldsymbol{v}\rangle\|_2$. Unlike the Frobenius norm, which measures the action of $\boldsymbol{A}$ on a complete basis, this seminorm measures only its action on $|\boldsymbol{v}\rangle$. The components satisfying $\delta\boldsymbol{D}_t|\boldsymbol{v}\rangle = 0$ are not penalized because they do not affect the instantaneous energy change of the target flow, which is computationally cheaper than the original residual R. This projected loss bounds the error in the instantaneous energy descent,

$$\left|\langle\boldsymbol{v}|\boldsymbol{D}_t|\boldsymbol{v}\rangle - \langle\boldsymbol{v}|[i\tilde{\boldsymbol{H}}_K, \boldsymbol{H}_t]|\boldsymbol{v}\rangle\right| \le \|\delta\boldsymbol{D}_t|\boldsymbol{v}\rangle\|_2 = \sqrt{2\mathcal{L}}. \tag{11}$$

Therefore, a sufficiently small residual therefore gives accurate matching of the instantaneous energy change. Minimizing the residual does not guarantee that its attained value is small for every selected support. In particular, the fitted direction decreases the represented energy to first order whenever $\mathcal{L} < 2\,\mathrm{Var}(\boldsymbol{H}_t)^2$, while avoiding the need to match operator components that are irrelevant to the chosen reference state.

Finding optimal $\boldsymbol{\lambda}$. For a sampled Pauli set $\mathbb{S}_{\boldsymbol{P}} = \{\boldsymbol{P}_1, \ldots, \boldsymbol{P}_K\}$ according to $P_{\boldsymbol{W}}$, only the real coefficients $\boldsymbol{\lambda} \in \mathbb{R}^K$ need to be optimized. Here, we define

$$\boldsymbol{B}_j = [i\boldsymbol{P}_j, \boldsymbol{H}_t], \qquad \boldsymbol{D}_t = -[[\rho, \boldsymbol{H}_t], \boldsymbol{H}_t], \tag{12}$$

Thus, we have $[i\tilde{\boldsymbol{H}}_K, \boldsymbol{H}_t] = \sum_{j=1}^{K} \lambda_j \boldsymbol{B}_j$ and the $\mathcal{L}$ becomes

$$\mathcal{L}(\boldsymbol{\lambda}, \mathbb{S}_{\boldsymbol{P}}) = \frac{1}{2} \left\| \left(\boldsymbol{D}_t - \sum_{j=1}^{K} \lambda_j \boldsymbol{B}_j \right) |\boldsymbol{v}\rangle \right\|_2^2 . \tag{13}$$

Define the state vectors $|\boldsymbol{d}_t\rangle = \boldsymbol{D}_t|\boldsymbol{v}\rangle$ and $|\boldsymbol{a}_j\rangle = \boldsymbol{B}_j|\boldsymbol{v}\rangle$. Then, the Eq. (13) can be rewritten as

$$\mathcal{L}(\boldsymbol{\lambda}, \mathbb{S}_{\boldsymbol{P}}) = \frac{1}{2} \left\| |\boldsymbol{d}_t\rangle - \sum_{j=1}^{K} \lambda_j |\boldsymbol{a}_j\rangle \right\|_2^2 , \tag{14}$$

Expanding the squared norm gives

$$\mathcal{L}(\boldsymbol{\lambda}, \mathbb{S}_{\boldsymbol{P}}) = \frac{1}{2}\langle \boldsymbol{d}_t|\boldsymbol{d}_t\rangle - \sum_{j=1}^{K} \lambda_j \ \mathrm{Re}\langle \boldsymbol{a}_j|\boldsymbol{d}_t\rangle + \frac{1}{2} \sum_{i,j=1}^{K} \lambda_i \lambda_j \ \mathrm{Re}\langle \boldsymbol{a}_i|\boldsymbol{a}_j\rangle. \tag{15}$$

Therefore, define the Gram matrix $\boldsymbol{G}$ and target vector $\boldsymbol{b}$ as

$$G_{ij} = \mathrm{Re}[\langle \boldsymbol{a}_i|\boldsymbol{a}_j\rangle] = \mathrm{Re}[\langle \boldsymbol{v}|\boldsymbol{B}_i\boldsymbol{B}_j|\boldsymbol{v}\rangle], \qquad b_i = \mathrm{Re}[\langle \boldsymbol{a}_i|\boldsymbol{d}_t\rangle] = \mathrm{Re}[\langle \boldsymbol{v}|\boldsymbol{B}_i\boldsymbol{D}_t|\boldsymbol{v}\rangle]. \tag{16}$$

Hence, $\mathcal{L}$ can then be written as

$$\mathcal{L}(\boldsymbol{\lambda}, \mathbb{S}_{\boldsymbol{P}}) = \frac{1}{2}\boldsymbol{\lambda}^{\mathsf{T}}\boldsymbol{G}\boldsymbol{\lambda} - \big(\boldsymbol{b}\big)^{\mathsf{T}}\boldsymbol{\lambda} + c, \tag{17}$$

where the constant $c = \frac{1}{2}\langle \boldsymbol{v}|\boldsymbol{D}_t^2|\boldsymbol{v}\rangle$. Since c does not depend on $\boldsymbol{\lambda}$, it can therefore be omitted during optimization, i.e. $\min_{\boldsymbol{\lambda}} \mathcal{L}$, yielding the quadratic program

$$\boldsymbol{\lambda}^* = \arg\min_{\boldsymbol{\lambda} \in \mathbb{R}^K} \left[\frac{1}{2}\boldsymbol{\lambda}^{\mathsf{T}}\boldsymbol{G}\boldsymbol{\lambda} - \boldsymbol{b}^{\mathsf{T}}\boldsymbol{\lambda} \right]. \tag{18}$$

Since $\boldsymbol{G}$ is positive semidefinite and the problem of Eq. (18) is convex, we can find the optimal $\boldsymbol{\lambda}$ by solving the following linear equation.

$$\boldsymbol{G}\boldsymbol{\lambda}^* = \boldsymbol{b} \tag{19}$$

If $\boldsymbol{G}^{(v)}$ is nonsingular, then we have $\boldsymbol{\lambda}^* = \boldsymbol{G}^{-1}\boldsymbol{b}$, otherwise $\boldsymbol{\lambda}$ can be obtained via the pseudoinverse or Tikhonov regularisation (Golub et al., 1999),

$$\boldsymbol{\lambda}^* = (\boldsymbol{G} + \mu \boldsymbol{I})^{-1} \boldsymbol{b}. \tag{20}$$

Here, we choose $\mu = 10^{-6}$.

Updating $\boldsymbol{W}$. At each flow step t, we train the autoregressive model using a batch of $N = 16K$ candidate Pauli strings, $\mathcal{B} = (\boldsymbol{P}_1, \ldots, \boldsymbol{P}_N)$, sampled independently from $P_{\boldsymbol{W}}$. Each candidate is evaluated through a separate one-dimensional coefficient fit. Specifically, define $\lambda_t^*(\boldsymbol{P}) \in \arg\min_{\lambda \in \mathbb{R}} \ell_t(\lambda, \boldsymbol{P})$, and $r_t(\boldsymbol{P}) = \ell_t\big(\lambda_t^*(\boldsymbol{P}), \boldsymbol{P}\big)$, where $\boldsymbol{D}_t = -[[\rho, \boldsymbol{H}_t], \boldsymbol{H}_t]$. The scalar fits used to evaluate candidates are separate from the joint coefficient fit on the selected K-string support.

We optimize the expected individual fitting loss,

$$\min_{\boldsymbol{W}} \mathcal{J}_{\mathrm{ind}}(\boldsymbol{W}), \qquad \mathcal{J}_{\mathrm{ind}}(\boldsymbol{W}) = \mathbb{E}_{\boldsymbol{P} \sim P_{\boldsymbol{W}}}[r_t(\boldsymbol{P})]. \tag{21}$$

With $\boldsymbol{H}_t$ held fixed during the policy update, $r_t(\boldsymbol{P})$ has no explicit dependence on $\boldsymbol{W}$ for a fixed candidate $\boldsymbol{P}$. The score-function identity therefore gives

$$\nabla_{\boldsymbol{W}} \mathcal{J}_{\mathrm{ind}} = \mathbb{E}_{\boldsymbol{P} \sim P_{\boldsymbol{W}}} \left[r_t(\boldsymbol{P}) \nabla_{\boldsymbol{W}} \log P_{\boldsymbol{W}}(\boldsymbol{P}) \right]. \tag{22}$$

We estimate this gradient using the sampled batch and update the network parameters as

$$\widehat{\boldsymbol{g}}_{\boldsymbol{W}} = \frac{1}{N} \sum_{n=1}^{N} r_t(\boldsymbol{P}_n) \nabla_{\boldsymbol{W}} \log P_{\boldsymbol{W}}(\boldsymbol{P}_n), \tag{23}$$

$$\boldsymbol{W} \leftarrow \boldsymbol{W} - \eta \widehat{\boldsymbol{g}}_{\boldsymbol{W}}. \tag{24}$$

The individual fitting objective provides a surrogate for training the sampler to favor candidates that approximate the target direction well individually. The generator coefficients are subsequently fitted jointly on the selected K-string support, accounting for interactions among the selected directions. After training the autoregressive model at flow step t, we draw a fresh batch of K Pauli strings from the updated distribution and jointly fit their coefficients to construct the generator.

C.4 PAULI PROPAGATION

For Hamiltonian $\boldsymbol{H}_t = \sum_j \alpha_j \boldsymbol{Q}_j$ at flow step t, to efficient evaluate $\boldsymbol{U}_t^\dagger \boldsymbol{H}_t \boldsymbol{U}_t$, where $\boldsymbol{U}_t = \exp(-i\Delta t \tilde{\boldsymbol{H}}_K)$, $\tilde{\boldsymbol{H}}_K = \sum_j \lambda_j \boldsymbol{P}_j$ and , we apply the Trotterized Pauli propagation (Rudolph et al., 2026). Concretely, we first utilize the first-order Trotterization to decompose the unitary $\boldsymbol{U}_t$ into a product of exponentials of individual terms in $\tilde{\boldsymbol{H}}_K$ (Childs et al., 2021),

$$\boldsymbol{U}_t = \left(\prod_{k=1}^{K} \boldsymbol{V}_k \right)^N, \tag{25}$$

where $\boldsymbol{V}_k = \exp(-\frac{i\Delta t}{N} \lambda_k \boldsymbol{P}_k), N \in \mathbb{N}^+$. Since any pair of Pauli strings $\boldsymbol{P}_k, \boldsymbol{Q}_j$ either commute ($[\boldsymbol{P}_k, \boldsymbol{Q}_j] = \boldsymbol{P}_k\boldsymbol{Q}_j - \boldsymbol{Q}_j\boldsymbol{P}_k = 0$) or anti-commute ($\{\boldsymbol{P}_k, \boldsymbol{Q}_j\} = \boldsymbol{P}_k\boldsymbol{Q}_j + \boldsymbol{Q}_j\boldsymbol{P}_k = 0$), we can analytically evaluate $\boldsymbol{U}_t^\dagger \boldsymbol{H}_t \boldsymbol{U}_t$ by applying $\{\boldsymbol{V}_k\}$ sequentially on $\boldsymbol{H}_t$,

$$\boldsymbol{U}_t^\dagger \boldsymbol{H}_t \boldsymbol{U}_t = \sum_j \alpha_j \mathcal{E}_{\boldsymbol{V}_K}[\cdots \mathcal{E}_{\boldsymbol{V}_1}[\mathcal{E}_{\boldsymbol{V}_K}[\cdots \mathcal{E}_{\boldsymbol{V}_1}[\boldsymbol{Q}_j]]]] \tag{26}$$

where we define $\mathcal{E}_{\boldsymbol{V}_k}[\cdot] = \boldsymbol{V}_k^\dagger(\cdot)\boldsymbol{V}_k$. Specifically, if $\boldsymbol{P}_k$ and $\boldsymbol{Q}_j$ anti-commute, it transforms as $\mathcal{E}_{\boldsymbol{V}_k}[\boldsymbol{Q}_j] = \cos(2\theta)\boldsymbol{Q}_j + i\sin(2\theta)\boldsymbol{Q}'_j$, where $\theta = \frac{\Delta t}{N}\lambda_k, \boldsymbol{Q}'_i = [\boldsymbol{P}_k, \boldsymbol{Q}_j]/2$ is also a Pauli string, whereas $\boldsymbol{Q}_j$ remains unchanged.

$$e^{i\theta \boldsymbol{P}_k} \boldsymbol{Q}_j e^{-i\theta \boldsymbol{P}_k} = \begin{cases} \boldsymbol{Q}_j, & [\boldsymbol{P}_k, \boldsymbol{Q}_j] = 0, \\ \cos(2\theta)\boldsymbol{Q}_j + i\sin(2\theta)[\boldsymbol{P}_k, \boldsymbol{Q}_j]/2, & \{\boldsymbol{P}_k, \boldsymbol{Q}_j\} = 0. \end{cases} \tag{27}$$

Here, to avoid operator proliferation, we merge the equal Pauli strings and select a reasonable threshold τ to perform truncation at each propagation step, as shown in Fig. 8. The operator $\tilde{\boldsymbol{H}}_t$ is obtained by propagating its Pauli expansion and discarding coefficients below the threshold τ after individual rotations. Its spectrum is generally not preserved exactly. Theorem 3 bounds its deviation using the total discarded coefficient mass along the same realized sequence.

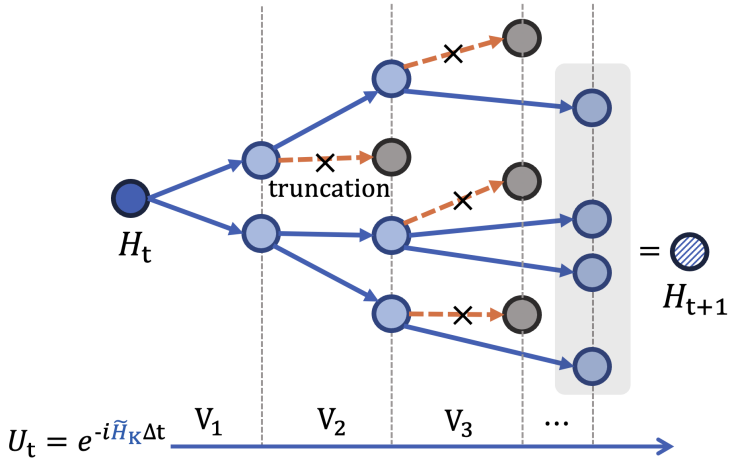


Figure 8: Illustration of Pauli propagation with truncation (Rudolph et al., 2026; Du et al., 2025). The evolution operator $\boldsymbol{U} = e^{-i\tilde{\boldsymbol{H}}_K \Delta t}$ is decomposed into a sequence of Pauli rotations, which are successively applied to the Pauli terms of $\boldsymbol{H}_t$. Each rotation may generate additional Pauli terms, producing a branching operator expansion. Blue nodes and solid arrows denote the retained terms and propagation paths, whereas gray nodes and dashed orange arrows denote generated terms discarded by the truncation criterion. After all rotations, the retained terms are collected to construct the approximate evolved Hamiltonian $\boldsymbol{H}_{t+1} \leftarrow \boldsymbol{U}_t^\dagger \boldsymbol{H}_t \boldsymbol{U}_t$, thereby controlling operator proliferation during the flow.

Remark. We release our source code at `https://anonymous.4open.science/r/NQO_flow-EF78`.

D Details of Experiments

D.1 Molecular Hamiltonians

The electronic Hamiltonian in quantum chemistry is expressed as

$$\boldsymbol{H} = E_{\text{nuc}}\boldsymbol{I} + \sum_{pq} h_{pq}\boldsymbol{a}_p^\dagger \boldsymbol{a}_q + \frac{1}{2}\sum_{pqrs} v_{pqrs}\boldsymbol{a}_p^\dagger \boldsymbol{a}_q^\dagger \boldsymbol{a}_s \boldsymbol{a}_r. \tag{28}$$

where $\boldsymbol{a}_p^\dagger$ and $\boldsymbol{a}_q$ are fermionic creation and annihilation operators, h_{pq}, v_{pqrs} represent one and two-electron integrals. Here, we select STO-3g as the basis set for all molecules and utilize the Jordan-Wigner transformation to map the above fermionic Hamiltonian into a qubit Hamiltonian (Jordan & Wigner, 1928). The size and geometries of molecules in our experiments are listed in Table 2.

Table 2: Molecular information summary used in the experiments.

System	qubits	electrons	bond data
LiH	12	4	$R = 1.60\,\text{Å}$
H_2O	14	10	$R_{\text{OH}} = 0.96\,\text{Å}$, 104.5°
NH_3	16	10	$R_{\text{NH}} = 1.01\,\text{Å}$
CH_4	18	10	$R_{\text{CH}} = 1.09\,\text{Å}$
N_2	20	14	$R_{\text{NN}} = 1.10\,\text{Å}$
Li_2O	30	14	$R_{\text{LiO}} = 1.00\,\text{Å}$, 120°

CH_4 has tetrahedral geometry, and NH_3 has a pyramidal geometry with nitrogen at the origin and hydrogen positions $(a, 0, z)$ and $(-a/2, \pm\sqrt{3}a/2, z)$, where $a = 0.8118954292\,\text{Å}$ and $z = 0.6007710146\,\text{Å}$. Chemical accuracy is a conventional target for computational chemistry: predicting an energy or energy difference with an error of no more than $1\,\text{kcal/mol}$, approximately 1.6×10^{-3} Hartree (Ha). For ground energy estimation, this means $|E_{\text{estimated}} - E_{\text{reference}}| \leq 1.6 \times 10^{-3}\,\text{Ha}$.

D.2 Many-body System Hamiltonian

For an open chain of n sites, the Hamiltonians of the transverse-field Ising and quantum Heisenberg models are presented as

$$H_{\text{TFIM}} = -J\sum_{i=1}^{n-1} \boldsymbol{Z}_i\boldsymbol{Z}_{i+1} - h\sum_{i=1}^{n} \boldsymbol{X}_i, \tag{29}$$

$$H_{\text{QHM}} = J\sum_{i=1}^{n-1} (\boldsymbol{X}_i\boldsymbol{X}_{i+1} + \boldsymbol{Y}_i\boldsymbol{Y}_{i+1} + \boldsymbol{Z}_i\boldsymbol{Z}_{i+1}). \tag{30}$$

In the experiments evaluating the performance of NQOF, for QHM, we select $J = 1$ without an external magnetic field, and select the initial state as $\bigotimes_{j=1}^{n/2}\left[(|01\rangle - |10\rangle)/\sqrt{2}\right]$. For the 20-qubit TFIM, we use $J = h = 1$ and $|+\rangle^{\otimes n}$ as the initial state. In the 8-qubit TFIM sensitivity study, we select $J = 1$ and $h = 0.9$. Besides, we obtain the reference ground energies of TFIM and QHM with different settings by DMRG with bond dimension $\chi = 128$ (White, 1993).

D.3 hyperparameters

Common hyperparameter settings across models are autoregressive transformer hidden dimension 64, head 4, FFN dimension 128, output dimension 4, and Adam optimizer with learning rate 10^{-4}, $(\beta_1, \beta_2) = (0.9, 0.999)$, $\epsilon = 10^{-8}$, zero weight decay, and gradient-norm clipping at 1.

E Benchmarks

Full Configuration Interaction (FCI). Full Configuration Interaction (FCI) provides the exact solution to the electronic Schrödinger equation within a fixed one-particle basis by diagonalizing the molecular Hamiltonian in the complete many-electron configuration space (Knowles & Handy, 1984).

Table 3: Hyperparameters in experiments

System	n	K	Cutoff	Steps	Adam updates flow step $t = 0 \mid t > 0$
Molecular energy and size comparisons					
LiH	12	64	10^{-7}	92	1,024 / 256
H_2O	14	32	10^{-7}	92	1,024 / 256
NH_3	16	64	10^{-8}	384	1,024 / 256
CH_4	18	128	10^{-8}	384	1,024 / 256
N_2	20	128	10^{-7}	94	1,024 / 256
Lattice energy comparison					
QHM	20	128	10^{-7}	64	1,024 / 256
TFIM	20	128	10^{-7}	256	1,024 / 256
TFIM size comparison					
TFIM	$10, 20, 30$	128	10^{-7}	256	1,024 / 256
Molecular convergence, and potential-energy curve					
LiH	12	16	10^{-8}	100	1024 / 256
N_2	20	16	10^{-7}	400	1024 / 256
Li_2O	30	128	10^{-7}	800	1,024 / 256
Pauli-term growth					
LiH	12	64	0	24	1,024 / 256
N_2	20	16	10^{-8}	256	1,024 / 256
Li_2O	30	256	10^{-7}	448	1,024 / 256
Sensitivity to K: $K \in \{8, 16, 32, 64\}$					
LiH	12	8–64	10^{-7}	24	1,024 / 256
QHM	8	8–64	10^{-7}	32	1,024 / 256
TFIM	8	8–64	10^{-7}	64	1,024 / 256
Density-matrix flow validation					
TFIM	4	8	10^{-7}	80	1,024 / 256

Consequently, FCI serves as the reference method for evaluating the accuracy of approximate ground state energy estimators. However, the dimension of the configuration space grows combinatorially with the number of orbitals and electrons, resulting in exponential computational and memory costs that restrict FCI to relatively small molecular systems.

Hartree-Fock (HF). Hartree-Fock method is a mean-field method that approximates the many-electron wavefunction by a single Slater determinant and variationally optimizes its constituent molecular orbitals (Roothaan, 1951). By treating electron–electron interactions through an effective self-consistent field, HF provides a computationally efficient approximation to the ground state energy. However, because a single-determinant representation neglects electron correlation beyond the mean-field description, HF is primarily used as a reference state and as an initialization for more accurate correlated methods.

Coupled Cluster with Singles and Doubles (CCSD). Coupled Cluster with Singles and Doubles (CCSD) improves upon the Hartree-Fock reference by parameterizing the correlated wavefunction as

$$|\Psi_{\mathrm{CCSD}}\rangle = e^{\boldsymbol{T}_1+\boldsymbol{T}_2}|\Psi_{\mathrm{HF}}\rangle, \tag{31}$$

where $\boldsymbol{T}_1$ and $\boldsymbol{T}_2$ generate single and double electronic excitations, respectively (Purvis & Bartlett, 1982). Through its exponential ansatz, CCSD incorporates important higher-order correlation effects while remaining substantially more tractable than FCI, making it a standard high-accuracy method in quantum chemistry. We therefore use CCSD as a representative structured approximation for comparison with NQOF on molecular systems.

Neural Quantum States (NQS). Neural Quantum States (NQS) parameterize the amplitudes and phases of a many-body wavefunction using a neural network and optimize its parameters variationally by minimizing the energy expectation value (Carleo & Troyer, 2017). Their flexible representations can capture complex correlations in both quantum many-body and molecular systems without explicitly storing the full wavefunction (Choo et al., 2020; Barrett et al., 2022; Shang et al., 2025). As a learning-based, state-centered approach, NQS provides a particularly relevant baseline for NQOF, enabling a direct comparison between neural optimization in the state space and the operator-centered learning strategy adopted by NQOF.

E.1 SETTINGS OF THE REFERENCE METHODS

Settings of NQS for molecular and many-body systems. All molecular Hamiltonians are constructed using the STO-3G basis set. The molecular NQS baseline adopts the autoregressive ansatz introduced by Barrett et al. (Barrett et al., 2022), with a hidden width of 64 for the amplitude network and two hidden layers of width 512 for the single phase network. The model is trained for $10{,}000$ updates using Adam with $(\beta_1, \beta_2) = (0.9, 0.99)$ and random seed 111. The learning rate is set to 10^{-3} for the first $5{,}000$ updates and reduced to 5×10^{-4} for the remaining updates. Adaptive count sampling begins with 10^7 nominal draws and increases up to a maximum of 10^{12}, targeting 10^4-10^5 unique configurations.

For the many-body systems, the NQS baseline employs a four-block Vision Transformer (ViT) (Dosovitskiy et al., 2021) with a hidden width of 60, ten attention heads, a patch size of 2, and a feed-forward width of 240. Each optimization update uses $4{,}096$ samples. The model is trained using stochastic gradient descent with a learning rate of 0.0075 and $\mu = 10^{-4}$ for the diagonal shift used in stochastic reconfiguration. For the QHM, sampling is restricted to the zero-magnetization sector using exchange moves, and we report the state obtained after 350 updates. For the TFIM, training consists of 399 updates using local Metropolis sampling, followed by 401 updates using exact sampling.

DBQA settings. The DBQA baseline performs three reduced group-commutator updates using $\boldsymbol{D} = \sum_i \boldsymbol{Z}_i$. At each update, it searches over seven equally spaced step sizes in the interval $[0.0005, 0.03]$ and implements the Hamiltonian evolution using a symmetric second-order product formula. For the molecular systems, DBQA is initialized from the Hartree-Fock state. For the lattice systems, it is initialized using a six-layer Hamiltonian variational ansatz (HVA). Specifically, the QHM experiments employ the archived XXZ circuit, whereas the TFIM experiments use alternating ZZ and X blocks initialized from a GHZ state. The TFIM circuit contains 12 parameters, each initialized to 0.25, and is optimized through five basin-hopping iterations with COBYLA. Each COBYLA run is limited to 700 iterations with a convergence tolerance of 10^{-4}.

F THEORETICAL ANALYSIS

In this section, we provide the related proof of Theorem 1, Proposition 1 and truncation error analysis.

F.1 PROOF OF THEOREM 1

Proof. The isospectrality property means that all the eigenvalues of $\boldsymbol{H}_t$ during the evolution driven by the Brockett-Wegner flow (Brockett, 1991) remain unchanged. Here, we define $\boldsymbol{G}_t = -[\rho, \boldsymbol{H}_t]$. As the density matrix ρ and $\boldsymbol{H}_t$ are both Hermitian, then we have $\boldsymbol{G}_t^\dagger = -\boldsymbol{G}_t$ and the Brockett-Wegner flow can be rewritten as

$$\frac{d\boldsymbol{H}_t}{dt} = -[[\rho, \boldsymbol{H}_t], \boldsymbol{H}_t] = [\boldsymbol{G}_t, \boldsymbol{H}_t]. \tag{32}$$

Let $\boldsymbol{U}_t$ solve $\frac{d}{dt}\boldsymbol{U}_t = -\boldsymbol{U}_t\boldsymbol{G}_t$ and $\boldsymbol{U}_0 = \boldsymbol{I}$, then we have

$$\frac{d}{dt}\boldsymbol{U}_t\boldsymbol{U}_t^\dagger = -\boldsymbol{U}_t(\boldsymbol{G}_t^\dagger + \boldsymbol{G}_t)\boldsymbol{U}_t^\dagger = 0. \tag{33}$$

Hence, $\boldsymbol{U}_t\boldsymbol{U}_t^\dagger = \boldsymbol{I}$, it implies $\boldsymbol{U}_t$ is unitary. Now, we define $\boldsymbol{H}'_t = \boldsymbol{U}_t^\dagger \boldsymbol{H}\boldsymbol{U}_t$ and have

$$\frac{d\boldsymbol{H}'_t}{dt} = \boldsymbol{G}_t\boldsymbol{U}_t^\dagger\boldsymbol{H}\boldsymbol{U}_t - \boldsymbol{U}_t^\dagger\boldsymbol{H}\boldsymbol{U}_t\boldsymbol{G}_t \tag{34}$$

$$= [\boldsymbol{G}_t, \boldsymbol{H}'_t]. \tag{35}$$

With $\boldsymbol{G}_t$ regarded as a prescribed continuous matrix-valued function of time, both $\boldsymbol{H}_t$ and $\boldsymbol{H}'_t$ solve the same linear initial-value problem

$$\dot{\boldsymbol{X}}_t = [\boldsymbol{G}_t, \boldsymbol{X}_t], \qquad \boldsymbol{X}_0 = \boldsymbol{H}. \tag{36}$$

Uniqueness therefore yields $\boldsymbol{H}_t = \boldsymbol{U}_t^\dagger\boldsymbol{H}\boldsymbol{U}_t$. Thus $\boldsymbol{H}_t$ remains on $\mathcal{M}(\boldsymbol{H})$ and has the same spectrum as $\boldsymbol{H}$.

For the monotonic decrease of the energy, we consider the energy function $E(t) = \langle \boldsymbol{v}|\boldsymbol{H}_t|\boldsymbol{v}\rangle = \mathrm{Tr}(\rho \boldsymbol{H}_t)$, where $\rho = |\boldsymbol{v}\rangle\langle\boldsymbol{v}|$. By differentiating it, we have

$$\frac{d}{dt}E(t) = \mathrm{Tr}\left(\rho \frac{d\boldsymbol{H}_t}{dt}\right) = -\mathrm{Tr}\left(\rho\left[[\rho, \boldsymbol{H}_t], \boldsymbol{H}_t\right]\right). \tag{37}$$

By the cyclic property of the trace, we obtain

$$\frac{d}{dt}E(t) = -\mathrm{Tr}\left([\rho, \boldsymbol{H}_t][\boldsymbol{H}_t, \rho]\right) = \mathrm{Tr}\left([\rho, \boldsymbol{H}_t]^2\right). \tag{38}$$

Since $[\rho, \boldsymbol{H}_t]$ is anti-Hermitian, we have

$$\mathrm{Tr}\left([\rho, \boldsymbol{H}_t]^2\right) = \mathrm{Tr}\left([\rho, \boldsymbol{H}_t]^\dagger[\boldsymbol{H}_t, \rho]\right) = -\|[\rho, \boldsymbol{H}_t]\|_F^2 \tag{39}$$

Consequently, the energy along the Brockett-Wegner flow decreases monotonically.

$$\frac{d}{dt}E(t) = -\|[\rho, \boldsymbol{H}_t]\|_F^2 \leq 0. \tag{40}$$

□

F.2 Proof of Proposition 1

Proof. Let $\boldsymbol{H}$ be a finite-dimensional Hermitian operator with distinct eigenvalues $E_0 < E_1 < \cdots < E_s$ and corresponding orthogonal spectral projectors $\boldsymbol{\Pi}_0, \boldsymbol{\Pi}_1, \ldots, \boldsymbol{\Pi}_s$. Thus $\boldsymbol{H} = \sum_{j=0}^{s} E_j \Pi_j$. We assume $s \geq 1$, a normalized reference state $|\boldsymbol{v}\rangle$, nonzero ground-space overlap $p_0 = \langle v|\Pi_0|v\rangle > 0$, and a positive gap $\Delta = E_1 - E_0 > 0$. Let $\boldsymbol{U}_t$ solve $\dot{\boldsymbol{U}}_t = \boldsymbol{U}_t[\rho, \boldsymbol{H}_t], \boldsymbol{U}_0 = \boldsymbol{I}$. Theorem 1 gives $\boldsymbol{H}_t = \boldsymbol{U}_t^\dagger \boldsymbol{H} \boldsymbol{U}_t$ and unitarity of $\boldsymbol{U}_t$. Set $|\boldsymbol{u}_t\rangle = \boldsymbol{U}_t|\boldsymbol{v}\rangle$ and $E_t = \langle \boldsymbol{u}_t|\boldsymbol{H}|\boldsymbol{u}_t\rangle = \langle \boldsymbol{v}|\boldsymbol{H}_t|\boldsymbol{v}\rangle$, we have

$$\frac{d|\boldsymbol{u}_t\rangle}{dt} = \boldsymbol{U}_t[\rho, \boldsymbol{H}_t]|\boldsymbol{v}\rangle = -(\boldsymbol{H} - E_t\boldsymbol{I})|\boldsymbol{u}_t\rangle. \tag{41}$$

This normalized imaginary-time equation has solution (Gluza et al., 2026)

$$|\boldsymbol{u}_t\rangle = \frac{e^{-t\boldsymbol{H}}|\boldsymbol{v}\rangle}{\|e^{-t\boldsymbol{H}}|\boldsymbol{v}\rangle\|}. \tag{42}$$

Write $p_j = \langle \boldsymbol{v}|\boldsymbol{\Pi}_j|\boldsymbol{v}\rangle$. The spectral decomposition gives

$$\epsilon_t = E_t - E_0 = \frac{\sum_{j=1}^{s} p_j(E_j - E_0)e^{-2(E_j - E_0)t}}{p_0 + \sum_{j=1}^{s} p_j e^{-2(E_j - E_0)t}}. \tag{43}$$

The denominator is at least p_0, and each excited energy satisfies $E_j - E_0 \geq \Delta$. Therefore,

$$\epsilon_t \leq (E_s - E_0)\frac{1 - p_0}{p_0}e^{-2t\Delta}. \tag{44}$$

Multiplying the exact expression by $e^{2t\Delta}$ and taking $t \to \infty$ yields $\Delta p_1/p_0$ when $p_1 > 0$. If $p_1 = 0$, the leading exponent is determined by the lowest excited eigenvalue with nonzero initial weight. If $p_0 = 1$, the energy error is identically zero. □

F.3 Proof of Pauli Coefficient Truncation Bound

Pauli coefficient truncation. To control operator proliferation, we discard Pauli terms whose coefficient magnitudes fall below a threshold. The accumulated error is controlled by the total absolute coefficient weight discarded during the evolution. We analyze the individual Pauli rotations in the implemented Trotter product. Let N_{T} denote the Trotter repetition number, distinct from the sampling batch size. For M outer iterations with K generators each, there are $R = MKN_{\mathrm{T}}$ rotations and truncation events. We index them by $t = 0, \ldots, R-1$ and write

$$\boldsymbol{U}_t = e^{-i\theta_t \boldsymbol{P}_t}, \qquad \theta_t = \frac{\Delta t}{N_{\mathrm{T}}}\lambda_{m,a}$$

for the corresponding generator at outer iteration m. Hamiltonians in this subsection are indexed by individual rotations, so index R corresponds to M outer iterations in Theorem 2.

At finite N_{T}, the Trotter product generally approximates the exponential update in the main text. The comparison below uses the implemented product itself; it isolates coefficient truncation error and does not include the Trotter error relative to the exact exponential. Spectrum preservation refers to the untruncated unitary propagation.

Theorem 3 (Accumulated coefficient truncation error)**.** *Let*

$$\boldsymbol{H}_0 = \sum_{\boldsymbol{Q}\in\mathcal{P}_n} \alpha_{0,\boldsymbol{Q}}\boldsymbol{Q}, \qquad S_0 = \sum_{\boldsymbol{Q}\in\mathcal{P}_n} |\alpha_{0,\boldsymbol{Q}}|^2,$$

where $\mathcal{P}_n = \{\boldsymbol{I},\boldsymbol{X},\boldsymbol{Y},\boldsymbol{Z}\}^{\otimes n}$ is the Hermitian Pauli basis. Consider

$$\boldsymbol{H}_{t+1} = \boldsymbol{U}_t^\dagger \boldsymbol{H}_t \boldsymbol{U}_t, \qquad t = 0,\dots,R-1.$$

Starting from $\widehat{\boldsymbol{H}}_0 = \boldsymbol{H}_0$, define

$$\boldsymbol{A}_t = \boldsymbol{U}_t^\dagger \widehat{\boldsymbol{H}}_t \boldsymbol{U}_t = \sum_{\boldsymbol{Q}\in\mathcal{P}_n} \beta_{t,\boldsymbol{Q}}\boldsymbol{Q},$$

and apply coefficient truncation with threshold $\tau_t > 0$:

$$\widehat{\boldsymbol{H}}_{t+1} = \sum_{\boldsymbol{Q}:|\beta_{t,\boldsymbol{Q}}|\ge\tau_t} \beta_{t,\boldsymbol{Q}}\boldsymbol{Q}.$$

Both propagations use the same realized rotations, including the generators and coefficients chosen from the truncated run. The comparison therefore does not involve rerunning the optimization without truncation.

Define

$$\mathcal{D}_t = \{\boldsymbol{Q} : |\beta_{t,\boldsymbol{Q}}| < \tau_t\}, \qquad \eta_t = \sum_{\boldsymbol{Q}\in\mathcal{D}_t} |\beta_{t,\boldsymbol{Q}}|.$$

Then

$$|E_0(\widehat{\boldsymbol{H}}_R) - E_0(\boldsymbol{H}_0)| \le \|\widehat{\boldsymbol{H}}_R - \boldsymbol{H}_R\|_{\mathrm{op}} \le \sum_{t=0}^{R-1} \eta_t. \tag{45}$$

Moreover, the number N_{t+1} of nonzero retained Pauli terms satisfies

$$N_{t+1} \le \frac{S_0}{\tau_t^2}. \tag{46}$$

For a normalized reference state $|v\rangle$, define

$$\widehat{E}_R = \langle v|\widehat{\boldsymbol{H}}_R|v\rangle, \qquad E_R^{\mathrm{prod}} = \langle v|\boldsymbol{H}_R|v\rangle.$$

Then

$$|\widehat{E}_R - E_R^{\mathrm{prod}}| \le \sum_{t=0}^{R-1} \eta_t, \qquad E_0(\boldsymbol{H}_0) \le E_R^{\mathrm{prod}} \le \widehat{E}_R + \sum_{t=0}^{R-1} \eta_t. \tag{47}$$

Proof. Let

$$\boldsymbol{D}_t = \boldsymbol{A}_t - \widehat{\boldsymbol{H}}_{t+1} = \sum_{\boldsymbol{Q}\in\mathcal{D}_t} \beta_{t,\boldsymbol{Q}}\boldsymbol{Q}$$

be the discarded operator. Since every Pauli string has operator norm one,

$$\|\boldsymbol{D}_t\|_{\mathrm{op}} \le \sum_{\boldsymbol{Q}\in\mathcal{D}_t} |\beta_{t,\boldsymbol{Q}}| = \eta_t. \tag{48}$$

Write $\mathcal{E}_t = \widehat{\boldsymbol{H}}_t - \boldsymbol{H}_t$. The two propagation rules imply

$$\mathcal{E}_{t+1} = \boldsymbol{U}_t^\dagger \mathcal{E}_t \boldsymbol{U}_t - \boldsymbol{D}_t.$$

By unitary invariance of the operator norm,

$$\|\mathcal{E}_{t+1}\|_{\mathrm{op}} \le \|\mathcal{E}_t\|_{\mathrm{op}} + \eta_t.$$

Since $\mathcal{E}_0 = 0$, iteration yields

$$\|\widehat{\boldsymbol{H}}_R - \boldsymbol{H}_R\|_{\mathrm{op}} \le \sum_{t=0}^{R-1} \eta_t. \tag{49}$$

The variational principle gives

$$|E_0(\widehat{\boldsymbol{H}}_R) - E_0(\boldsymbol{H}_R)| \le \|\widehat{\boldsymbol{H}}_R - \boldsymbol{H}_R\|_{\mathrm{op}}.$$

Because the untruncated propagation preserves the spectrum, $E_0(\boldsymbol{H}_R) = E_0(\boldsymbol{H}_0)$. This proves equation 45.

To establish the sparsity bound, write

$$\widehat{\boldsymbol{H}}_t = \sum_{\boldsymbol{Q}} \alpha_{t,\boldsymbol{Q}} \boldsymbol{Q}, \qquad S_t = \sum_{\boldsymbol{Q}} |\alpha_{t,\boldsymbol{Q}}|^2.$$

Pauli orthogonality and unitary invariance imply

$$\begin{aligned}\sum_{\boldsymbol{Q}} |\beta_{t,\boldsymbol{Q}}|^2 &= 2^{-n} \operatorname{Tr}(\boldsymbol{A}_t^\dagger \boldsymbol{A}_t) \\ &= 2^{-n} \operatorname{Tr}(\widehat{\boldsymbol{H}}_t^\dagger \widehat{\boldsymbol{H}}_t) = S_t.\end{aligned}$$

Truncation removes a subset of these coefficients, so

$$S_{t+1} = S_t - \sum_{\boldsymbol{Q} \in \mathcal{D}_t} |\beta_{t,\boldsymbol{Q}}|^2 \le S_t. \tag{50}$$

Consequently, $S_{t+1} \le S_0$. Every retained coefficient has magnitude at least τ_t, and hence

$$N_{t+1} \tau_t^2 \le \sum_{\boldsymbol{Q}:|\beta_{t,\boldsymbol{Q}}| \ge \tau_t} |\beta_{t,\boldsymbol{Q}}|^2 = S_{t+1} \le S_0.$$

This proves equation 46.

Finally,

$$|\widehat{E}_R - E_R^{\mathrm{prod}}| \le \|\widehat{\boldsymbol{H}}_R - \boldsymbol{H}_R\|_{\mathrm{op}}.$$

Since $\boldsymbol{H}_R$ is unitarily equivalent to $\boldsymbol{H}_0$, $E_R^{\mathrm{prod}} \ge E_0(\boldsymbol{H}_0)$. This proves equation 47. □

Thus, $\widehat{E}_R + \sum_t \eta_t$ is a certified upper bound on the ground-state energy, even when $\widehat{E}_R$ itself is not variational. The quantity $\sum_t \eta_t$ can be accumulated during propagation and provides an a posteriori certificate without a short-time or locality assumption.

Short-time accuracy of coefficient truncation. Let $w(\boldsymbol{Q})$ be the number of non-identity factors in a Pauli string $\boldsymbol{Q}$, and let

$$J_0 = \sum_{\boldsymbol{Q}} |\alpha_{0,\boldsymbol{Q}}| > 0, \qquad \alpha_{0,\boldsymbol{Q}} = 0 \quad \text{whenever } w(\boldsymbol{Q}) > k.$$

At outer iteration m, write

$$\boldsymbol{K}_m = \sum_{a=1}^{K} \lambda_{m,a} \boldsymbol{P}_{m,a},$$

where $\lambda_{m,a} \in \mathbb{R}$, $w(\boldsymbol{P}_{m,a}) \le q$, $q \ge 2$, and

$$\sup_m \max_x \sum_{a: x \in \operatorname{supp} \boldsymbol{P}_{m,a}} |\lambda_{m,a}| \le g.$$

We use the individual-rotation propagation and truncation defined above. Its angles satisfy

$$\max_x \sum_{t: x \in \operatorname{supp} \boldsymbol{P}_t} |\theta_t| = \max_x \sum_{m=0}^{M-1} \Delta t \sum_{a: x \in \operatorname{supp} \boldsymbol{P}_{m,a}} |\lambda_{m,a}| \le gT, \qquad T = M\Delta t. \tag{51}$$

Here the N_{T} repetitions cancel the factor $1/N_{\mathrm{T}}$ in each rotation angle.

Proposition 2 (Short-time coefficient truncation bound)**.** *Set $r = q - 1$ and $t^\star = 1/(2egr)$. For $T = M\Delta t < t^\star$, define*

$$\mathcal{A}_T = J_0 e^{k/r}, \qquad c_T = \frac{1}{r} - 2egT > 0, \qquad \nu_n = \frac{c_T}{c_T + \log(3n)}.$$

Then

$$\eta_t \le \mathcal{A}_T^{1-\nu_n} (e\tau_t)^{\nu_n}. \tag{52}$$

Furthermore, we have

$$\sum_{t=0}^{R-1} \eta_t \le \mathcal{A}_T^{1-\nu_n} \left(e \sum_{t=0}^{R-1} \tau_t \right)^{\nu_n}. \tag{53}$$

Consequently,

$$|E_0(\widehat{\boldsymbol{H}}_R) - E_0(\boldsymbol{H}_0)| \le \|\widehat{\boldsymbol{H}}_R - \boldsymbol{H}_R\|_{\mathrm{op}} \le \mathcal{A}_T^{1-\nu_n} \left(e \sum_{t=0}^{R-1} \tau_t \right)^{\nu_n}. \tag{54}$$

The same upper bound controls $|\widehat{E}_R - E_R^{\mathrm{prod}}|$.

Proof. Interpolate rotation t as $\boldsymbol{U}_t(u) = e^{-iu\theta_t \boldsymbol{P}_t}$, $0 \le u \le 1$, and use $s = t + u$ as an interpolation parameter. Define the cumulative absolute angle at site x by

$$h_x(t+u) = \sum_{\substack{0 \le j < t \\ x \in \operatorname{supp} \boldsymbol{P}_j}} |\theta_j| + u|\theta_t| \, \mathbf{1}_{\{x \in \operatorname{supp} \boldsymbol{P}_t\}}.$$

By equation 51, $h_x(s) \le gT$. Set

$$\mu_x(s) = \frac{1}{r} - 2e\, h_x(s), \qquad \omega_s(\boldsymbol{Q}) = \exp\left(\sum_{x \in \operatorname{supp} \boldsymbol{Q}} \mu_x(s) \right).$$

Then

$$c_T \le \mu_x(s) \le \frac{1}{r}, \qquad \omega_s(\boldsymbol{Q}) \ge e^{c_T w(\boldsymbol{Q})}.$$

Follow the Pauli coefficients along the actual interpolated trajectory, including every truncation, and define

$$W(s) = \sum_{\boldsymbol{Q}} \omega_s(\boldsymbol{Q}) |\alpha_{\boldsymbol{Q}}(s)|.$$

During rotation t, a nonzero commutator maps $\boldsymbol{Q}$ to a Pauli string $\boldsymbol{Q}'$ proportional to $\boldsymbol{P}_t\boldsymbol{Q}$, with coefficient rate of magnitude $2|\theta_t|$. The supports must overlap, and at most $q - 1 = r$ new sites can be introduced. Since $0 \le \mu_x(s) \le 1/r$,

$$\omega_s(\boldsymbol{Q}') \le e\, \omega_s(\boldsymbol{Q}).$$

Moreover,

$$\frac{d}{ds} \omega_s(\boldsymbol{Q}) = -2e|\theta_t| \, |\operatorname{supp} \boldsymbol{Q} \cap \operatorname{supp} \boldsymbol{P}_t| \, \omega_s(\boldsymbol{Q}).$$

The coefficient equation therefore implies, almost everywhere between truncations,

$$\frac{dW}{ds} \le 2e|\theta_t| \sum_{\boldsymbol{Q}} \omega_s(\boldsymbol{Q}) |\alpha_{\boldsymbol{Q}}(s)| \Big(\mathbf{1}_{\{[\boldsymbol{P}_t, \boldsymbol{Q}] \ne 0\}} - |\operatorname{supp} \boldsymbol{Q} \cap \operatorname{supp} \boldsymbol{P}_t| \Big) \le 0.$$

At truncation t, the decrease in W is exactly

$$\Gamma_t = \sum_{\boldsymbol{Q} \in \mathcal{D}_t} \omega_{t+1}(\boldsymbol{Q}) |\beta_{t,\boldsymbol{Q}}|.$$

Hence

$$\sum_{t=0}^{R-1} \Gamma_t \le W(0) \le J_0 e^{k/r} = \mathcal{A}_T. \tag{55}$$

Write $\nu = \nu_n$ and $b = \log(3n)$, so that $(1-\nu)c_T = \nu b$. Pauli counting gives

$$\sum_{\boldsymbol{Q}\in\mathcal{P}_n} e^{-bw(\boldsymbol{Q})} = (1+3e^{-b})^n = \left(1+\frac{1}{n}\right)^n \le e.$$

Applying Hölder's inequality to the discarded coefficients,

$$\begin{aligned}\eta_t &= \sum_{\boldsymbol{Q}\in\mathcal{D}_t} \left(e^{c_T w(\boldsymbol{Q})}|\beta_{t,\boldsymbol{Q}}|\right)^{1-\nu}\left(e^{-bw(\boldsymbol{Q})}|\beta_{t,\boldsymbol{Q}}|\right)^{\nu} \\ &\le \left(\sum_{\boldsymbol{Q}\in\mathcal{D}_t} e^{c_T w(\boldsymbol{Q})}|\beta_{t,\boldsymbol{Q}}|\right)^{1-\nu}\left(\sum_{\boldsymbol{Q}\in\mathcal{D}_t} e^{-bw(\boldsymbol{Q})}|\beta_{t,\boldsymbol{Q}}|\right)^{\nu} \\ &\le \Gamma_t^{1-\nu}(e\tau_t)^{\nu}.\end{aligned}$$

Equation equation 52 follows from $\Gamma_t \le \mathcal{A}_T$. A second application of Hölder's inequality gives

$$\sum_{t=0}^{R-1}\eta_t \le \left(\sum_{t=0}^{R-1}\Gamma_t\right)^{1-\nu}\left(e\sum_{t=0}^{R-1}\tau_t\right)^{\nu} \le \mathcal{A}_T^{1-\nu}\left(e\sum_{t=0}^{R-1}\tau_t\right)^{\nu},$$

which proves equation 53. The remaining claims follow from Theorem 3. □

For a uniform threshold τ, Theorem 2 follows by setting

$$C_n = \mathcal{A}_T^{1-\nu_n}(eKN_{\mathrm{T}})^{\nu_n},$$

since $R = MKN_{\mathrm{T}}$.

Choice of threshold. Fix M, K, N_{T} and uniform bounds k, q, g satisfying the proposition for the runs under consideration. For a target error $\varepsilon > 0$, the uniform threshold

$$\tau_t = \tau = \frac{\mathcal{A}_T}{eR}\left(\frac{\varepsilon}{\mathcal{A}_T}\right)^{1/\nu_n}$$

ensures

$$\|\widehat{\boldsymbol{H}}_R - \boldsymbol{H}_R\|_{\mathrm{op}} \le \varepsilon.$$

The spectral and expectation-value errors above are therefore also at most ε.

The same weighted coefficient bound controls the number N_{t+1} of retained terms. Writing $\nu = \nu_n$, Hölder's inequality gives

$$N_{t+1} \le \tau^{-(1-\nu)} \sum_{|\beta_{t,\boldsymbol{Q}}|\ge\tau} |\beta_{t,\boldsymbol{Q}}|^{1-\nu} \le e^{\nu}\left(\frac{\mathcal{A}_T}{\tau}\right)^{1-\nu}.$$

With the prescribed threshold,

$$N_{t+1} \le eR^{1-\nu_n}\left(\frac{\mathcal{A}_T}{\varepsilon}\right)^{\log(3n)/c_T}.$$

For fixed k, q, g, T with $T < t^\star$, if J_0, R, and ε^{-1} are polynomially bounded in n, this retained-term bound is $\exp[O((\log n)^2)]$.

Applicability and relation to energy convergence. The bounded-weight and local-strength conditions are assumptions of Proposition 2; they are not automatically enforced by the learned Pauli sampler or the coefficient solve. In particular, fixed k, q and $J_0 = O(n)$ are not assumed for all molecular Hamiltonians. The a posteriori bound in Theorem 3 remains valid without these assumptions or the short-time restriction. Its numerical value must be evaluated from the coefficients discarded in the actual run.

The truncation results compare the same realized Trotter sequence with and without coefficient deletion. They do not establish that this learned sequence follows the ideal Brockett-Wegner trajectory.

Combining them with the ideal-flow convergence result requires control of the learned-direction and Trotter errors, as well as verification that the required evolution lies within the applicable short-time regime.

For completeness, the state-projected loss $\mathcal{L}$ in Appendix B.3 gives the following local descent criterion. For a fixed represented Hamiltonian $\widehat{\boldsymbol{H}}$, Eq. (11) and the Brockett-Wegner variance identity imply

$$\frac{d}{ds}\langle v|e^{is\tilde{\boldsymbol{H}}_K}\widehat{\boldsymbol{H}}e^{-is\tilde{\boldsymbol{H}}_K}|v\rangle\Big|_{s=0} \leq -2\,\mathrm{Var}_v(\widehat{\boldsymbol{H}}) + \sqrt{2\mathcal{L}}.$$

Thus the learned direction decreases the represented energy to first order whenever $\mathcal{L} < 2\,\mathrm{Var}_v(\widehat{\boldsymbol{H}})^2$. This local criterion alone does not imply convergence of the finite-step, truncated optimization to the ground state.

G Additional Experiments

In this section, we present additional numerical experiments to further evaluate the performance and generality of NQOF. Specifically, we examine its timing and convergence performance on representative molecular and many-body quantum systems, and we extend its application to the MaxCut problem to assess its capability for combinatorial optimization. Taken together, these results demonstrate the broad applicability of NQOF on both physical and discrete quantum optimization tasks.

G.1 Timing comparison

We investigate the wall-clock time of NQOF compared with NQS. Here, we select the mask and no-mask neural autoregressive quantum state (NAQS) (Barrett et al., 2022) as the reference method. As shown in Fig. 9, NQOF achieves the fastest initial energy reduction and consistently obtains a lower energy than NAQS. Despite being trained from a fresh initialization, NQOF steadily approaches the FCI reference and ultimately reaches chemical accuracy. In comparison, NAQS without the mask converges more rapidly than its masked counterpart. Both NAQS variants remain above the FCI energy within the reported runtime. These results demonstrate that the compact, operator-centered optimization of NQOF provides greater wall-clock efficiency and more stable convergence than the state-centered NAQS baselines under the same computational setting.

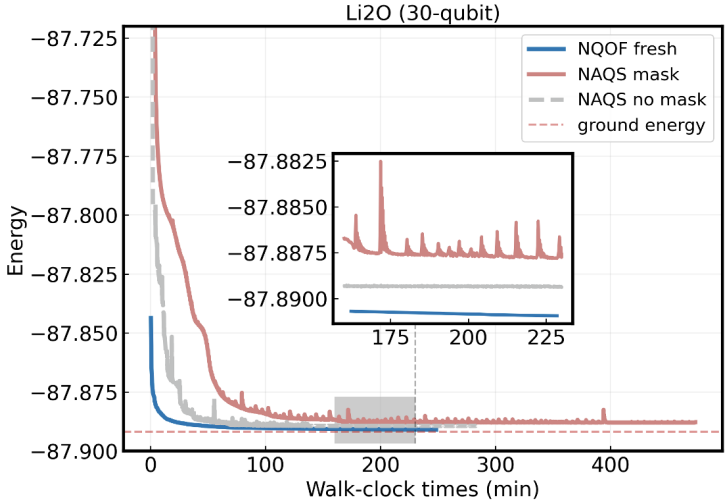


Figure 9: Wall-clock time comparison between NQOF and NAQS (Barrett et al., 2022) for the 30-qubit Li_2O. The horizontal dashed line denotes the ground state energy.

G.2 Performance of NQOF for molecular and many-body quantum systems

Here, we investigate the convergence of NQOF across molecular and many-body quantum systems.

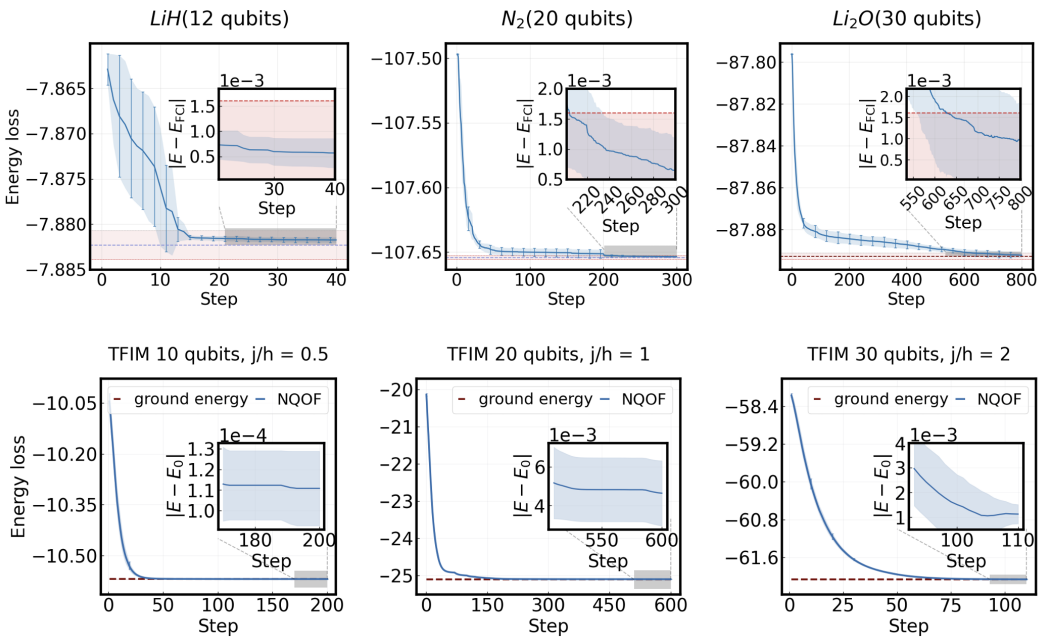


Figure 10: Convergence of NQOF across molecular and many-body quantum systems. The top row shows the energy loss for LiH (12 qubits), N_2 (20 qubits), and Li_2O (30 qubits). The horizontal reference lines indicate the corresponding FCI energies, while the red shaded regions denote the chemical accuracy window of (1.6×10^{-3} (Ha)). The bottom row presents the convergence of NQOF for TFIMs with $10 \sim 30$ qubits and $J/h \in \{0.5, 1.0, 2.0\}$, where the dark-red dashed lines indicate the exact ground energies. Solid blue curves show the mean energies over independent runs, and the error bars or shaded regions represent one standard deviation.

As shown in Fig. 10, NQOF consistently exhibits a rapid initial decrease in energy followed by a slower refinement toward the reference ground energy. For all three molecular systems, the final mean errors fall within chemical accuracy, including the 30-qubit Li_2O system, although larger systems generally require more flow steps to converge. The decreasing variation across independent runs further indicates that the learned operator flow becomes increasingly stable near convergence. A similar trend is observed for the TFIMs: NQOF closely approaches the exact ground energy across different system sizes and coupling regimes. The $J/h = 1$ case converges more slowly and retains a larger residual error than the $J/h = 0.5$ and $J/h = 2$ cases, which is consistent with the increased optimization difficulty near the critical regime. Overall, the results show that NQOF remains accurate and stable as the number of qubits increases, while its convergence rate depends more strongly on the physical complexity of the system than on system size alone.

G.3 Application to Combinational Optimization

We further examine whether NQOF can be extended from ground energy estimation to combinatorial optimization. Here, we consider unweighted MaxCut instances on sparse graphs with 10, 12, and 14 vertices, each constrained to have maximum degree $\deg(v) \leq 3$. For a graph $G = (V, E)$, the MaxCut problem is to partition the vertices V into two disjoint subsets S and its complement S^c and maximize the number of edges that have one endpoint in S and one endpoint in S^c. Generally, the objective function of MaxCut can be written as $C(\mathbf{z}) = \frac{1}{2}\sum_{(i,j)\in E}(1 - z_i z_j)$ where $z_i \in \{\pm 1\}$. It can be equivalently encoded into a diagonal Pauli-Z Hamiltonian $\boldsymbol{H}_c$. Thus, finding the solution of the MaxCut problem is equivalent to preparing the ground state of the encoded Hamiltonian $\boldsymbol{H}_c$. NQOF is then used to learn an operator flow that concentrates the reference state toward the low-energy state of the corresponding Hamiltonian $\boldsymbol{H}_c$, from which candidate bit strings are sampled and decoded as graph bipartitions.

As shown in Fig. 11(a), NQOF reaches the optimal cut with high probability across all tested instances, achieving 0.995, 0.990, and 0.991 for the 10, 12, and 14 qubits cases, respectively. As shown in Fig. 11(b), we track both the probability P_{opt} of sampling an optimal solution and the

corresponding optimization loss during the flow. P_{opt} increases rapidly with the number of flow steps for all tested graph sizes, reaching near-unity values for $n = 10$, 12, and 14. The late-stage inset further confirms that the final sampling probability remains above 0.98, indicating that the optimized distribution is strongly concentrated on globally optimal cuts rather than merely high-quality suboptimal configurations.

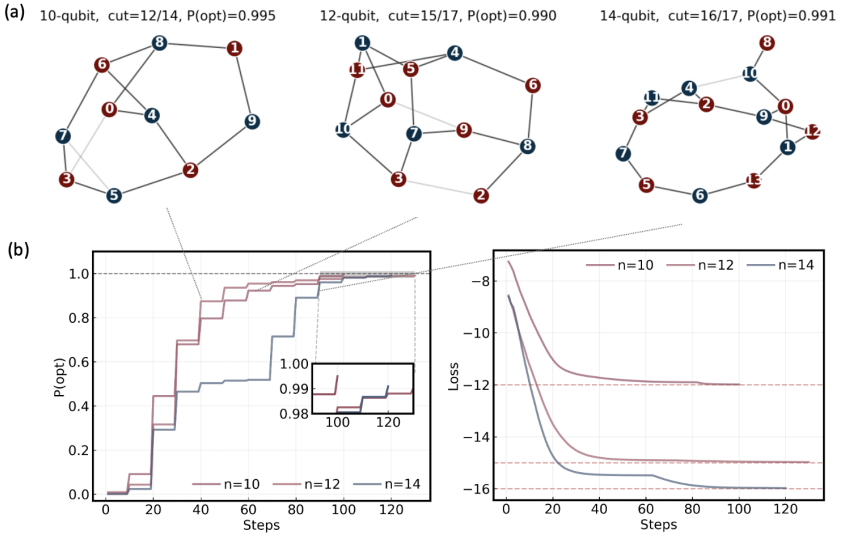


Figure 11: NQOF for MaxCut optimization. (a). Instance graphs on 10, 12, and 14 vertices, encoded as Ising cost Hamiltonians and optimized with NQOF. All graphs have bounded degree $\deg(v) \leq 3$. Red and blue vertices indicate the two partitions S and S^c obtained from the bit string sampled after the NQOF evolution, and edges connecting vertices in different partitions contribute to the cut value. (b). In the left panel, the probability P_{opt} of obtaining an optimal MaxCut solution with NQOF as a function of the flow step for graph sizes $n = 10$, 12, and 14. The corresponding loss curve during the NQOF evolution is shown in the right panel.

The loss curves provide a complementary view of the optimization dynamics as shown in the right panel of Fig. 11(b). For all three graph sizes, NQOF decreases the MaxCut loss quickly during the early stage and then saturates near the exact optimum, indicated by the dashed horizontal reference lines. The convergence behaviour is stable across increasing problem size: although the larger instances require a longer refinement stage, the loss remains well aligned with the optimal value while the probability of sampling an optimal cut continues to improve. Together, these results show that NQOF not only lowers the optimization objective, but also reshapes the output distribution so that optimal bit strings are sampled with high probability. This suggests that the NQOF can also serve as a general optimization primitive for discrete problems, extending its applicability beyond molecular and many-body system ground energy estimation.