VIBE-QC: an open-source, general-purpose electronic-structure code for molecules and solids developed with generative AI

Michael F. Peintinger*

Mulliken Center for Theoretical Chemistry, University of Bonn, Bonn, Germany

Abstract

We present VIBE-QC, a general-purpose open-source electronic-structure code for molecules and solids, developed almost entirely through generative-AI-assisted programming ("vibe-coding") under expert human direction. The molecular stack provides Hartree–Fock and Kohn–Sham DFT over the full LIBXC library, analytic gradients and Hessians, Møller–Plesset theory and double hybrids, coupled cluster with local-correlation variants, multireference methods, excited states, dispersion, continuum solvation, effective core potentials, and semiempirical methods. The periodic stack treats one-, two-, and three-dimensional systems at multiple $\boldsymbol{k}$ through three Fock-build routes. Production claims are validated by out-of-process cross-code comparisons against established reference programs, reaching the nanohartree scale for mean-field methods and $3.3 \times 10^{-7}\,E_\mathrm{h}$ for conventional CCSD(T); external parity for periodic total energies is open work at this release, tracked in the public issue tracker. VIBE-QC adds a container-based results-storage and visualization format (QVF), an interactive viewer (VIBE-VIEW), and a job queue (VIBE-QUEUE), under the MPL-2.0 license.

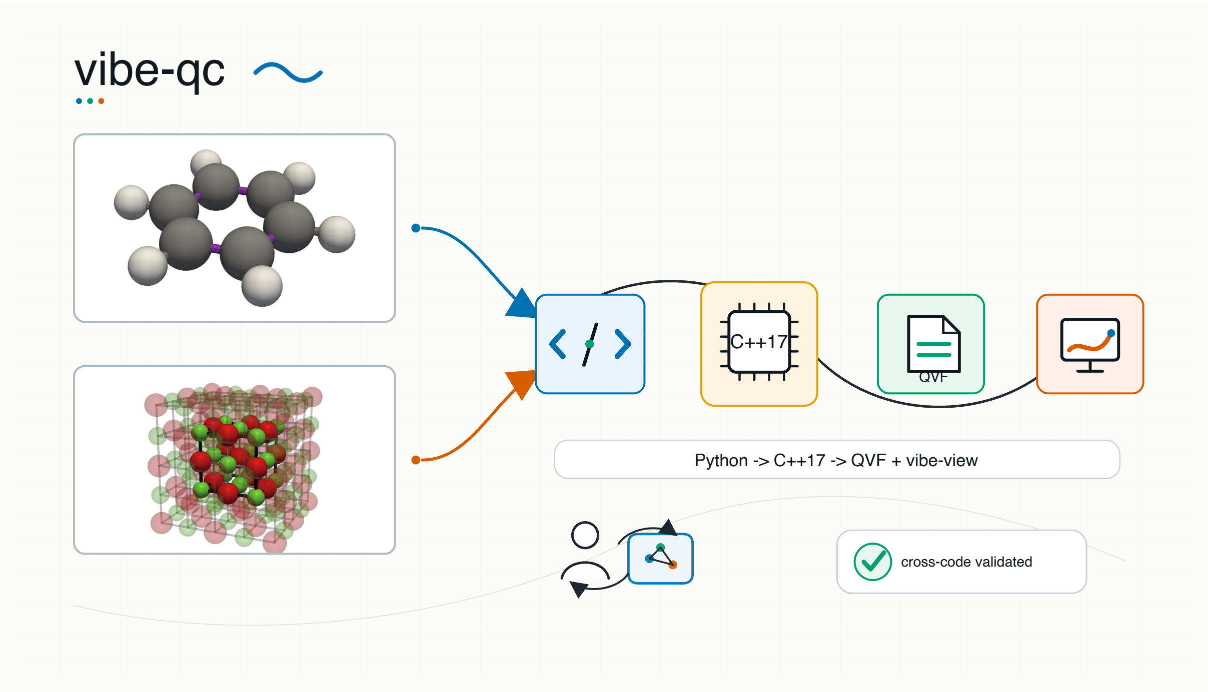

*Corresponding author. Emails: mpei@vibe-qc.com, mpei@uni-bonn.de

Graphical TOC summary. vibe-qc is an open-source, general-purpose electronic-structure code for molecules and solids that was built almost entirely through generative-AI-assisted programming under expert direction, and validated capability-by-capability where production status is claimed. It spans Hartree–Fock through coupled cluster, multireference, and excited states for molecules, and three periodic Fock-build routes for solids, and ships its own modern file format (qvf), viewer (vibe-view), and cluster-aware queue (vibe-queue).

Keywords. quantum chemistry, electronic-structure software, periodic systems, generative-AI-assisted development, vibe-coding.

1 Introduction

vibe-qc began as an experiment to discover how far a current large language model could be pushed, under expert direction, toward a working quantum-chemistry program. An initial prototype, reported on the project blog as a one-day effort [1], performed cross-code-validated Hartree–Fock; development continued from there, and the code reported in this paper is the result. The name records the method: vibe-qc is "vibe-coding" applied to quantum chemistry, and the development methodology, generative-AI programming under human scientific control, is the lens through which we present it.

We make one claim about that methodology and defend it throughout: a code produced this way can be held to exactly the same cross-code validation discipline as any conventionally written quantum-chemistry package, and meeting that discipline is what makes the resulting numerics trustworthy. Accordingly the release distinguishes production-validated capabilities from implemented but still-gated features, and backs the production claims with out-of-process reference runs (Section 4). The development workflow is documented in the open: working sessions were logged on the project blog [1–3], and the documentation includes a chapter on driving the code with a language model [4]. Section 6.1 returns to the methodology and names the specific models used.

What the experiment produced is a general-purpose code. vibe-qc is not a single-method demonstrator but a broad electronic-structure package that treats molecules and solids on the same footing, in the spirit of the established general-purpose suites Orca [5], PySCF [6], Crystal [7, 8], GPAW [9], and OpenMolcas [10, 11]. For molecules it spans Hartree–Fock, the full Kohn–Sham DFT ladder, Møller–Plesset theory and double hybrids, canonical coupled cluster and pilot local-correlation routes, complete-active-space and other multireference methods, and excited states. For solids it provides Gaussian-basis periodic HF and DFT in one, two, and three dimensions through three complementary Fock-build routes, with space-group symmetry, effective core potentials, and bonding analysis. The breadth is itself a product of the development method: generative-AI assistance makes the many parsers, drivers, and file-format writers that a general-purpose code needs cheap enough for one person to build and maintain. The code rests on the standard open-source numerical libraries of the field (libint, libxc, spglib, libecpint, and FFTW), reads and writes the established quantum-chemistry file formats, and adds a modern container-based, general-purpose quantum-chemistry results-storage and visualization format (qvf), an interactive viewer (vibe-view), and a cluster-aware job queue (vibe-queue).

The end goal is a feature-complete, rigorously validated general-purpose electronic-structure code that treats molecules and solids on an equal footing, up to and including correlated periodic methods. A detailed development roadmap was laid out at the start and was revised several times as capabilities were completed earlier than planned; the version history records that pace. Each release carries a three-part number `A.B.C`, in which `A` is the generation, `B` the minor release, and `C` the patch level. The minor releases, their dates, and the foundational work each principal feature builds on are listed in Table 4.

This paper describes VIBE-QC at the 0.15 release series. Because the series is under active development, each reported result is pinned by the `.system` manifest archived with it rather than by a version string alone. Section 2 catalogs the implemented theory and notes where analytic gradients are available. Section 3 describes the C++17 backend, the `pybind11` frontend, the file-format and visualization surface, and the documentation, which doubles as a self-contained course in quantum chemistry. Section 4 reports the per-method cross-code validation, each method against the most appropriate of five reference programs: ORCA for molecular methods, CRYSTAL 23 for the periodic bipolar-expansion route, PySCF for periodic Gaussian density fitting, GPAW for the Gaussian-and-plane-waves route, and OPENMOLCAS for complete-active-space methods. Section 5 demonstrates representative calculations, and Section 6 returns to the development methodology, the release history, the scope of the claims, and the outlook.

2 Theory and methods

This section catalogs the theory VIBE-QC implements in the 0.15 release series. Standard methods are summarized briefly; the description weights toward the parts particular to VIBE-QC, namely the three periodic Fock-build routes, the multi-$\boldsymbol{k}$ machinery, and the periodic exchange–correlation and symmetry infrastructure. Every production claim made here is assigned either an archived cross-code parity run in Section 4 or an explicit gated status. The status vocabulary used in this manuscript is summarized in Table 1. The development methodology is described in Section 6.1. All integrals, functionals, symmetry, ECP integrals, and fast transforms are provided by the standard open-source libraries of the field, LIBINT, LIBXC, SPGLIB, LIBECPINT, and FFTW respectively, which VIBE-QC orchestrates through its own C++17 and Python code.

The methods implemented in the 0.15 series are, in summary:

- **Mean-field SCF:** RHF, UHF, RKS, UKS, ROHF, and ROKS, for molecules and for one-, two-, and three-dimensional periodic systems (two-dimensional slabs currently restricted to RHF, RKS, and UKS), with SAD, core-Hamiltonian, Hückel, minimal-basis, density-restart, broken-symmetry, and fragment (FRAGMO) initial guesses and a convergence stack of DIIS, EDIIS, level shifting, and a second-order fallback.
- **Density-functional theory:** the full LIBXC ladder: LDA, GGA, τ-meta-GGA (including r^{2}SCAN and r^{2}SCAN01), global hybrids, and range-separated hybrids (ωB97X, ωB97X-V, ωB97M-V) with the VV10 nonlocal correlation kernel.
- **Perturbation theory and double hybrids:** MP2, RI-MP2, SCS-MP2, and SOS-MP2; B2PLYP, DSD-PBEP86, PWPB95, and revDSD-PBEP86-D4.
- **Coupled cluster:** CCSD and CCSD(T) with RHF, UHF, and ROHF references and frozen-natural-

Table 1: Method-maturity vocabulary used for the 0.15 series release paper. A feature is promoted only as far as the archived evidence supports.

Status	Meaning	Examples in this release line
`production`	Supported for release use, exercised on production-sized inputs, and matched against an out-of-process reference where one exists.	Core molecular SCF, DFT, and molecular ECP routes.
`verified`	Carries targeted cross-code evidence on specific cases, but not yet across enough of its input surface to be called production.	Selected correlated and property routes.
`under-review`	Implemented and runnable, but a convention or a reference comparison is unresolved, so no quantitative claim is made.	Pilot local-correlation, multireference, BIPOLE, periodic-ECP, COOP/COHP, and periodic DFT+U routes.
`experimental`	Available in the tree but excluded from validation and from every quantitative claim in this paper.	GFN2-xTB and other isolated research paths.

orbital acceleration, and the local DLPNO-MP2, DLPNO-CCSD, and DLPNO-CCSD(T) hierarchy.

- **Multireference and non-mean-field:** CASCI, CASSCF (state-averaged), NEVPT2, and CASPT2, with full configuration interaction, selected CI, and experimental DMRG active-space solvers.
- **Excited states and propagators:** time-dependent DFT (full Casida/RPA and the Tamm–Dancoff approximation) and CIS, with natural transition orbitals and UV/Vis spectra, and GF2/OVGF ionization energies.
- **Corrections:** DFT-D3(BJ) and DFT-D4 dispersion, composite hf-3c and r^2SCAN-3c methods, CPCM and COSMO continuum solvation, and molecular and periodic effective core potentials.
- **Semiempirical and machine-learning:** MSINDO (including a periodic cyclic-cluster treatment at the semiempirical level), GFN2-xTB, PM6, the OMx family, and density-functional tight binding, plus the MACE machine-learning interatomic potential.
- **Periodic Fock builds:** Gaussian density fitting (GDF, $\boldsymbol{\Gamma}$ and multi-$\boldsymbol{k}$, with mixed density fitting), a bipolar-expansion route (BIPOLE; multi-$\boldsymbol{k}$, all spin cases), and a Gaussian-and-plane-waves route (GPW), with route-specific periodic ECP and DFT+U support.
- **Analysis and properties:** Mulliken, Löwdin, and Hirshfeld charges, Mayer bond orders (molecular and periodic), COOP/COHP, QTAIM topology, band structures, fat bands, total and projected densities of states, and Fermi surfaces.
- **Potential-energy surfaces:** geometry optimization, nudged-elastic-band searches, molecular dynamics, metadynamics, and minimum-energy conical-intersection optimization.

Implementation origin is recorded separately from maturity. In particular, an optional external engine can be production, verified, under-review, or experimental; the origin label does not itself assert validation status.

Analytic nuclear gradients are available for RHF, UHF, RKS, UKS, and ROHF, for single-state CASSCF, for single-state closed-shell IC-CASPT2 (unshifted, on the small-space engine), for the periodic GDF and bipolar-expansion routes (the latter at multiple $\boldsymbol{k}$ and including fractional occupations), and for GPW at $\boldsymbol{\Gamma}$; the RHF Hessian is analytic. NEVPT2, the remaining CASPT2 variants, and periodic meta-GGA forces are finite-difference on the same driver surface. The remainder of this section describes the parts particular to VIBE-QC, weighting toward the three periodic Fock-build routes and the multi-$\boldsymbol{k}$, symmetry, and exchange–correlation infrastructure.

2.1 Molecular self-consistent field

Restricted (RHF, RKS), unrestricted (UHF, UKS), and restricted-open-shell (ROHF, ROKS) references are dispatched from a common SCF driver solving the Roothaan–Hall [12], Pople–Nesbet [13], and open-shell [14] equations. In the AO basis the closed-shell Fock matrix is

$$F_{\mu\nu} = h_{\mu\nu} + \sum_{\lambda\sigma} D_{\lambda\sigma} \left[(\mu\nu|\lambda\sigma) - \tfrac{1}{2}(\mu\sigma|\lambda\nu) \right], \tag{1}$$

with the two-electron repulsion integrals (ERIs) evaluated through LIBINT; for KS-DFT the exact-exchange term is scaled by the hybrid fraction and the exchange–correlation potential is added (Section 2.2). The two-electron build is selected from a polymorphic family: a conventional four-index build, a resolution-of-the-identity build (RIJK), and the chain-of-spheres RIJCOSX build [15] for exact exchange on larger systems.

Convergence is accelerated by Pulay DIIS [16, 17] on the commutator residual $\boldsymbol{e} = \boldsymbol{FDS} - \boldsymbol{SDF}$, backed by EDIIS [18], a Saunders–Hillier level shift [19], and a second-order fallback. Initial guesses include superposition of atomic densities (SAD, the default), core Hamiltonian, Hückel, projected minimal-basis, density restart (READ), broken-symmetry spin-flip guesses, and a fragment molecular-orbital (FRAGMO) guess that assembles a block-diagonal density from converged fragments. Near-linear dependence of the AO basis is removed by canonical orthogonalization of the overlap below a configurable eigenvalue threshold.

2.2 Exchange–correlation functionals

All functionals are evaluated through LIBXC [20] on atom-centered Becke fuzzy-Voronoi grids [21] with Treutler–Ahlrichs radial [22] and Lebedev angular [23] quadrature. The full ladder is available: LDA; GGAs (PBE [24], BLYP); global hybrids (PBE0 [25], B3LYP [26, 27]); τ-dependent meta-GGAs including TPSS, SCAN, r^2SCAN [28] and its constraint-restored variant r^2SCAN01 [29]; range-separated hybrids (ωB97X [30]); and the VV10 nonlocal correlation kernel [31], which completes ωB97X-V and the meta-GGA hybrid ωB97M-V [32]. The B3LYP keyword maps to the VWN5 convention; the legacy VWN-RPA form is reachable explicitly.

2.3 Analytic gradients and Hessians

Analytic nuclear gradients are implemented for HF and KS-DFT across the LDA, GGA, hybrid, and τ-meta-GGA rungs. For a closed-shell reference,

$$\frac{\partial E}{\partial \boldsymbol{R}_A} = \sum_{\mu\nu} D_{\mu\nu} \frac{\partial h_{\mu\nu}}{\partial \boldsymbol{R}_A} + \frac{1}{2} \sum_{\mu\nu\lambda\sigma} D_{\mu\nu} D_{\lambda\sigma} \frac{\partial(\mu\nu\|\lambda\sigma)}{\partial \boldsymbol{R}_A} - \sum_{\mu\nu} W_{\mu\nu} \frac{\partial S_{\mu\nu}}{\partial \boldsymbol{R}_A} + \frac{\partial E_{\text{nuc}}}{\partial \boldsymbol{R}_A}, \tag{2}$$

with W the energy-weighted density and the ∂S term the Pulay force; integral derivatives come from Libint. Analytic gradients are also available for ROHF and, at this release, for single-state CASSCF, verified against the full-energy finite-difference force to tight tolerance and driving all three geometry optimizers (L-BFGS, Brent line search, and the delocalized-internal-coordinate provider). Single-state, unshifted, closed-shell IC-CASPT2 also carries an analytic gradient, evaluated from the coupled orbital/CI Lagrangian and validated against the fully relaxed total-energy finite difference on H_2/6-31G and LiH/STO-3G CAS(2,2) to within $1 \times 10^{-7}\, E_\text{h}\, a_0{}^{-1}$. The other CASPT2 variants (shifted, open-shell, and large active-space cases) and NEVPT2 remain on finite-difference gradients on the same unified surface. The gradient round-trips through the ASE Calculator interface [33] for geometry optimization, which offers rational-function (RFO/P-RFO) [34], GDIIS, and FIRE [35] steppers in delocalized internal coordinates [36]; the RHF analytic Hessian gives harmonic frequencies and IR intensities, and a finite-difference Hessian covers the remaining references.

2.4 Møller–Plesset theory and double hybrids

MP2 [37] is implemented in canonical and resolution-of-the-identity (RI-MP2) forms for closed- and open-shell references, together with the spin-component-scaled (SCS-MP2) [38] and scaled-opposite-spin (SOS-MP2) [39] variants. The double hybrids B2PLYP [40], DSD-PBEP86, PWPB95, and the GMTKN55-trained revDSD-PBEP86-D4 [41] add a perturbative RI-MP2 correlation correction to a hybrid KS reference. The canonical and RI kernels are validated on small molecules; their memory envelope at production scale has not yet been characterized, and no production-scale claim is made in this release.

2.5 Coupled cluster: canonical, local, and frozen-natural-orbital

Coupled-cluster singles and doubles (CCSD) [42] and its perturbative-triples extension CCSD(T) [43] are available with RHF, UHF (UCCSD/UCCSD(T)), and ROHF references through a density-fitted implementation. The 0.15-series local-correlation surface includes the domain-based local pair-natural-orbital hierarchy DLPNO-MP2/DLPNO-CCSD/DLPNO-CCSD(T) [44, 45], with the scaling-preserving iterative-triples correction [46] as the default, within the current pilot size cap. These local-correlation routes are included as pilot implementations in this release; the validation tables do not claim Orca TightPNO parity for them. Open-shell DLPNO-MP2 is exposed through the UHF-reference DLPNO-UMP2 route [47]; open-shell DLPNO-CCSD(T) remains a correctness pilot. Frozen natural orbitals [48] accelerate closed-shell CCSD(T) with a ΔMP2 correction; the no-truncation path reproduces the canonical result. Canonical CCSD(T) remains a dense small-molecule validation route in this release; production-size bounded-memory coverage is not claimed for canonical or density-fitted triples.

2.6 Multireference and non-mean-field methods

The active-space stack provides CASCI and CASSCF [49] (single-state and state-averaged, spin-pure), the n-electron valence perturbation theory NEVPT2 [50], and internally-contracted, multistate, and extended-multistate CASPT2 [51]. For the active-space solver itself VIBE-QC offers exact full configuration interaction (FCI) [52], selected CI of the heat-bath [53] and CIPSI [54] families, and the density-matrix renormalization group [55]; a transcorrelated-Hamiltonian pre-step is also available. Exact-in-basis FCI energies serve as an internal oracle for the approximate solvers.

2.7 Excited states and electron propagators

Excited-state methods are a central capability of this release. Vertical excitation energies are obtained from linear-response time-dependent DFT, in both the full Casida (random-phase-approximation) form and the Tamm–Dancoff approximation [56], for restricted and unrestricted references, and from configuration-interaction singles [57]. The driver prints a UV/Vis spectrum, writes natural transition orbitals [58] (including for unrestricted references) to the QVF archive, and exposes finite-difference excited-state gradients; TD-DFT is reachable both from `run_job`/`run_periodic_job` with `tddft=True` and from a `vibeqc tddft` command-line entry point, and is available for periodic systems as well as molecules. Ionization energies and electron affinities are computed directly from the one-particle Green's function through second-order (GF2) and renormalized OVGF-type propagator methods [59].

2.8 Eigensolvers

The iterative diagonalizations underneath the SCF, the active-space solvers, and the excited-state response are routed through a unified eigensolver framework selected by a `solver=` keyword. The production solver is the Davidson method [60]; a locally optimal block preconditioned conjugate-gradient solver (LOBPCG) [61], a periodic Davidson variant, and a semiempirical solver are also provided, with Jacobi–Davidson and GPLHR present as experimental options.

2.9 Dispersion and composite methods

London dispersion is treated by DFT-D3 with Becke–Johnson damping [62, 63] for molecular and periodic systems, and by the charge-dependent D4 model [64] for molecules (reference `dftd4` backend by default, with a native MPL-2.0 backend including its electronegativity-equilibration charges for H–Ne). The composite "-3c" methods [65, 66] combine a small-basis SCF with a geometrical counterpoise correction [67], a short-range basis-incompleteness term, and dispersion; the counterpoise and short-range terms reproduce the reference `mctc-gcp` implementation to better than $1\,\mathrm{n}E_\mathrm{h}$. Molecular D3(BJ) remains available as an output capability, but no quantitative claim is made for it in this release; the open component-level comparison is tracked in the issue tracker.

2.10 Implicit solvation and effective core potentials

Continuum solvation uses CPCM [68] and COSMO [69] within the polarizable-continuum family [70], with the smooth continuous-surface-charge switching cavity of Scalmani and Frisch [71] and an analytic gradient. ECP integrals are evaluated by LIBECPINT [72] for the Stuttgart/Cologne and def2 families [73] and enter the SCF and gradient paths. A distinguishing feature of this release line is *periodic* ECP support: the CRYSTAL-format inline ECP Gaussian expansions are routed through LIBECPINT, auto-attached for the pob-TZVP-REV2 family, and lattice-summed into the periodic core Hamiltonian, extending the periodic stack to heavy-element (4d/5d) crystals; the corresponding reference comparisons are open work at this release.

2.11 Semiempirical methods and machine-learning potentials

The semiempirical platform spans five method families: DFTB0 and SCC-DFTB [74], GFN2-xTB [75], PM6 [76], the OMx orthogonalization-corrected methods [77], and MSINDO [78] (including a COSMO-solvated variant). Validation status for the semiempirical engines is given in Table 1; several have no direct counterpart among the reference programs of Section 4, and no quantitative semiempirical claim is made in this paper. The NDDO methods PM6 and OMx, molecular-only in their original implementations, are extended to one-, two-, and three-dimensional crystals by real-space lattice summation of the Fock matrix; GFN2-xTB is present but remains gated experimental for quantitative work in this release line; and MSINDO gains a cyclic cluster model with Ewald embedding and analytic nuclear gradients. A dedicated article on the semiempirical platform, including a comparison with machine-learning interatomic potentials, is in preparation; here we record the implemented surface and reserve numerical claims for validated rows.

The pre-trained machine-learning interatomic potential MACE [79] is available as an external engine behind an optional extra: VIBE-QC marshals geometry in and reads energy and forces out, names the model, and never presents the result as its own electronic energy.

2.12 Potential-energy-surface exploration

Beyond single points and optimization, VIBE-QC drives nudged-elastic-band transition-state searches, velocity-Verlet molecular dynamics, well-tempered metadynamics [80, 81], and a minimum-energy conical-intersection optimizer, all on the same energy and gradient surfaces and available with the semiempirical and MLIP engines for affordable exploration.

2.13 Periodic SCF: crystal orbitals and $\boldsymbol{k}$ sampling

VIBE-QC implements the Bloch-summed Gaussian crystal-orbital ansatz [82] for HF and KS-DFT in one, two, and three dimensions. The periodic AO is the lattice sum

$$\phi_{\mu\boldsymbol{k}}(\boldsymbol{r}) = \sum_{\boldsymbol{R}} e^{i\boldsymbol{k}\cdot\boldsymbol{R}}\,\chi_\mu(\boldsymbol{r}-\boldsymbol{R}), \tag{3}$$

and the Fock and density matrices carry a crystal-momentum label. The Brillouin zone is sampled on Monkhorst–Pack [83] or generalized regular [84] meshes, with the irreducible set reduced by SPGLIB [85]; an automatic recommender selects the mesh density, optionally through a gated starter machine-learning predictor following the $\boldsymbol{k}$-spacing model of Choudhary and Tavazza [86]. Metallic occupations use Fermi–Dirac, Methfessel–Paxton [87], or Marzari–Vanderbilt cold smearing [88], or parameter-free Gilat–Raubenheimer tetrahedron integration, with the Mermin free energy as the variational functional. Three Fock-build routes share this infrastructure and are selected by one keyword.

2.14 GDF: periodic Gaussian density fitting

The GDF route fits the periodic density in an auxiliary Gaussian basis using the range-separated Gaussian-density-fitting construction [89]. The three-index integrals are fitted against the two-index lattice metric, Cholesky-decomposed, and contracted with the density to assemble the Hartree and exchange matrices; the $\boldsymbol{k} = 0$ exchange divergence is removed by the Born–von-Kármán Madelung correction. A mixed-density-fitting variant [90] supplements the Gaussian fit with a plane-wave correction for heavy-core systems. GDF is the default route and supports $\boldsymbol{\Gamma}$ and multi-$\boldsymbol{k}$ RHF, RKS, and (with per-spin smearing) UHF and UKS, with analytic gradients. External total-energy parity for the route is open work at this release (Section 4.2).

2.15 Bipolar-expansion Fock build

The second periodic route evaluates the two-electron lattice sum directly in real space under an exact Ewald split of the Coulomb operator,

$$\frac{1}{r_{12}} = \frac{\mathrm{erfc}(\omega r_{12})}{r_{12}} + \frac{\mathrm{erf}(\omega r_{12})}{r_{12}}, \tag{4}$$

whose split point ω is a pure gauge parameter to which the converged energy is invariant. The short-range erfc channel of both the Coulomb and exchange builds is summed exactly over direct-lattice cells; the long-range erf channel is evaluated in reciprocal space together with the neutralizing background, so the reciprocal channel plays the role of the far-field model of the density. This reproduces the bielectronic-zone and monoelectronic-zone partition of the CRYSTAL Coulomb machinery [82, 91] in the Ewald-split gauge. An optional restricted exact zone bounds the expensive exact-ERI zone below the one-electron fold-convergence cutoff, with the omitted tail a measured per-case quantity. A quartet-level bipolar multipole far field is implemented but not certified: it does not preserve the three-translation periodic Fock domain of the exact contraction, and is refused before SCF at this release.

Both periodic Gaussian routes default to the same corrected Ewald electrostatic gauge, so their total energies are directly comparable in that configuration; the corrected multi-$\boldsymbol{k}$ gauge requires a complete Monkhorst–Pack mesh, and a legacy non-unified gauge remains reachable for diagnostics. Route-to-route agreement is reported per case rather than assumed (Section 4.2). The bipolar-expansion route supports multi-$\boldsymbol{k}$ RHF, UHF, RKS, and UKS with space-group-reduced direct ERIs. Production atomic forces are exact finite differences of the converged route energy, and fixed-cell atomic relaxation is available on that path; analytic BIPOLE gradients, including the fractional-occupation (Mermin) free-energy form, are implemented as a research preview pinned against finite differences, and are not the production force path at this release. Variable-cell optimization fails closed pending a certified strain convention. The route is deliberately narrow at its boundaries: configurations whose Hamiltonian it cannot execute exactly, including effective-core-potential bases, periodic force constants, linear response, explicit crystal-orbital population analysis, non-Fermi–Dirac occupation broadening, variable-cell and restart-preserving optimization, low-dimensional cells, and non-Monkhorst–Pack $\boldsymbol{k}$ sets under the corrected gauge, fail before the SCF starts rather than silently substituting a surrogate operator, and an unconverged SCF is a hard failure in both the manifest and the process exit code. It is the route benchmarked against CRYSTAL 23 [8] in Section 4.2.

2.16 GPW: the Gaussian-and-plane-waves route

The third route represents the density on a regular real-space grid and solves the Hartree problem by fast Fourier transform, the Gaussian-and-plane-waves (GPW) scheme [92, 93]. VIBE-QC's GPW path covers $\mathbf{\Gamma}$ RHF/UHF/RKS/UKS and multi-$\boldsymbol{k}$ pure-DFT RKS, with periodic D3(BJ), an analytic $\mathbf{\Gamma}$ gradient, and smearing. The GPW route as a whole carries an experimental runtime gate in the 0.15 series and is excluded from the validation claims of this paper; its open stability defects are tracked in the issue tracker. The Gaussian-augmented-plane-waves variant (GAPW) [94] extends the same route to all-electron treatment by restoring the hard core density near the nuclei. GAPW is not a fourth Fock-build route but an augmentation of this one; it is present at this release but gated experimental.

2.17 Multi-$\boldsymbol{k}$ convergence aids and density mixing

At multiple $\boldsymbol{k}$ the DIIS residual is the Frobenius-norm-weighted concatenation of the per-$\boldsymbol{k}$ commutators. Beyond DIIS and EDIIS [18], the periodic driver offers reciprocal-space density mixing: Anderson/Pulay mixing [95], a limited-memory Broyden second-method mixer, and a Kerker preconditioner [96] for long-wavelength charge sloshing in metals. The mixers operate on the per-$\boldsymbol{k}$ density matrices and preserve the electron count exactly.

2.18 Periodic Becke partition for XC grids

For periodic KS-DFT the molecular Becke partition [21] is extended by partitioning each home-cell grid point among the union of the home-cell atom and its periodic images within a cutoff. The weights sum to unity per grid point and the partition-weighted density integrated over the home cell recovers the electron count per cell, avoiding both the double counting of a naive scheme and the site-symmetry breaking a molecular grid imposes on a high-symmetry cell. The periodic XC quadrature assembles the *full* cross-cell density at each grid point, rather than the home-cell contribution alone; on dense crystals this reduces the real-space XC error by more than two orders of magnitude relative to a home-cell-only assembly (Section 4.2), for both the energy and the basis-set optimization XC gradient.

2.19 Space-group symmetry

SPGLIB [85] identifies the space group; for each operation the AO permutation matrix and the real spherical-harmonic Wigner rotation are built and verified against explicit rotations, and the crystallographic character tables project the AO basis into irreducible representations (SALCs). The lattice sum over cell pairs in the Fock and overlap builds is reduced to one representative per orbit and the rest reconstructed, giving a point-group-order compression. Geometry symmetrization, primitive-cell reduction, and irreducible-$\boldsymbol{k}$ generation use the same machinery.

2.20 DFT+U, bonding analysis, and properties

A rotationally invariant Dudarev DFT+U [97] correction,

$$E_U = \frac{U_{\text{eff}}}{2} \sum_I \operatorname{Tr}\big[\boldsymbol{n}_I(\mathbf{1} - \boldsymbol{n}_I)\big], \tag{5}$$

is available for molecular references and for periodic calculations through route-specific drivers: **Γ** molecular-style references, BIPOLE UHF/UKS, and explicit closed-shell RKS on the GPW route and its GAPW augmentation, including the multi-***k*** GPW RKS branch. Open-shell multi-***k*** GPW +U is not part of this release and fails closed rather than silently switching route; at **Γ** the GPW +U driver also covers UHF and UKS. The periodic DFT+U validation entry records which Fock-build route and +U driver were actually executed. Population analysis (Mulliken, Löwdin, Hirshfeld), Mayer bond orders, dipole moments, and the static polarizability are computed from the converged density. For the validated BIPOLE bulk cases, Mulliken and Löwdin charges and Mayer bond orders are evaluated from the same lattice density and overlap convention as the SCF energy; ordinary electric dipole moments for three-dimensional bulk cells remain explicitly gated until a polarization convention is implemented. Periodic density volumes are evaluated from the density the SCF returned, with the Born–von-Kármán transform verified against every source block and against the primitive-cell electron count before export; a volume that cannot be produced truthfully is omitted rather than approximated. Molecular Löwdin and Mayer values are exposed as output capabilities but remain unvalidated pending a rerun at the released version. The bonding-analysis surface adds crystal-orbital overlap and Hamilton populations (COOP [98] and COHP [99], including spin-polarised), Mayer bond orders for periodic systems [100], and a quantum-theory-of-atoms-in-molecules topological analysis [101] with critical-point search and bond-path tracing from the analytic density Hessian. For solids, VIBE-QC produces band structures along crystallographic ***k*** paths [102], fat (orbital-projected) bands, total and projected densities of states by tetrahedron or Gaussian broadening (with a Gilat–Raubenheimer variant for metals), and Fermi-surface exports.

3 Implementation

3.1 Architecture

VIBE-QC is a two-layer code. The numerical core is C++17: it owns the hot kernels (integral driving through LIBINT, the Fock builds, the lattice sums, the gradient primitives, and the XC evaluation on Becke grids), uses EIGEN [103] for dense linear algebra, and parallelizes with OpenMP. The orchestration layer is Python: it owns job setup, I/O, the property and visualization surface, and the workflow tooling. The two layers meet at a `pybind11` boundary [104] with zero-copy EIGEN↔NumPy transfer. Calculations are configured by Python dataclasses (`RHFOptions`, `UHFOptions`, `RKSOptions`, `UKSOptions`, and their periodic counterparts) passed across the boundary, and are driven by two high-level entry points, `run_job` for molecules and `run_periodic_job` for crystals, which wrap the low-level per-method drivers. Figure 1 sketches the layering and Listing 1 shows the surface.

Listing 1: The high-level API. The same `Molecule` drives every molecular method; the periodic route is selected by one keyword.

```
import vibeqc as vq

# Molecule: a closed-shell DFT single point with dispersion.
mol = vq.Molecule([vq.Atom(8, [0.0, 0.0, 0.0]),
 vq.Atom(1, [0.0, 1.43, -0.98]),
 vq.Atom(1, [0.0, -1.43, -0.98])])
vq.run_job(mol, basis="def2-tzvp", method="rks", functional="pbe0",
 dispersion="d3bj", output_qvf=True)
```

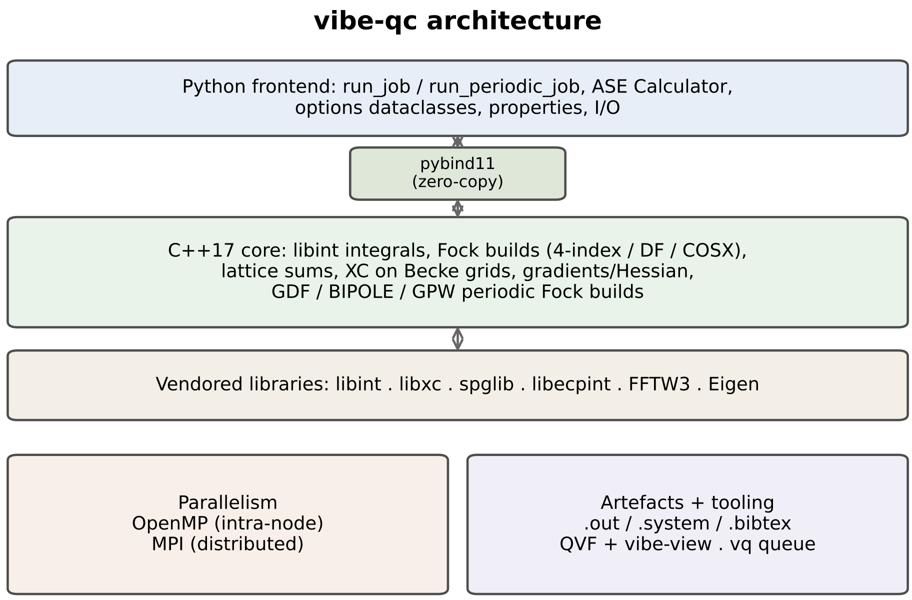

Figure 1: Two-layer architecture: a Python orchestration frontend over a C++17 numerical core meeting at a `pybind11` boundary, with the vendored libraries beneath and the artifact bundle, VIBE-VIEW, and `vq` queue alongside.

```
# Crystal: multi-k periodic HF via Gaussian density fitting.
import numpy as np
a = 7.72
lat = (a/2) * np.array([[0,1,1],[1,0,1],[1,1,0]])
sysp = vq.PeriodicSystem(3, lat, [vq.Atom(3, [0,0,0]),
 vq.Atom(1, [a/2, a/2, a/2])])
basis = vq.BasisSet(sysp.unit_cell_molecule, "sto-3g")
vq.run_periodic_job(sysp, basis, method="RHF",
 jk_method="gdf", kpoints=(2, 2, 2))
```

3.2 Vendored dependencies and the source install

The vendored native dependencies are LIBINT 2.13.1 (Gaussian integrals; LGPL-3.0) [105], LIBXC 7.0.0 (XC functionals; MPL-2.0) [20], SPGLIB (crystal symmetry; BSD-3-Clause) [85], LIBECPINT 1.0.7 (ECP integrals; MIT) [72], and FFTW3 (reciprocal-space transforms; GPL-2.0) [106]; EIGEN and `pybind11` are header-only. The combined binary is effectively GPL-2.0 through the FFTW3 link; an FFT-backend abstraction for commercial-friendly builds is on the roadmap. The supported install is a source clone plus a bootstrap script,

```
git clone https://gitlab.peintinger.com/mpei/vibeqc.git
cd vibeqc && ./scripts/install.sh
```

which accepts `--dev` (track `main`) and `--branch NAME`, and builds every native dependency in tree under `third_party/`. The default `release` branch fast-forwards from each tag, so a bare clone lands on the latest release. There is no binary wheel at the 0.15 series; the source install trades binary-size convenience for build-time reproducibility. The basis library bundles 239 Gaussian basis sets, of which 90 are taken unmodified from the LIBINT distribution and 149 carry curated or custom sources, including our own pob-* solid-state family [107, 108].

3.3 Reproducibility artifacts and automatic citations

Every job writes a small, self-describing artifact bundle. The `.out` file carries the banner, the SCF trace, the energy breakdown, the orbital table with the HOMO–LUMO gap, Mulliken, Löwdin, and Hirshfeld charges, Mayer bond orders, the dipole, and wall-clock timings. The `.system` file is a TOML manifest pinning the host, the build (version and Git SHA, vendored-library versions), and a run fingerprint, so a reported wall time is interpretable and a result is traceable to an exact build. Output writers neutralize bidirectional-override, zero-width, and control characters in cube headers and TOML manifests, closing the Trojan-Source class of attacks (CVE-2021-42574 / CWE-150).

Two artifacts directly serve the citation discipline this journal expects. Each job emits a `.bibtex` and a `.references` file listing every method, functional, basis set, ECP, and linked library it actually used, assembled at run time from a single citation database that is the project's source of truth for what the code claims to cite. A reference block is also appended to the `.out` file. The basis-set and functional citations propagate from the per-file headers in the bundled basis library and from the LIBXC convention, so a user who runs a calculation receives, without extra effort, the list of papers the Methods section of their own manuscript should cite.

3.4 File formats and the QVF format

VIBE-QC reads and writes the established quantum-chemistry file formats, including Molden, extended XYZ, POSCAR, CIF, XSF and BXSF, and Gaussian cube, so that it interoperates with existing structure builders, viewers, and analysis tools. Implementing and maintaining this breadth of parsers and writers is itself a benefit of the development method: generative-AI assistance makes the many small, well-specified input/output paths that a general-purpose code needs cheap to build and to keep correct.

Alongside these, VIBE-QC defines a modern format of its own, and we regard it as a deliberate design improvement over the legacy text formats. Those formats are single-purpose and text-based: each carries one kind of data, so the outputs of one calculation scatter across a cube file, an XYZ file, a Molden file, an XSF file, a log file, and program-specific sidecars, with no shared schema, no typing, and no integrity checks, and many of the associated single-purpose tools and viewers are no longer maintained. VIBE-QC instead defines the QVF (quantum visualization format) [109], a modern container-based, general-purpose quantum-chemistry results-storage and visualization format: a versioned, ZIP-based, random-access archive (v1.2 at this release) that keeps a calculation's structure, scalar fields, spectra, bands, trajectories, provenance, and viewer hints in a single typed, checksummed file. It is designed to be producer-neutral: the format names the writing program in a mandatory provenance field, its section kinds are not VIBE-QC-specific, and a consumer may read only the sections it understands and still open the file. A single archive can bundle the structure and bonds, volumetric data (density, density difference, electrostatic potential, the electron localization function, and the reduced density gradient for non-covalent-interaction analysis), the Gaussian wavefunction, canonical and localized orbitals, natural transition orbitals, band structures and fat (orbital-projected) bands, total and projected densities of states, Fermi surfaces, phonon bands and densities of states, equation-of-state curves, optical/ECD and NMR spectra, vibrational modes, optimization and reaction-path trajectories, bond-order and COOP/COHP tables, QTAIM critical points and bond paths, and the provenance, SCF history, and citation records. It is written by default: `output_qvf` defaults to true on both `run_job` and `run_periodic_job`, so an ordinary calculation produces a QVF archive without being asked.

QVF is a published, versioned, and fully specified format intended for adoption by other quantum-chemistry codes, not just VIBE-QC. The normative specification (version 1.2 at this release) is maintained in the repository at `qvf-writer/spec/qvf-format-spec.md` and published with the documentation; the companion design document at `docs/design_qvf_format.md` records the rationale and the per-section reference. A machine-readable contract is provided as a JSON-schema manifest, shipped both with the code at `python/vibeqc/output/formats/qvf_manifest.schema.json` and with the toolkit described below.

To make adoption practical rather than merely permitted, the reference implementation is distributed as a standalone, self-contained toolkit (`qvf-writer/`) licensed under Apache 2.0, so that it can be vendored into any code including a proprietary one, independently of VIBE-QC's own MPL 2.0 license. It contains the specification and JSON schema, reference writers in Python and in zero-dependency C++17 that carry their own ZIP, CRC-32, SHA-256, and JSON emitters, a C ABI for callers in C and Fortran, a validator that checks schema conformance, member checksums, binary byte sizing, cross-reference resolution, and extension governance, and worked integration examples. Correctness is made testable rather than asserted: a conformance suite ships a frozen corpus of reference archives, each paired with the values a correct implementation must decode from it, together with a language-agnostic runner, so a third-party reader or writer can be certified against the standard instead of merely agreeing with one other implementation.

Extension is governed rather than ad hoc. A producer with data that no canonical section kind covers writes it under a reserved `x_<vendor>` namespace with a declaration in the archive root; consumers report such sections rather than misinterpreting them, and a vendor kind that two independent codes emit with a stable shape becomes a candidate for promotion to a canonical kind through a documented registry process. The manifest integer stays at 1 across the `v1.x` specification line: an optional addition is the additive case and does not bump it, and a proposed bump to 2 for periodic reaction paths was withdrawn by governance ruling in favor of an optional member, so that no consumer is forced to reject an archive it can in fact read.

3.5 The VIBE-VIEW interactive viewer

VIBE-VIEW is the companion interactive three-dimensional viewer for QVF archives, shipped as a separate package behind an optional `[viewer]` extra and versioned independently of VIBE-QC, since it is a consumer of the format rather than part of the code. It is built on `PyVista`, `Trame`, and `VTK`, and presents a local browser interface with no JavaScript build step and no external GPU daemon. Reading one or more QVF files directly, it renders structures and periodic cells, volumetric isosurfaces of the density or wavefunction with clip planes, difference and spin-density maps, atom-resolved charge overlays, band structures and densities of state, optical, vibrational, and NMR spectra, vibrational-mode and reaction-path animations, the SCF-convergence history, and an embedded citation panel. A compare mode overlays

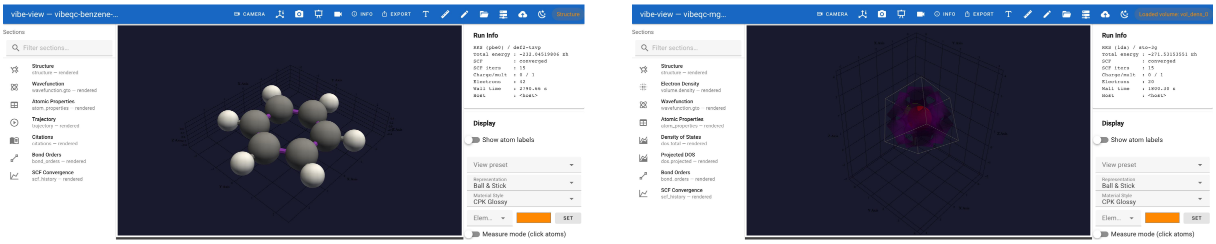

Figure 2: The VIBE-VIEW browser viewer. Left: the optimized structure of benzene (M20, RKS PBE0/def2-TZVP) with CPK spheres and inferred bonds. Right: the electron density of MgO (KRKS LDA/STO-3G, GDF, $k = (2, 2, 2)$) as an isosurface inside the periodic cell. Both are unmodified screenshots captured from archives produced by the release-paper calculation set.

structures from several archives with Kabsch alignment, and biomolecular systems are drawn as cartoon ribbons from the optional chain and secondary-structure metadata that the QVF `structure` section may carry. Figure 2 shows two representative panels; further screenshots and the full feature list are given in the Supporting Information.

3.6 Documentation, examples, and tutorials

VIBE-QC ships an extensive documentation, example, and tutorial corpus at `https://vibe-qc.com`. Beyond the API and user-guide reference, the tutorials are written to be read in sequence as a self-contained introduction to quantum chemistry: they progress from Hartree–Fock and basis sets through DFT, perturbation theory, coupled cluster, and multireference methods to periodic systems, each step pairing the theory with a runnable VIBE-QC input and an interpretation of the output. Producing and maintaining teaching material at this depth alongside the code is, again, made tractable by the development method.

3.7 Memory pre-flight, parallelism, and remote execution

A static pre-flight estimator sizes the AO and MO buffers, the ERI and auxiliary storage, the grid memory, the DIIS history, and the periodic L_{pq} tensors before any allocation, and aborts with an actionable `InsufficientMemoryError` rather than thrashing mid-SCF; the estimate can be overridden. Intra-node parallelism is provided by OpenMP across the heavy kernels and is the production parallel surface of this release. An optional `mpi4py` substrate and low-level collectives exist in the code base, together with an experimental GPW Hartree–J z-slab overlay, but production distributed execution through `mpirun` or MPI for the HF, DFT, or MP2 drivers is not part of this release; campaigns that need more than one node are distributed at the `vq` workflow and queue level. GPU acceleration of the compute kernels remains on the roadmap (the viewer is separately GPU-accelerated). The BLAS/LAPACK backend (Apple Accelerate on macOS, OpenBLAS or MKL on Linux) is selectable at build time. For longer jobs VIBE-QC ships its own job queue, VIBE-QUEUE (invoked as `vq`), a single-host supervisor that dispatches work to a workstation or server over SSH with cgroup-enforced limits on memory, wall time, and CPU starvation, a per-job scratch directory, and a read-only web dashboard, and reads results back through the same artifact bundle. Jobs are described by a declarative, backend-agnostic specification that persists to disk, so a running job survives a daemon or machine restart rather than being lost. On a shared cluster the same specification can be handed to an external batch scheduler rather than run directly, so one submission interface spans a laptop, a workstation, and an HPC system. Shipped scheduler dialects include PBS/Torque and Slurm; both have been exercised end to end on university clusters, submitting and retrieving VIBE-QC and reference-program jobs through the same interface used locally. The queue architecture is described in the Supporting Information.

3.8 Development toolchain

The code was written with heavy generative-AI assistance under direct human supervision; the rationale and the division of labor are the subject of Section 6.1. For the record, the models were used in distinct roles: Anthropic Claude Opus 4.8 [110] (and Opus 4.7 earlier) [111], later joined by Claude Opus 5 and Claude Fable 5 [112, 113], as the main development chats, DeepSeek [114, 115] and later DeepSeek v4 Pro [116] for high-volume routine generation, and OpenAI Codex 5.5 [117] together with ChatGPT 5.6 Sol [118] for debugging and refactoring, with the models frequently cross-checking one another's output.

Release-paper validation additionally used a bounded agentic controller rather than an unconstrained multi-agent search. The controller applies the method-maturity policy of Table 1 to a release-paper inventory and runs calculation, result retrieval, scientific judgment, defect repair, and paper/SI promotion as concurrent lanes; a lane pauses only for a direct dependency, so a row with a confirmed defect waits for its fixed release while unaffected rows continue. The controller is closed by default, requires an explicit human activation, and pauses itself with a recorded blocker when no evidence-gated work remains. Scientific authority remains human throughout.

Scientific judgment separates numerical, provenance, and performance gates. Before comparing energies, the controller audits method, basis, ECP, charge, multiplicity, frozen-core policy, grids, fitting, integral mode, symmetry, $\boldsymbol{k}$ sampling, and solver controls. This audit prevented, for example, an apparent water MP2 discrepancy from being attributed to VIBE-QC accuracy: the reference calculation had frozen core orbitals whereas the VIBE-QC calculation was all-electron, so the comparison was held and a matched reference was computed instead. Provenance checks require full output, manifest and QVF records, and the exact immutable release identity recorded before submission and recovered from the job artifacts; corrected inputs receive new identities rather than in-place edits, and automatic retries are disabled. Performance acceptance compares matched wall and CPU times, memory, and algorithmic modes, and an efficiency gap blocks acceptance only when those are comparable. Literature values are resolved through a maintained local catalog, and missing sources enter a request queue rather than being reconstructed from memory. Reference programs remain out-of-process oracles; their source is neither copied nor treated as publication authority. Figure 3 shows the lanes and their direct dependencies.

Standard build-time tooling (CMake, scikit-build-core, `pytest`, Sphinx) and a GitLab continuous-integration pipeline complete the toolchain. Each development milestone is required to pass the affected local test lanes and production-shaped evidence where applicable. The release pipeline then provides the clean-build and release-candidate integration gate; broader regression lanes, too memory-intensive for the shared runner, execute on a development workstation before a tagged release.

4 Validation

Every capability promoted to production status in this paper is compared, on a canonical system, against an independent implementation in an established program, executed out-of-process so that the compari-

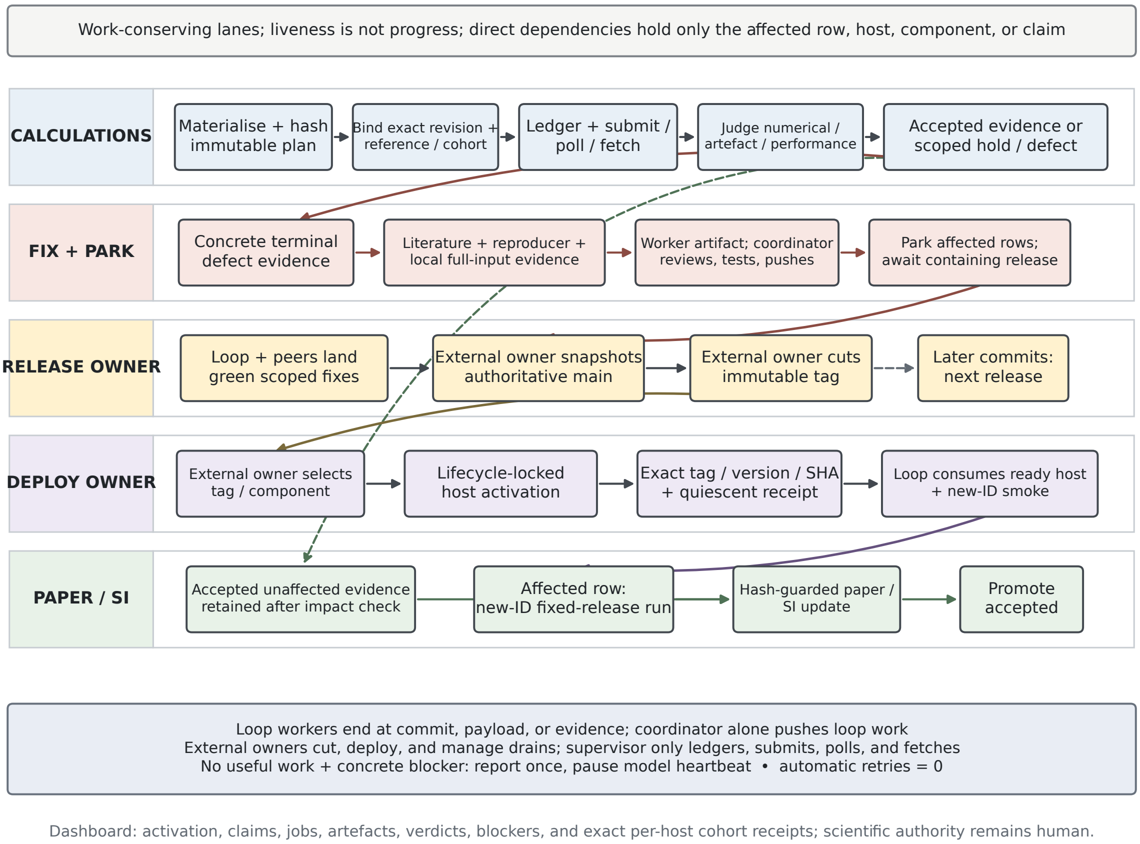

Figure 3: Concurrent, evidence-gated validation loop. Horizontal swimlanes run independently; arrows show only direct dependencies. Evidence for unaffected capabilities can proceed to the paper while a confirmed defect follows the fix, release, and rerun path. Scientific authority remains human: the controller requires explicit activation and pauses with a recorded blocker when no evidence-gated work remains.

son is between two codes and not between a wrapper and itself. One reference is used per method family, matched to the formalism: ORCA 6 [5] for molecular methods, OPENMOLCAS [10, 11] for complete-active-space and multireference methods, PySCF [6] for the periodic Gaussian-density-fitting route, CRYSTAL 23 [8] for the periodic bipolar-expansion route, and GPAW [9] for the Gaussian-and-plane-waves route. Capabilities that ship but are gated, or that lack a suitable oracle, are outside the scope of this paper's claims; each is labeled in the code and recorded in the project's public issue tracker (Section 8). The full per-method tables, the systems, the settings, and the input and output archives are in the Supporting Information.

Agreement is limited by different approximations in different method classes, and the results in Table 2 group accordingly. Mean-field energies are SCF-limited and agree with ORCA at the nanohartree scale, with a mean absolute deviation of 7.7×10^{-10} E_h over the mean-field rows of the table. Correlated and density-functional energies are limited by density fitting and by the integration grid rather than by SCF convergence, and agree at the 7.2×10^{-5} E_h level, the mean absolute deviation over the six density-fitted and grid-limited rows, with a largest deviation of 1.5×10^{-4} E_h. The separation is informative: the conventional, non-density-fitted CCSD(T) route agrees with ORCA to at most 3.3×10^{-7} E_h, which localizes the larger residual of the density-fitted CCSD comparison to the fitting approximation rather than to the coupled-cluster implementation. Re-execution of fourteen correlated cases on the corrected density-fitting

Table 2: Molecular-method validation, one reference per method. Test IDs refer to the archived input set (Section 5); the full per-method tables are in the Supporting Information.

Method	System / basis	Reference	Agreement
RHF, UHF, ROHF	H_2O, OH, NH_2 / def2-TZVP	Orca	$1.2 \times 10^{-9}\ E_h$
KS-DFT (LDA→meta→RSH)	H_2O / def2-TZVP	Orca	$1.5 \times 10^{-4}\ E_h$
MP2 / SCS / SOS	H_2O / cc-pVTZ	Orca	MP2 $2.5 \times 10^{-5}\ E_h$; SCS $5.2 \times 10^{-5}\ E_h$; SOS $6.5 \times 10^{-5}\ E_h$
Double hybrid (B2PLYP)	H_2O / cc-pVTZ	Orca	$1.4 \times 10^{-6}\ E_h$ (SCF half $1.3 \times 10^{-6}\ E_h$, scaled MP2 $7.7 \times 10^{-8}\ E_h$)
CCSD(T) (all-e., non-DF)	H_2O, CO / cc-pVXZ, def2-XVP	Orca conv.	$\leq 3.3 \times 10^{-7}$ E_h
DF-CCSD (all-e.)	H_2O / cc-pVDZ	Orca conv. (non-DF)	$1.4 \times 10^{-4}\ E_h$ (fitting-limited)
DLPNO-MP2/CCSD(T)	S22 subset (TightPNO)	canonical	truncation error 0.17 kcal mol^{-1} (Section 4.1)
CASSCF	N_2 / cc-pVDZ (6e,6o)	OpenMolcas	$\leq 3.0 \times 10^{-10}\ E_h$
FCI	H_2 / 6-31G	internal (cross-release)	$1.4 \times 10^{-11}\ E_h$
TDDFT (TDA)	CH_2O / def2-SVP	Orca	4.4×10^{-3} eV
OVGF orbital energies	9 molecules / def2-TZVP	Orca	1×10^{-4} eV
Composite HF-3c	$(H_2O)_2$ / MINIX	Orca/ `mctc-gcp`	$1.7 \times 10^{-9}\ E_h$
Dispersion D4, E_{disp} only	$(H_2O)_2$ / def2-TZVP, PBE0	Orca	$2.0 \times 10^{-9}\ E_h$
Geometry optimization	$(H_2O)_2$ / def2-TZVP, PBE0-D3(BJ)	Orca	$3.1 \times 10^{-4}\ E_h$; r_{OO} 3.5×10^{-3} Å
ECP (dhf-ECP)	I_2 / dhf-TZVP	Orca	$1.32 \times 10^{-8}\ E_h$
RHF analytic Hessian / IR	H_2O / def2-TZVP	Orca	0.11 cm^{-1}
Interaction energies, S22 set	22 complexes / def2-SVP, RKS-PBE-D3(BJ)	Orca	0.016 kcal mol^{-1} MAD, 22/22 within tolerance

route reproduces every earlier default-route total energy to within $1.1 \times 10^{-9}\ E_h$. The local DLPNO hierarchy is reported as a pilot implementation rather than as a TightPNO parity claim, and the complete-active-space energies reproduce a matched OpenMolcas reference on a tagged release (CASSCF to $3.0 \times 10^{-10}\ E_h$, CASCI to $8.2 \times 10^{-9}\ E_h$), with Orca as an independent second oracle for CASSCF.

Two rows deserve comment because the agreement holds term by term. The B2PLYP double hybrid agrees to $1.4 \times 10^{-6}\ E_h$ once the auxiliary-basis and scaling conventions are matched on both sides, $1.3 \times 10^{-6}\ E_h$ in the Kohn–Sham half and $7.7 \times 10^{-8}\ E_h$ in the scaled MP2 contribution, so the small residual is not a cancellation of larger errors. The composite HF-3c reproduces Orca to $1.7 \times 10^{-9}\ E_h$, the tightest molecular agreement in the table, and its gCP and short-range-basis terms match to the same order.

4.1 Published benchmark sets

Cross-code parity establishes that we solve the same equations as an established program. It says nothing about whether the resulting chemistry is of the quality the method is known to deliver. For that we compare against published benchmark data, and Table 3 reports three such sets. These are accuracy statements about the method as implemented here, not parity statements about the implementation, and the two must not be read as the same claim.

The excitation energies are the clearest case. Over 332 excitation-energy comparisons from QUESTDB [119]

(94 B3LYP, 105 PBE0, 99 PBE, and 34 CIS entries), spanning 32 molecules and compared against the theoretical best estimates at aug-cc-pVTZ, the mean absolute deviation is 0.153 eV for B3LYP, 0.172 eV for PBE0, 0.184 eV for PBE, and 0.238 eV for CIS. That ordering, and those magnitudes, are what the linear-response formalism is documented to give for global hybrids and for CIS, so the implementation reproduces the known behavior of the methods it provides. A further 234 QUESTDB entries are excluded for reasons of principle rather than of numerics: doubly excited states, which the adiabatic approximation cannot describe, and fluorescence entries, which refer to a different geometry.

The barrier heights require the same care in the opposite direction. Against BH9 [120], PBE gives a mean absolute deviation of 6.96 kcal mol^{-1} on forward barriers, 9.05 kcal mol^{-1} on reverse barriers, and 4.85 kcal mol^{-1} on reaction energies over 942 unfiltered comparisons, retaining 20 outlier rows whose deviations exceed three times the published reference-method error. Those are large numbers, and they are the expected ones: a semilocal GGA systematically underestimates barrier heights by several kcal mol^{-1}, and BH9 is referenced to a DLPNO-CCSD(T_0)/CBS focal point. The result validates that our PBE reproduces PBE, not that PBE is accurate. We report only PBE here because that is the functional the present campaign covered; a ladder study across BH9 belongs with the functional-assessment literature rather than with a code release paper.

The local-correlation truncation error is measured directly against the canonical result on the S22 subset used by Liakos *et al.* [121], per method and PNO setting. At the TightPNO setting the DLPNO-CCSD(T) interaction-energy truncation error has a mean absolute value of 0.17 kcal mol^{-1} and a maximum of 0.46 kcal mol^{-1} over the five complexes computed, and the DLPNO-MP2 truncation error has a mean absolute value of 0.75 kcal mol^{-1} and a maximum of 2.76 kcal mol^{-1} over 22 complexes, with the largest errors on dispersion-dominated complexes. NormalPNO and LoosePNO results are recorded in the database but not claimed, and DLPNO-CCSD is excluded from this statistic under a tracked convergence defect. The density-fitting error on the corresponding total energies is below the reporting threshold. This quantifies the truncation error left open by the pilot designation of the DLPNO hierarchy in Table 2.

The electron propagator illustrates the distinction between the two kinds of claim directly. The underlying Hartree–Fock orbital energies match Orca to 1×10^{-4} eV, which is implementation parity. The OVGF-corrected first ionization energies, measured against IP-EOM-CCSD on the same nine molecules at def2-TZVP, deviate by 0.60 eV on average with a maximum of 1.22 eV, which is method accuracy and is the error the diagonal second-order approximation is expected to carry. The same calculation is therefore a tight parity result and a loose accuracy result at once, depending on which question is asked of it.

4.2 Periodic methods

External total-energy parity for the periodic routes is open work at this release, and no periodic total-energy comparison is promoted as a parity claim in this paper. The residuals against the matched references (PySCF, Crystal 23, and GPAW) are characterized rather than closed: a matched PySCF cross-check localizes the density-fitting route's offset to a fitting-completeness deficit in the Coulomb block on dense-core cells, with the exchange–correlation increments themselves consistent across LDA, PBE, and r^{2}SCAN, and the bipolar-expansion comparison differs by an exchange finite-size convention. Completing the fitted Coulomb tails is the principal open task of the periodic stack. The comparisons, their current residuals, and every periodic capability that ships without validation at this release are recorded

Table 3: Validation against published benchmark data. These are accuracy measures for the method as implemented, not cross-code parity measures. Reference values are the published ones cited in each row.

Set	Quantity	Method	n	MAD
QUESTDB [119]	vertical excitation	B3LYP	94	0.153 eV
		PBE0	105	0.172 eV
		PBE	99	0.184 eV
		CIS	34	0.238 eV
BH9 [120]	barrier, forward	PBE	304	6.96 kcal mol^{-1}
	barrier, reverse	PBE	331	9.05 kcal mol^{-1}
	reaction energy	PBE	307	4.85 kcal mol^{-1}
S22 subset [121, 122]	DLPNO truncation error (TightPNO)	DLPNO-CCSD(T)	5	0.17 kcal mol^{-1}
		DLPNO-MP2	22	0.75 kcal mol^{-1}

in the public issue tracker (Section 8); representative periodic calculations within the release's scope are shown in Section 5.

4.3 Regression suite and route coverage

Cross-code parity establishes whether the reported numbers are right; it does not establish whether the code runs across its whole surface. For an implementation produced the way this one was, breadth of execution deserves separate evidence, because that is where a generated code is most likely to fail quietly. Two mechanisms provide it. The repository carries a repo-wide `pytest` suite that collects more than 15 500 parametrized test items across 726 test files, covering molecular and periodic SCF, exchange–correlation evaluation, gradients and Hessians, post-Hartree–Fock and multireference methods, density fitting, periodic Coulomb treatments, basis and ECP handling, and output and provenance behavior. Separately, a curated end-to-end parity runner holds 44 scientific cases, 40 molecular and 4 periodic, that execute the matched reference program alongside VIBE-QC. Beyond these, a separately maintained input library is drawn on as a route-coverage stress test. From it, 859 registered cases across 57 molecular systems were drawn for this release; 819 executed and 652 produced a terminal energy, yielding 935 comparisons, and a curated subset of 26 of those, spanning 14 systems, carries matched ORCA references. The library exercises input paths far beyond the validated set, and it is reported in the Supporting Information as coverage evidence rather than as accuracy evidence: of its 935 comparisons, 301 pass, 298 miss their tolerance, and the remainder did not reach a number.

Beyond one-oracle parity, a two-code consensus reference set is maintained with the same input library: canonical RHF total energies for which VIBE-QC and ORCA, executed independently at identical geometry, basis, and protocol, must agree within a per-atom tolerance before a value is stored, so each stored energy is backed by two independently written implementations, and no value is created when the codes disagree. At the 2026-08-14 snapshot the set holds 44 consensus energies across 21 systems and seven basis sets, with a largest cross-code spread of 3.1×10^{-9} E_h, together with one exact FCI energy (H_2/STO-3G). Two cross-code contracts in the same method family remain open, are excluded, and are tracked in the issue tracker; the complete set and its open contracts are tabulated in the Supporting Information.

5 Representative calculations

This section demonstrates VIBE-QC on calculations that exercise each centerpiece capability. The inputs are archived with the manuscript as a runnable test set (thirty scripts, twenty molecular and ten periodic, driven entirely through the high-level API); the manifest maps each script to its method, basis, and reference. Every run also emits its own `.system` manifest and citation files. Capabilities outside the validated scope of this release are not illustrated here (Section 6.2).

5.1 Molecular optimization and vibrational frequencies

The water dimer is optimized at PBE0-D3(BJ)/def2-TZVP through the ASE bridge (script M07), exercising analytic KS gradients and the post-SCF dispersion contribution to the forces under a force-only convergence criterion (maximum force 0.02 eV Å^{-1}). The optimized O$\cdots$O separation is 3.0873 Å for the symmetric double-donor arrangement preserved from the starting geometry; the case validates the gradient and dispersion forces on a fixed hydrogen-bond topology and is not a conformational search for the global-minimum dimer structure. Harmonic frequencies and IR intensities for the water monomer are obtained from the analytic RHF Hessian (script M14); the symmetric and antisymmetric stretches and the bend are reproduced to within 0.11 cm^{-1} of the ORCA reference. Both results are embedded in the job's QVF archive and can be animated in VIBE-VIEW.

5.2 A representative organic molecule

To exercise the molecular stack at a realistic size, in the spirit of the benchmark systems used in peer release papers, benzene is optimized at PBE0-D3(BJ)/def2-TZVP (script M20). The optimized C–C bond length is 1.387 452 Å.

Timing is reported only from colocated controls, in which both programs run on the same node with the same allocated CPU count and memory. For the H_2O_2/DLPNO-CCSD(T)/cc-pVTZ comparison with one CPU and 6144 MiB each, VIBE-QC required 699.88 s wall time against 156.67 s for ORCA 6.1.1, a factor of 4.47, with peak resident memory of 1015.5 MiB and 2007.2 MiB respectively. We make no claim that VIBE-QC is faster than the established codes; on the local-correlation routes it is demonstrably slower, and the open efficiency defects are tracked in the issue tracker.

5.3 Heavy elements through effective core potentials

I_2 at RHF/dhf-TZVP (script M09) exercises the shipped dhf small-core ECP through LIBECPINT on a heavy main-group element. The total energy differs from the matched ORCA calculation by $1.3 \times 10^{-8}\, E_\mathrm{h}$, placing the molecular ECP path close to the mean-field parity floor of the all-electron routes.

5.4 Periodic dispersion and property export

The face-centered-cubic neon crystal at RKS-PBE with periodic D3(BJ) on a $4 \times 4 \times 4$ mesh (script P06) exercises the periodic lattice-sum dispersion path on a system whose cohesion is dispersion-dominated; the corrected total energy is archived in the QVF. Densities of states are exported from the archive and rendered through the VIBE-VIEW plotting helper for the two-dimensional graphene and hexagonal boron nitride cells, demonstrating the two- and three-dimensional periodic property surface end to end. Bulk observables derived from these exports, including band gaps and equation-of-state quantities, are not reported at this release (Section 6.2).

6 Discussion

6.1 The development methodology

The lens of this paper is the development methodology: VIBE-QC is a general-purpose quantum-chemistry code written almost entirely through generative-AI programming under human scientific control. It began as an experiment with a single question, how far a current large language model could be pushed toward a correct, non-trivial scientific code when a domain expert supplies the direction, reported on the project blog as a one-day prototype [1]. The prototype did real cross-code-validated Hartree–Fock, development continued, and the molecular and periodic stacks reported here grew from that start over the following months. The name records the method. The project has been transparent about it throughout: the documentation includes a chapter on driving the code with a language model [4], the development was logged publicly [1–3], and the project funding statement names the model subscription as the primary recurring cost. The substance that makes the methodology credible is the validated science of the preceding sections; the two are inseparable.

The division of labor was consistent throughout, and it used several models in distinct roles, including against one another. We set the architecture and scope, chose the methods and the scientific trade-offs, decided what counted as validated, read the diffs, and signed each release. The main development chat was Anthropic Claude Opus 4.8 [110] (with Opus 4.7 earlier in the project) [111], joined later by Claude Opus 5 and Claude Fable 5 [112, 113]; these chats translated equations into C++ kernels, scaffolded the Python surface and the test suite, and drafted documentation. DeepSeek [114, 115], later DeepSeek v4 Pro [116], was used for high-volume routine generation at low cost, and OpenAI Codex 5.5 [117], later ChatGPT 5.6 Sol [118], handled debugging and refactoring. A productive pattern was to have the models check one another: a kernel written in one chat was reviewed, stress-tested, or re-derived in another, which caught mistakes before they reached the cross-code parity gate. The models did not choose the physics, decide what was correct, or release anything.

The most notable outcome of this arrangement is the development pace. One person produced the release line of Table 4 in roughly ten weeks, spanning Hartree–Fock to coupled cluster, excited states, and three periodic Fock-build routes.

We report the cadence as a descriptive fact of the methodology; the cross-code validation discipline of Section 4 is what makes results produced at this pace defensible.

What makes the result trustworthy is not the workflow but the discipline imposed on it. Production status is gated on cross-code parity against an independent implementation (Section 4), and the repo-wide automated regression suite of Section 4.3 is the mechanism that keeps the gate shut. This discipline caught the failure modes the workflow produced: off-by-one errors in lattice sums, caught by parity; wrong index permutations in symmetry operations, caught by the Wigner-D verification; a periodic Coulomb gauge error that converged to an impossible energy, caught by translation invariance and by comparison to PySCF; and plausible-but-wrong citations in early documentation, caught by DOI resolution. In one case, a set of

Table 4: Release history of VIBE-QC (patch releases omitted), with the foundational work each release's principal capability builds on.

Release	Date	Work the principal capability builds on
0.1.0	2026-04-18	Roothaan–Hall equations (C. C. J. Roothaan)
0.2.0†	—	Ewald summation (P. P. Ewald)
0.2.5†	—	Wigner D-matrices and group theory (E. Wigner)
0.3.0†	—	Molecular-orbital theory (R. S. Mulliken)
0.4.0	2026-04-27	Hellmann–Feynman theorem (H. Hellmann)
0.5.0	2026-04-30	FG-matrix vibrational analysis (E. B. Wilson)
0.6.0	2026-04-30	DIIS and analytic gradients (P. Pulay)
0.7.0	2026-05-02	Canonical orthogonalization (P.-O. Löwdin)
0.8.0	2026-05-17	DFT-D3/D4 dispersion (S. Grimme)
0.9.0	2026-05-23	Coupled cluster (P. J. Knowles)
0.10.0	2026-05-27	Periodic Hartree–Fock (C. Pisani)
0.11.0	2026-06-03	Range-separated density fitting (Q. Sun)
0.12.0	2026-06-11	Algorithmic analysis (D. E. Knuth)
0.13.0	2026-06-17	Generalized-regular $\boldsymbol{k}$ grids (P. Wisesa)
0.14.0	2026-06-22	CCSD(T) (R. J. Bartlett)
0.15.0	2026-06-26	DLPNO-CCSD(T) (F. Neese)

† `0.2.0`, `0.2.5`, and `0.3.0` were internal milestones that were not separately tagged.

exchange–correlation residuals reported at `v0.9.0` did not reproduce the parity floor on re-examination and was retracted at `v0.10.0`; the error was detected by the validation gate rather than by the generation step.

6.2 Scope of the claims at the 0.15 series

We draw the scope line plainly. The claims of this paper are the validated ones of Section 4: the molecular mean-field and molecular-ECP routes are production-ready within the documented envelope; the correlated MP2 and CCSD(T) routes carry small-molecule numerical evidence, with production-size memory coverage still open; and the DLPNO hierarchy is a pilot local-correlation surface, not a TightPNO parity claim. External total-energy parity for the periodic routes is open work at this release (Section 4.2), and no periodic total-energy comparison is promoted as a parity claim; the periodic stack converges cleanly in the molecular-limit regime and is being hardened for tight-cell bulk work.

Every further capability that ships in the tree without validation at this release is gated under the vocabulary of Table 1, is excluded from the claims of this paper, and is recorded as an open issue in the project's public issue tracker (Section 8), with its reproducer where one exists. The tracker, not this paper, is the canonical record of known defects and their status.

Compute-kernel GPU acceleration and production MPI execution through `mpirun` are on the roadmap; the current MPI surface is an optional `mpi4py` substrate plus an experimental GPW Hartree–J z-slab overlay, while the validated parallel surface in this release is intra-node OpenMP plus queue-level distribution of independent jobs through `vq`. There is no binary wheel; the supported channel is the source install.

6.3 Reproducibility and provenance

Every result in this paper is reproducible by cloning the source tree, running the bootstrap script, and executing the archived input at the build recorded in that row's `.system` manifest. Because the 0.15 series is under active development, the manifest rather than the version string is the object that pins a result: each job writes a `.system` file recording the exact build, the vendored-library versions, and the hardware, together with an automatic citation manifest naming every method invoked, so a result carries its own provenance. We are careful about one boundary: a language-model-driven development session is not itself reproducible run-to-run, so what is reproducible is the released code at a named commit and its outputs, not the construction process that produced it. That is the appropriate object of reproducibility for any scientific code, however it was written.

6.4 Outlook

Near-term work hardens the periodic stack: all-electron $1.0 \times 10^{-6}\, E_\mathrm{h}$ accuracy on heavy-core crystals through mixed density fitting, analytic periodic meta-GGA forces, and broader multi-$\boldsymbol{k}$ coverage of the correlated and excited-state methods. The longer-range direction reuses the validated machinery of this paper for the ab initio cyclic cluster model (AICCM) [123], a $\boldsymbol{\Gamma}$ real-space approach in which a finite, periodic-image-aware cluster is treated with molecular correlated wavefunction theory. The plan is to climb a Jacob's ladder of correlation treatments within the cyclic cluster, starting from the semiempirical cyclic-cluster treatment that already exists in the tree (Section 2.11) and ascending through Hartree–Fock, MP2, and local DLPNO-MP2 to DLPNO-CCSD(T), with the pob-* basis sets [107, 108] as the periodic-side default. The AICCM work is under active development in the tree as gated, experimental code, is deliberately out of scope here, and will be reported in its own dedicated papers; the present release is the validated numerical foundation on which it is being built.

7 Conclusions

We have introduced VIBE-QC at the 0.15 series, a general-purpose open-source electronic-structure code that treats molecules and solids on the same footing: molecular Hartree–Fock, the full DFT ladder, Møller–Plesset theory and double hybrids, canonical coupled cluster, pilot local-correlation routes, complete-active-space and other multireference methods, and excited states, together with periodic HF and DFT in one, two, and three dimensions through three Fock-build routes (Gaussian density fitting, a bipolar-expansion route, and Gaussian-and-plane-waves) under one Python interface. Production-status claims are validated out-of-process against the reference matched to each method family, ORCA for molecular methods and OPENMOLCAS for complete-active-space methods; external total-energy parity for the periodic routes against PYSCF, CRYSTAL 23, and GPAW is open work at this release, and capabilities outside the validated scope are gated and tracked in the public issue tracker. The code rests on the standard open-source libraries of the field, reads and writes the established file formats, and adds the QVF format, the VIBE-VIEW viewer, and the cluster-aware VIBE-QUEUE job queue.

The lens through which we have presented VIBE-QC is its construction: it is, to our knowledge, the first general-purpose quantum-chemistry code built almost entirely through generative-AI programming under expert human direction. What started as a one-day experiment in vibe-coding became a broad, cross-code-validated package, and the cross-code validation, not the novelty of the workflow, is what makes the promoted numerics trustworthy. That same gate is what keeps the periodic total energies, the local-correlation hierarchy, and several property outputs outside the claims of this release; their open items live in the public issue tracker rather than in this paper. The longer-range aim is to reuse this validated foundation for an ab initio cyclic cluster model that climbs from the semiempirical level to DLPNO-CCSD(T), to be reported separately. We invite the community to use the code, to check it against their own references, and to repeat the experiment.

8 Software availability and computational details

Availability. VIBE-QC is openly developed at `https://gitlab.peintinger.com/mpei/vibeqc` with documentation at `https://vibe-qc.com`, licensed under the Mozilla Public License, version 2.0. This paper describes the 0.15 release series. Each reported calculation carries a `.system` manifest recording the exact build used, distributed with the calculation archive. Known defects and every capability that ships without validation at this release are tracked in the public issue tracker at `https://gitlab.peintinger.com/mpei/vibeqc/-/issues`.

Installation. Installation is from source:

```
git clone https://gitlab.peintinger.com/mpei/vibeqc.git
cd vibeqc && ./scripts/install.sh
```

The bootstrap script builds the vendored native dependencies in tree under `third_party/`: LIBINT 2.13.1, LIBXC 7.0.0, SPGLIB, LIBECPINT 1.0.7, and FFTW3. There is no binary wheel at the 0.15 series.

Computational details. The reported calculations were run through `vq` on x86_64 Linux workstations (Intel Core i9 and AMD Ryzen 9 class, 8–16 physical cores, 62–192 GB), on Apple silicon under macOS, and on two university clusters, one PBS/Torque and one Slurm. The hardware and operating systems are tabulated in the Supporting Information; node-level hostnames are pseudonymized in the archived manifests, and the OpenMP thread count of each run is recorded in its `.system` manifest. Shared-memory OpenMP is the production parallel surface of this release; the reported timings should not be read as distributed-memory MPI scaling, which is limited here to an optional `mpi4py` substrate and an experimental GPW z-slab overlay. The functional, basis, grid, and convergence settings for each calculation are those in the archived input scripts and are echoed in full in each `.out` file and `.system` manifest; the per-job `.bibtex` manifests list the corresponding method, functional, basis-set, and ECP citations. Platforms tested for the release are macOS on Apple silicon and Linux x86_64 (Arch/Manjaro, Debian, Ubuntu).

Acknowledgments

VIBE-QC rests on the open-source work of others. The author thanks E. F. Valeev (LIBINT), S. Lehtola and the LIBXC developers, A. Togo (SPGLIB), R. A. Shaw and J. G. Hill (LIBECPINT), the EIGEN and `pybind11` teams, the FFTW developers, and A. Hjorth Larsen and the ASE developers, whose libraries VIBE-QC links against. The author is grateful to T. Bredow at the Mulliken Center for Theoretical Chemistry, University of Bonn, for the original cyclic cluster model and pob-* basis-set work that motivates this code, and to the ORCA, OPENMOLCAS, PYSCF, and CRYSTAL teams, whose programs serve as the cross-code parity oracles. The code was developed with the assistance of the generative-AI models named in Section 6.1.

Author contributions

Michael F. Peintinger: conceptualization, methodology, software architecture, validation, formal analysis, writing: original draft, and writing: review and editing. In line with the development methodology of Section 6.1, generative-AI tools (Anthropic Claude Opus 4.7, Opus 4.8, Opus 5, and Fable 5; DeepSeek Flash, Pro, and v4 Pro; and OpenAI Codex 5.5 and ChatGPT 5.6 Sol) were used for code generation, test scaffolding, and documentation drafting under the direct supervision, review, and scientific control of the author. The same tools were used to assist in drafting and editing this manuscript. The author reviewed, verified, and edited all AI-assisted output, is responsible for all design decisions, validation, and interpretation, and takes full responsibility for the content of this manuscript. No AI tool is listed as an author, in accordance with the publisher's authorship criteria.

Funding

Development of VIBE-QC is fully self-financed by the author and received no external, institutional, or commercial funding.

Conflicts of interest

The author declares no conflicts of interest.

Data availability

Source code, documentation, and tutorials are openly available at `https://gitlab.peintinger.com/mpei/vibeqc` and `https://vibe-qc.com`, licensed under the Mozilla Public License, version 2.0. This paper describes the 0.15 release series. Each reported calculation carries a `.system` manifest recording the exact build it was produced with, and the manifests are distributed with the calculation archive in the Supporting Information. Known defects are recorded in the repository's public issue tracker. The default `release` branch fast-forwards from each tag, so a bare `git clone` lands the user on the latest published snapshot.

References

[1] M. F. Peintinger. *I vibe-coded a quantum chemical program in one day*. `https://peintinger.com/2026/04/22/i-vibe-coded-a-quantum-chemical-program-in-one-day/`. Apr. 2026.

[2] M. F. Peintinger. *The giants vibe-qc stands on: a guided tour of the stack*. `https://peintinger.com/2026/04/23/the-giants-vibe-qc-stands-on-a-guided-tour-of-the-stack/`. Apr. 2026.

[3] M. F. Peintinger. *vibe-qc documentation as educational artefact*. `https://peintinger.com/2026/04/26/vibe-qc-documentation-as-educational-artefact/`. Apr. 2026.

[4] M. F. Peintinger. *Using Claude to drive vibe-qc*. `https://vibe-qc.com/using_claude.html`. Part of the vibe-qc documentation, release branch. 2026.

[5] F. Neese. "Software update: the ORCA program system—version 6.0". In: *WIREs Computational Molecular Science* 15.1 (2025), e70019. DOI: `10.1002/wcms.70019`.

[6] Q. Sun, X. Zhang, S. Banerjee, P. Bao, M. Barbry, N. S. Blunt, N. A. Bogdanov, G. H. Booth, J. Chen, Z.-H. Cui, et al. "Recent developments in the PySCF program package". In: *The Journal of Chemical Physics* 153.2 (2020), p. 024109. DOI: `10.1063/5.0006074`.

[7] R. Dovesi, R. Orlando, A. Erba, C. M. Zicovich-Wilson, B. Civalleri, S. Casassa, L. Maschio, M. Ferrabone, M. De La Pierre, P. D'Arco, Y. Noël, M. Causà, M. Rérat, and B. Kirtman. "CRYSTAL14: a program for the ab initio investigation of crystalline solids". In: *International Journal of Quantum Chemistry* 114.19 (2014), pp. 1287–1317. DOI: `10.1002/qua.24658`.

[8] A. Erba, J. K. Desmarais, S. Casassa, B. Civalleri, L. Donà, I. J. Bush, B. Searle, L. Maschio, L. Edith-Daga, A. Cossard, C. Ribaldone, E. Ascrizzi, N. L. Marana, J.-P. Flament, and B. Kirtman. "CRYSTAL23: a program for computational solid state physics and chemistry". In: *Journal of Chemical Theory and Computation* 19.20 (2023), pp. 6891–6932. DOI: `10.1021/acs.jctc.2c00958`.

[9] J. Enkovaara, C. Rostgaard, J. J. Mortensen, J. Chen, M. Dułak, L. Ferrighi, J. Gavnholt, C. Glinsvad, V. Haikola, H. A. Hansen, H. H. Kristoffersen, M. Kuisma, A. H. Larsen, L. Lehtovaara, M. Ljungberg, O. Lopez-Acevedo, P. G. Moses, J. Ojanen, T. Olsen, V. Petzold, N. A. Romero, J. Stausholm-Møller, M. Strange, G. A. Tritsaris, M. Vanin, M. Walter, B. Hammer, H. Häkkinen, G. K. H. Madsen, R. M. Nieminen, J. K. Nørskov, M. Puska, T. T. Rantala, J. Schiøtz, K. S. Thygesen, and K. W. Jacobsen. "Electronic structure calculations with GPAW: a real-space implementation of the projector augmented-wave method". In: *Journal of Physics: Condensed Matter* 22.25 (2010), p. 253202. DOI: `10.1088/0953-8984/22/25/253202`.

[10] I. Fdez. Galván, M. Vacher, A. Alavi, C. Angeli, F. Aquilante, J. Autschbach, J. J. Bao, S. I. Bokarev, N. A. Bogdanov, R. K. Carlson, et al. "OpenMolcas: from source code to insight". In: *Journal of Chemical Theory and Computation* 15.11 (2019), pp. 5925–5964. DOI: `10.1021/acs.jctc.9b00532`.

[11] G. Li Manni, I. F. Galván, A. Alavi, F. Aleotti, F. Aquilante, J. Autschbach, D. Avagliano, A. Baiardi, J. J. Bao, S. Battaglia, et al. "The OpenMolcas web: a community-driven approach to advancing computational chemistry". In: *Journal of Chemical Theory and Computation* 19.20 (2023), pp. 6933–6991. DOI: `10.1021/acs.jctc.3c00182`.

[12] C. C. J. Roothaan. "New developments in molecular orbital theory". In: *Reviews of Modern Physics* 23.2 (1951), pp. 69–89. DOI: `10.1103/RevModPhys.23.69`.

[13] J. A. Pople and R. K. Nesbet. "Self-consistent orbitals for radicals". In: *The Journal of Chemical Physics* 22.3 (1954), pp. 571–572. DOI: `10.1063/1.1740120`.

[14] C. C. J. Roothaan. "Self-consistent field theory for open shells of electronic systems". In: *Reviews of Modern Physics* 32.2 (1960), pp. 179–185. DOI: 10.1103/RevModPhys.32.179.

[15] F. Neese, F. Wennmohs, A. Hansen, and U. Becker. "Efficient, approximate and parallel Hartree–Fock and hybrid DFT calculations. A "chain-of-spheres" algorithm for the Hartree–Fock exchange". In: *Chemical Physics* 356.1-3 (2009), pp. 98–109. DOI: 10.1016/j.chemphys.2008.10.036.

[16] P. Pulay. "Convergence acceleration of iterative sequences. The case of SCF iteration". In: *Chemical Physics Letters* 73.2 (1980), pp. 393–398. DOI: 10.1016/0009-2614(80)80396-4.

[17] P. Pulay. "Improved SCF convergence acceleration". In: *Journal of Computational Chemistry* 3.4 (1982), pp. 556–560. DOI: 10.1002/jcc.540030413.

[18] K. N. Kudin, G. E. Scuseria, and E. Cancès. "A black-box self-consistent field convergence algorithm: one step closer". In: *The Journal of Chemical Physics* 116.19 (2002), pp. 8255–8261. DOI: 10.1063/1.1470195.

[19] V. R. Saunders and I. H. Hillier. "A "level-shifting" method for converging closed shell Hartree–Fock wave functions". In: *International Journal of Quantum Chemistry* 7.4 (1973), pp. 699–705. DOI: 10.1002/qua.560070407.

[20] S. Lehtola, C. Steigemann, M. J. T. Oliveira, and M. A. L. Marques. "Recent developments in libxc — a comprehensive library of functionals for density functional theory". In: *SoftwareX* 7 (2018), pp. 1–5. DOI: 10.1016/j.softx.2017.11.002.

[21] A. D. Becke. "A multicenter numerical integration scheme for polyatomic molecules". In: *The Journal of Chemical Physics* 88.4 (1988), pp. 2547–2553. DOI: 10.1063/1.454033.

[22] O. Treutler and R. Ahlrichs. "Efficient molecular numerical integration schemes". In: *The Journal of Chemical Physics* 102.1 (1995), pp. 346–354. DOI: 10.1063/1.469408.

[23] V. I. Lebedev. "Quadratures on a sphere". In: *USSR Computational Mathematics and Mathematical Physics* 16.2 (1976), pp. 10–24. DOI: 10.1016/0041-5553(76)90100-2.

[24] J. P. Perdew, K. Burke, and M. Ernzerhof. "Generalized gradient approximation made simple". In: *Physical Review Letters* 77.18 (1996), pp. 3865–3868. DOI: 10.1103/PhysRevLett.77.3865.

[25] C. Adamo and V. Barone. "Toward reliable density functional methods without adjustable parameters: the PBE0 model". In: *The Journal of Chemical Physics* 110.13 (1999), pp. 6158–6170. DOI: 10.1063/1.478522.

[26] A. D. Becke. "Density-functional thermochemistry. III. The role of exact exchange". In: *The Journal of Chemical Physics* 98.7 (1993), pp. 5648–5652. DOI: 10.1063/1.464913.

[27] P. J. Stephens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch. "Ab initio calculation of vibrational absorption and circular dichroism spectra using density functional force fields". In: *The Journal of Physical Chemistry* 98.45 (1994), pp. 11623–11627. DOI: 10.1021/j100096a001.

[28] J. W. Furness, A. D. Kaplan, J. Ning, J. P. Perdew, and J. Sun. "Accurate and numerically efficient r^{2}SCAN meta-generalized gradient approximation". In: *The Journal of Physical Chemistry Letters* 11.19 (2020), pp. 8208–8215. DOI: 10.1021/acs.jpclett.0c02405.

[29] J. W. Furness, A. D. Kaplan, J. Ning, J. P. Perdew, and J. Sun. "Construction of meta-GGA functionals through restoration of exact constraint adherence to regularized SCAN functionals". In: *The Journal of Chemical Physics* 156.3 (2022), p. 034109. DOI: 10.1063/5.0073623.

[30] J.-D. Chai and M. Head-Gordon. "Systematic optimization of long-range corrected hybrid density functionals". In: *The Journal of Chemical Physics* 128.8 (2008), p. 084106. DOI: 10.1063/1.2834918.

[31] O. A. Vydrov and T. Van Voorhis. "Nonlocal van der Waals density functional: the simpler the better". In: *The Journal of Chemical Physics* 133.24 (2010), p. 244103. DOI: 10.1063/1.3521275.

[32] N. Mardirossian and M. Head-Gordon. "ωB97M-V: a combinatorially optimized, range-separated hybrid, meta-GGA density functional with VV10 nonlocal correlation". In: *The Journal of Chemical Physics* 144.21 (2016), p. 214110. DOI: 10.1063/1.4952647.

[33] A. Hjorth Larsen, J. Jørgen Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. Bjerre Jensen, J. Kermode, J. R. Kitchin, E. Leonhard Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. Bergmann Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, and K. W. Jacobsen. "The atomic simulation environment — a Python library for working with atoms". In: *Journal of Physics: Condensed Matter* 29.27 (2017), p. 273002. DOI: 10.1088/1361-648X/aa680e.

[34] A. Banerjee, N. Adams, J. Simons, and R. Shepard. "Search for stationary points on surfaces". In: *The Journal of Physical Chemistry* 89.1 (1985), pp. 52–57. DOI: 10.1021/j100247a015.

[35] E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, and P. Gumbsch. "Structural relaxation made simple". In: *Physical Review Letters* 97.17 (2006), p. 170201. DOI: 10.1103/PhysRevLett.97.170201.

[36] J. Baker, A. Kessi, and B. Delley. "The generation and use of delocalized internal coordinates in geometry optimization". In: *The Journal of Chemical Physics* 105.1 (1996), pp. 192–212. DOI: 10.1063/1.471864.

[37] C. Møller and M. S. Plesset. "Note on an approximation treatment for many-electron systems". In: *Physical Review* 46.7 (1934), pp. 618–622. DOI: 10.1103/PhysRev.46.618.

[38] S. Grimme. "Improved second-order Møller–Plesset perturbation theory by separate scaling of parallel- and antiparallel-spin pair correlation energies". In: *The Journal of Chemical Physics* 118.20 (2003), pp. 9095–9102. DOI: 10.1063/1.1569242.

[39] Y. Jung, R. C. Lochan, A. D. Dutoi, and M. Head-Gordon. "Scaled opposite-spin second order Møller–Plesset correlation energy: an economical electronic structure method". In: *The Journal of Chemical Physics* 121.20 (2004), pp. 9793–9802. DOI: 10.1063/1.1809602.

[40] S. Grimme. "Semiempirical hybrid density functional with perturbative second-order correlation". In: *The Journal of Chemical Physics* 124.3 (2006), p. 034108. DOI: 10.1063/1.2148954.

[41] G. Santra, N. Sylvetsky, and J. M. L. Martin. "Minimally empirical double-hybrid functionals trained against the GMTKN55 database: revDSD-PBEP86-D4, revDOD-PBE-D4, and DOD-SCAN-D4". In: *The Journal of Physical Chemistry A* 123.24 (2019), pp. 5129–5143. DOI: 10.1021/acs.jpca.9b03157.

[42] G. D. Purvis and R. J. Bartlett. "A full coupled-cluster singles and doubles model: the inclusion of disconnected triples". In: *The Journal of Chemical Physics* 76.4 (1982), pp. 1910–1918. DOI: 10.1063/1.443164.

[43] K. Raghavachari, G. W. Trucks, J. A. Pople, and M. Head-Gordon. "A fifth-order perturbation comparison of electron correlation theories". In: *Chemical Physics Letters* 157.6 (1989), pp. 479–483. DOI: 10.1016/S0009-2614(89)87395-6.

[44] C. Riplinger and F. Neese. "An efficient and near linear scaling pair natural orbital based local coupled cluster method". In: *The Journal of Chemical Physics* 138.3 (2013), p. 034106. DOI: 10.1063/1.4773581.

[45] C. Riplinger, B. Sandhoefer, A. Hansen, and F. Neese. "Natural triple excitations in local coupled cluster calculations with pair natural orbitals". In: *The Journal of Chemical Physics* 139.13 (2013), p. 134101. DOI: 10.1063/1.4821834.

[46] Y. Guo, C. Riplinger, U. Becker, D. G. Liakos, Y. Minenkov, L. Cavallo, and F. Neese. "Communication: an improved linear scaling perturbative triples correction for the domain based local pair-natural orbital based singles and doubles coupled cluster method [DLPNO-CCSD(T)]". In: *The Journal of Chemical Physics* 148.1 (2018), p. 011101. DOI: 10.1063/1.5011798.

[47] M. Saitow, U. Becker, C. Riplinger, E. F. Valeev, and F. Neese. "A new near-linear scaling, efficient and accurate, open-shell domain-based local pair natural orbital coupled cluster singles and doubles theory". In: *The Journal of Chemical Physics* 146.16 (2017), p. 164105. DOI: 10.1063/1.4981521.

[48] A. E. DePrince and C. D. Sherrill. "Accuracy and efficiency of coupled-cluster theory using density fitting / Cholesky decomposition, frozen natural orbitals, and a t_1-transformed Hamiltonian". In: *Journal of Chemical Theory and Computation* 9.6 (2013), pp. 2687–2696. DOI: 10.1021/ct400250u.

[49] B. O. Roos, P. R. Taylor, and P. E. M. Siegbahn. "A complete active space SCF method (CASSCF) using a density matrix formulated super-CI approach". In: *Chemical Physics* 48.2 (1980), pp. 157–173. DOI: 10.1016/0301-0104(80)80045-0.

[50] C. Angeli, R. Cimiraglia, S. Evangelisti, T. Leininger, and J.-P. Malrieu. "Introduction of n-electron valence states for multireference perturbation theory". In: *The Journal of Chemical Physics* 114.23 (2001), pp. 10252–10264. DOI: 10.1063/1.1361246.

[51] K. Andersson, P.-Å. Malmqvist, and B. O. Roos. "Second-order perturbation theory with a complete active space self-consistent field reference function". In: *The Journal of Chemical Physics* 96.2 (1992), pp. 1218–1226. DOI: 10.1063/1.462209.

[52] P. J. Knowles and N. C. Handy. "A new determinant-based full configuration interaction method". In: *Chemical Physics Letters* 111.4-5 (1984), pp. 315–321. DOI: 10.1016/0009-2614(84)85513-X.

[53] A. A. Holmes, N. M. Tubman, and C. J. Umrigar. "Heat-bath configuration interaction: an efficient selected configuration interaction algorithm inspired by heat-bath sampling". In: *Journal of Chemical Theory and Computation* 12.8 (2016), pp. 3674–3680. DOI: 10.1021/acs.jctc.6b00407.

[54] B. Huron, J. P. Malrieu, and P. Rancurel. "Iterative perturbation calculations of ground and excited state energies from multiconfigurational zeroth-order wavefunctions". In: *The Journal of Chemical Physics* 58.12 (1973), pp. 5745–5759. DOI: 10.1063/1.1679199.

[55] S. R. White. "Density matrix formulation for quantum renormalization groups". In: *Physical Review Letters* 69.19 (1992), pp. 2863–2866. DOI: 10.1103/PhysRevLett.69.2863.

[56] R. Bauernschmitt and R. Ahlrichs. “Treatment of electronic excitations within the adiabatic approximation of time dependent density functional theory”. In: *Chemical Physics Letters* 256.4-5 (1996), pp. 454–464. DOI: 10.1016/0009-2614(96)00440-X.

[57] J. B. Foresman, M. Head-Gordon, J. A. Pople, and M. J. Frisch. “Toward a systematic molecular orbital theory for excited states”. In: *The Journal of Physical Chemistry* 96.1 (1992), pp. 135–149. DOI: 10.1021/j100180a030.

[58] R. L. Martin. “Natural transition orbitals”. In: *The Journal of Chemical Physics* 118.11 (2003), pp. 4775–4777. DOI: 10.1063/1.1558471.

[59] W. von Niessen, J. Schirmer, and L. S. Cederbaum. “Computational methods for the one-particle Green’s function”. In: *Computer Physics Reports* 1.2 (1984), pp. 57–125. DOI: 10.1016/0167-7977(84)90002-9.

[60] E. R. Davidson. “The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices”. In: *Journal of Computational Physics* 17.1 (1975), pp. 87–94. DOI: 10.1016/0021-9991(75)90065-0.

[61] A. V. Knyazev. “Toward the optimal preconditioned eigensolver: locally optimal block preconditioned conjugate gradient method”. In: *SIAM Journal on Scientific Computing* 23.2 (2001), pp. 517–541. DOI: 10.1137/S1064827500366124.

[62] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg. “A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H–Pu”. In: *The Journal of Chemical Physics* 132.15 (2010), p. 154104. DOI: 10.1063/1.3382344.

[63] S. Grimme, S. Ehrlich, and L. Goerigk. “Effect of the damping function in dispersion corrected density functional theory”. In: *Journal of Computational Chemistry* 32.7 (2011), pp. 1456–1465. DOI: 10.1002/jcc.21759.

[64] E. Caldeweyher, S. Ehlert, A. Hansen, H. Neugebauer, S. Spicher, C. Bannwarth, and S. Grimme. “A generally applicable atomic-charge dependent London dispersion correction”. In: *The Journal of Chemical Physics* 150.15 (2019), p. 154122. DOI: 10.1063/1.5090222.

[65] R. Sure and S. Grimme. “Corrected small basis set Hartree–Fock method for large systems”. In: *Journal of Computational Chemistry* 34.19 (2013), pp. 1672–1685. DOI: 10.1002/jcc.23317.

[66] S. Grimme, A. Hansen, S. Ehlert, and J.-M. Mewes. “r^2SCAN-3c: a “Swiss army knife” composite electronic-structure method”. In: *The Journal of Chemical Physics* 154.6 (2021), p. 064103. DOI: 10.1063/5.0040021.

[67] H. Kruse and S. Grimme. “A geometrical correction for the inter- and intra-molecular basis set superposition error in Hartree–Fock and density functional theory calculations for large systems”. In: *The Journal of Chemical Physics* 136.15 (2012), p. 154101. DOI: 10.1063/1.3700154.

[68] V. Barone and M. Cossi. “Quantum calculation of molecular energies and energy gradients in solution by a conductor solvent model”. In: *The Journal of Physical Chemistry A* 102.11 (1998), pp. 1995–2001. DOI: 10.1021/jp9716997.

[69] A. Klamt and G. Schüürmann. “COSMO: a new approach to dielectric screening in solvents with explicit expressions for the screening energy and its gradient”. In: *Journal of the Chemical Society, Perkin Transactions 2* 5 (1993), pp. 799–805. DOI: 10.1039/P29930000799.

[70] J. Tomasi, B. Mennucci, and R. Cammi. “Quantum mechanical continuum solvation models”. In: *Chemical Reviews* 105.8 (2005), pp. 2999–3094. DOI: 10.1021/cr9904009.

[71] G. Scalmani and M. J. Frisch. "Continuous surface charge polarizable continuum models of solvation. I. General formalism". In: *J. Chem. Phys.* 132.11 (2010), p. 114110. DOI: 10.1063/1.3359469.

[72] R. A. Shaw and J. G. Hill. "Prescreening and efficiency in the evaluation of integrals over ab initio effective core potentials". In: *The Journal of Chemical Physics* 147.7 (2017), p. 074108. DOI: 10.1063/1.4986887.

[73] B. Metz, H. Stoll, and M. Dolg. "Small-core multiconfiguration-Dirac–Hartree–Fock-adjusted pseudopotentials for post-d main group elements: application to PbH and PbO". In: *The Journal of Chemical Physics* 113.7 (2000), pp. 2563–2569. DOI: 10.1063/1.1305880.

[74] M. Elstner, D. Porezag, G. Jungnickel, J. Elsner, M. Haugk, T. Frauenheim, S. Suhai, and G. Seifert. "Self-consistent-charge density-functional tight-binding method for simulations of complex materials properties". In: *Physical Review B* 58.11 (1998), pp. 7260–7268. DOI: 10.1103/PhysRevB.58.7260.

[75] C. Bannwarth, S. Ehlert, and S. Grimme. "GFN2-xTB — an accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions". In: *Journal of Chemical Theory and Computation* 15.3 (2019), pp. 1652–1671. DOI: 10.1021/acs.jctc.8b01176.

[76] J. J. P. Stewart. "Optimization of parameters for semiempirical methods V: modification of NDDO approximations and application to 70 elements". In: *Journal of Molecular Modeling* 13.12 (2007), pp. 1173–1213. DOI: 10.1007/s00894-007-0233-4.

[77] P. O. Dral, X. Wu, L. Spörkel, A. Koslowski, W. Weber, R. Steiger, M. Scholten, and W. Thiel. "Semiempirical quantum-chemical orthogonalization-corrected methods: theory, implementation, and parameters". In: *Journal of Chemical Theory and Computation* 12.3 (2016), pp. 1082–1096. DOI: 10.1021/acs.jctc.5b01046.

[78] B. Ahlswede and K. Jug. "Consistent modifications of SINDO1: I. Approximations and parameters". In: *Journal of Computational Chemistry* 20.6 (1999), pp. 563–571. DOI: 10.1002/(SICI)1096-987X(19990430)20:6<563::AID-JCC1>3.0.CO;2-2.

[79] I. Batatia, D. P. Kovács, G. N. C. Simm, C. Ortner, and G. Csányi. "MACE: higher order equivariant message passing neural networks for fast and accurate force fields". In: *Advances in Neural Information Processing Systems (NeurIPS)*. Vol. 35. 2022, pp. 11423–11436.

[80] A. Laio and M. Parrinello. "Escaping free-energy minima". In: *Proceedings of the National Academy of Sciences* 99.20 (2002), pp. 12562–12566. DOI: 10.1073/pnas.202427399.

[81] A. Barducci, G. Bussi, and M. Parrinello. "Well-Tempered Metadynamics: A Smoothly Converging and Tunable Free-Energy Method". In: *Phys. Rev. Lett.* 100.2 (2008), p. 020603. DOI: 10.1103/PhysRevLett.100.020603.

[82] C. Pisani, R. Dovesi, and C. Roetti. *Hartree–Fock Ab Initio Treatment of Crystalline Systems*. Vol. 48. Lecture Notes in Chemistry. Berlin, Heidelberg: Springer, 1988. DOI: 10.1007/978-3-642-93385-1.

[83] H. J. Monkhorst and J. D. Pack. "Special points for Brillouin-zone integrations". In: *Physical Review B* 13.12 (1976), pp. 5188–5192. DOI: 10.1103/PhysRevB.13.5188.

[84] P. Wisesa, K. A. McGill, and T. Mueller. "Efficient generation of generalized Monkhorst-Pack grids through the use of informatics". In: *Physical Review B* 93.15 (2016), p. 155109. DOI: 10.1103/PhysRevB.93.155109.

[85] A. Togo and I. Tanaka. "Spglib: a software library for crystal symmetry search". In: *arXiv preprint* (2018). arXiv: 1808.01590.

[86] K. Choudhary and F. Tavazza. "Convergence and machine learning predictions of Monkhorst-Pack k-points and plane-wave cut-off in high-throughput DFT calculations". In: *Computational Materials Science* 161 (2019), pp. 300–308. DOI: 10.1016/j.commatsci.2019.02.006.

[87] M. Methfessel and A. T. Paxton. "High-precision sampling for Brillouin-zone integration in metals". In: *Physical Review B* 40.6 (1989), pp. 3616–3621. DOI: 10.1103/PhysRevB.40.3616.

[88] N. Marzari, D. Vanderbilt, A. De Vita, and M. C. Payne. "Thermal contraction and disordering of the Al(110) surface". In: *Physical Review Letters* 82.16 (1999), pp. 3296–3299. DOI: 10.1103/PhysRevLett.82.3296.

[89] H.-Z. Ye and T. C. Berkelbach. "Fast periodic Gaussian density fitting by range separation". In: *J. Chem. Phys.* 154.13 (2021), p. 131104. DOI: 10.1063/5.0046617.

[90] Q. Sun, T. C. Berkelbach, J. D. McClain, and G. K.-L. Chan. "Gaussian and plane-wave mixed density fitting for periodic systems". In: *The Journal of Chemical Physics* 147.16 (2017), p. 164119. DOI: 10.1063/1.4998644.

[91] R. Dovesi, C. Pisani, C. Roetti, and V. R. Saunders. "Treatment of Coulomb interactions in Hartree-Fock calculations of periodic systems". In: *Phys. Rev. B* 28.10 (1983), pp. 5781–5792. DOI: 10.1103/PhysRevB.28.5781.

[92] G. Lippert, J. Hutter, and M. Parrinello. "A hybrid Gaussian and plane wave density functional scheme". In: *Molecular Physics* 92.3 (1997), pp. 477–488. DOI: 10.1080/00268979709482119.

[93] J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter. "Quickstep: fast and accurate density functional calculations using a mixed Gaussian and plane waves approach". In: *Computer Physics Communications* 167.2 (2005), pp. 103–128. DOI: 10.1016/j.cpc.2004.12.014.

[94] M. Krack and M. Parrinello. "All-electron ab-initio molecular dynamics". In: *Physical Chemistry Chemical Physics* 2.10 (2000), pp. 2105–2112. DOI: 10.1039/b001167n.

[95] D. G. Anderson. "Iterative procedures for nonlinear integral equations". In: *Journal of the ACM* 12.4 (1965), pp. 547–560. DOI: 10.1145/321296.321305.

[96] G. P. Kerker. "Efficient iteration scheme for self-consistent pseudopotential calculations". In: *Physical Review B* 23.6 (1981), pp. 3082–3084. DOI: 10.1103/PhysRevB.23.3082.

[97] S. L. Dudarev, G. A. Botton, S. Y. Savrasov, C. J. Humphreys, and A. P. Sutton. "Electron-energy-loss spectra and the structural stability of nickel oxide: an LSDA+U study". In: *Physical Review B* 57.3 (1998), pp. 1505–1509. DOI: 10.1103/PhysRevB.57.1505.

[98] T. Hughbanks and R. Hoffmann. "Chains of trans-edge-sharing molybdenum octahedra: metal-metal bonding in extended systems". In: *Journal of the American Chemical Society* 105.11 (1983), pp. 3528–3537. DOI: 10.1021/ja00349a027.

[99] R. Dronskowski and P. E. Blöchl. "Crystal orbital Hamilton populations (COHP): energy-resolved visualization of chemical bonding in solids based on density-functional calculations". In: *The Journal of Physical Chemistry* 97.33 (1993), pp. 8617–8624. DOI: 10.1021/j100135a014.

[100] I. Mayer. "Charge, bond order and valence in the AB initio SCF theory". In: *Chemical Physics Letters* 97.3 (1983), pp. 270–274. DOI: 10.1016/0009-2614(83)80005-0.

[101] R. F. W. Bader. *Atoms in Molecules: A Quantum Theory*. Vol. 22. International Series of Monographs on Chemistry. Oxford: Oxford University Press, 1990.

[102] Y. Hinuma, G. Pizzi, Y. Kumagai, F. Oba, and I. Tanaka. "Band structure diagram paths based on crystallography". In: *Computational Materials Science* 128 (2017), pp. 140–184. DOI: 10.1016/j.commatsci.2016.10.015.

[103] G. Guennebaud, B. Jacob, et al. *Eigen v3*. https://eigen.tuxfamily.org. 2010.

[104] W. Jakob, J. Rhinelander, and D. Moldovan. *pybind11 — seamless operability between C++11 and Python*. https://github.com/pybind/pybind11. 2017.

[105] E. F. Valeev. *Libint: a library for the evaluation of molecular integrals of many-body operators over Gaussian functions*. https://github.com/evaleev/libint. Version 2.13.1. 2025.

[106] M. Frigo and S. G. Johnson. "The design and implementation of FFTW3". In: *Proceedings of the IEEE* 93.2 (2005), pp. 216–231. DOI: 10.1109/JPROC.2004.840301.

[107] M. F. Peintinger, D. Vilela Oliveira, and T. Bredow. "Consistent Gaussian basis sets of triple-zeta valence with polarization quality for solid-state calculations". In: *Journal of Computational Chemistry* 34.6 (2013), pp. 451–459. DOI: 10.1002/jcc.23153.

[108] D. Vilela Oliveira, J. Laun, M. F. Peintinger, and T. Bredow. "BSSE-correction scheme for consistent Gaussian basis sets of double- and triple-zeta valence with polarization quality for solid-state calculations". In: *Journal of Computational Chemistry* 40.27 (2019), pp. 2364–2376. DOI: 10.1002/jcc.26013.

[109] M. F. Peintinger. *Quantum chemistry needs a modern file format (QVF)*. https://peintinger.com/2026/05/21/qc-needs-a-modern-file-format-qvf/. May 2026.

[110] Anthropic. *Claude Opus 4.8*. https://www.anthropic.com/claude. Large language model. Used as a code-generation collaborator under human supervision. 2026.

[111] Anthropic. *Claude Opus 4.7*. https://www.anthropic.com/claude. Large language model. Used as a code-generation collaborator under human supervision. 2026.

[112] Anthropic. *Claude Opus 5*. https://www.anthropic.com/claude. Large language model. Used as a code-generation collaborator under human supervision. 2026.

[113] Anthropic. *Claude Fable 5*. https://www.anthropic.com/news/claude-fable-5-mythos-5. Large language model. Used as a code-generation collaborator under human supervision. 2026.

[114] DeepSeek. *DeepSeek Flash*. https://www.deepseek.com. Large language model used as a secondary code-generation and review aid. 2026.

[115] DeepSeek. *DeepSeek Pro*. https://www.deepseek.com. Large language model used as a secondary code-generation and review aid. 2026.

[116] DeepSeek. *DeepSeek v4 Pro*. https://www.deepseek.com. Large language model used as a secondary code-generation and review aid. 2026.

[117] OpenAI. *Codex 5.5*. https://openai.com. Code-generation model used as a secondary implementation aid. 2026.

[118] OpenAI. *ChatGPT 5.6 Sol*. https://openai.com. Large language model used as a secondary debugging and refactoring aid. 2026.

[119] M. Véril, A. Scemama, M. Caffarel, F. Lipparini, M. Boggio-Pasqua, D. Jacquemin, and P.-F. Loos. "QUESTDB: A database of highly accurate excitation energies for the electronic structure community". In: *WIREs Comput. Mol. Sci.* 11.5 (2021), e1517. DOI: 10.1002/wcms.1517.

[120] V. K. Prasad, A. Otero-de-la-Roza, and G. A. DiLabio. "BH9, a New Comprehensive Benchmark Data Set for Barrier Heights and Reaction Energies: Assessment of Density Functional Approximations and Basis Set Incompleteness Potentials". In: *J. Chem. Theory Comput.* 18.1 (2021), pp. 151–166. DOI: 10.1021/acs.jctc.1c00694.

[121] D. G. Liakos, Y. Guo, and F. Neese. "Comprehensive Benchmark Results for the Domain Based Local Pair Natural Orbital Coupled Cluster Method (DLPNO-CCSD(T)) for Closed- and Open-Shell Systems". In: *J. Phys. Chem. A* 124.1 (2019), pp. 90–100. DOI: 10.1021/acs.jpca.9b05734.

[122] P. Jurečka, J. Šponer, J. Černý, and P. Hobza. "Benchmark database of accurate (MP2 and CCSD(T) complete basis set limit) interaction energies of small model complexes, DNA base pairs, and amino acid pairs". In: *Phys. Chem. Chem. Phys.* 8.17 (2006), pp. 1985–1993. DOI: 10.1039/B600027D.

[123] M. F. Peintinger and T. Bredow. "The cyclic cluster model at Hartree–Fock level". In: *Journal of Computational Chemistry* 35.11 (2014), pp. 839–846. DOI: 10.1002/jcc.23550.