

Addressing Quantum Many-Body Problems via Analog Devices

Vinícius Mohr,¹ Tristan Zerweck,¹ Luis Mickeler,² Rachel Grange,² Juan Carrasquilla¹

¹Institute for Theoretical Physics, ETH Zürich

²Institute for Quantum Electronics, ETH Zürich

Abstract

Neural quantum states are among the most expressive variational ansätze for quantum many-body problems, yet evaluating them incurs a digital computational cost that grows with network size. We propose realizing the ansatz directly in analog hardware, using experimentally demonstrated devices as variational wavefunctions within the variational Monte Carlo (VMC) framework. We study two platforms: the microring resonator meshgrid (MRM) and the optical reservoir (OR), a disordered lithium niobate slab that maps phase-encoded spin configurations to high-dimensional speckle features. Benchmarked against a restricted Boltzmann machine (RBM) on frustrated two-dimensional spin models and the H_2 molecule, both platforms match the digital baseline. These results establish analog hardware as a viable substrate for neural quantum states.

1 Introduction

The simulation of strongly correlated quantum many-body systems is one of the foremost challenges in modern physics. State-of-the-art numerical methods each face characteristic limitations: density functional theory fails in strongly correlated regimes [1, 2], tensor network methods scale exponentially with entanglement and become intractable beyond quasi-one-dimensional systems [3, 4], quantum Monte Carlo [5] suffers from the sign problem in frustrated or fermionic systems, and neural quantum states (NQS) within variational Monte Carlo (VMC) [6] are increasingly limited by the throughput and power consumption of conventional digital hardware.

We introduce a paradigm in which analog physical devices serve as trainable variational wavefunction ansätze, optimized within the VMC framework. Rather than instantiating the wavefunction on digital hardware, we let a device’s own physical dynamics carry out the computation, using it as a physical neural network (PNN) [7]. This approach is compelling because many physical processes perform the core operations of computation directly in hardware, executing linear-algebra primitives at a throughput and energy efficiency that conventional processors cannot match [8–10] and, in the optical case, even acting as universal function approximators [11]. These are precisely the properties a variational ansatz demands: enough expressivity to capture nontrivial quantum correlations, and cheap enough evaluation to facilitate the variational optimization. We adapt two experimentally realized platforms (Fig. 1), both previously demonstrated on image-classification tasks: the microring resonator meshgrid (MRM) [12] and the optical reservoir (OR) [13].

2 Analog Wavefunction Architectures

2.1 Microring Resonator Meshgrid

The MRM [12], as illustrated in Fig. 1a, is a layered network of optically coupled resonator modes in which the input spin configuration $\sigma \in \{-1, +1\}^N$ is encoded directly into the resonator detunings rather than the propagating light. This causes the scattering to produce a nonlinear input-output map from a fully linear physical process. Setting $x_0 = \text{diag}(\sigma)$, successive layers effectively apply

$$x_{n+1} = f_{\theta_n}(x_n) = (A_{\theta_n}^\dagger x_n A_{\theta_n} + b_{\theta_n})^{-1}, \quad (1)$$




Figure 1: (a) Microring resonator meshgrid: spin inputs modulate resonator detunings, the summed scattering output yields $\log \psi$. Adapted from Ref. [12]. (b) Optical reservoir: spin configurations are phase-encoded and scattered through a $\chi^{(2)}$ medium, speckle features at ω and 2ω are read out by a trainable linear layer. Adapted from Ref. [13].

where A_{θ_n} encodes interlayer couplings and b_{θ_n} encodes decay rates and reference frequencies. The variational wavefunction amplitude is then the composition of these layers

$$\log \psi_{\theta}(\sigma) = \sum_i [(f_{\theta_n} \circ \dots \circ f_{\theta_1})(x_0)]_{i,i}, \quad (2)$$

which nicely mirrors a deep neural network with matrix inversion as a nonlinear activation. Another beneficial feature of this setup is also that gradients can be extracted from local scattering measurements [12], enabling efficient training.

2.2 Optical Reservoir

The OR [13] uses a polycrystalline lithium niobate slab as a fixed, untrained feature extractor. Input spins are phase-encoded via a spatial light modulator ($\sigma_i = +1$: zero phase, $\sigma_i = -1$: π phase) and the phase-modulated wavefront propagates through the slab, where random quasi-phase-matching and $\chi^{(2)}$ nonlinearity produce high-dimensional speckle patterns at both ω and 2ω . The resulting intensity spectra are then captured using a camera and passed to a single trainable linear layer that maps the speckle features to the variational log-amplitude:

$$\log \psi_{\theta}(\sigma) = A_{\theta} \log \tilde{\psi}_{\text{FH/SH}}(\sigma) + b_{\theta}, \quad (3)$$

where $\log \tilde{\psi}_{\text{FH/SH}}(\sigma)$ denotes the fundamental- and second-harmonic speckle intensity. Because the linear readout is the only component evaluated digitally, any expressivity beyond product states must arise from the reservoir's physical nonlinear mapping.

3 Ground-State Benchmarks

3.1 Lattice Spin Hamiltonians

We benchmarked both architectures against a restricted Boltzmann machine (RBM) on three two-dimensional spin Hamiltonians:

$$H_{\text{TFIM}} = -J \sum_{\langle i,j \rangle} Z_i Z_j - h \sum_i X_i \quad (J = 1, h = 3.044), \quad (4)$$

$$H_{\text{Heis}} = J \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{S}_j \quad (J = 1), \quad (5)$$

$$H_{J_1 J_2} = J_1 \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{S}_j + J_2 \sum_{\langle\langle i,j \rangle\rangle} \mathbf{S}_i \cdot \mathbf{S}_j \quad (J_1 = 1, J_2 = 0.5). \quad (6)$$

All three architectures used comparable parameter counts (1648 RBM, 608 MRM, 1601 OR) and were optimized via stochastic reconfiguration [16] with spatial symmetries imposed. On a




Figure 2: Ground-state benchmarks and quantum chemistry results. (a) Relative energy errors on a symmetrized 4×4 lattice. (b) Spin structure factor $S(\mathbf{q})$ for the 4×4 lattice. (c) Potential energy surface of H_2 (6-31G basis, $N = 8$ spins via Jordan-Wigner) traced by the OR versus exact diagonalization, the equilibrium bond length $d_{\text{eq}} = 0.74 \text{\AA}$ is accurately recovered. (d) Energy errors on an unsymmetrized 10×10 lattice (MPS at bond dimension 2000 and QMC as baselines). (e) Structure factor for the 10×10 lattice. (f) Relative ground-state energy errors at equilibrium, the OR and RBM reach $\sim 10^{-7}$, while the MRM reaches $\sim 2 \times 10^{-3}$.

4×4 lattice, both analog architectures achieve relative energy errors matching the RBM baseline across all three Hamiltonians (Fig. 2a). Both architectures reproduce the spin structure factor $S(\mathbf{q})$, from exact diagonalization calculations (Fig. 2b), capturing the antiferromagnetic ordering of the Heisenberg model and the frustrated stripe order of the J_1 - J_2 model at $J_2/J_1 = 0.5$, which is a particularly challenging frustrated regime in which sign-problem-free QMC is unavailable. Extending the MRM (unsymmetrized) 10×10 lattice, it remains competitive against the RBM, especially for the Heisenberg model (Fig. 2d,e), demonstrating that analog wavefunctions can be scaled up to larger, two-dimensional systems. Here, the "exact" energy and structure factor results were generated using DMRG for the TFIM and QMC for the Heisenberg model.

3.2 Quantum Chemistry: the H_2 Molecule

To demonstrate that the approach extends beyond lattice spin models, we applied the OR to the hydrogen molecule in the 6-31G basis. The H_2 Hamiltonian is mapped to an $N = 8$ spin problem via the Jordan-Wigner transformation using molecular integrals from PySCF [17, 18]. As can be seen in Fig. 2c, the OR traces the full potential energy surface (PES) in near-exact agreement with exact diagonalization across interatomic distances from 0.1 to $2.5 \text{\AA}$, correctly locating the equilibrium geometry at $d_{\text{eq}} = 0.74 \text{\AA}$. At equilibrium, the OR achieves a relative ground-state energy error of $\sim 2 \times 10^{-7}$, comparable to the RBM ($\sim 5 \times 10^{-7}$) and nearly four orders of magnitude lower than the MRM (Fig. 2f). This demonstrates that analog variational ansätze extend beyond spin systems to problems in quantum chemistry.

4 Conclusion

We have introduced analog physical devices as variational wavefunction ansätze for quantum many-body problems. Both the MRM and the OR integrate seamlessly into the VMC framework and match or approach RBM accuracy on spin models in two dimensions, with the OR further achieving near-exact energies for the H_2 molecule. Future work will address theoretical expressivity bounds, hardware noise tolerance, and the regimes where analog devices offer a genuine computational advantage over digital neural quantum states.

References

- [1] Pierre Hohenberg and Walter Kohn. Inhomogeneous electron gas. *Physical Review*, 136(3B):B864–B871, 1964. doi: 10.1103/PhysRev.136.B864. URL <https://doi.org/10.1103/PhysRev.136.B864>.
- [2] Walter Kohn and Lu Jeu Sham. Self-consistent equations including exchange and correlation effects. *Physical Review*, 140(4A):A1133–A1138, 1965. doi: 10.1103/PhysRev.140.A1133. URL <https://doi.org/10.1103/PhysRev.140.A1133>.
- [3] Steven R. White. Density matrix formulation for quantum renormalization groups. *Physical Review Letters*, 69(19):2863–2866, November 1992. doi: 10.1103/PhysRevLett.69.2863. URL <https://doi.org/10.1103/PhysRevLett.69.2863>.
- [4] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of physics*, 326(1):96–192, 2011. URL <https://doi.org/10.1016/j.aop.2010.09.012>.
- [5] Federico Becca and Sandro Sorella. Quantum Monte Carlo Approaches for Correlated Systems by Federico Becca. /core/books/quantum-monte-carlo-approaches-for-correlated-systems/EB88C86BD9553A0738BDAE400D0B2900, November 2017. URL <https://doi.org/10.1017/9781316417041>.
- [6] Matija Medvidović and Javier Robledo Moreno. Neural-network quantum states for many-body physics. *The European Physical Journal Plus*, 139(7):631, July 2024. ISSN 2190-5444. doi: 10.1140/epjp/s13360-024-05311-y. URL <https://doi.org/10.1140/epjp/s13360-024-05311-y>.
- [7] Logan G. Wright, Tatsuhiko Onodera, Martin M. Stein, Tianyu Wang, Darren T. Schachter, Zoey Hu, and Peter L. McMahon. Deep physical neural networks trained with backpropagation. *Nature*, 601(7894):549–555, January 2022. ISSN 1476-4687. doi: 10.1038/s41586-021-04223-6. URL <https://doi.org/10.1038/s41586-021-04223-6>.
- [8] Hailong Zhou, Jianji Dong, Junwei Cheng, Wenchao Dong, Chaoran Huang, Yichen Shen, Qiming Zhang, Min Gu, Chao Qian, Hongsheng Chen, Zhichao Ruan, and Xinliang Zhang. Photonic matrix multiplication lights up photonic accelerator and beyond. *Light: Science & Applications*, 11(1):30, February 2022. ISSN 2047-7538. doi: 10.1038/s41377-022-00717-8. URL <https://doi.org/10.1038/s41377-022-00717-8>.
- [9] Bo Peng et al. An integrated large-scale photonic accelerator for high-performance computing. *Nature*, 2025. doi: 10.1038/s41586-025-08786-6. URL <https://doi.org/10.1038/s41586-025-08786-6>. Demonstrates an integrated photonic accelerator targeting energy-efficient, high-throughput compute (including linear-algebra primitives).
- [10] Arash Shafiee, Aniruddha Nag, Naveen Muralimanohar, Rajeev Balasubramanian, John Paul Strachan, Miao Hu, R. Stanley Williams, and Vivek Srikumar. ISAAC: A Convolutional Neural Network Accelerator with In-Situ Analog Arithmetic in Crossbars. In *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, pages 14–26, 2016. doi: 10.1109/ISCA.2016.12. URL <https://doi.org/10.1109/ISCA.2016.12>.
- [11] Benjamin Savinson, David J. Norris, Siddhartha Mishra, and Samuel Lanthaler. Universality of physical neural networks with multivariate nonlinearity, 2025. URL <https://arxiv.org/abs/2509.05420>.
- [12] Clara C. Wanjura and Florian Marquardt. Fully nonlinear neuromorphic computing with linear wave scattering. *Nature Physics*, 20(9):1434–1440, September 2024. ISSN 1745-2481. doi: 10.1038/s41567-024-02534-9. URL <https://doi.org/10.1038/s41567-024-02534-9>.
- [13] Hao Wang, Jianqi Hu, Andrea Morandi, Alfonso Nardi, Fei Xia, Xuanchen Li, Romolo Savo, Qiang Liu, Rachel Grange, and Sylvain Gigan. Large-scale photonic computing with nonlinear disordered media. *Nature Computational Science*, 4(6):429–439, June 2024. ISSN 2662-8457. doi: 10.1038/s43588-024-00644-1. URL <https://doi.org/10.1038/s43588-024-00644-1>.
- [14] Luca Leone, Simone F. E. Oliviero, and Alioscia Hamma. Stabilizer rényi entropy. *Physical*

- Review Letters*, 128:050402, 2022. doi: 10.1103/PhysRevLett.128.050402. URL <https://link.aps.org/doi/10.1103/PhysRevLett.128.050402>.
- [15] Sergey Bravyi, Dan Browne, Padraic Calpin, Earl Campbell, David Gosset, and Mark Howard. Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum*, 3:181, September 2019. ISSN 2521-327X. doi: 10.22331/q-2019-09-02-181. URL <https://doi.org/10.22331/q-2019-09-02-181>.
 - [16] Filippo Vicentini, Damian Hofmann, Attila Szabó, Dian Wu, Christopher Roth, Clemens Giuliani, Gabriel Pescia, Jannes Nys, Vladimir Vargas-Calderón, Nikita Astrakhantsev, and Giuseppe Carleo. Codebase release 3.4 for NetKet. *SciPost Phys. Codebases*, pages 7–r3.4, 2022. doi: 10.21468/SciPostPhysCodeb.7-r3.4. URL <https://scipost.org/10.21468/SciPostPhysCodeb.7-r3.4>.
 - [17] Qiming Sun, Xing Zhang, Samraghi Banerjee, Peng Bao, Marc Barbry, Nick S. Blunt, Nikolay A. Bogdanov, George H. Booth, Jia Chen, Zhi-Hao Cui, Janus J. Eriksen, Yang Gao, Shengguo Li, Tianyu Li, Jing Liu, Weijie Luo, Chao Ma, Mario Motta, Jean Moussa, Guangjin Nan, Anders M. N. Niklasson, Juan E. Peralta, Benjamin Pritchard, Wibowo A. Purwanto, Paul J. Robinson, Edoardo Ronca, Elvira R. Sayfutyarova, Markus Scheurer, Kanav Setia, Sukin Sim, Alexei Y. Sokolov, Shuyu Sun, Sheng Sun, Marcel Swart, Chia-Nan Tsai, Wen-Tao Ting, Vijay Voora, Xinyu Zhang, Yi-Qi Zhang, Zheng Zhang, Timothy C. Berkelbach, Garnet Kin-Lic Chan, C. David Sherrill, Shuaicheng Wang, Takeshi Yanai, Yihan Xu, Huan Zhai, Yuzhong He, and Jialiang Zhao. Recent developments in the PySCF program package. *The Journal of Chemical Physics*, 153(2):024109, 2020. doi: 10.1063/5.0006074. URL <https://doi.org/10.1063/5.0006074>.
 - [18] Qiming Sun, Timothy C. Berkelbach, Nick S. Blunt, George H. Booth, Sheng Guo, Zhi Li, Junzi Liu, James D. McClain, Elvira R. Sayfutyarova, Sandeep Sharma, Sebastian Wouters, and Garnet Kin-Lic Chan. PySCF: the Python-based simulations of chemistry framework. *WIREs Computational Molecular Science*, 8(1):e1340, 2018. doi: 10.1002/wcms.1340. URL <https://doi.org/10.1002/wcms.1340>.