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.

SlaKoNet-VQD: A universal Slater-Koster tight-binding Hamiltonian for variational quantum band-structure calculations on near-term hardware

Akshaya Ajith
Johns Hopkins University
Jaehyung Lee
Johns Hopkins University
Charles Rhys Campbell
Johns Hopkins University
Kamal Choudhary
drkamal@jhu.edu

Johns Hopkins University

Article

Keywords:

Posted Date: October 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11020957/v1

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

Additional Declarations: No competing interests reported.

SlaKoNet-VQD: A universal Slater-Koster tight-binding Hamiltonian for variational quantum band-structure calculations on near-term hardware

Akshaya Ajith,[1] Jaehyung Lee,[1] Charles Rhys Campbell,[1,2] and Kamal Choudhary[1,2,*]

[1] *Department of Materials Science and Engineering, Johns Hopkins University, Baltimore, MD 21218, USA*

[2] *Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD 21218, USA*

(Dated: September 20, 2026)

Abstract

Variational quantum algorithms such as the variational quantum eigensolver (VQE) and variational quantum deflation (VQD) have emerged as leading candidates for computing electronic structure on near-term quantum hardware. For periodic solids, however, their reach has been limited not by the quantum algorithm itself but by the cost of constructing a faithful second-quantized Hamiltonian for each material, typically via density-functional theory (DFT) followed by Wannierization or by hand-fitted empirical tight-binding parameters. SlaKoNet, a recently developed method, significantly reduces the computational cost of electronic Hamiltonian construction with a small accuracy trade-off. It integrates modern deep-learning methods and the Slater-Koster tight-binding formalism to fit thousands of onsite and distance-dependent hopping and overlap tight-binding parameters across 65 elements of the periodic table. After fitting, these parameters can be used to deterministically construct an electronic Hamiltonian for any crystal made up of any of the 65 elements in the training set, and the crystal's band structure can subsequently be computed. In this work, we combine a SlaKoNet model trained on the JARVIS-TBmBJ dataset with a Qiskit-based variational quantum deflation algorithm, replacing the traditional and costly Hamiltonian construction step with a universal neural Hamiltonian generator. Within the 65-element set, the resulting SlaKoNet-VQD workflow is structure-agnostic, fully differentiable, fast, and suited to high-throughput screening of band structures. We benchmark the workflow on silicon, recovering the full eight-band band structure along the standard high-symmetry k-path with a mean absolute deviation of 1.78 meV from exact diagonalization on a 3-qubit statevector simulator. We further compute the full band structure for five additional conventional superconductors, Al, Ta, Nb, V, and ZrN, achieving a mean absolute deviation within approximately 1.8 meV, and demonstrate full execution of SlaKoNet-VQD on IBM Quantum hardware for a representative k-point ground-state calculation on aluminum with a mean absolute error of roughly 0.37 eV. Beyond the single-particle level, we promote the same SlaKoNet Hamiltonian to a correlated Hubbard model and solve it with dynamical mean-field theory, recovering the weakening of correlations across the group-5 metals and strong quasiparticle renormalization in the cuprate La_2CuO_4, and we identify the resulting impurity problem as the natural target for a quantum solver. The pipeline opens a path to high-throughput benchmarking of variational quantum algorithms across the periodic table and to gradient-based co-optimization of ansatz and Hamiltonian for materials discovery. An interactive web app is available at website: https://atomgpt.org/quantum.

I. INTRODUCTION

Quantum computing has been heralded as a natural platform for electronic-structure problems, with the variational quantum eigensolver (VQE) [1, 2] and its excited-state extension, variational quantum deflation (VQD) [3], emerging as the leading near-term algorithms. While the molecular case has matured into a well-mapped landscape [4, 5], periodic solids remain substantially less developed. Recent work has begun to bridge this gap [6–10], but the input to the quantum solver, a faithful second-quantized Hamiltonian, has invariably been built by hand for each material. Two routes dominate: density-functional theory (DFT) followed by Wannierization to obtain a localized tight-binding (TB) Hamiltonian [11, 12], or hand-fitted empirical Slater-Koster parameters [14]. Both are accurate when properly executed, but both demand human expertise per material and are difficult to compose into a high-throughput or differentiable pipeline.

This per-material setup cost has become the dominant bottleneck in scaling quantum band-structure studies. Constructing a maximally localized Wannier Hamiltonian requires careful selection of disentanglement windows [13], initial projections, and convergence checks; empirical TB models require fitting against a target band structure or experimental data. Neither workflow is differentiable, neither transfers across chemistries without re-fitting, and neither composes naturally with the gradient-based optimization that underpins modern variational quantum algorithms.

In parallel, machine-learning approaches to electronic-

* drkamal@jhu.edu

structure Hamiltonians have matured to the point of replacing these hand-built inputs. End-to-end graph neural networks such as DeepH-E3 [16], HamGNN [17], and equivariant wavefunction predictors [18] learn the Hamiltonian matrix elements directly from the local environment; SlaKoNet [19] pursues a complementary route that keeps the physically interpretable Slater-Koster (SK) tight-binding parametrization and instead uses automatic differentiation to learn the distance-dependent SK hopping and overlap integrals themselves, across 65 elements of the periodic table, against TBmBJ-level JARVIS-DFT data [21, 22]. Once trained, SlaKoNet constructs the tight-binding Hamiltonian $H(\mathbf{k})$ and overlap matrix $S(\mathbf{k})$ for any composition within its training domain by polynomial interpolation of the learned parameters at runtime. Critically for our purposes, the entire mapping from atomic coordinates to $H(\mathbf{k})$ is differentiable, runs on the GPU, and requires no system-specific manual parametrization.

In this work we close the loop by feeding SlaKoNet's tight-binding Hamiltonian directly into a Qiskit [28] VQD pipeline. The resulting SlaKoNet-VQD workflow is, to our knowledge, the first quantum band-structure calculation in which the Hamiltonian is produced by a pretrained universal TB model rather than by per-material DFT, Wannierization, or hand-fitting. We make five specific contributions. First, we build a clean pipeline: an interface between the SlaKoNet Hamiltonian generator and the Qiskit `Estimator` primitive, including a Hermitization step that handles finite-precision numerical noise from the Hamiltonian generation step (Sec. II). Second, we provide a validation: a full-Brillouin-zone benchmark on silicon along the standard high-symmetry path, comparing 3-qubit VQD eigenvalues to exact diagonalization at every k-point (Sec. III B). Third, we report hardware execution: running on IBM Quantum hardware with error mitigation and quantifying the descrepancy between simulated hardware and real hardware in this context (Sec. III E). Fourth, we demonstrate generalization: with no retraining, the same workflow produces band structures for a panel of conventional superconductors spanning elemental free-electron and transition metals and a transition-metal nitride (Sec. III F). Fifth, we sketch a correlated extension, promoting the same neural Hamiltonian to a Hubbard model solved by dynamical mean-field theory, recovering known correlation trends across the group-5 metals and strong quasiparticle renormalization in the cuprate La_2CuO_4, and identifying the resulting impurity problem as the natural target for a quantum solver (Sec. III G).

The remainder of the paper is organized as follows. Section II describes the SlaKoNet model, the Hermitian solver, the VQD ansatz, and the hardware-execution pipeline. Section III presents single-k-point benchmarks, full band structures, hardware results, multi-material generalization, and the correlated DMFT extension. Section IV discusses what this enables for high-throughput quantum materials screening and the limitations of the current implementation. Section V closes with an outlook on phonons, the path to an ML-accelerated quantum dynamical mean-field theory, and fault-tolerant quantum band-structure calculations.

II. METHODS

The complete workflow, from a crystal structure to validated quantum eigenvalues, is summarized in Fig. 1 and detailed in the subsections below.

A. The SlaKoNet universal Slater-Koster framework

SlaKoNet [19] is a parameter optimization framework that produces a non-orthogonal tight-binding (TB) Hamiltonian and overlap in the Slater-Koster (SK) two-center form [14, 15] for an arbitrary periodic crystal containing elements with atomic number $Z \leq 65$. In contrast to graph-neural-network approaches that predict Hamiltonian matrix elements as a function of the local environment, SlaKoNet directly parametrizes the onsite energies and hopping and overlap integrals themselves. The model learns them end-to-end via automatic differentiation against DFT-generated tuples of crystals and their calculated bandgaps and densities of states.

For a minimal sp^3 basis, the on-site energies $\varepsilon_s, \varepsilon_p$ and the SK hopping integrals $V_{ss\sigma}, V_{sp\sigma}, V_{pp\sigma}, V_{pp\pi}$ (with corresponding overlap integrals) are stored as one-dimensional vectors on a uniform interatomic-distance grid extending to 7 Å about each atom. For elements requiring d orbitals, the basis is extended by $\varepsilon_d, V_{sd\sigma}, V_{pd\sigma}, V_{pd\pi}, V_{dd\sigma}, V_{dd\pi}, V_{dd\delta}$ together with their overlap counterparts. The subscripts for the onsite energies ε indicate which orbital is being represented, the first two subscripts of the SK hopping parameters V which two-orbital interaction is being represented, and the third subscript on V indicates whether a σ or π bond is being represented. These distance-tabulated vectors are stored as trainable `nn.ParameterDict` entries, one per element pair (Z_I, Z_J), and are interpolated at runtime by local polynomial fitting over eight neighboring grid points, with polynomial tail functions enforcing smooth decay to zero beyond the cutoff. The two-body parameter dictionary is initialized from the three-body tight-binding (TB3Py) parametrization of Garrity and Choudhary [20] (originally fit to PBEsol DFT) and is then refined by gradient descent against the higher-fidelity TBmBJ [21] dataset.

For a crystal with atoms I, J at positions $\mathbf{R}_I, \mathbf{R}_J$ and atomic species Z_I, Z_J, the Slater-Koster formalism yields real-space Hamiltonian and overlap matrix elements

$$H_{I\alpha,J\beta}(\mathbf{R}) \;=\; h_{\alpha\beta}(\mathbf{R}_J + \mathbf{R} - \mathbf{R}_I,\, Z_I, Z_J)\,, \qquad (1)$$

where α, β label angular characters of the atomic orbitals, $\mathbf{r}$ runs over lattice vectors within the cutoff. The Bloch

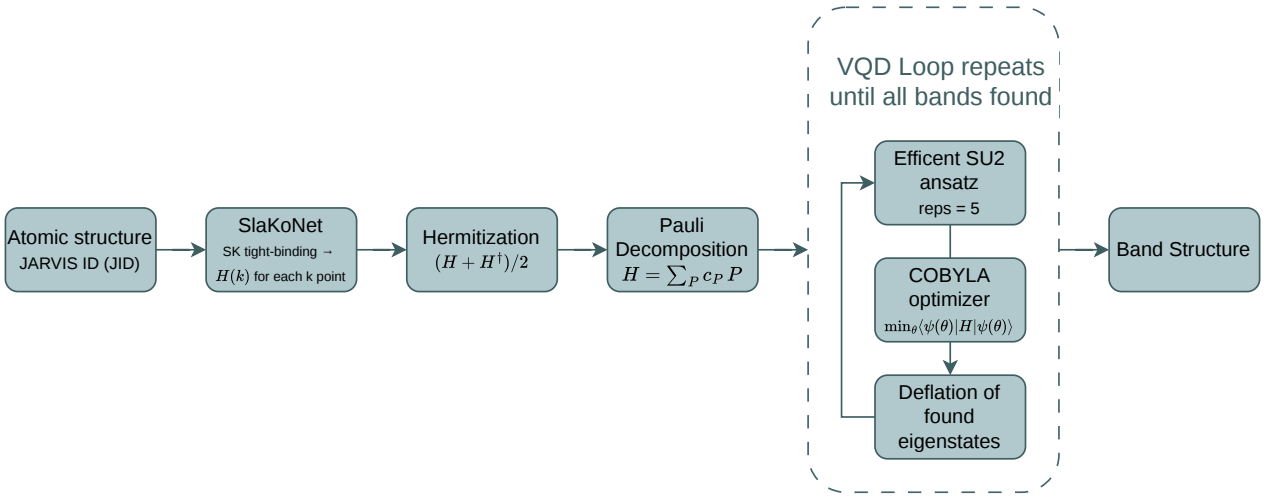

FIG. 1. End-to-end SlaKoNet-VQD pipeline to determine full band structure. For a given material, SlaKoNet takes the atomic structure from JARVIS and produces the k-resolved tight-binding Hamiltonian $H(k)$ for each k-point, each of which is Hermitized via $(H + H^\dagger)/2$ then decomposed into the Pauli basis. The VQD loop (dashed box) is applied iteratively at each k-point. After each eigenstate is found, it is deflated from the spectrum so that the next application returns the following eigenvalue, repeating until all bands are recovered. The resulting eigenvalues are validated against exact diagonalization of the same Hamiltonian.

sum

$$H_{\alpha\beta}(\mathbf{k}) \;=\; \sum_{\mathbf{R}} e^{i\mathbf{k}\cdot\mathbf{R}}\, H_{I\alpha,J\beta}(\mathbf{R}) \tag{2}$$

yields the $\mathbf{k}$-resolved Hamiltonian $H(\mathbf{k})$ used as input to the quantum algorithm. Finally, the orbital overlap matrix in reciprocal space is computed in the same fashion.

SlaKoNet is trained on the JARVIS-DFT database [22] using the Tran-Blaha modified Becke-Johnson (TBmBJ) meta-GGA functional [21], which is known to yield significantly improved bandgaps over standard GGA functionals while remaining computationally tractable for the high-throughput setting. The training set comprises approximately 20,000 materials with $Z \leq 65$, spanning oxides, carbides, nitrides, chalcogenides, halides, and intermetallics. The training objective combines a density-of-states reconstruction loss with a bandgap loss (rather than direct Hamiltonian regression); optimization uses AdamW with learning-rate scheduling and gradient clipping, with early stopping at 200 epochs. On a held-out set of 52 semiconductors and insulators, SlaKoNet attains a mean absolute error of 0.74 eV against experimental bandgaps, compared with 1.14 eV for OptB88vdW DFT and 0.38 eV for the underlying TBmBJ reference [19]. The current release does not include spin-orbit coupling or repulsive pair potentials, limitations that propagate directly to the present work and are addressed in Sec. IV B. For all results below we use the publicly released `slakonet_v0` weights without fine-tuning.

B. From neural Hamiltonian to qubit operator

For each $\mathbf{k}$ along the high-symmetry path, SlaKoNet returns a real symmetric matrix $H(\mathbf{k}) \in \mathbb{R}^{N\times N}$, where N is the number of orbitals in the unit cell ($N = 8$ for the silicon sp^3 basis used in Sec. III). In finite precision the GPU tensor occasionally breaks matrix symmetry at the 10^{-9} level due to non-symmetric accumulation of the Bloch sum in Eq. (2). We restore matrix symmetry using a symmetric projection

$$H(\mathbf{k}) \;\leftarrow\; \tfrac{1}{2}\big[H(\mathbf{k}) + H(\mathbf{k})^\dagger\big]\,, \tag{3}$$

which we have verified introduces a per-element deviation below 10^{-7} relative to the raw output and is essential for the Qiskit `SparsePauliOp` decomposition to succeed.

The symmetrized matrix is decomposed into a sum of weighted Pauli strings,

$$H(\mathbf{k}) = \sum_{P\in\{I,X,Y,Z\}^{\otimes n}} c_P(\mathbf{k})\, P, \tag{4}$$

with $n = \lceil \log_2 N \rceil$ qubits ($n = 3$ for the silicon basis). The decomposition is performed analytically via tensor-product Pauli projectors implemented in the `HermitianSolver` interface of the Jarvis-tools/Qiskit module [22].

C. Variational quantum deflation

Excited states of $H(\mathbf{k})$ are obtained by VQD [3]; the ground state corresponds to the zero-deflation limit of

VQD and is equivalent to a standard VQE run, the form we use for the single-**k**-point and hardware benchmarks below. Starting from a parameterized ansatz $|\psi(\boldsymbol{\theta})\rangle$, the ground state minimizes

$$E_0(\boldsymbol{\theta}_0) = \min_{\boldsymbol{\theta}} \langle\psi(\boldsymbol{\theta})| H(\mathbf{k}) |\psi(\boldsymbol{\theta})\rangle . \tag{5}$$

For the j-th excited state, the cost function is augmented with overlap penalties against the previously obtained eigenstates,

$$F_j(\boldsymbol{\theta}_j) = \langle\psi(\boldsymbol{\theta}_j)| H(\mathbf{k}) |\psi(\boldsymbol{\theta}_j)\rangle + \sum_{i<j} \beta_i \left| \langle\psi(\boldsymbol{\theta}_i)|\psi(\boldsymbol{\theta}_j)\rangle \right|^2, \tag{6}$$

with β_i chosen large enough to enforce orthogonality but small enough to leave the eigenvalue spectrum undistorted. We use a common value $\beta_i = 4$ for all deflated states; the recovered eigenvalues are empirically robust to this choice over a broad range for the eight-band silicon problem.

D. Ansatz and classical optimizer

Six hardware-efficient [2] ansatz topologies are implemented in `QuantumCircuitLibrary`; we report results with `circuit6` (`EfficientSU2`) at five repetitions (`reps`=5), which consists of 36 trainable parameters arranged as alternating R_Y/R_Z single-qubit rotation layers on all 3 qubits, interlaced with a linear entanglement chain of CNOT gates ($q_0 \to q_1$, $q_1 \to q_2$), as shown in Fig. 2. The classical outer loop uses COBYLA [24] with a maximum of 400 iterations following [8], which found it to be the fastest among common gradient-free optimizers for related TB Hamiltonians; an ansatz and optimizer ablation is reported in Sec. III C.

E. Brillouin-zone sampling

The high-symmetry **k**-path is generated using the Setyawan-Curtarolo convention [23] as implemented in Jarvis-tools [22]. For silicon we use a line density of 10, yielding 161 **k**-points along the path $\Gamma-X-W-K-\Gamma-L-U-W-L-K|U-X$. SlaKoNet evaluates $H(\mathbf{k})$ for the entire path in a single batched GPU call.

F. Hardware execution

Hardware runs use IBM Quantum [29] via the `qiskit-ibm-runtime Estimator` primitive on the `ibm_boston` backend, accessed through the IBM Quantum CHM232 instance. The ansatz is transpiled to the device basis gate set $\{\sqrt{X}, X, R_z, CZ\}$ at optimization level 2, with dynamical decoupling and TREX error mitigation enabled. The `Estimator` returns each energy as an expectation value together with its ensemble standard error, providing shot-noise error bars on the reported eigenvalues. To bound runtime cost we restrict hardware execution to one representative **k**-point and report both raw and mitigated estimates.

TABLE I. Ground-state energy of the silicon SlaKoNet Hamiltonian at the L point, $\mathbf{k} = (0.5, 0.5, 0.5)$, on a 3-qubit statevector simulator with the `circuit6` ansatz at `reps` = 1. Energies are reported relative to the Fermi level E_F obtained from SlaKoNet.

Method	$E - E_\mathrm{F}$ (eV)	Relative error
Exact diagonalization	-0.8477	N/A
VQE (COBYLA, statevector)	-0.8398	1.10×10^{-2}

G. Software stack and reproducibility

The full pipeline is built on Jarvis-tools [22] for crystal structure handling and **k**-path generation, SlaKoNet for neural Hamiltonian evaluation, and Qiskit [28] (`qiskit`≥`2.0`, `qiskit-aer`, `qiskit-algorithms`, `qiskit-ibm-runtime`) for the variational quantum algorithm. Simulator runs use the noiseless `statevector_simulator` backend. Random seeds, package versions, and exact ansatz parameters are documented in the companion repository at `https://github.com/atomgptlab/atomqc`.

III. RESULTS

A. Single-k-point benchmark

We first validate the pipeline at a single **k**-point. SlaKoNet evaluates the silicon sp^3 Hamiltonian $H(\mathbf{k})$ at the default $\mathbf{k} = (0.5, 0.5, 0.5)$ on the GPU; the resulting 8×8 Hermitian matrix is mapped to a 3-qubit operator following Sec. II B. The ansatz `circuit6` with `reps` = 1 is optimized by COBYLA on a noiseless statevector simulator.

Table I compares the VQE ground-state energy to the exact eigenvalue of the same matrix obtained by direct diagonalization. The relative error is $|E_\mathrm{VQE} - E_\mathrm{exact}|/|E_\mathrm{VQE}| \approx 1.10 \times 10^{-2}$, well within the regime where the VQD overlap penalty in Eq. (6) can resolve excited states without distortion.

B. Silicon band structure across the Brillouin zone

We next compute the full silicon band structure by running VQD at every **k**-point along the high-symmetry path described in Sec. II E, comprising 161 **k**-points. At each **k**, all eight eigenvalues of the 8×8 Hamiltonian are extracted by sequential VQD with the overlap penalty in Eq. (6).

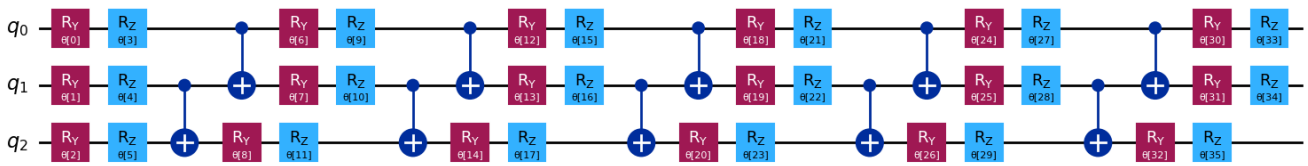

FIG. 2. EfficientSU2 ansatz circuit (reps=5) used for VQD band structure calculations on the 3-qubit SlaKoNet Hamiltonian of Si (JVASP-1002). The circuit consists of alternating layers of parameterized RY and RZ single-qubit rotation gates and CNOT entangling gates, yielding 36 variational parameters. Circuit diagram generated using Qiskit.

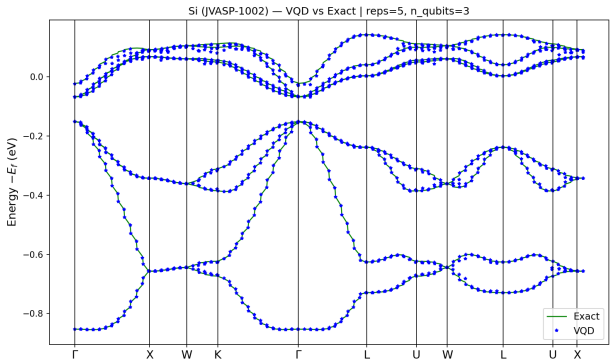

FIG. 3. Silicon band structure from SlaKoNet-VQD (blue markers) compared to exact diagonalization of the same neural Hamiltonian (green lines), along the standard high-symmetry path. The Fermi level E_F obtained from SlaKoNet is shown as the dashed red line. All eight bands are reproduced from a 3-qubit ansatz with `reps` = 5, 400 COBYLA iterations per state, noiseless statevector backend.

Figure 3 overlays the VQD eigenvalues (markers) on the exact band structure (solid lines) obtained by direct diagonalization of the same SlaKoNet Hamiltonian. The eight bands are reproduced across the full path, including the indirect gap between the valence-band maximum at Γ and the conduction-band minimum near X. Quantitatively, the mean absolute error across all bands and all **k**-points is 1.78 meV, with the largest deviations of 29.32 meV occurring on the highest excited states where the deflation penalty has been applied seven times in succession.

C. Convergence and ansatz sensitivity

Having established accuracy at fixed ansatz settings, we now examine how the VQD accuracy depends on the choice of circuit ansatz, its depth, and the classical optimizer. Using the silicon Hamiltonian (`JVASP-1002`) as a fixed target, we sweep all six ansatz topologies in `QuantumCircuitLibrary` over repetitions `reps` = 1 to 10 and three optimizers, COBYLA, SPSA, and L_BFGS_B, reporting in each case the MAE of the VQD eigenvalues relative to exact diagonalization.

Figure 4 shows the depth dependence. Circuit 6 (`EfficientSU2`) reaches the 40 meV chemical-accuracy threshold under all three optimizers, with L_BFGS_B converging fastest in the number of repetitions. Circuit 3 attains a comparable final MAE with fewer variational parameters at equivalent depth, although its convergence with increasing `reps` is slower. Figure 5 fixes the depth at `reps` = 5 and compares all six circuits across the three optimizers. Circuits 3 and 6 are the only ansatze that fall below the chemical-accuracy threshold under every optimizer, reaching MAE several orders of magnitude below it ($\sim 10^{-4}$ meV) under COBYLA and L_BFGS_B, whereas circuits 1 and 2 fail to reach chemical accuracy under any optimizer. These sweeps motivate the `circuit6`/COBYLA combination used for the band-structure results above, which balances accuracy, parameter count, and robustness across **k**-points.

D. Comparison to a DFT+Wannier baseline

To place the SlaKoNet Hamiltonian on the same footing as prior quantum band-structure work [6, 8, 9], we compare it to a DFT+Wannier90 [12] reference for silicon drawn from the JARVIS Wannier tight-binding database [22]. A direct element-wise comparison of the two Hamiltonians is complicated by the differing bases, the Wannier model uses a disentangled valence-plus-conduction set while SlaKoNet employs a non-orthogonal sp^3d^5 basis, so the band manifolds do not correspond one-to-one. Comparing basis-independent features, the SlaKoNet and Wannier band structures agree on the silicon valence bandwidth (11.3 versus 11.7 eV) and on the binding energy of the deep s-derived band at Γ (to within ~ 0.4 eV), confirming that SlaKoNet captures the gross valence electronic structure without any Wannierization. A fully basis-matched, element-wise Hamiltonian comparison, which requires regenerating the Wannier functions with sp^3d^5 projections matched to SlaKoNet, is a target for future quantitative validation.

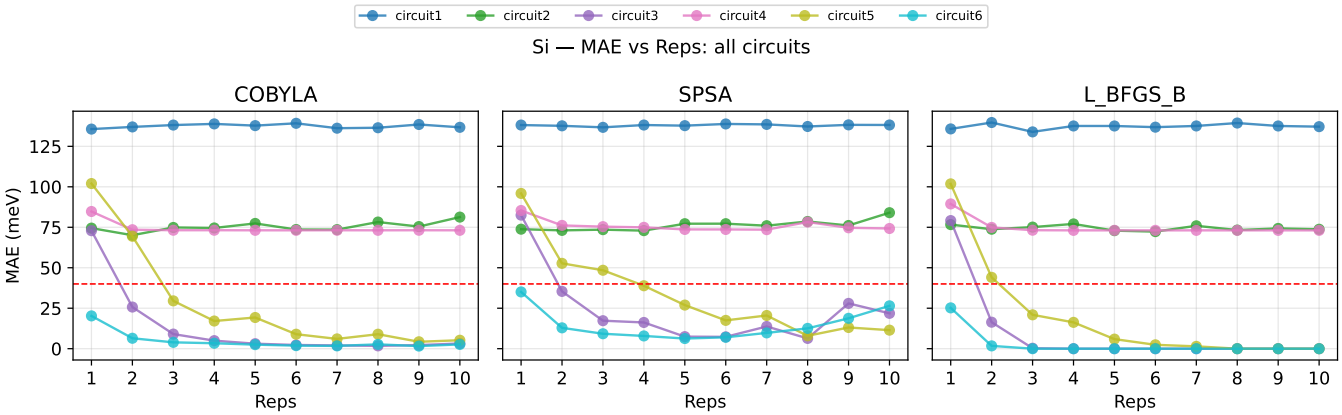

FIG. 4. Mean absolute error (MAE, meV) of VQD eigenvalues relative to exact diagonalization (NumPy) of the SlaKoNet Hamiltonian for Si (JVASP-1002), as a function of ansatz repetitions (reps=1 to 10) for six circuit ansatzes and three optimizers: COBYLA, SPSA, and L_BFGS_B. The dashed red line indicates the 40 meV chemical accuracy threshold. Circuit 6 (EfficientSU2) achieves sub-threshold accuracy under all three optimizers, with L_BFGS_B converging most rapidly. Circuit 3 reaches comparable final MAE values despite requiring fewer variational parameters at equivalent reps, though convergence is slower.

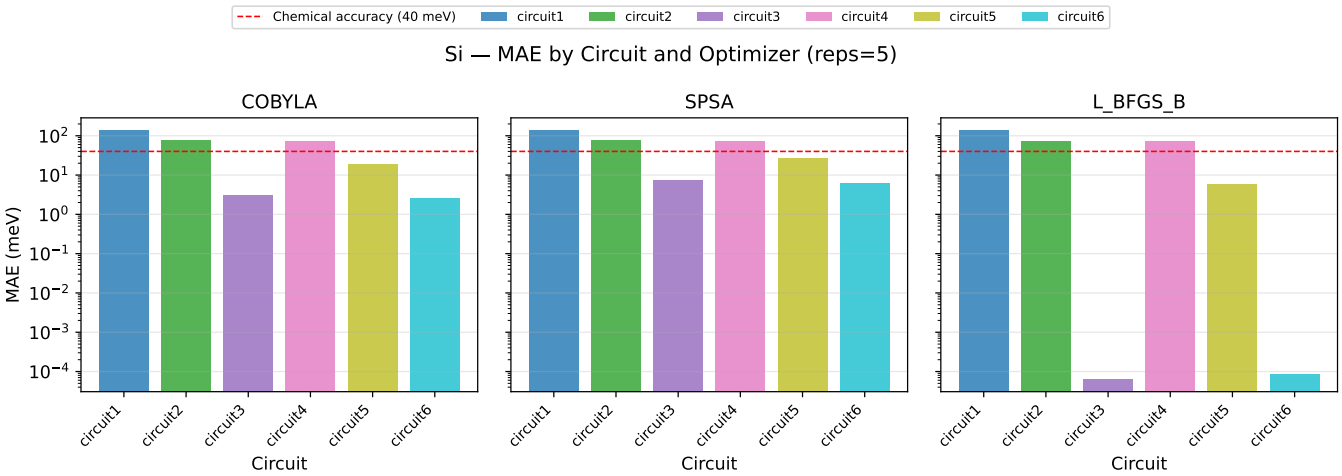

FIG. 5. MAE (meV) for all six circuit ansatzes at reps=5 across three optimizers, for Si (JVASP-1002). The dashed red line indicates the 40 meV chemical accuracy threshold. Circuits 3 and 6 (EfficientSU2) achieve sub-threshold accuracy under all optimizers. While circuit 6 slightly underperforms circuit 3 on L_BFGS_B, it outperforms circuit 3 on COBYLA. Both reach MAE values several orders of magnitude below the chemical accuracy threshold ($\sim 10^{-4}$ meV) under COBYLA and L_BFGS_B. Circuits 1 and 2 fail to reach chemical accuracy under any optimizer.

E. Execution on IBM Quantum hardware

We next move the same VQD calculation to physical hardware via the IBM Quantum runtime [29]. To keep queue and shot budgets manageable we select $\mathbf{k} = (0.5, 0.5, 0.5)$ for aluminum and run VQE on `ibm_boston` with 20480 shots per energy estimate. The transpiled circuit has a single two-qubit CZ gate after optimization level 2. The job (IBM Quantum workload `d7fjjt21u7fs739m3geg`) consumed 25 minutes of quantum-processor time and evaluated the energy at 50 ansatz-parameter settings along the optimization trajectory.

Figure 6 shows the 50 energy evaluations of the VQE optimization executed on `ibm_boston`, each carrying a shot-noise (ensemble) standard error of ≈ 0.06 eV returned directly by the `Estimator`. The best variational estimate reached on hardware is $\langle H \rangle = 9.17 \pm 0.06$ eV, reported here on the absolute energy scale including the identity term of the Pauli decomposition; the ≈ 0.5 eV spread across evaluations reflects the COBYLA search on a noisy device rather than the statistical uncertainty of any single estimate. Referred to the exact ground state of the same Hamiltonian, the error-mitigated estimate deviates by approximately 0.37 eV, comparable to typical DFT-functional differences and larger than chemical ac-

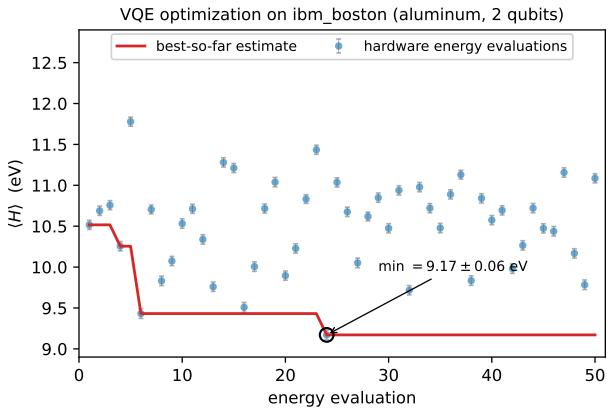

FIG. 6. Aluminum SlaKoNet-VQE on IBM Quantum hardware (`ibm_boston`, workload `d7fjjt21u7fs739m3geg`) at the representative **k**-point $(0.5, 0.5, 0.5)$. Blue markers are the 50 hardware energy evaluations of the variational optimization with their ensemble (shot-noise) standard errors; the red curve is the best-so-far variational estimate, reaching 9.17 ± 0.06 eV. Energies are absolute expectation values including the identity term of the Pauli decomposition.

curacy. This run exercises only the ground state on hardware; resolving the deflated excited-state spectrum on a physical device, where the overlap penalty compounds gate noise, remains more demanding and is left to future work.

F. Generalization across the periodic table

The strength of the SlaKoNet-VQD pipeline is that the same workflow applies to any composition covered by SlaKoNet's training distribution, with no retraining or per-material setup. We illustrate this on a panel of conventional (BCS) superconductors drawn from the Jarvis-DFT database, spanning elemental free-electron and transition metals and a transition-metal nitride. Aluminum (`JVASP-816`) is a face-centered-cubic free-electron superconductor ($T_c \approx 1.2$ K) that tests behavior at the Fermi level and is also the system used for hardware execution (Sec. III E). Vanadium (`JVASP-1041`), tantalum (`JVASP-1014`), and niobium (`JVASP-934`) are body-centered-cubic transition-metal superconductors with $3d$, $5d$, and $4d$ valence manifolds respectively ($T_c \approx 5.4$, 4.5, and 9.3 K, the last being the highest elemental critical temperature). Zirconium nitride (`JVASP-19679`) is a rock-salt transition-metal nitride superconductor ($T_c \approx 10$ K) that tests transferability to a binary compound with mixed metallic and covalent bonding.

Across all five materials the qubit count varies from $n = 2$ to $n = 3$ depending on the orbital basis size, with the resulting band structures shown in Fig. 7. Mean absolute errors per band are summarized in Table II and remain below 1.8 meV for every material tested. The elemental metals (Al, Ta, V, Nb), represented in a two-qubit basis, are reproduced essentially exactly, with per-band MAEs at or below 0.05 meV; the bands crossing the Fermi level in these free-electron and transition-metal superconductors are therefore resolved more accurately than the wider-energy-window states of gapped silicon. The three-qubit rock-salt nitride ZrN reaches 1.74 meV, comparable to silicon (1.78 meV) and consistent with the larger Hilbert space and denser spectrum that the sequential deflation must climb.

TABLE II. Generalization of SlaKoNet-VQD across material classes. All runs use the `circuit6` ansatz at `reps` = 5, COBYLA optimizer, and statevector backend.

Material	Class	n_{qubits}	MAE per band (meV)
Al	metal	2	0.05
Ta	metal	2	0.02
V	metal	2	0.01
Nb	metal	2	0.02
Si	semiconductor	3	1.78
ZrN	metal nitride	3	1.74

G. Toward correlated electronic structure: SlaKoNet + DMFT and a quantum impurity solver

The benchmarks above all concern the *single-particle* band structure, for which exact diagonalization of $H(\mathbf{k})$ is trivial and the variational solver claims no advantage over a classical eigensolver. The scientifically interesting and classically hard regime for the superconductors studied here is the *many-body* problem obtained once electron-electron interactions are restored. We therefore promote the SlaKoNet Hamiltonian to a correlated lattice model by adding a local Hubbard repulsion on the transition-metal d manifold,

$$\hat{H} = \sum_{\mathbf{k},\,\alpha\beta,\,\sigma} H_{\alpha\beta}(\mathbf{k})\, c^{\dagger}_{\mathbf{k}\alpha\sigma} c_{\mathbf{k}\beta\sigma} + U \sum_{i} \hat{n}_{i\uparrow}\hat{n}_{i\downarrow}, \tag{7}$$

and treat it within dynamical mean-field theory (DMFT) [30]. DMFT maps the lattice problem onto a self-consistent Anderson impurity model whose local Green's function reproduces that of the lattice,

$$\begin{aligned} G(\mathbf{k},\omega) &= \left[\omega + \mu - H(\mathbf{k}) - \Sigma(\omega)\right]^{-1}, \\ G_{\text{loc}}(\omega) &= \frac{1}{N_k}\sum_{\mathbf{k}} G(\mathbf{k},\omega), \end{aligned} \tag{8}$$

following the lattice-Green's-function construction of Choudhary [6]. We take the non-interacting density of states directly from the SlaKoNet eigenvalue grid, with no Wannierization, close the DMFT loop on the real axis with an iterated-perturbation-theory (IPT) impurity solver, and extract the interacting spectral function $A(\omega) = -\pi^{-1}\mathrm{Im}\, G_{\text{loc}}(\omega)$ and the quasiparticle weight

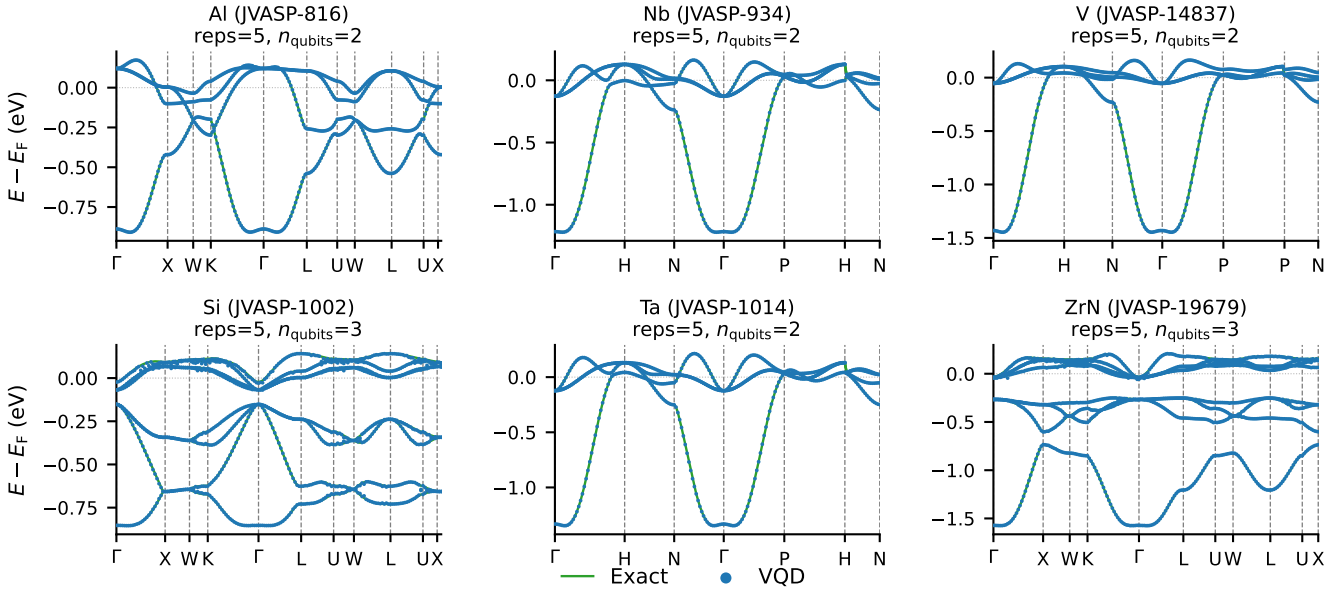

FIG. 7. SlaKoNet-VQD band structures for a panel of conventional superconductors (the elemental metals Al, Ta, V, and Nb and the transition-metal nitride ZrN). In each panel, VQD eigenvalues (markers) are overlaid on exact diagonalization of the same SlaKoNet Hamiltonian (lines). All panels use the same ansatz and optimizer; only the input crystal structure changes.

$Z = [1 - \partial_\omega \mathrm{Re}\,\Sigma(\omega)|_0]^{-1}$. IPT is quantitatively controlled near particle-hole symmetry (half filling), so the quasiparticle weights reported below are intended to capture correlation *trends* rather than converged, filling-accurate values; the cuprate case treated below, a single half-filled correlated band, is the regime in which IPT is most reliable.

As the Hubbard U is increased, the SlaKoNet-derived spectral function of vanadium develops a renormalized quasiparticle peak at E_F flanked by incipient lower and upper Hubbard sidebands, and the quasiparticle weight is progressively reduced from the non-interacting value $Z = 1$ to $Z = 0.65$ at $U = 6$ eV as spectral weight transfers away from the Fermi level. None of this physics is present in the bare band structure of Sec. III F; it is generated entirely by the self-energy $\Sigma(\omega)$. Tracking $Z(U)$ across the isoelectronic group-5 sequence V ($3d$), Nb ($4d$), Ta ($5d$) shows that correlations weaken monotonically down the column as the d orbitals become more extended, so for a fixed U the heavier elements retain a larger Z ($Z = 0.65$, 0.70, 0.74 at $U = 6$ eV for V, Nb, Ta), in agreement with the established chemical trend and with the larger bandwidths SlaKoNet assigns to the $4d$ and $5d$ metals. The free-electron control Al stays essentially unrenormalized ($Z = 0.96$), and the low-$N(E_F)$ nitride ZrN is only weakly correlated; both remain metallic over the full U range, as expected for these weak-to-moderate-coupling superconductors. Table III collects the Fermi-level density of states $N(E_F)$, which sets the bare (BCS) pairing strength, together with Z at representative U.

TABLE III. SlaKoNet + DMFT(IPT) summary for the superconductor panel: Fermi-level density of states $N(E_F)$ (states/eV per cell) from the bare SlaKoNet bands, and the quasiparticle weight Z from the interacting solver at $U = 2$ and 4 eV. $Z \to 1$ is weakly correlated; smaller Z indicates stronger correlation.

Material	orbital	$N(E_F)$	$Z\,(U{=}2)$	$Z\,(U{=}4)$
V	$3d$	0.79	0.95	0.81
Nb	$4d$	0.74	0.96	0.84
Ta	$5d$	0.85	0.96	0.87
ZrN	nitride	0.38	0.99	0.97
Al	sp	0.23	0.99	0.98

a. A strongly correlated cuprate. The elemental metals above are weakly to moderately correlated. As a more demanding test we apply the same pipeline to La_2CuO_4, the parent compound of the first high-T_c cuprate superconductor and a textbook strongly correlated material. SlaKoNet returns a metallic band structure for the seven-atom La_2CuO_4 cell (consistent with the density-functional starting point, which misses the Mott gap), with a Cu-d derived band crossing the Fermi level. Applying the Hubbard interaction to the *full* multi-band density of states leaves the quasiparticle weight essentially unrenormalized ($Z > 0.97$ up to $U = 4$ eV), because the interaction is diluted across the many O-p, La, and Cu states. Restricting the DMFT bath to the low-energy manifold that crosses the Fermi level, a proxy for the single-band Cu-$d_{x^2-y^2}$ downfolding used in rigorous treatments [31], recovers strong corre-

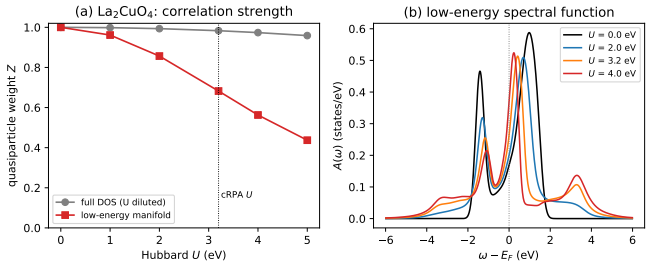

FIG. 8. SlaKoNet + DMFT(IPT) for the cuprate La_2CuO_4 (`JVASP-7907`). (a) Quasiparticle weight $Z(U)$: applied to the full multi-band density of states the Hubbard interaction is diluted and Z stays near unity (gray); restricting the DMFT bath to the low-energy manifold crossing E_F (red), a proxy for the single-band Cu-d downfolding, recovers strong correlation, with $Z = 0.69$ at the constrained-RPA value $U = 3.2$ eV (dotted line). (b) Corresponding low-energy spectral function: as U increases the quasiparticle peak at E_F is suppressed and weight transfers to incipient Hubbard bands. Energies are referenced to E_F.

lation: the quasiparticle weight falls to $Z = 0.69$ at the constrained-RPA value $U = 3.2$ eV and to $Z = 0.57$ at $U = 4$ eV (Fig. 8), with spectral weight transferring from the Fermi level toward incipient Hubbard bands. This manifold restriction overestimates Z relative to a full orbital-projected downfolding (which yields $Z \approx 0.27$ for a related cuprate [31]), so our values are best read as an upper bound on the correlation strength rather than a quantitative prediction; a rigorous Wannier downfolding of the SlaKoNet Hamiltonian is left for future work. The methodological point stands: SlaKoNet supplies the differentiable non-interacting Hamiltonian for a strongly correlated cuprate with no DFT or Wannierization, leaving only the interaction strength U to be supplied (here taken from the constrained-RPA literature value).

b. Why this is the natural target for the quantum solver. Within DMFT the lattice self-consistency in Eq. (**??**) is cheap; the entire computational cost, and the exponential scaling in the number of correlated orbitals, lives in the impurity solver, which must compute the ground- and excited-state Green's function of an interacting fermionic Hamiltonian. This is exactly the object a variational quantum algorithm is suited to, and it has recently been demonstrated on hardware: Selisko *et al.* [31] solve the DMFT impurity problem for the cuprate superconductor $Ca_2CuO_2Cl_2$ on IBM Quantum devices with up to 14 qubits, using VQE for the ground state and a quantum equation-of-motion expansion for the impurity Green's function, and recover the quasiparticle weight to within 2.2% of exact diagonalization and ARPES. Earlier proposals established the hybrid quantum-classical DMFT loop on small systems [32–35]. Crucially, in all of these the impurity model is built from a *per-material* DFT+Wannier+cRPA front end. The contribution of the present pipeline is to replace that front end: SlaKoNet supplies the material-specific, differentiable $H(\mathbf{k})$ and its non-interacting density of states for any composition in a single forward pass, feeding the same DMFT self-consistency and the same quantum impurity solver without DFT, Wannierization, or constrained-RPA. The IPT results in Fig. 8 are the classical reference against which such a quantum impurity solver is validated, in the same role that exact diagonalization and continuous-time quantum Monte Carlo play in Ref. [31].

IV. DISCUSSION

A. What this enables

The SlaKoNet-VQD pipeline removes a structural bottleneck that has limited quantum band-structure studies to one material at a time. Three concrete consequences follow.

a. High-throughput VQA benchmarking. Because SlaKoNet emits $H(\mathbf{k})$ in milliseconds and requires no per-material human input, it becomes feasible to run VQE/VQD on thousands of Jarvis-DFT entries in a single sweep. The closest prior effort [6] reported 307 Wannier-derived TB Hamiltonians as test inputs; with SlaKoNet-VQD the same scale is reachable without any Wannierization, opening a route to systematic studies of how ansatz expressibility, qubit count, and material chemistry interact. The superconductor panel in Sec. III F is a first step in this direction: a single ansatz and optimizer setting transfers across five chemically distinct systems at two and three qubits with no per-material tuning.

b. Differentiable end-to-end pipeline. SlaKoNet is implemented as a PyTorch `nn.Module`; its output $H(\mathbf{k})$ carries gradients with respect to atomic positions, lattice parameters, and the trainable SK integrals themselves. Combined with parameter-shift gradients of the VQD cost function [26], this exposes the entire chain from crystal structure to variational quantum eigenvalue as a single differentiable graph. Concrete consequences include gradient-based co-design of ansatz and Hamiltonian for hardware-aware compilation, and direct optimization of crystal structure under a quantum-evaluated objective. A concrete numerical demonstration of this end-to-end co-optimization is left to future work.

c. A clean substrate for algorithmic research. The SlaKoNet Hamiltonian is reproducible to floating-point precision, has fixed structure across all materials, and ships with the model weights. This makes it a natural benchmark target for new variational algorithms (subspace-search VQE [10, 25], ADAPT-VQE, quantum filter diagonalization), in much the same role that the FCIDUMP format plays in molecular quantum-chemistry quantum-algorithm research.

B. Limitations

We take pains to be explicit about what this pipeline does and does not deliver in its current form.

a. Accuracy is bounded by SlaKoNet, not by VQD. The mean absolute errors reported in Sec. III reflect two distinct sources of deviation: (i) the variational solver's departure from exact diagonalization of the SlaKoNet Hamiltonian, which is well below 10 meV in our setting, and (ii) the deviation of SlaKoNet itself from the underlying TBmBJ reference, which on a 52-material benchmark is reported as 0.74 eV in mean absolute bandgap error [19], improving substantially over OptB88vdW DFT (1.14 eV) but trailing the TBmBJ upper bound of 0.38 eV. The dominant contribution to absolute electronic-structure accuracy is therefore (ii), not the quantum solver. Reported errors on ultrawide-gap insulators (> 6 eV) are particularly large (~ 2.3 eV [19]), and our pipeline inherits this behavior.

b. Basis, spin-orbit coupling, and repulsive potentials. SlaKoNet employs a non-orthogonal Slater-Koster basis with sp^3 orbitals for light elements and an extended sp^3d^5 basis where d orbitals are required, but the released model does not include spin-orbit coupling and does not provide optimized repulsive pair potentials [19]. Heavy-element systems where SOC is essential (e.g., Bi_2Se_3, Pb halide perovskites) therefore lie outside the present demonstration, and total-energy or structural quantities that rely on the repulsive part of the potential are not accessible. These are model limitations, not pipeline limitations; an SOC extension to SlaKoNet, or a paired repulsive potential, would slot in unchanged.

c. VQD overlap-penalty scaling. The cost in Eq. (6) grows linearly in the number of previously found states, and excited-state fidelity degrades as the spectrum is climbed. For the $N = 8$ silicon Hamiltonian we observe 30.3 times larger errors on the highest band than on the ground state. Larger orbital bases (e.g., transition-metal d shells) will compound this; subspace-search variants such as SSVQE [25] or filter-diagonalization approaches may be preferable as the system grows.

d. Hardware noise. Even with error mitigation, current superconducting-qubit hardware introduces deviations of ~ 372 meV on the aluminum eigenvalues, comparable to typical DFT-functional differences and larger than chemical accuracy. This is an architectural property of present-day devices, not a flaw of the pipeline; the same $H(\mathbf{k})$ inputs will benefit automatically from improvements in coherence times, gate fidelities, and mitigation protocols.

C. Relation to prior quantum band-structure work

The conceptual closest neighbor to this work is Choudhary [6], which applied VQE/VQD to 307 Wannier tight-binding Hamiltonians for solids; the present paper retains the variational solver and replaces the Wannierization step with a universal, differentiable Slater-Koster TB framework. Sherbert *et al.* [7] and the recent Sci. Reports study [8] pursue similar DFT+Wannier inputs. The IOP study [9] applies the VQD algorithm to crystalline band structures along similar lines.

The GaAs-focused work of Miháliková *et al.* [10] is the closest methodological sibling: it also applies VQD (together with SSVQE) to the tight-binding band structure of a solid and studies ansatz and optimizer choice, paralleling our convergence analysis (Sec. III C). It differs from the present work in three respects that sharpen our contribution. First, its Hamiltonian is an empirical, hand-fitted sp^3s^* parametrization specific to GaAs, whereas SlaKoNet supplies the Hamiltonian for any composition with no per-material fitting. Second, it assigns one qubit per orbital, requiring 10 qubits for the ten-orbital cell, whereas our compact Pauli decomposition of the $N \times N$ matrix uses $\lceil \log_2 N \rceil$ qubits (3 for the eight-band silicon basis); the two encodings trade qubit count against the particle-conserving gate structure theirs exploits. Third, its results are confined to a noiseless simulator, while we additionally execute on IBM Quantum hardware and extend the workflow to a panel of superconductors and to the correlated DMFT regime (Sec. III E and Sec. III G).

The very recent measurement-protocol work (arXiv:2511.04389), which reduces the Pauli-measurement count for TB Hamiltonians, is complementary: it improves the quantum side of the same problem we attack from the Hamiltonian-construction side, and the two optimizations stack.

V. CONCLUSION AND OUTLOOK

We have demonstrated that a pretrained universal Slater-Koster tight-binding framework, SlaKoNet, can serve as the Hamiltonian source for variational quantum band-structure calculations, removing the per-material DFT or Wannierization step that has constrained prior work in this space. On silicon, the resulting SlaKoNet-VQD pipeline reproduces all eight bands along the standard high-symmetry path with a mean absolute error of 1.78 meV on a 3-qubit statevector simulator. A representative ground-state calculation on aluminum runs end-to-end on IBM Quantum hardware, deviating by approximately 0.37 eV from exact diagonalization after error mitigation. The same workflow extends without modification across a panel of conventional superconductors spanning elemental free-electron and transition metals and a transition-metal nitride.

Two natural extensions point the way forward. First, SlaKoNet can be generalized to emit dynamical matrices in addition to electronic Hamiltonians; feeding these into the same VQD machinery yields phonon band structures, in the spirit of the original Wannier-based phonon work of Choudhary [6], but again without the finite-difference setup cost per material. Sec-

ond, the correlated extension sketched in Sec. III G points to a fully ML-accelerated dynamical mean-field theory on quantum hardware. Here, SlaKoNet supplies the differentiable $H(\mathbf{k})$ that defines the DMFT bath, while the impurity Green's function is evaluated by the same variational machinery, replacing the per-material DFT+Wannier+cRPA construction used in recent quantum-DMFT hardware experiments [31]. Because the impurity problem carries the entire many-body cost, this is where a quantum solver is expected to matter for the correlated superconductors studied here.

Looking further out, removing the Hamiltonian-construction bottleneck has implications beyond the variational era. Once fault-tolerant resources permit quantum phase estimation [27] on solids, the input to QPE will still need to be a structurally faithful second-quantized Hamiltonian; a universal neural model that can produce one for any composition is at least as valuable in that regime as in this one. The pipeline demonstrated here is a concrete step toward that capability.

DATA AND CODE AVAILABILITY

The SlaKoNet model weights and training data are available through the AtomGPT ecosystem at `https://github.com/atomgptlab/slakonet`. The Qiskit-based `HermitianSolver` and `QuantumCircuitLibrary` interfaces used here are distributed with Jarvis-tools [22]. A reproducibility repository containing the notebooks, configuration files, and scripts that generate every figure in this work will be deposited at `https://github.com/atomgptlab/atomqc` upon publication. The IBM Quantum hardware run reported in Sec. III E is archived under workload `d7fjjt21u7fs739m3geg` and can be retrieved through the `qiskit-ibm-runtime` service.

FUNDING DECLARATION

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

ACKNOWLEDGMENTS

This research used resources of the Oak Ridge Leadership Computing Facility (OLCF), a U.S. Department of Energy Office of Science User Facility at Oak Ridge National Laboratory, through allocation CHM217, "Cross-Platform Benchmarking of Quantum Hardware for Materials Discovery and Design." We acknowledge the use of IBM Quantum services for the hardware experiments reported in this work; the views expressed are those of the authors and do not reflect the official policy or position of IBM or the IBM Quantum team.

[1] A. Peruzzo *et al.*, *A variational eigenvalue solver on a photonic quantum processor*, Nat. Commun. **5**, 4213 (2014).

[2] A. Kandala *et al.*, *Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets*, Nature **549**, 242 (2017).

[3] O. Higgott, D. Wang, and S. Brierley, *Variational quantum computation of excited states*, Quantum **3**, 156 (2019).

[4] Y. Cao *et al.*, *Quantum chemistry in the age of quantum computing*, Chem. Rev. **119**, 10856 (2019).

[5] S. McArdle *et al.*, *Quantum computational chemistry*, Rev. Mod. Phys. **92**, 015003 (2020).

[6] K. Choudhary, *Quantum computation for predicting electron and phonon properties of solids*, J. Phys.: Condens. Matter **33**, 385501 (2021).

[7] K. Sherbert, F. T. Cerasoli, and M. Buongiorno Nardelli, *A systematic variational approach to band theory in a quantum computer*, RSC Adv. **11**, 39438 (2021).

[8] N. Khandelwal, N. Verma, P. Jamdagni *et al.*, *Quantum computation of the electronic structure of some prototype solids*, Sci. Rep. **15**, 44074 (2025); doi:10.1038/s41598-025-27675-6.

[9] M. Ďuriška, I. Miháliková, and M. Friák, *Quantum computing of the electronic structure of crystals by the variational quantum deflation algorithm*, Phys. Scr. **100**, 045111 (2025); doi:10.1088/1402-4896/adbb29.

[10] I. Miháliková, M. Krejčí, and M. Friák, *The impact of quantum circuit architecture and hyperparameters on variational quantum algorithms exemplified in the electronic structure of the GaAs crystal*, Sci. Rep. **15**, 15746 (2025); doi:10.1038/s41598-025-00151-x.

[11] N. Marzari, A. A. Mostofi, J. R. Yates, I. Souza, and D. Vanderbilt, *Maximally localized Wannier functions: Theory and applications*, Rev. Mod. Phys. **84**, 1419 (2012).

[12] G. Pizzi *et al.*, *Wannier90 as a community code: new features and applications*, J. Phys.: Condens. Matter **32**, 165902 (2020).

[13] I. Souza, N. Marzari, and D. Vanderbilt, *Maximally localized Wannier functions for entangled energy bands*, Phys. Rev. B **65**, 035109 (2001).

[14] J. C. Slater and G. F. Koster, *Simplified LCAO method for the periodic potential problem*, Phys. Rev. **94**, 1498 (1954).

[15] D. Porezag, T. Frauenheim, T. Köhler, G. Seifert, and R. Kaschner, *Construction of tight-binding-like potentials on the basis of density-functional theory: Application to carbon*, Phys. Rev. B **51**, 12947 (1995).

[16] H. Li *et al.*, *Deep-learning density functional theory Hamiltonian for efficient ab initio electronic-structure calculation*, Nat. Comput. Sci. **2**, 367 (2022).

[17] Y. Zhong, H. Yu, X. Gong, and H. Xiang, *Transferable equivariant graph neural networks for the Hamilto-*

nians of molecules and solids, npj Comput. Mater. **9**, 182 (2023).

[18] O. T. Unke *et al.*, *SE(3)-equivariant prediction of molecular wavefunctions and electronic densities*, Adv. Neural Inf. Process. Syst. **34**, 14434 (2021).

[19] K. Choudhary, *SlaKoNet: A unified Slater-Koster tight-binding framework using neural network infrastructure for the periodic table*, J. Phys. Chem. Lett. **16**, 11109 (2025); doi:10.1021/acs.jpclett.5c02456.

[20] K. F. Garrity and K. Choudhary, *Fast and accurate prediction of material properties with three-body tight-binding model for the periodic table*, Phys. Rev. Mater. **7**, 044603 (2023).

[21] K. Choudhary, Q. Zhang, A. C. E. Reid, S. Chowdhury, N. Van Nguyen, Z. Trautt, M. W. Newrock, F. Y. Congo, and F. Tavazza, *Computational screening of high-performance optoelectronic materials using OptB88vdW and TB-mBJ formalisms*, Sci. Data **5**, 180082 (2018).

[22] K. Choudhary *et al.*, *The joint automated repository for various integrated simulations (JARVIS) for data-driven materials design*, npj Comput. Mater. **6**, 173 (2020).

[23] W. Setyawan and S. Curtarolo, *High-throughput electronic band structure calculations: Challenges and tools*, Comput. Mater. Sci. **49**, 299 (2010).

[24] M. J. D. Powell, *A direct search optimization method that models the objective and constraint functions by linear interpolation*, in *Advances in Optimization and Numerical Analysis* (Springer, 1994), pp. 51-67.

[25] K. M. Nakanishi, K. Mitarai, and K. Fujii, *Subspace-search variational quantum eigensolver for excited states*, Phys. Rev. Research **1**, 033062 (2019).

[26] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, *Evaluating analytic gradients on quantum hardware*, Phys. Rev. A **99**, 032331 (2019).

[27] A. Yu. Kitaev, *Quantum measurements and the Abelian stabilizer problem*, arXiv:quant-ph/9511026 (1995).

[28] A. Javadi-Abhari *et al.*, *Quantum computing with Qiskit*, arXiv:2405.08810 (2024); doi:10.48550/arXiv.2405.08810.

[29] IBM Quantum, https://quantum.ibm.com, accessed September 20, 2026.

[30] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, *Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions*, Rev. Mod. Phys. **68**, 13 (1996).

[31] J. Selisko, M. Amsler, C. Wever, Y. Kawashima, G. Samsonidze, R. Ul Haq, F. Tacchino, I. Tavernelli, and T. Eckl, *Dynamical mean field theory for real materials on a quantum computer*, npj Comput. Mater. **11**, 325 (2025); doi:10.1038/s41524-025-01772-6.

[32] B. Bauer, D. Wecker, A. J. Millis, M. B. Hastings, and M. Troyer, *Hybrid quantum-classical approach to correlated materials*, Phys. Rev. X **6**, 031045 (2016).

[33] J. M. Kreula, S. R. Clark, and D. Jaksch, *Few-qubit quantum-classical simulation of strongly correlated lattice fermions*, EPJ Quantum Technol. **3**, 11 (2016).

[34] I. Rungger *et al.*, *Dynamical mean field theory algorithm and experiment on quantum computers*, arXiv:1910.04735 (2019).

[35] T. Keen, T. Maier, S. Johnston, and P. Lougovski, *Quantum-classical simulation of two-site dynamical mean-field theory on noisy quantum hardware*, Quantum Sci. Technol. **5**, 035001 (2020).