Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Gradient-estimator design overcomes trainability barriers in neural-network-based variational optimization

Rui Wang

rwang89@nju.edu.cn

Nanjing University https://orcid.org/0000-0002-3636-1399

Yi-Ran Xue

Nanjing University

Baigeng Wang

Nanjing University https://orcid.org/0009-0006-0931-8527

Chenan Wei

University of Massachusetts Amherst https://orcid.org/0000-0002-1263-7354

Article

Keywords:

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11582001/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

Gradient-estimator design overcomes trainability barriers in neural-network-based variational optimization

Yi-Ran Xue,[1,2] Rui Wang,[1,3,4,5,*] Baigeng Wang,[1,3,4,†] and Chenan Wei[1,2,6,7,‡]

[1] *National Laboratory of Solid State Microstructures and Department of Physics, Nanjing University, Nanjing 210093, China*
[2] *Department of Physics, University of Massachusetts, Amherst, Massachusetts 01003, USA*
[3] *Collaborative Innovation Center of Advanced Microstructures, Nanjing University, Nanjing 210093, China*
[4] *Jiangsu Physical Science Research Center, Nanjing 210093, China*
[5] *Hefei National Laboratory, Hefei 230088, People's Republic of China*
[6] *A. Alikhanyan National Science Laboratory, Br. Alikhanian 2, Yerevan 0036, Armenia*
[7] *Institute of Theoretical Physics, Faculty of Physics, University of Warsaw, Warsaw, Poland*

(Dated: November 3, 2026)

Neural networks provide expressive representations for scientific computing. However, even sufficiently expressive networks can suffer training failure in weak-gradient regimes, limiting their practical use in quantum many-body physics and *ab initio* quantum chemistry. Here we derive an unbiased direct gradient estimator and introduce the adaptive minimum-variance phase (AMVP) estimator for neural-network variational optimization. By improving the signal-to-noise ratio of weak gradients, these methods enable reliable scientific calculations where training previously failed, while substantially reducing computational cost. The framework enables compact networks to outperform larger and fine-tuned default standard-estimator models with over an order of magnitude less GPU time on correlated flux models, and ultimately exceed the density matrix renormalization group (DMRG) accuracy. It further achieves chemical accuracy in N_2 bond breaking and, for the first time, in heavy-element I_2 with explicit spin–orbit coupling. These results demonstrate that gradient-estimator design expands the capabilities of neural-network variational methods for accurate scientific computing.

Introduction. Neural networks provide expressive representations for scientific computing[1]. Across quantum many-body physics and *ab initio* quantum chemistry, neural-network variational methods have shown great potential in representing complex wavefunctions in challenging systems[2–6]. However, greater network expressivity does not by itself ensure successful optimization: training can still fail despite extensive fine-tuning and substantial computational investment, limiting the practical application of these methods[7–9].

Advances in sampling and optimization have helped neural-network variational Monte Carlo (NN-VMC) better exploit the expressive power of wavefunctions. For example, autoregressive models ease sampling bottlenecks through independent Born sampling[10], while minimum-step stochastic reconfiguration (MinSR) enables accurate training of deep, expressive networks[11]. Despite these advances, training can still fail in complex systems when weak descent signals are overwhelmed by gradient noise[12, 13]. In particular, one of the challenging cases is learning the signs and phases that govern quantum interference in strongly correlated and spin–orbit-coupled (SOC) systems[7, 14–16].

In this work, we tackle such training failures in NN-VMC through estimator design. This principle has shown its ability to recover effective optimization signals in machine learning[17–20], suggesting an overlooked route to successful NN-VMC calculations. Specifically, we derive an unbiased direct gradient estimator and further introduce the adaptive minimum-variance phase (AMVP) estimator, which improves the signal-to-noise ratio of weak phase gradients. This framework broadens the practical scope of NN-VMC, enabling reliable and accurate calculations for systems previously inaccessible at significantly lower computational cost.

We first demonstrate these capabilities in synthetic-flux quantum systems motivated by ultracold-atom realizations[21–23]. On the ladder, a compact network reaches subpercent energy error, far outperforming standard-estimator ResNet and multilayer perceptron (MLP) baselines across learning-rate, width and depth sweeps, with an order of magnitude less GPU time. On the two-dimensional torus, AMVP outperforms our density matrix renormalization group (DMRG)[24] benchmark in median variational energy. Furthermore, we overcome intrinsic sign and phase barriers from N_2 bond breaking to SOC[16]. For the heavy-element molecule I_2, we achieve, to our knowledge, the first NN-VMC optimization with explicit SOC: most runs succeed in reaching chemical accuracy, while all standard runs fail.

To explain these gains, we combine theoretical analysis with matched-state numerical diagnostics, linking reduced gradient variance to successful sign and phase learning in correlated matter and relativistic quantum chemistry. More broadly, these findings establish gradient-estimator design as a core component of scientific computation with neural networks.

* Corresponding author: rwang89@nju.edu.cn
† Corresponding author: bgwang@nju.edu.cn
‡ Corresponding author: chenanwei@umass.edu

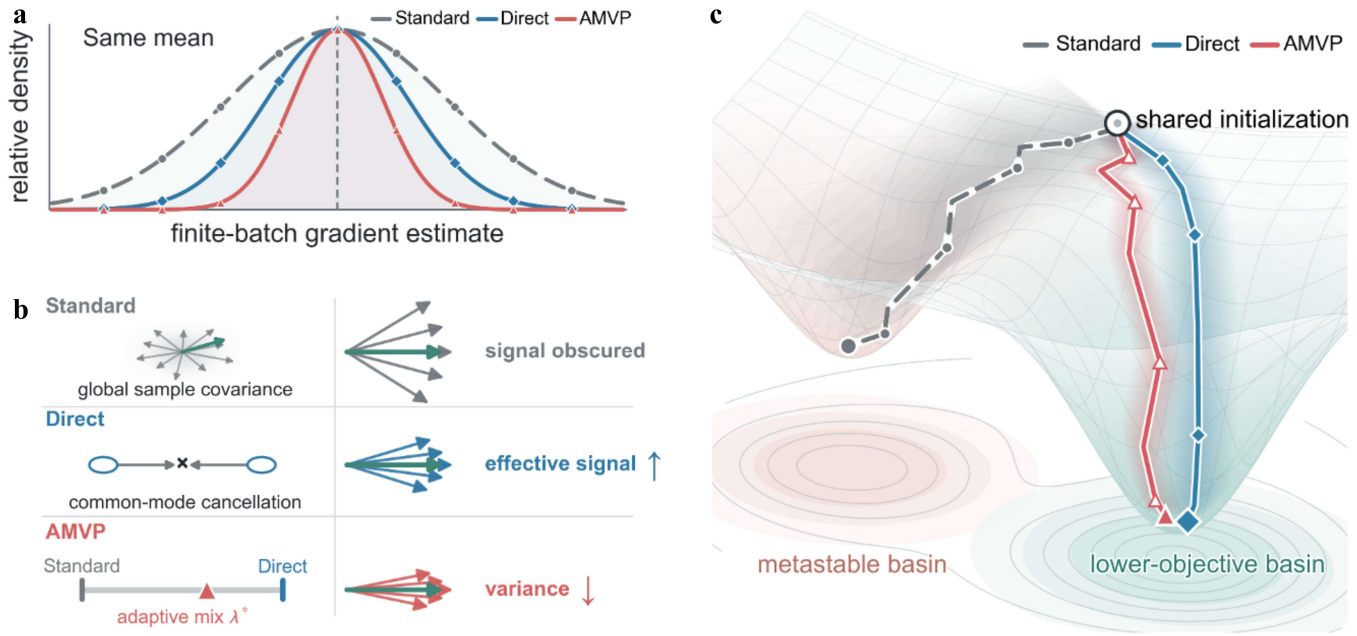


FIG. 1. **Estimator design changes the stochastic optimization dynamics.** (a) The standard, direct, and adaptive minimum-variance phase (AMVP) estimators share the same population gradient but exhibit different finite-sample fluctuations. (b) The standard estimator first aggregates the imaginary local-energy contributions before coupling them to the phase score, whereas the direct estimator retains connection-resolved phase differences that remove common-mode fluctuations. AMVP adaptively combines unbiased estimators according to their stochastic properties. (c) Different estimator noise structures lead to different optimization trajectories. The standard estimator can leave optimization trapped in metastable regimes, whereas the direct and AMVP estimators allow neural states to better exploit their available capacity and approach target states.

Results.

Gradient estimators define distinct stochastic optimization dynamics. At finite sample sizes, statistical fluctuations can obscure weak descent signals and impede optimization. Training dynamics depend on both the population gradient and the finite-sample fluctuations of its estimator. For an unbiased estimator,

$$\hat{g} = \nabla E + \xi, \tag{1}$$

where ∇E denotes the population gradient and ξ represents finite-sample fluctuations. Different estimators can therefore target the same variational gradient while generating different optimization trajectories[25].

This distinction becomes particularly important in neural-network variational methods. Specifically, in neural-network quantum states (NQS), this challenge is especially pronounced in the phase sector, where effective optimization signals can be weak, while estimator fluctuations remain substantial. Here, we compare the standard score-function estimator and a direct phase-gradient estimator while keeping the ansatz, sampling procedure, and optimization protocol unchanged. Although both estimators are unbiased for the same variational force, $\langle \hat{g}_{\mathrm{std}} \rangle = \langle \hat{g}_{\mathrm{dir}} \rangle = \nabla E$, they generate different finite-sample stochastic dynamics.

For $\Psi_\theta(R) = \exp[u_\theta(R) + \mathrm{i}\varphi_\theta(R)]$, a phase-only parameter β satisfies $\partial_\beta u_\theta = 0$, and its conventional score estimator is $\hat{g}_\beta^{\mathrm{std}}(R) = 2\,\mathrm{Im}\,E_{\mathrm{loc}}(R)\,\partial_\beta\varphi_\theta(R)$. When amplitude and phase share parameters, the bare local-energy derivative omits the amplitude contribution. We therefore retain the standard amplitude force and replace only the phase channel by its direct form,

$$\hat{\boldsymbol{g}}_\theta^{\mathrm{dir,cpl}}(R) = 2\,[\mathrm{Re}\,E_{\mathrm{loc}}(R) - E]\,\nabla_\theta u_\theta(R) + \sum_{R'} \mathrm{Im}[W_{RR'}]\,[\nabla_\theta\varphi_\theta(R) - \nabla_\theta\varphi_\theta(R')]\,, \tag{2}$$

where $W_{RR'} = H_{RR'}\Psi_\theta(R')/\Psi_\theta(R)$, and gradients through u_θ are stopped in the direct phase term. This construction preserves the full population gradient while allowing connection-resolved phase-score fluctuations to cancel before summation.

The AMVP dynamically combines the two unbiased phase estimators,

$$\hat{g}_\lambda = (1-\lambda)\hat{g}_{\mathrm{std}} + \lambda\hat{g}_{\mathrm{dir}}, \tag{3}$$

where the mixing coefficient is selected to minimize the measured gradient variance[26, 27] (see Methods and Supplementary Information Sec. 4),

$$\lambda^\star = \arg\min_{\lambda\in[0,1]} \mathrm{Var}(\hat{g}_\lambda). \tag{4}$$

As illustrated in Fig. 1, the AMVP estimator exploits the freedom of unbiased estimator design[28] to select

the stochastic channel with lower gradient fluctuations. Here, Hermiticity yields a direct local-energy derivative estimator of the phase force that shares the population mean of the standard score estimator but has different finite-sample fluctuations. In coupled amplitude–phase networks, parameter sharing further introduces amplitude–phase cross-covariances, so AMVP minimizes the variance of the full stochastic force rather than that of an isolated phase estimator.

Furthermore, by suppressing estimator-induced stochastic fluctuations, AMVP stabilizes optimization, allowing neural states to realize more of their available expressive capacity and more accurately approximate optimal states without increasing architecture complexity or extensive hyperparameter tuning.

Overcoming phase-optimization barriers in quantum flux models. To examine whether NQS with complex phase structures can fully exploit their available variational capacity, we first consider a flux ladder model. The Peierls phases generate chiral currents and introduce non-trivial phase structures into the wavefunction[21, 29, 30], making the system sensitive to the quality of phase-gradient estimation. For the flux ladder benchmark, we use separated amplitude–phase networks, which allow the phase-gradient estimator to be isolated directly. This setting provides a controlled test of whether the estimator choice alone can alter the optimization outcome. The Hamiltonian is given by

$$H_{\text{ladder}} = \sum_{\langle ij\rangle\in\text{legs}} \left[J_z S_i^z S_j^z + \frac{J_{\text{leg}}}{2} \left(e^{\mathrm{i}A_{ij}} S_i^+ S_j^- + \text{h.c.} \right) \right] + \sum_{\langle ij\rangle\in\text{rungs}} \left[J_z S_i^z S_j^z + \frac{J_{\text{rung}}}{2} \left(S_i^+ S_j^- + \text{h.c.} \right) \right], \tag{5}$$

where $A_{ij} = +\Phi/2$ on one leg and $A_{ij} = -\Phi/2$ on the other leg in the symmetric gauge. We use $J_{\text{leg}} = 1$, $J_{\text{rung}} = 0.8$, and $J_z = 0.5$.

For a 50-rung flux ladder at $\Phi = 0.3\pi$, we keep the amplitude network, sampler, optimizer and training protocol fixed while varying only the phase-gradient estimator. Despite extensive learning-rate optimization and increases in network capacity, the standard estimator remains unable to access the low-error regime. Changing only the estimator instead brings the same compact neural-network representation to subpercent error. The direct and AMVP estimators achieve median errors of 0.50% and 0.43%, respectively (Fig. 2a).

The improvement extends beyond the variational energy. Both direct and AMVP recover the DMRG chiral-current profile, with $j_c/j_c^{\text{DMRG}} \approx 1.01$, whereas the standard result remains near 0.79 (Fig. 2b). Estimator choice therefore determines access to the correct phase-sensitive observable, rather than merely lowering the energy.

Notably, increasing neural-network complexity does not remove this barrier: a larger ResNet and wider or deeper MLPs trained with the standard estimator remain trapped in a higher-error regime (Fig. 2a,c). Moreover, the compact AMVP-estimator MLP reaches the low-error state with more than an order of magnitude less GPU time than the standard-estimator ResNet, showing that estimator design can substitute for brute-force increases in model size and computational cost.

To test whether this behavior extends beyond separated amplitude–phase optimization, we further consider an 8×8 flux square with periodic boundary conditions using a coupled amplitude–phase neural-network architecture. The shared parameters introduce amplitude–phase correlations into the full stochastic force, which are incorporated by the AMVP variance minimization. In two dimensions, the AMVP estimator enables a compact NQS ansatz to reach energies below the DMRG results obtained up to $\chi = 2048$, whereas the standard estimator remains far above this low-error regime (Fig. 2d). Furthermore, early in training, most AMVP and direct runs reach sub-percent error, whereas no standard run reaches this threshold at any point during optimization (Fig. 2e).

Beyond flux models, we further verify the robustness of the estimator advantage in quantum systems with different origins of complex phase structures (See Supplementary Information Fig. S1,2), including chiral spin chains[31] and continuum fractional quantum Hall (FQH) systems[32, 33]. The same trend is observed across these different systems, indicating that the optimization limitation is associated with stochastic phase-gradient estimation rather than a particular model class.

Estimator-limited chemical accuracy. Moving beyond flux-driven phases, molecular electronic structure changes the origin of the phase problem. Relative determinant signs arise from correlation-driven interference[34, 35] rather than an externally imposed gauge field; in heavy-element systems, spin–orbit coupling[16, 36] further extends the problem to a broad distribution of complex phases in the working determinant basis. Because this information is embedded in the many-electron state, the associated gradients are less explicit and more easily obscured by finite-sample fluctuations, yet accurate phase learning remains essential for correlation-energy recovery and chemical accuracy [35].

For N_2, we perform joint amplitude–phase optimization in a selected-CI determinant space using a shared MLP,

$$\Psi_\theta(I) = \exp[u_\theta(I) + \mathrm{i}\varphi_\theta(I)], \qquad I \in \mathcal{S}. \tag{6}$$

For spin-free molecular Hamiltonians, phase optimization primarily learns relative determinant signs. With explicit spin–orbit coupling, the I_2 wavefunction instead develops a nontrivial complex phase structure in the determinant basis used by the ansatz: no single global phase renders all coefficients real, and their phases span a broad distribution rather than being restricted to 0 and π.

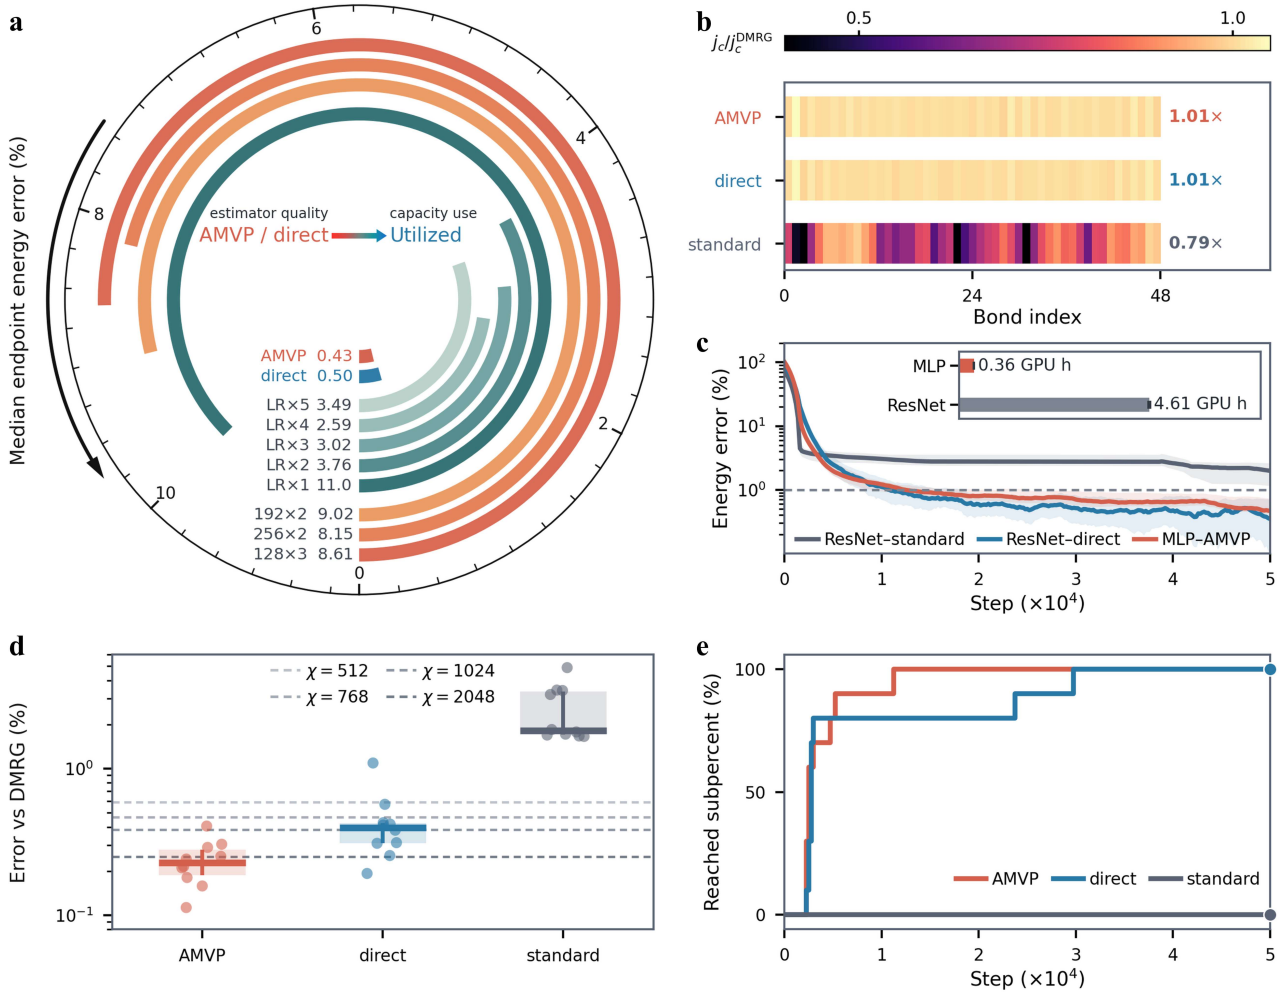


FIG. 2. **Gradient-estimator choice overcomes a phase-optimization barrier.** (a) On a 50-rung flux ladder, the standard estimator remains trapped in a high-error regime despite variations in learning rate and network capacity. In contrast, the direct and adaptive minimum-variance phase (AMVP) estimators allow the same neural-network representation to reach subpercent errors, with median errors of 0.50% and 0.43%, respectively. (b) The chiral current profile obtained with direct and AMVP agrees with the DMRG reference ($j_c/j_c^{\mathrm{DMRG}} \approx 1.01$), whereas the standard estimator shows a marked deviation ($j_c/j_c^{\mathrm{DMRG}} \approx 0.79$). (c) Training trajectories and computational cost. Both a direct-trained ResNet and a compact AMVP-trained MLP reach energy errors below 0.5%, whereas the standard-trained ResNet remains in a higher-error regime. The inset compares allocated A100 GPU times for AMVP-trained MLP and standard-trained ResNet runs: 0.36 GPU h and 4.61 GPU h, respectively. (d) Energy errors for an 8×8 flux square with periodic boundary conditions. Energies obtained with AMVP estimators are lower than the finite-χ DMRG energies obtained up to $\chi = 2048$, whereas the standard estimator remains in a higher-error regime. (e) Cumulative fraction of runs that have reached subpercent energy error.

Across all six N_2 geometries, direct and AMVP place every paired run within chemical accuracy, whereas the standard estimator fails in every run at the first five geometries and reaches only 3/10 at the longest bond length (Fig. 3a). This separation persists as the independently calculated wrong-spin gap collapses from approximately 292 to 1.4 mHa, demonstrating robust access to the target state throughout bond stretching.

The I_2 calculation tests the same principle for a heavy-element molecule with explicit spin–orbit coupling (Fig. 3b,c). With joint amplitude–phase optimization, direct and AMVP reach median full configuration interaction (FCI) errors of 1.13 and 1.19 mHa, respectively, and chemical accuracy in 9/10 paired runs; standard gives 2.39 mHa and 0/10 (Fig. 3b). Under fixed-amplitude phase optimization, at batch size 4096, direct and AMVP reach median FCI errors near 1.3 mHa, compared with 2.4 mHa for standard; at the largest batch, both redesigned estimators reach chemical accuracy in all paired runs (Fig. 3c). The accompanying phase-infidelity reduction shows that the energy gain reflects more accurate complex phase learning.

Furthermore, the C_2 width–batch map separates sampling budget from neural-network capacity (Fig. 3e). At

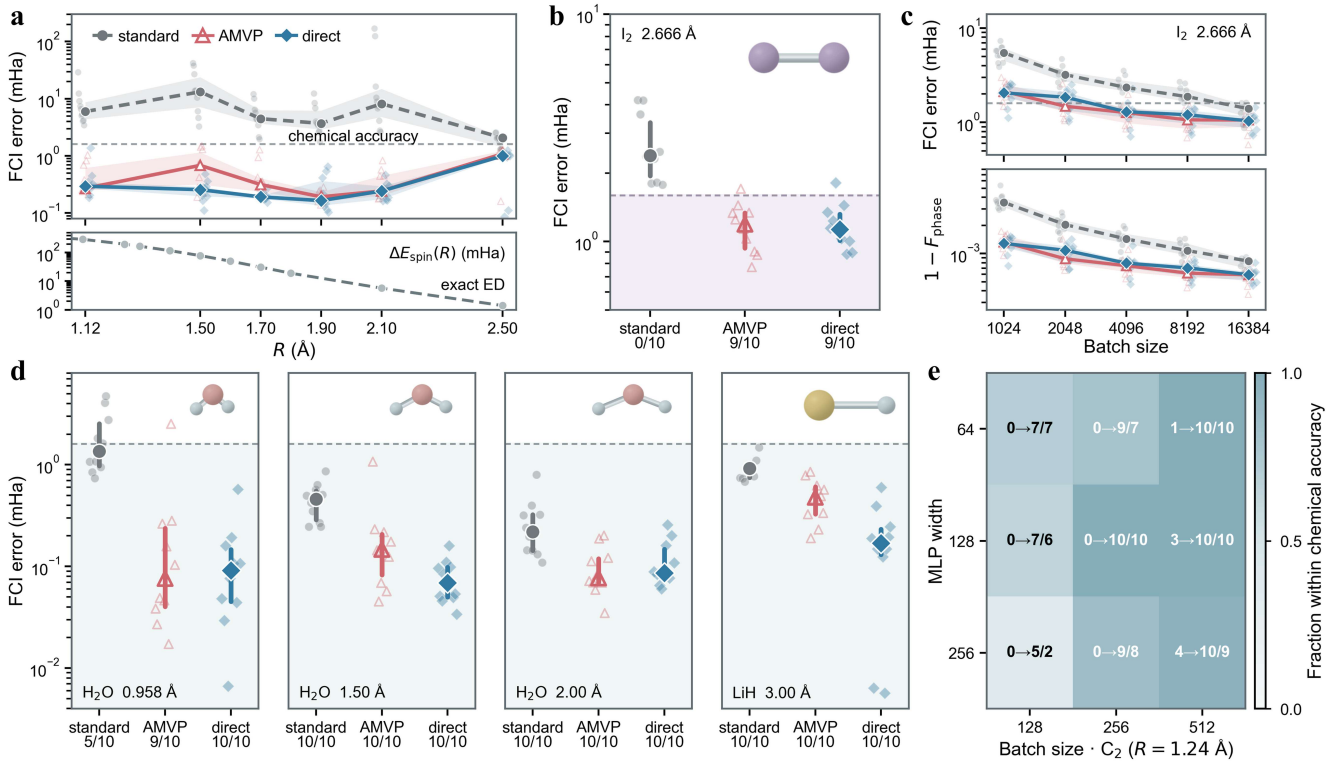


FIG. 3. **Estimator-limited trainability of molecular wavefunctions.** (a) FCI errors along the N_2 bond stretch for the jointly optimized amplitude–phase MLP. The lower strip gives the independently calculated excitation gap from the target-singlet ground state to the first wrong-spin state. (b) FCI errors for I_2 at $R = 2.666$ Å with explicit spin–orbit coupling under joint amplitude–phase optimization. (c) Batch-size dependence of fixed-amplitude phase optimization for I_2 at $R = 2.666$ Å with explicit spin–orbit coupling. The upper and lower panels show the FCI error and phase infidelity $1 - F_{\text{phase}}$, respectively. (d) FCI errors for fixed-amplitude phase-learning controls on H_2O at three geometries and LiH. (e) Fraction of paired runs within chemical accuracy for C_2 at $R = 1.24$ Å across MLP widths and batch sizes. Each cell reports success counts in the order standard → direct/AMVP; the shading gives the mean direct–AMVP success fraction. In (b,d), the left-to-right order is standard, AMVP, direct. Small symbols denote individual runs; large symbols show medians; shaded bands and vertical bars indicate interquartile ranges.

batch size 256, the standard estimator produces no chemically accurate run at any tested width, whereas direct and AMVP are reliable across the map. Increasing width or batch size alone does not systematically close this gap.

In addition, the same behavior extends across chemically distinct molecular systems (Fig. 3d). Direct reaches chemical accuracy for all ten runs in every displayed condition and AMVP does so for nine or ten, while standard is less reliable at the equilibrium H_2O geometry. Even where standard crosses the threshold, the redesigned estimators further lower the median error. The common pattern across molecules shows that estimators with the same expectation value can differ sharply in the precision and reliability of finite-sample training.

Estimator design therefore reduces the reliance on ever-larger sampling budgets and neural networks in *ab initio* quantum chemistry.

Matched-state diagnostics. The accuracy improvements in the lattice and molecular calculations raise a mechanistic question: do the estimators possess different stochastic properties at the same neural-network state, or do their measured statistics merely reflect the fact that training has already carried them to different regions of parameter space? We explain this distinction through matched-state diagnostics.

For an estimator m, we denote its total gradient variance by $V_m = \mathrm{tr}[\mathrm{Cov}(\hat{g}_m)]$. The variance gain and shared-signal batch signal-to-noise ratio (SNR) are

$$\mathcal{R}_V^{(m)} = \frac{V_{\text{std}}}{V_m}, \qquad \mathrm{SNR}_m = \frac{\sqrt{B}\,\|\mu_{\text{common}}\|}{\sqrt{V_m}}, \tag{7}$$

where B is the batch size. The common population signal μ_{common} is estimated symmetrically and debiased for finite-batch noise, so its norm is identical for the standard, direct and AMVP estimators. A variance ratio above unity denotes lower variance than the standard estimator.

At initialization in the 8×8 flux-torus calculation, direct and AMVP reduce the gradient variance by roughly ninefold (Fig. 4a). The relative ordering changes dur-

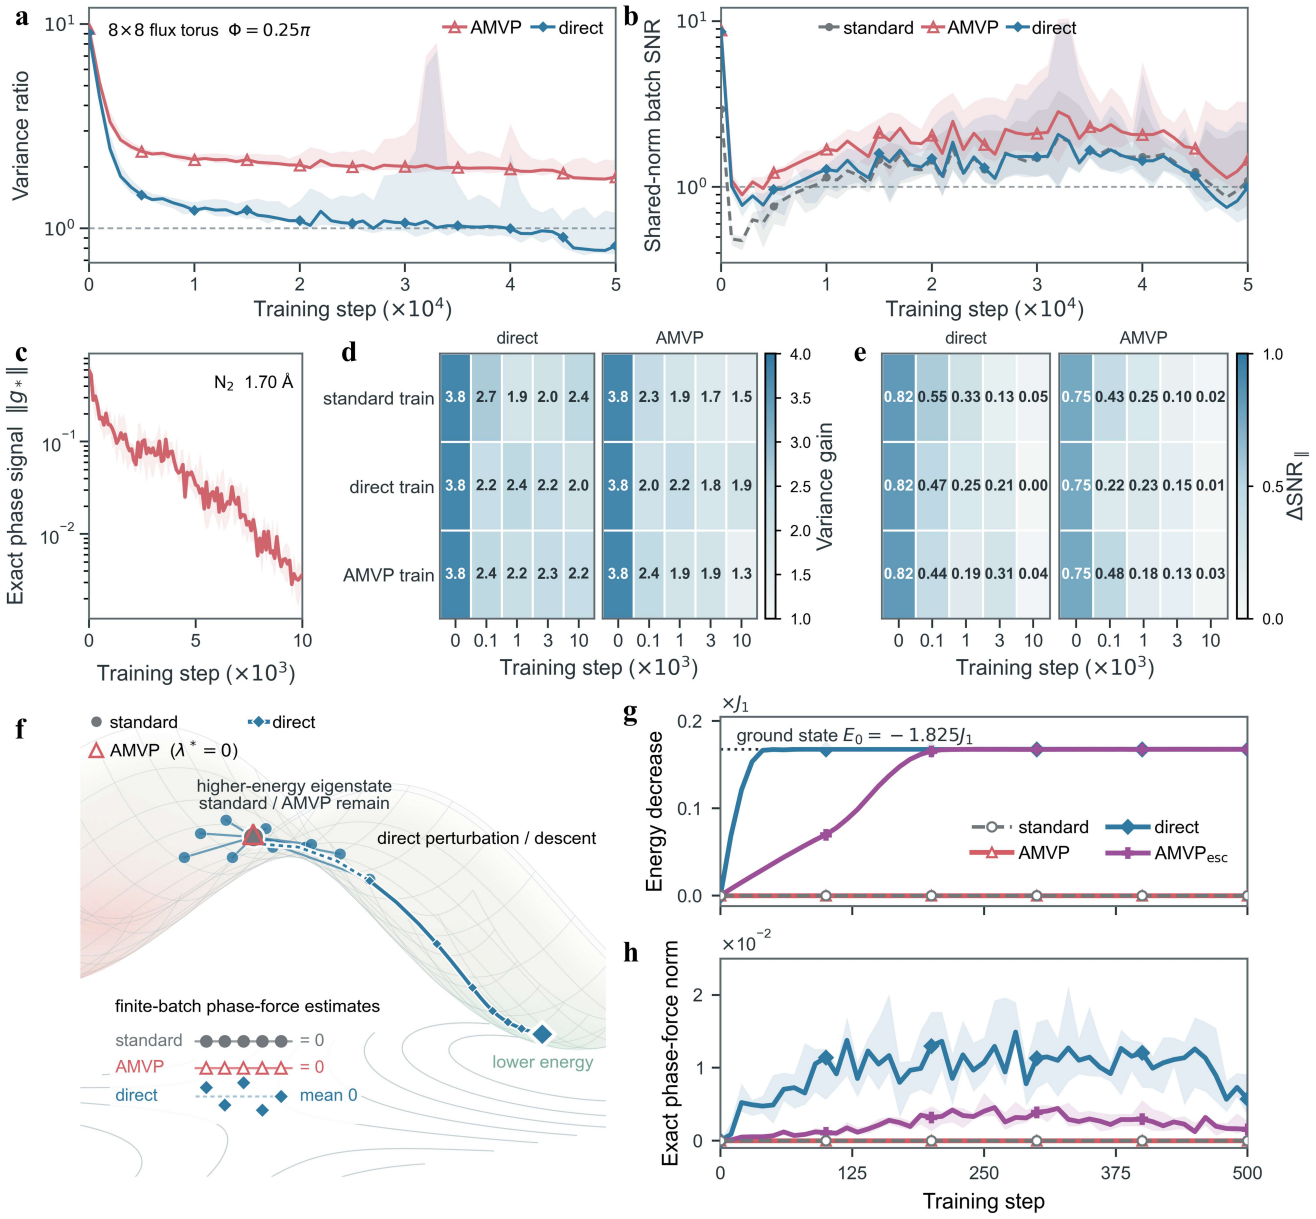


FIG. 4. **Matched-state phase-gradient statistics in quantum many-body physics and molecular calculations.** (a) Variance ratios V_{std}/V_m for the direct and AMVP estimators along the training process of an 8×8 flux torus. (b) Shared-norm batch SNR for the same flux-torus states, using an identical debiased population-signal estimate for all estimators. (c) Exact phase-gradient norm for the N_2 fixed-amplitude phase-channel diagnostic at $R = 1.70$ Å. (d) Rows give the estimator that produced the state, columns the checkpoint in thousands of updates; the gains do not depend on which estimator produced the state, showing that the variance difference is a property of the estimator itself rather than of the states. (e) Corresponding gains in descent-aligned SNR relative to the standard estimator. (f) Exact-eigenstate control for the periodic honeycomb J_1–J_2 model. At the higher-energy eigenstate, the standard score-function phase-force estimate vanishes for every finite batch, and AMVP consequently selects $\lambda^* = 0$. The direct estimator instead has nonzero finite-batch realizations despite having the same zero population mean, allowing a stochastic displacement toward lower energy. (g,h) Evolution from the same excited eigenstate: energy decrease (g) and exact phase-force norm (h). $\mathrm{AMVP}_{\mathrm{esc}}$ denotes the escape-aware extension with a temporary lower bound on the direct mixing weight. Curves and bands in (a,b,c,g,h) show medians and interquartile ranges; heatmap cells in (d,e) show medians.

ing training and the direct estimator eventually becomes slightly noisier than standard, while AMVP retains the lower-variance channel through its adaptive interpolation.

The shared-signal SNR shows the same trend (Fig. 4b): at initialization, it increases from 2.95 for the standard

estimator to 8.83 for AMVP, which maintains the highest SNR as the relative estimator variances evolve during training. Because all estimators use the same signal norm, the SNR differences arise entirely from their finite-sample fluctuations.

The phase sector of molecular models is a progressively weaker-signal regime. In the N_2 matched-state fixed-amplitude phase-channel diagnostic, the exact phase-gradient norm decreases from 0.590 at initialization to 0.00364 (Fig. 4c). To measure whether a stochastic gradient remains effective as this signal collapses, we define the descent-aligned SNR

$$\mathrm{SNR}_{\parallel}^{(m)} = \frac{\overline{g}_m \cdot g_\star}{\|g_\star\|\sqrt{V_m}}, \tag{8}$$

where $g_\star$ is the exact full-support phase gradient and $\overline{g}_m$ is the minibatch-mean gradient. This quantity rewards both low variance and alignment with the exact descent direction.

Cross-evaluation at fixed checkpoints shows consistently lower direct and AMVP variance across states produced by all three training trajectories (Fig. 4d). Similar values across the three rows show that the improvement is not restricted to states generated by direct or AMVP training.

The descent-aligned SNR gains are largest during early and intermediate training and decrease as the exact phase force approaches stationarity (Fig. 4e). Their seed medians remain non-negative across all three trajectory sources. Evaluating all estimators at the same network checkpoints shows that the direct and AMVP estimators have lower variance. The gain therefore arises from estimator-specific statistics rather than from differences among the states reached during training.

The connection between these statistics and optimization comes from the standard smoothness bound [37]. For an unbiased estimator with variance V and an update of size η on an L_{sm}-smooth objective,

$$\mathbb{E}[E_{t+1}] - E_t \le -\eta\|\nabla E\|^2 + \frac{L_{\mathrm{sm}}\eta^2}{2}\left(\|\nabla E\|^2 + V\right). \tag{9}$$

In the noise-dominated regime $V \gg \|\nabla E\|^2$, estimator fluctuations constrain the updates that produce reliable descent. The lower variance therefore raises the descent-aligned SNR at fixed batch size, while leaving the population force and the variational objective unchanged.

Quantum many-body physics and molecular diagnostics reveal the same finite-sample mechanism in two different neural-network state parameterizations. The direct estimator is most effective when weak phase-gradient signals are dominated by standard-estimator noise. AMVP maintains lower variance as the relative ordering changes, adapting between the two unbiased stochastic channels. These diagnostics establish the gradient-quality mechanism underlying the accuracy and reliability improvements in Figs. 2 and 3.

The distinction in Eq. (1) between ξ_{std} and ξ_{dir} becomes most transparent at an exact eigenstate, where the population phase force vanishes and only estimator-specific finite-batch fluctuations remain. We demonstrate this in the periodic honeycomb J_1–J_2 model (Fig. 4f–h). At the excited Bloch eigenstate, $\mathrm{Im}\, E_{\mathrm{loc}} = 0$ pointwise, so every finite-batch standard phase-force estimate vanishes. The variance-minimizing AMVP weight therefore collapses to the standard endpoint, $\lambda^\star = 0$. The direct estimator has the same zero population mean but has nonzero connection-resolved fluctuations on individual batches. A batch realization with a component along a descending direction can perturb the state away from this point. The exact phase force then becomes nonzero and supports continued energy descent. This control reveals a complementary role of estimator fluctuations: variance reduction improves descent in weak-signal regimes, whereas nonzero direct fluctuations can also help optimization leave an excited-state zero-force point.

This analysis motivates an escape-aware AMVP (Fig. 4g,h), in which a temporary lower bound on the direct weight is activated when the standard phase force vanishes but direct finite-batch fluctuations remain. An independent energy evaluation can guide the decision to continue the escape phase or return to variance-minimizing mixing.

Discussion. This work establishes gradient-estimator design as a practical route to more accurate and resource-efficient neural-network variational simulations. The origin of phase information changes across the systems studied. Flux phases are externally induced, whereas determinant signs arise from electronic correlations and spin–orbit coupling produces complex phases in heavy-element molecules. Estimator redesign remains effective throughout this progression, allowing compact neural-network representations to reach low-energy states and chemical accuracy.

Matched-state diagnostics identify the finite-sample origin of these gains. Although the standard and direct estimators share the same population force, estimator redesign lowers gradient variance and raises the signal-to-noise ratio of weak phase directions. The exact-eigenstate analysis reveals a further effect: alternative unbiased estimators can possess different finite-batch zero sets, changing whether stochastic optimization can leave an excited-state zero-force point. Gradient-estimator design therefore shapes both the precision of descent and the stochastic directions available to optimization.

Estimator performance changes during optimization: the direct form is advantageous in weak-signal regimes, whereas the standard estimator becomes zero-variance at exact eigenstates, with AMVP adapting to this crossover. Their distinct finite-batch zero sets further motivate escape-aware extensions beyond variance minimization alone.

More broadly, these results elevate gradient estimation from a numerical implementation choice to a component of neural-network variational design. The same perspective may apply wherever the same variational gradient admits alternative finite-sample representations with distinct stochastic properties[25, 38]. Non-Hermitian variational problems provide a further direction, where the Hermiticity identity underlying the standard form no longer holds and finite-sample design can affect both the construction and the efficiency of the variational force.

Methods. Architectural and optimization details are provided in Supplementary Information Sec. 1.

Variational objective and phase-force estimators. For a basis configuration R, the Born probability, local energy and variational energy are [39, 40]

$$\begin{aligned} p_\theta(R) &= \frac{|\Psi_\theta(R)|^2}{\langle\Psi_\theta|\Psi_\theta\rangle}, \\ E_{\mathrm{loc}}(R) &= \sum_{R'} H_{RR'} \frac{\Psi_\theta(R')}{\Psi_\theta(R)}, \\ E(\theta) &= \sum_R p_\theta(R)\,\mathrm{Re}\,E_{\mathrm{loc}}(R). \end{aligned} \tag{10}$$

We write $\Psi_\theta(R) = \exp[u_\theta(R) + \mathrm{i}\varphi_\theta(R)]$ and use $E \equiv E(\theta)$. For a phase-only parameter β, one configuration contributes the amplitude row, the standard phase row and the direct phase row

$$\begin{aligned} \hat{\boldsymbol{g}}_{\mathrm{amp}}(R) &= 2[\mathrm{Re}\,E_{\mathrm{loc}}(R) - E]\nabla_\theta u_\theta(R), \\ \hat{g}^{\mathrm{std}}_\beta(R) &= 2\,\mathrm{Im}\,E_{\mathrm{loc}}(R)\,\partial_\beta\varphi_\theta(R), \\ \hat{g}^{\mathrm{dir}}_\beta(R) &= \mathrm{Re}\,\partial_\beta E_{\mathrm{loc}}(R) \\ &= \sum_{R'} \mathrm{Im}[W_{RR'}][\partial_\beta\varphi_\theta(R) - \partial_\beta\varphi_\theta(R')]. \end{aligned} \tag{11}$$

where $W_{RR'} = H_{RR'}\Psi_\theta(R')/\Psi_\theta(R)$ and gradients through u_θ are stopped in the direct phase term. Hermiticity gives

$$\langle\hat{g}^{\mathrm{dir}}_\beta\rangle_{p\theta} = 2\langle\mathrm{Im}\,E_{\mathrm{loc}}\,\partial_\beta\varphi_\theta\rangle_{p\theta} = \langle\hat{g}^{\mathrm{std}}_\beta\rangle_{p\theta}, \tag{12}$$

under the usual differentiability and support conditions. Thus the raw estimators have the same population force but different configuration-level fluctuations. For a shared real two-head network, the amplitude row in Eq. (11) is retained and only the phase row is replaced, yielding Eq. (2).

Adaptive minimum-variance phase estimation. Let $\hat{\boldsymbol{g}}^{\mathrm{ph}}_{\mathrm{std}}$ and $\hat{\boldsymbol{g}}^{\mathrm{ph}}_{\mathrm{dir}}$ denote the two phase rows. On the calibration sample we evaluate $V_s = \mathrm{tr}\,\mathrm{Cov}(\hat{\boldsymbol{g}}^{\mathrm{ph}}_{\mathrm{std}})$, $V_d = \mathrm{tr}\,\mathrm{Cov}(\hat{\boldsymbol{g}}^{\mathrm{ph}}_{\mathrm{dir}})$, their cross-covariance trace C, and $A_m = \mathrm{tr}\,\mathrm{Cov}(\hat{\boldsymbol{g}}_{\mathrm{amp}}, \hat{\boldsymbol{g}}^{\mathrm{ph}}_m)$. Minimizing the trace covariance of the complete force gives

$$\begin{aligned} \lambda^\star &= \mathrm{clip}_{[0,1]}\,\frac{V_s - C + A_s - A_d}{V_s + V_d - 2C}, \\ \overline{\boldsymbol{g}}_{\lambda^\star} &= \overline{\boldsymbol{g}}_{\mathrm{amp}} + (1-\lambda^\star)\overline{\boldsymbol{g}}^{\mathrm{ph}}_{\mathrm{std}} + \lambda^\star\overline{\boldsymbol{g}}^{\mathrm{ph}}_{\mathrm{dir}}. \end{aligned} \tag{13}$$

Algorithm 1. AMVP: one optimization update.

1. Draw a production batch for the parameter update. On recalibration steps, also obtain a calibration batch for estimating $\lambda^\star$.
2. Evaluate E_{loc}, the conventional amplitude force, and the standard and direct phase forces, stopping gradients through u_θ only when computing the direct phase force.
3. On calibration updates, estimate V_s, V_d, C and, for shared parameters, A_s, A_d; evaluate Eq. (13) and retain the coefficient until the next calibration.
4. Form the complete force using Eq. (13). Fixed endpoints use $\lambda = 0$ (standard) or $\lambda = 1$ (direct).
5. Apply the common preconditioner or optimizer and update θ.

Here $A_s = A_d = 0$ for disjoint amplitude and phase parameter blocks, and overbars are production-batch means. Hence $\lambda = 0$ corresponds to the standard endpoint and $\lambda = 1$ corresponds to the direct endpoint.

Quantum many-body physics calculations. The principal lattice system is the 50-rung flux ladder of Eq. (5). Independent real networks represent u_θ and φ_θ. The standard controls vary only the phase-component multiplier after the common preconditioned solve; the MLP and ResNet capacity controls otherwise use matched estimator protocols.

The ladder reference, $E_{\mathrm{ref}} = -43.302963$, is obtained using finite MPS[24, 41] in TeNPy[42] with $\chi = 64$ and truncation cutoff 10^{-10}; doubling χ changes the energy by 7×10^{-6}. For an oriented bond $i \to j$, the current operator is $\widehat{j}_{ij} = -\partial H/\partial A_{ij}$, and $j_c(x) = [j_{\mathrm{bottom}}(x) - j_{\mathrm{top}}(x)]/2$. The scalar current ratio is $\mathrm{median}_x\left[\mathrm{median}_{\mathrm{runs}}\, j_c(x)/j_c^{\mathrm{DMRG}}(x)\right]$ over the displayed bonds.

The square-flux system contains 8×8 sites with periodic boundary condition and $S^z_{\mathrm{tot}} = 0$,

$$H_\square = \sum_{\langle ij\rangle}\left[J_z S^z_i S^z_j + \frac{J_{xy}}{2}\left(e^{\mathrm{i}A_{ij}} S^+_i S^-_j + \mathrm{h.c.}\right)\right], \tag{14}$$

with $J_{xy} = 1$, $J_z = 0$ and flux $\Phi = 0.25\pi$ per plaquette. In Landau gauge, positive vertical bonds carry $A_{(x,y),(x,y+1)} = \Phi x$; horizontal bonds are real except for the periodic-x crossing bond, which carries $-\Phi L_x y$. From the DMRG energies at $\chi = 512, 1024, 2048$, we set $r = (E_{2048} - E_{1024})/(E_{1024} - E_{512})$ and $E_{\mathrm{ref}} = E_{2048} + (E_{2048} - E_{1024})r/(1-r) \simeq -28.3931$.

The supplementary chiral-chain Hamiltonian is $H_{\mathrm{ch}} = J\sum_i \mathbf{S}_i\cdot\mathbf{S}_{i+1} + \alpha\sum_i \mathbf{S}_i\cdot(\mathbf{S}_{i+1}\times\mathbf{S}_{i+2})$[31], with $J = \alpha = 1$, 50 periodic sites and $S^z_{\mathrm{tot}} = 0$. A shared real two-head network is optimized with the same pipeline for all three estimators; the DMRG reference is -23.748041. The continuum calculation uses the DeepHall[33] complex

Psiformer[43] on the Haldane sphere[44] at $\nu = 1/3$ for $N = 6$ ($2Q = 15$) and $N = 8$ ($2Q = 21$). The corrected energy per particle is $\frac{E_c}{N} = \sqrt{\frac{2Q\nu}{N}}\,\frac{E - N/2 - N^2/(2\sqrt{Q})}{N}$ in units of $e^2/(\epsilon\ell_B)$.

Molecular phase optimization. For N_2 in Fig. 3a and Supplementary Information Fig. S3a,c, amplitude and phase are optimized jointly on a fixed determinant support. For I_2, Fig. 3c isolates phase optimization at fixed amplitude, whereas Fig. 3b optimizes amplitude and phase jointly under the same explicitly spin–orbit-coupled Hamiltonian.

The spin gap[45] in Fig. 3a is the energy difference between the target-singlet ground state and the lowest excited state with a different total spin.

Matched-state diagnostics and exact-state control. The square-torus diagnostics evaluate all phase-force estimators on the same B configurations using the variance and shared-signal SNR definitions in Eq. (7).

The exact-eigenstate analysis in Fig. 4f–h uses the periodic one-particle honeycomb J_1–J_2 model,

$$\hat{H} = -J_1 \sum_{\langle i,j\rangle} \left(\hat{c}_i^\dagger \hat{c}_j + \text{h.c.}\right) + J_2 \sum_{\langle\langle i,j\rangle\rangle} \left(\hat{c}_i^\dagger \hat{c}_j + \text{h.c.}\right). \quad (15)$$

Statistical analysis. All comparisons use ten paired independent runs and are summarized by the median and interquartile range.

Data availability. The data supporting the findings of this study are publicly available in the AMVP repository[46].

Code availability. The associated source code and calculation workflows are publicly available in the AMVP repository[46].

Acknowledgements. We acknowledge Yubing Qian, Xiaoyong Ni, Xin Jin, Ji-Zhong Jiang and Khachatur Nazaryan for fruitful discussions. We also thank Ji Chen's group for making the DeepHall code available, parts of which served as a reference for our implementation. This work was supported by the National R&D Program of China (2024YFA1410500, 2022YFA1403601), the Innovation Program for Quantum Science and Technology (Grant No. 2021ZD0302800), the National Natural Science Foundation of China (No. 12322402, No. 12274206), the Natural Science Foundation of Jiangsu Province (No. BK20233001), the Fundamental Research Funds for the Central Universities (No. KG202501), the Armenian Higher Education and Science Committee ARPI Remote Laboratory program 24RL-1C024, research projects 21AG-1C024 and 25Post-Doc1C003.

Competing interests. The authors declare no competing interests.

Contributions. Y.-R.X., R.W., and C.W. conceived this project. Y.-R.X. designed the algorithm and performed the calculation supervised by R.W., B.W., and C.W. All authors contributed to the analysis of the data, the discussion of the results and the writing of the manuscript.

[1] X. Gao and L.-M. Duan, Efficient representation of quantum many-body states with deep neural networks, Nature Communications **8**, 662 (2017).
[2] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, Science **355**, 602 (2017).
[3] D. Pfau, J. S. Spencer, A. G. D. G. Matthews, and W. M. C. Foulkes, Ab initio solution of the many-electron Schrödinger equation with deep neural networks, Physical Review Research **2**, 033429 (2020).
[4] J. Hermann, Z. Schätzle, and F. Noé, Deep-neural-network solution of the electronic Schrödinger equation, Nature Chemistry **12**, 891 (2020).
[5] M. Medvidović and J. Robledo Moreno, Neural-network quantum states for many-body physics, The European Physical Journal Plus **139**, 631 (2024).
[6] Z. Tang, H. Chen, Y. Li, Y. Qian, Y. Wang, W. Fu, J. Li, C. Si, W. Duan, J. Chen, and Y. Xu, Deep-learning electronic structure calculations, Nature Computational Science **5**, 1133 (2025).
[7] M. Bukov, M. Schmitt, and M. Dupont, Learning the ground state of a non-stoquastic quantum Hamiltonian in a rugged neural network landscape, SciPost Physics **10**, 147 (2021).
[8] J. Hermann, J. Spencer, K. Choo, A. Mezzacapo, W. M. C. Foulkes, D. Pfau, G. Carleo, and F. Noé, Ab initio quantum chemistry with neural-network wavefunctions, Nature Reviews Chemistry **7**, 692 (2023).
[9] E. Ledinauskas and E. Anisimovas, Universal performance gap of neural quantum states applied to the Hofstadter–Bose–Hubbard model, SciPost Physics **18**, 011 (2025).
[10] O. Sharir, Y. Levine, N. Wies, G. Carleo, and A. Shashua, Deep autoregressive models for the efficient variational simulation of many-body quantum systems, Physical Review Letters **124**, 020503 (2020).
[11] A. Chen and M. Heyl, Empowering deep neural quantum states through efficient optimization, Nature Physics **20**, 1476 (2024).
[12] A. Sinibaldi, C. Giuliani, G. Carleo, and F. Vicentini, Unbiasing time-dependent Variational Monte Carlo by projected quantum evolution, Quantum **7**, 1131 (2023).
[13] A. Misery, L. Gravina, A. Santini, and F. Vicentini, Looking elsewhere: improving variational Monte Carlo gradients by importance sampling, Machine Learning: Science and Technology **7**, 015035 (2026), arXiv:2507.05352.
[14] T. Westerhout, N. Astrakhantsev, K. S. Tikhonov, M. I. Katsnelson, and A. A. Bagrov, Generalization properties of neural network approximations to frustrated magnet ground states, Nature Communications **11**, 1593 (2020).
[15] A. Szabó and C. Castelnovo, Neural network wave functions and the sign problem, Physical Review Research **2**, 033075 (2020).
[16] C. A. Melton, M. C. Bennett, and L. Mitas, Quantum Monte Carlo with variable spins, The Journal of Chemical Physics **144**, 244113 (2016).
[17] R. Ranganath, S. Gerrish, and D. M. Blei, Black box variational inference, in *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics (AISTATS)*, Proceedings of Machine Learning Research,

Vol. 33 (2014) pp. 814–822.
[18] P. Parmas, C. E. Rasmussen, J. Peters, and K. Doya, PIPPS: Flexible model-based policy search robust to the curse of chaos, in *Proceedings of the 35th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 80 (2018) pp. 4065–4074.
[19] W. Grathwohl, D. Choi, Y. Wu, G. Roeder, and D. Duvenaud, Backpropagation through the void: Optimizing control variates for black-box gradient estimation, in *International Conference on Learning Representations (ICLR)* (2018) arXiv:1711.00123.
[20] G. Tucker, D. Lawson, S. Gu, and C. J. Maddison, Doubly reparameterized gradient estimators for Monte Carlo objectives, in *International Conference on Learning Representations (ICLR)* (2019).
[21] M. Atala, M. Aidelsburger, M. Lohse, J. T. Barreiro, B. Paredes, and I. Bloch, Observation of chiral currents with ultracold atoms in bosonic ladders, Nature Physics **10**, 588 (2014).
[22] Q. Liang, Z. Dong, J.-S. Pan, H. Wang, H. Li, Z. Yang, W. Yi, and B. Yan, Chiral dynamics of ultracold atoms under a tunable SU(2) synthetic gauge field, Nature Physics **20**, 1738 (2024).
[23] A. Impertro, S. Huh, S. Karch, J. F. Wienand, I. Bloch, and M. Aidelsburger, Strongly interacting Meissner phases in large bosonic flux ladders, Nature Physics **21**, 895 (2025).
[24] S. R. White, Density matrix formulation for quantum renormalization groups, Physical Review Letters **69**, 2863 (1992).
[25] S. Mohamed, M. Rosca, M. Figurnov, and A. Mnih, Monte Carlo gradient estimation in machine learning, Journal of Machine Learning Research **21**, 1 (2020).
[26] E. Greensmith, P. L. Bartlett, and J. Baxter, Variance reduction techniques for gradient estimates in reinforcement learning, Journal of Machine Learning Research **5**, 1471 (2004).
[27] L. Metz, N. Maheswaranathan, J. Nixon, D. Freeman, and J. Sohl-Dickstein, Understanding and correcting pathologies in the training of learned optimizers, in *Proceedings of the 36th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 97 (2019) pp. 4556–4565.
[28] R. Assaraf and M. Caffarel, Zero-variance principle for Monte Carlo algorithms, Physical Review Letters **83**, 4682 (1999).
[29] E. Orignac and T. Giamarchi, Meissner effect in a bosonic ladder, Physical Review B **64**, 144515 (2001).
[30] D. Hügel and B. Paredes, Chiral ladders and the edges of quantum Hall insulators, Physical Review A **89**, 023619 (2014).
[31] C. Wei, V. V. Mkhitaryan, and T. A. Sedrakyan, Unveiling chiral states in the XXZ chain: Finite-size scaling probing symmetry-enriched $c = 1$ conformal field theories, Journal of High Energy Physics **2024**, 125 (2024).
[32] Y. Teng, D. D. Dai, and L. Fu, Solving the fractional quantum hall problem with self-attention neural network, Physical Review B **111**, 205117 (2025).
[33] Y. Qian, T. Zhao, J. Zhang, T. Xiang, X. Li, and J. Chen, Describing landau level mixing in fractional quantum hall states with deep learning, Physical Review Letters **134**, 176503 (2025).
[34] K. Choo, A. Mezzacapo, and G. Carleo, Fermionic neural-network states for ab-initio electronic structure, Nature Communications **11**, 2368 (2020).
[35] T. D. Barrett, A. Malyshev, and A. I. Lvovsky, Autoregressive neural-network wavefunctions for Ab Initio quantum chemistry, Nature Machine Intelligence **4**, 351 (2022).
[36] H.-J. Flad, M. Dolg, and A. Shukla, Spin-orbit coupling in variational quantum Monte Carlo calculations, Physical Review A **55**, 4183 (1997).
[37] L. Bottou, F. E. Curtis, and J. Nocedal, Optimization methods for large-scale machine learning, SIAM Review **60**, 223 (2018).
[38] D. P. Kingma and M. Welling, Auto-encoding variational Bayes, in *International Conference on Learning Representations (ICLR)* (2014) arXiv:1312.6114.
[39] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Quantum Monte Carlo simulations of solids, Reviews of Modern Physics **73**, 33 (2001).
[40] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, Cambridge, 2017).
[41] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, Annals of Physics **326**, 96 (2011).
[42] J. Hauschild and F. Pollmann, Efficient numerical simulations with tensor networks: Tensor network python (TeNPy), SciPost Physics Lecture Notes , 5 (2018).
[43] I. von Glehn, J. S. Spencer, and D. Pfau, A self-attention ansatz for ab-initio quantum chemistry, in *International Conference on Learning Representations (ICLR)* (2023) arXiv:2211.13672.
[44] F. D. M. Haldane, Fractional quantization of the hall effect: A hierarchy of incompressible quantum fluid states, Physical Review Letters **51**, 605 (1983).
[45] Z. Li, Z. Lu, R. Li, X. Wen, X. Li, L. Wang, J. Chen, and W. Ren, Spin-symmetry-enforced solution of the many-body schrödinger equation with a deep neural network, Nature Computational Science **4**, 910 (2024).
[46] Chenan-Phys, NN-VMC training scripts: adaptive minimum-variance phase (AMVP) estimator, GitHub repository (2026).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SI.pdf