

Fermionic Shadow Regression: Amortized Learning of Molecular Electron Dynamics

Aniket Deshpande¹, Karthik Panicker^{2,3}, Luis Mantilla Calderon^{3,4,5}, Mohsen Bagherimehrab^{2,4},
and Alán Aspuru-Guzik^{2,3,4,5,6,7,8,*}

¹*Department of Physics, University of Illinois Urbana-Champaign, USA*

²*Department of Chemistry, University of Toronto, Canada*

³*Vector Institute for Artificial Intelligence, Toronto, Canada*

⁴*Department of Computer Science, University of Toronto, Canada*

⁵*NVIDIA Corporation, Santa Clara, USA*

⁶*Canadian Institute for Advanced Research (CIFAR), Toronto, Canada*

⁷*Department of Materials Science & Engineering, University of Toronto, Canada*

⁸*Department of Chemical Engineering & Applied Chemistry, University of Toronto, Canada*

*Correspondence: alan@aspuru.com

Abstract

Algorithmic shadow spectroscopy extracts molecular transition frequencies from randomized measurements of many observables over time at a fixed molecular geometry. We ask whether the underlying object, the time-dependent fermionic-shadow *signal matrix* of degree-two Majorana observables, can instead be *learned across a geometry scan*, amortizing shadow-spectroscopy observables so that the inputs to spectral post-processing at held-out geometries are predicted classically. We introduce the *Fermionic Shadow Regressor* (FSR), an explicit amplitude-phase network whose oscillation frequencies are conditioned on Hartree-Fock orbital energies and bounded, geometry by geometry, by an exact *line-spectrum screen*. On four simulated eight-qubit molecules a single recipe transfers with high held-out fidelity, outperforms generic Gaussian-process and Fourier-feature surrogates on the harder scans, and degrades in the spectral-congestion region identified by the screen; a LiH finite-shot test matches the infinite-shot reference at $S = 1000$. The targets are classically simulated and the screen is an exact-diagonalization oracle: this is an amortization-and-sampling-design study, *not* a demonstration of quantum advantage.

1 Problem and contribution

High-accuracy quantum spectroscopy is resource intensive, and classical shadows [2] reduce the measurement cost of estimating many observables: “algorithmic shadow spectroscopy” [1] estimates transition energies from the principal components of a bank of shadow-estimated observables sampled over time, without ever reconstructing the state. But a chemistry workflow rarely wants one spectrum; it wants spectra *across a geometry scan*, and re-acquiring shadow data independently at every geometry discards the strong structure shared between nearby geometries.

FSR amortizes this shared structure by learning the *observable signal matrix* $D_{\mu k}(\mathbf{x}) = \langle \Gamma_{\mu} \rangle(\mathbf{x}, t_k)$, the time-dependent expectations of the $N_o = 120$ degree-two Majorana observables of an eight-qubit active space at geometry $\mathbf{x}$. We train a model to predict $D(\mathbf{x}, \cdot)$ at held-out geometries from shadow targets on a training set, turning repeated per-geometry acquisition into a supervised regression problem over geometry and time. Where algorithmic shadow spectroscopy processes measured traces at a fixed instance, the FSR learns the signal matrix across a geometry scan. The targets come from fermionic/matchgate classical shadows [5–7]; the shadow-learning theory of Huang et al. [3, 4] motivates this amortization regime under smoothness and bounded-shadow-norm assumptions.

Scope. Targets are classically simulated, the systems are eight-qubit active spaces, and the screen is an exact-diagonalization oracle: this is an architecture-and-sampling-design study complementary to algorithmic shadow spectroscopy, not a quantum speedup.

2 Method

Harmonic model. Under $H(\mathbf{x})$, the geometry-parameterized active-space Hamiltonian, a time-evolved expectation is a sum of sinusoids at the Bohr frequencies $\omega_{jk} = E_j - E_k$. The FSR encodes this structure directly rather than asking a generic network to synthesize oscillations:

$$\hat{D}_\mu(\mathbf{x}, t) = \sum_{k=1}^K \left[a_{\mu k}(\mathbf{x}) \cos(\omega_k(\mathbf{x}) t) + b_{\mu k}(\mathbf{x}) \sin(\omega_k(\mathbf{x}) t) \right] + c_\mu(\mathbf{x}), \quad (1)$$

where every geometry-dependent quantity is a function of standardized Hartree–Fock orbital energies $\epsilon(\mathbf{x})$ (a chemically motivated coordinate that tracks the gap spectrum far better than the bare bond length). The frequencies $\omega_k(\mathbf{x})$ are confined, geometry by geometry, to an *adaptive bandwidth* (below); the amplitudes are produced by a low-rank trunk so the 120 channels share a common pool of temporal modes. Because (1) is linear in the harmonics, the model extrapolates in time by construction.

Learning and resolution criteria. (i) *Learnability*: specializing Huang et al. [4] to fermionic shadows, the map $\mathbf{x} \mapsto D(\mathbf{x}, \cdot)$ is learnable on average *inside* the sampled geometry box given bounded shadow norm and smoothness; it promises nothing about extrapolation. (ii) *Resolution*: the sampling step must satisfy Nyquist, $\Delta t \leq \pi/\Omega$ for bandwidth Ω , and a horizon T resolves only line spacings $\gtrsim 2\pi/T$.

Line-spectrum screen. Because the active space is small, we diagonalize $H(\mathbf{x})$ *exactly* at each geometry, read off the Bohr lines, and extract two numbers: the operational bandwidth $\omega_{\text{op}}(\mathbf{x})$ and the smallest co-dominant line spacing $\delta\omega(\mathbf{x})$. The resolution criterion then prescribes the grid $(\Delta t, N_T)$ and the FSR’s per-geometry frequency ceiling. The same spectrum localizes congested geometries where the surrogate is expected to fail, so the screen serves both as a grid-design rule and as a failure diagnostic.

3 Results

Cross-molecule transfer. The suite is H_4 , LiH , BeH_2 , and N_2 , all reduced to eight-qubit/120-channel active spaces, with CAS(4,4) for H_4 , BeH_2 , and N_2 , and CAS(2,4) for LiH . With the architecture and optimizer held fixed and *only* the screen-set grid changing, the recipe transfers. LiH is heteronuclear and has dominant frequencies roughly an order of magnitude slower than H_4 , requiring a four-times longer horizon. It reaches a mean held-out Pearson correlation of 0.992 (median 0.997), staying coherent across its full training box and horizon (Fig. 1b); BeH_2 matches it at 0.997. N_2 is accurate everywhere except a narrow band at $R \approx 2.1\text{--}2.25 \text{ \AA}$, the avoided-crossing congestion the screen flags beforehand, so the failure is localized by an independent diagnostic rather than discovered only after training.

Baselines and ablations. Against three surrogates trained on the same data (Fig. 1c), the FSR is decisively better on the harder H_4/N_2 scans; a fixed-frequency-bank regressor with the *same*




Figure 1: The Fermionic Shadow Regressor. (a) HF orbital features condition a frequency network and a low-rank amplitude trunk; harmonic synthesis at time t gives the 120-channel signal matrix $\hat{D}$. (b) One recipe, retrained only with LiH’s screen-set grid, stays coherent (green) across the box and full horizon. (c) Held-out accuracy vs. kernel-ridge, a parameter-matched Fourier-feature MLP, and a fixed-frequency-bank linear surrogate.

oracle ceiling matches FSR only on the smoothest molecule, isolating the value of *learning* the geometry-conditioned frequencies rather than tiling them. Ablating the adaptive bandwidth, HF features, or low-rank sharing each degrades accuracy.

Frequency and smoothness diagnostics. Predicted dominant frequencies sit below the co-dominant line spacing for all four molecules, so the model resolves Bohr lines rather than a smooth envelope. The screen, however, tests spectral *resolution* but not the *smoothness* of the geometry-to-signal map: a fifth molecule, H_2O , passes the resolution screen everywhere yet fails almost uniformly, indicating a distinct smoothness-failure mode and motivating a future screen for both resolution and geometry smoothness.

Finite-shot stress test. Replacing the infinite-shot targets by finite Born-sample averages over S outcomes per matchgate label, the LiH surrogate improves monotonically with S and matches the infinite-shot reference at $S = 1000$, so the regression is not an artifact of idealized targets.

4 Outlook

The two explicit limits are simulated targets and an exact-diagonalization oracle screen. Scaling beyond eight-qubit active spaces needs a sampling diagnostic that estimates bandwidth, line spacing, and geometry smoothness from shadow data rather than exact spectra, together with symmetry-aware architectures.

References

- [1] H. H. S. Chan, R. Meister, M. L. Goh, and B. Koczor. Algorithmic shadow spectroscopy. *PRX Quantum*, 6(1), Mar. 2025. ISSN 2691-3399. doi: 10.1103/prxquantum.6.010352. URL <http://dx.doi.org/10.1103/PRXQuantum.6.010352>.
- [2] H.-Y. Huang, R. Kueng, and J. Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020.
- [3] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill. Provably efficient machine learning for quantum many-body problems. *Science*, 377(6613):eabk3333, 2022.
- [4] H.-Y. Huang, S. Chen, and J. Preskill. Learning to predict arbitrary quantum processes. *PRX Quantum*, 4(4):040337, 2023.
- [5] G. H. Low. Classical shadows of fermions with particle number symmetry, 2022. URL <https://arxiv.org/abs/2208.08964>.
- [6] K. Wan, W. J. Huggins, J. Lee, and R. Babbush. Matchgate shadows for fermionic quantum simulation. *Communications in Mathematical Physics*, 404(2):629–700, 2023.
- [7] A. Zhao, N. C. Rubin, and A. Miyake. Fermionic partial tomography via classical shadows. *Physical Review Letters*, 127(11):110504, 2021.