

How to Use Quantum Computers for Biomolecular Free Energies

Jakob Günther, Thomas Weymuth, Moritz Bensberg, Freek Witteveen, Matthew S. Teynor, F. Emil Thomasen, Valentina Sora, William Bro-Jørgensen, Raphael T. Husistein, Mihael Erakovic, Marek Miller, Leah Weisburn, Minsik Cho, Marco Eckhoff, Aram W. Harrow,* Anders Krogh,* Troy Van Voorhis,* Kresten Lindorff-Larsen,* Gemma Solomon,* Markus Reiher,* and Matthias Christandl*

Cite This: *J. Chem. Theory Comput.* 2026, 22, 4329–4345


Read Online

ACCESS |

Metrics & More

Article Recommendations

ABSTRACT: Free energy calculations are at the heart of physics-based analyses of biochemical processes. They allow us to quantify molecular recognition mechanisms, which determine a wide range of biological phenomena, from how cells send and receive signals to how pharmaceutical compounds can be used to treat diseases. Quantitative and predictive free energy calculations require computational models that accurately capture both the varied and intricate electronic interactions between molecules as well as the entropic contributions from the motions of these molecules and their aqueous environment. However, accurate quantum-mechanical energies and forces can be obtained only for small atomistic models and not for large biomacromolecules. Here, we demonstrate how to consistently link accurate quantum-mechanical data obtained for substructures to the overall potential energy of biomolecular complexes using machine learning in an integrated algorithm. We do so using a two-fold quantum embedding strategy where the innermost quantum cores are treated at a very high level of accuracy. We demonstrate the viability of this approach for the molecular recognition of a ruthenium-based anticancer drug by its protein target by applying traditional quantum chemical methods. As such methods scale unfavorably with system size, we analyze the requirements for quantum computers to provide highly accurate energies that affect the resulting free energies. Once the requirements are met, our computational pipeline, FreeQuantum, is able to make efficient use of the quantum-computed energies, thereby enabling quantum computing-enhanced modeling of biochemical processes. This approach combines the exponential speedups of quantum computers for simulating interacting electrons with modern classical simulation techniques that incorporate machine learning to model large molecules.



1. INTRODUCTION

Free energy calculations represent the state-of-the-art method for predicting the affinity between biomolecules in aqueous solution. These methods provide an atomistic understanding of molecular recognition and ligand binding.^{1–3} This understanding is key to deciphering the mechanisms that underlie processes in cells, with far-reaching consequences for biological function and hence for disease and health. For these reasons, free energy calculations are increasingly implemented as part of strategies for discovering and improving small molecule modulators of biological function.^{4,5}

For biomolecular recognition driven by a change in free energy, the prototypical example is the interaction of a ligand (such as a small-molecule drug) with a biomacromolecular target, typically a protein. The ability to predict such free energy differences with high accuracy and speed—and for a wide range of types of molecules—would enable a revolutionary ability to describe biological processes at the molecular level, engineer proteins for predefined functions, and improve the treatment of diseases via improved drugs. The accurate calculation of the free energy difference of such a binding process, however, requires one to deal with the enormous size

of the configuration space that must be sampled for a proper representation of the thermodynamical ensemble. For each sampled atomistic structure, furthermore, an accurate estimate of the interaction energy needs to be computed.

While the sampling is a computational problem that in principle can be dealt with by well-established approaches,^{6,7} the interaction energy computation would be more accurate when based on a quantum-mechanical description of the electrons, in particular for complex molecular systems that are difficult to handle with standard force fields. Machine learning-based potentials provide a potential route to bridge the accuracy of quantum chemical calculations with the sampling needed for accurate biological free energy calculations. Previous work in this area has, for example, used various flavors of such potentials—typically trained on density

Received: December 15, 2025

Revised: February 24, 2026

Accepted: February 24, 2026

Published: April 19, 2026



functional theory (DFT) calculations across a broad set of fragment molecules—to perform free energy calculations.^{8–12}

Accurate quantum-mechanical calculations carried out by traditional computers, however, scale highly unfavorably with system size and are, therefore, almost always out of reach for routine studies of ligand binding. The basic reason for this can be found in the exponential overhead in the memory requirement for the description of a quantum-mechanical system on a traditional classical computer, even if only the protein–ligand interface is treated in a fully quantum-mechanical way. A future quantum computer, by contrast, is in principle able to represent the entire wave function, since the number of qubits required scales linearly with the number of electrons, opening the potential to overcome this ‘curse of dimensionality’.^{13–15}

To realize this potential, energy-sampling methods must leverage highly accurate quantum-mechanical calculations; however, current computational pipelines are not set up to handle such high-level methods—neither for calculations that can be carried out with traditional quantum chemical methods today nor for those carried out with future quantum computers. Without this capability, the promise of quantum computing for broad application in biochemistry remains unclear.

We note that in different biomolecular situations, there are different levels of importance of the energetic contributions when compared to the entropic contributions. In Figure 1, we sketch this ‘biomolecular quantum simulation quadrangle’ with low versus high complexity of energy when compared to entropy and give key examples. The situation with low entropic complexity ranges from simple closed-shell molecules (such as a single water molecule), which can be tackled with traditional

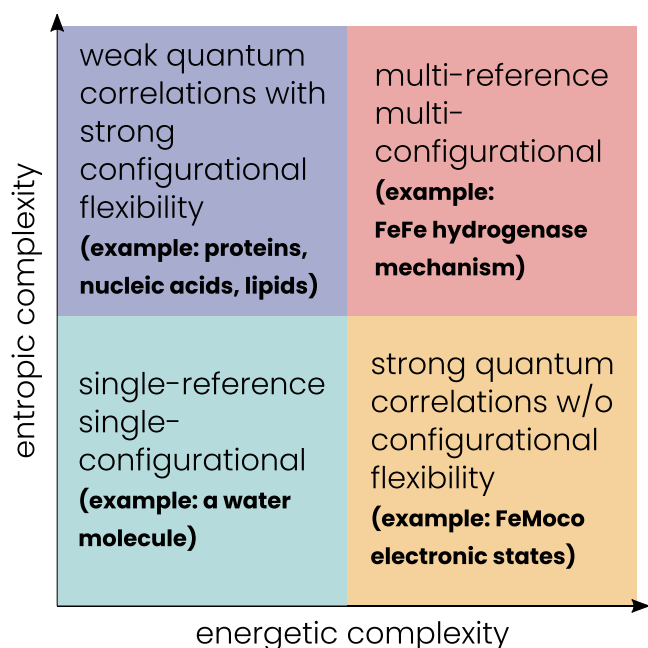


Figure 1. Biomolecular quantum simulation quadrangle: the complexity of the electronic structure problem is indicated on the horizontal axis. The importance and complexity of the sampling problem are indicated on the vertical axis. Considering the four paradigmatic situations where the energetic and entropic complexities are low or high gives rise to four quadrants. In each quadrant, we indicate a characteristic example.

methods, to open-shell 3d transition-metal clusters (such as the FeMo-cofactor of nitrogenase) for which traditional methods struggle but quantum computing is promising.¹⁶ The situation with high entropic complexity ranges from the case of comparatively weak electronic correlations (including most biochemical situations) to the case where high electronic complexity meets large configurational freedom (e.g., the FeFe hydrogenase, which involves an iron–sulfur cluster as an active site whose mechanism depends on the configurational flexibility of the surrounding protein scaffold.^{17–19} We note that for many systems in the top-left corner of the quadrangle, standard force fields can, in principle, perform well, especially if the system only has top-row elements from the periodic table. However, in the presence of transition metals, accuracy is often poor.

In previous work,^{20,21} we developed an approach exploiting the concept of hybrid machine learning/molecular mechanics models,^{22–24} including nonequilibrium switching corrections for free energy calculations,^{25–28} to predict binding free energies accurately. There, we already demonstrated that (i) we can efficiently represent a hybrid low-cost-quantum–classical model of a protein–guest complex by a machine learning potential from which a free energy of binding can be obtained²⁰ and that (ii) this machine learning model can be improved by transfer learning with local high-accuracy quantum energies obtained in a second embedding by traditional computations,²¹ where small quantum cores are embedded in the large, low-cost quantum region of the hybrid model.

In this work, we present a framework and implementation of an end-to-end pipeline for the calculation of free energies that makes efficient use of expensive, high-accuracy quantum-mechanical calculations. We demonstrate the accuracy, reliability, and plasticity of our algorithmic workflow using the example of a ruthenium-based anticancer compound binding to a target protein (GRP78/NKP1339), where the quantum data are obtained by traditional quantum chemical wave function-based methods. Our pipeline has the flexibility to allow for the direct replacement of this traditional computing engine with a quantum computing engine. We provide requirements for when this replacement will impact the free energy calculation. This is an important step toward charting and realizing the conditions for quantum advantage in computing free energies for large and complex biological systems.

For a graphical representation of our FreeQuantum pipeline, see Figure 2.

We demonstrate the FreeQuantum pipeline by calculating the free energy of binding of the aforementioned open-shell, transition-metal-containing anticancer drug to its protein target using a series of traditional correlated electronic structure methods. Based on our earlier work developing qubit-efficient quantum algorithms for ground state energy estimation,^{29–31} as well as algorithms based on the qubitization framework,^{32–34} we can exploit this example to arrive at general conclusions for the constraints on future quantum computers to achieve a quantum advantage in the general field of atomistic modeling for free energy calculations in biochemistry.

The remainder of this paper is organized as follows. In Section 2, we describe the scientific aspects of the computational pipeline as well as the results from a pipeline run for the anticancer drug–host complex obtained with traditional

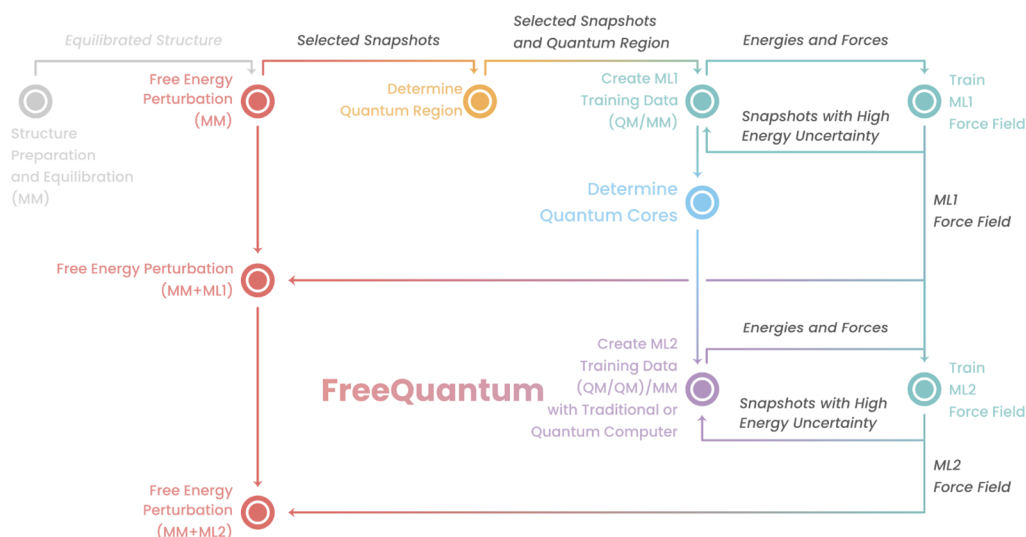


Figure 2. Workflow of the FreeQuantum pipeline: after a structure preparation step for the host protein and bound guest molecule by equilibration (gray), structural sampling with classical force fields starts in an initial alchemical free energy perturbation (FEP) step (red, top). The structures are then forwarded to the hybrid QM/MM modeling step, which defines the quantum region that covers the parts of the protein–ligand system that must be subjected to a quantum-mechanical description (orange). For the resulting QM/MM structures with associated energies and forces, a first machine learning (ML) potential (ML1) is trained with an active learning loop to interlace the QM subsystem data with the MM force field data (turquoise, top). The loop involves nonequilibrium (NEQ) switching calculations, the final result of which is output as the MM + ML1 FEP (red, middle). In a further step, the QM/MM hybrid model is refined by introducing a second type of embedding, where quantum-in-quantum subregions (quantum cores) are defined (blue), for which highly accurate quantum-mechanical calculations are then carried out, either by traditional quantum chemical calculations or by future quantum computation (violet). The resulting QM/QM/MM data are fed into a refining process of the earlier ML model (ML1) by transfer learning (turquoise, bottom). Sampling from this refined model (ML2) yields the third tier of free energy results (red, bottom). With the active learning and the FEP calculations, we automated all parts of the pipeline, which otherwise would require frequent manual intervention and substantial human time.

computational methods. In Section 3, we discuss the prospects of future quantum computing methods. Emerging from an analysis of the quantum resources required for impacting the free energy calculation for the ruthenium-based system, we offer general insights into quantum advantages in computational molecular biology. A conclusion is offered in Section 4. Additional technical information is provided in the Appendix.

2. THE COMPUTATIONAL PIPELINE FREEQUANTUM

The partition function Z of a molecular system in the Born–Oppenheimer approximation is given by

$$Z = \int d^{3N_{\text{nuc}}}R \text{tr} \exp(-\beta H(R)) \quad (1)$$

where $\beta = \frac{1}{k_{\text{B}}T}$ is the inverse temperature with k_{B} the Boltzmann constant and T the temperature. $H(R)$ is the electronic Hamiltonian for given nuclear coordinates R of the N_{nuc} nuclei. We include in $H(R)$ the electrostatic interaction of the nuclei. At ambient temperature in biological systems, higher electronic states can be neglected, resulting in the approximation $\text{tr} \exp(-\beta H(R)) \approx \exp(-\beta E_{\text{g}}(R))$, where $E_{\text{g}}(R)$ is the electronic ground state energy for a fixed nuclear configuration; varying over R gives the potential energy surface (PES).

Thermodynamic quantities, such as the Gibbs free energy G can be obtained from Z and therefore involve a sampling problem over a vast PES. We do this by traditional molecular dynamics methods, which carve out a discretized trajectory of nuclear positions with the help of a force field. We can restrict to states with specific properties, e.g., bound or unbound for our protein–ligand system, and indicate this in the subscript.

Our main goal is the computation of the difference

$$\Delta G_{\text{binding}} = G_{\text{bound}} - G_{\text{unbound}} \quad (2)$$

to high accuracy as this is key to any design attempt that would rely on differences of $\Delta G_{\text{binding}}$ obtained for different host molecules (for instance, within a drug discovery campaign). For each step in the trajectory, we therefore have to compute the ground state energy, i.e., the corresponding point on the PES with high accuracy. Ideally, we would employ a full quantum-mechanical energy calculation. As the system consists of thousands of atoms, this is infeasible, and we resort to a two-fold quantum embedding strategy resulting in a three-layer quantum-in-quantum-in-classical embedding (QM/QM/MM).

As in refs 20 and 21, we begin by considering the initial force field (MM) that has been used for the trajectory computation. We then choose the quantum region, which will be treated with DFT, giving rise to the QM/MM embedding. As the computation of the DFT data is more expensive, we do not compute it for every point on the trajectory but only for some, which we then use to train an ML potential in an active learning loop (ML1). To obtain further refinement, we identify one or more quantum cores within the quantum region, giving rise to the QM/QM/MM embedding. The energies of the quantum cores can be computed with traditional quantum chemical methods (currently) or quantum computing (in the future). In either case, these data points are used to obtain a refined ML potential (ML2), replacing ML1 via transfer learning.

In the following, we describe the main aspects of the implementation: first, regarding the free energy perturbation

method; second, regarding the ML potentials; and third, regarding the quantum embedding.

2.1. Free Energy Perturbation

The calculation of the binding free energy $\Delta G_{\text{binding}} = G_{\text{bound}} - G_{\text{unbound}}$ by direct computation of the free energies as outlined above is challenging, as the molecular dynamics simulation used in order to sample the phase space is slow and would therefore only explore a very small part. This leads to large uncompensated errors when computing free energy differences, since the sampling for G_{bound} and G_{unbound} is over distributions with little overlap.

Free energy perturbation (FEP) offers a way of avoiding this difficulty by telescoping the expression

$$G_{\text{bound}} - G_{\text{unbound}} = \sum_{k=1}^{s-1} \frac{G_{\lambda_k} - G_{\lambda_{k+1}}}{\Delta G_{\lambda_k}} \quad (3)$$

with s steps along a path (e.g., a reaction coordinate) parametrized by λ (see ref 15 and Figure 7). Here, $1 = \lambda_1 > \dots > \lambda_s = 0$, where 0 and 1 correspond to the unbound and bound situation,⁶ (Section 1.3), and where G_{λ_k} and $G_{\lambda_{k+1}}$ involve very similar sampling when suitable paths are chosen, giving rise to error compensation and thus to an effective algorithm.⁷

An elegant alternative strategy that connects the thermodynamically stable states with a path is to introduce an unphysical parameter into the Hamiltonian and to change this parameter so as to walk along the resulting 'alchemical' path. In protein–ligand binding, this parameter simply switches off the interaction between protein and ligand:

$$H_{\lambda} = H_{\text{protein}} + H_{\text{ligand}} + \lambda H_{\text{interaction}} \quad (4)$$

This defines G_{λ} with $G_{\text{bound}} = G_1$ and $G_{\text{unbound}} = G_0 = G_{\text{protein}} + G_{\text{ligand}}$.

This idea is especially useful, as it can be easily adapted to tackle the solvated situation, where it is hard to find a reaction coordinate path due to the presence of water molecules. Here, one first switches off the protein–ligand interaction in the solvated complex. The result corresponds to the sum of the solvated protein and the ligand in vacuum. Considering now the ligand separately, one can switch on the interaction with surrounding water molecules, resulting in

$$\begin{aligned} \Delta G_{\text{solvated binding}} &= \frac{G_{\text{solvated complex}} - (G_{\text{solvated protein}} + G_{\text{ligand}})}{\Delta G_{\text{partially solvated binding}}} \\ &+ \frac{(G_{\text{ligand}} - G_{\text{solvated ligand}})}{-\Delta G_{\text{ligand solvation}}} \end{aligned} \quad (5)$$

The telescoping sum (3) can be used for each of the two free energy differences on the right, leading to the alchemical FEP method that we implemented.^{35–37} Here, the free energy differences ΔG_{λ_k} are determined jointly using the multistate Bennett acceptance ratio (MBAR),³⁸ which has low variance and bias. We carefully choose the step sizes to ensure sufficient overlap of the potential energy distributions along the steps taken on the path, so that the free energies computed with MBAR are reliable.

2.2. Machine Learning Potentials

Training a system-specific ML potential from scratch requires a training set with a large number of representative structures. It is common to use both energy and force information (analytical derivatives), and this is done within an active learning loop in order to obtain ML1. Starting with 2000 QM/MM data points, as is done for similar ML potentials, we use 90% of the random structures for training and the rest for validation and train the potential until the root-mean-square error has converged. We then use this potential for nonequilibrium switching (NEQ), a step in our FEP computation during which new structures with high uncertainties are encountered. At this point, the active learning process begins, new QM/MM reference data for these structures are generated, and the training is repeated as described until convergence.


For the second embedding leading to ML2, active learning is a challenge, as the forces are not directly available in our embedding approaches (see below), and training solely on energies requires significantly more training data, as noted in the training of ML1.²¹ We address this challenge using transfer learning, which allows us to refine ML1 directly with sparse high precision data, a strategy that has previously shown good results.^{39,20} To select the structures, we calculate the median energy of the QM/MM energies and then calculate QM/QM/MM energies for all structures whose QM/MM energies are within 400 kJ/mol of the median. This eliminates artifacts from unphysical structures. In total, we used 4570 reference conformers for the protein–ligand complex and 5441 reference conformers for the solvated ligand in order to refine the ML1 potential energy surface in the transfer learning step to ML2.

2.3. Quantum Embedding

In the following, we will explain the construction and choices made in the quantum-mechanically refined PESs. To begin with, the free energy sampling is conducted via molecular dynamics simulations on the PES corresponding to a standard force field, with the equilibrated structure as a starting point. The energies and forces are expected to be inaccurate in certain areas of the PES due to quantum-mechanical effects. It is, therefore, desirable to switch to a quantum-mechanical description of the protein–ligand system. Since the system is too large, this is out of reach, even in principle, with traditional, high-accuracy classical computation. Such systems could, in principle, be represented on a quantum computer, but they are expected to remain out of reach in practice for at least the coming decades. A full quantum-mechanical representation of the system, however, is also not needed as the relevant quantum-mechanical effects in biochemical systems are localized. Furthermore, the errors in some localized regions are likely to be the same throughout all of the calculations, therefore leading to error cancellation since we are looking at a free energy difference. In the case of protein–ligand binding, the most relevant region is the interface of protein and ligand, in our case including the ligand, as it is small.

Quantum embedding methods allow us to define and incorporate a quantum region within a larger classical environment. The total energy is then the sum of the energy of the outer part (E_{MM}) and the ground state energy of an electronic structure Hamiltonian of the quantum region in an effective MM potential from the outer part

$$H_{\text{quantum region}} = H_{\text{electronic, quantum region}} + V_{\text{interaction, MM}} \quad (6)$$



Method	MCL-1/19G	GRP78/NKP-1339
MM	-37.5 ± 0.4 [20]	-19.1 ± 1.5 [20]
PBE/MM	-35.3 ± 1.8 [20]	-17.0 ± 2.6 [20]
UMP2/PBE/MM	—	-9.6 ± 2.2 this work
LCCSD(T)/PBE/MM	-37.2 ± 1.0 [21]	—
UCCSD(T)/PBE/MM	—	-10.8 ± 2.4 this work
NEVPT2/PBE/MM	—	-11.3 ± 2.9 this work
Experiment	-37.3 ± 0.1 [51]	—

Figure 3. (a) Host–guest complex of the chaperone BiP (binding immunoglobulin protein) (the protein host) with the ruthenium transition-metal complex (the small-molecule guest, highlighted in the blue circle in ball–stick representation, with carbon atoms in gray, chlorine atoms in green, nitrogen atoms in blue, hydrogen atoms in white, and Ru in orange; the Lewis structure is given in the lower left corner). The QM region in the QM/MM model is the complete Ru–drug molecule, whereas the quantum core for the QM-in-QM embedding is highlighted in the blue hexagon of the Lewis structure (a larger version of this figure is available in the appendix in Figure 9). (b) Overview of the binding free energies (in kJ/mol) obtained at different stages of the FreeQuantum pipeline and with different quantum chemical methods to treat the quantum core (see Methods for an explanation of all acronyms). Next to the results for the ruthenium drug–protein complex, we provide our previous results for another protein–ligand system, the myeloid cell leukemia 1 (MCL-1) protein, dysregulation of which is associated with various cancers, inhibited by the small organic closed-shell molecule 19G, for which classical force fields (MM) work very well and an experimental reference is available. The error bars are determined as in ref 21, eq 14.

As the required quantum region would still involve thousands of orbitals, an exact quantum-mechanical treatment remains out of the question for traditional methods and also for the foreseeable future on quantum computers. An approximate treatment with DFT is possible but leaves us with large and uncontrolled errors, which are not systematically improvable. As those DFT errors are expected to be further localized, we resort to an additional quantum-in-quantum (QM/QM) embedding. Here, we identify one or more smaller ‘quantum cores’ inside the quantum region, which we treat with accurate wave function methods embedded in the larger quantum region, which we treat with DFT. The Hamiltonian for the core consists of the electronic part plus an additional external potential coming from the quantum-in-quantum embedding:

$$H_{\text{quantum core}} = H_{\text{electronic, quantum core}} + V_{\text{interaction, DFT}} \quad (7)$$

Thus, we obtain a quantum-in-quantum-in-classical embedding (QM/QM/MM).

While it is desirable to let the quantum region cover the entire protein–ligand interface, which includes the small ligand, we here focused the quantum region on the ligand only in our ruthenium-compound pipeline run. We note that this still incorporates some of the quantum-mechanical interaction effects between host and ligand due to the external potential on the ligand (see eq 6). We have also chosen the quantum region and the quantum cores manually, but note that an automated algorithm for quantum region selection is implemented and can be used in the future.^{40,41}

2.4. Quantum Engines

The pipeline is set up such that the size of the quantum cores, to be treated with wave function-based methods, can be adjusted both to the protein–ligand complex at hand and to the computational resources available. At present, the computational resources consist of high-performance computers with which we have carried out quantum chemical calculations, as we discuss below. We highlight again that the curse of dimensionality will ultimately limit the ability of these computations to capture the correlations of large quantum cores, and at this point, it would be advantageous to replace them with quantum computations. The development of both

quantum algorithms and quantum hardware is rapid and has the potential to be a game changer within our approach, as we detail below.

In the full run of our pipeline on the ruthenium drug–protein complex, we have used a Huzinaga-type projection embedding.²¹ Here, a single core is created inside the quantum region. Since not all of its orbitals can be considered in a single energy calculation due to its large size, we have resorted to a complete active space (CAS) approach, which can deliver the static correlation with the help of a full configuration interaction (FCI) wave function in a restricted space of orbitals.⁴² The choice of active orbitals has been made with the autoCAS algorithm.⁴³ This algorithm enables a straightforward and fully automated selection of active spaces in a rigorous way.⁴⁴ We rely on the density matrix renormalization group (DMRG) as a proxy for the full configuration interaction (FCI) calculation.⁴⁵ To account for the dynamic correlations that cannot be captured within the CAS realm, we use second-order N-electron valence state perturbation theory (NEVPT2).⁴⁵

To achieve high flexibility for the application of quantum algorithms, FreeQuantum can employ a second, complementary embedding strategy based on bootstrap embedding.^{46,41} Here, several smaller overlapping quantum subsystems are chosen within one quantum core, each of which is subjected to an energy calculation before being embedded in a consistent manner into the (larger) quantum core. The size of these subsystems, however, affects the accuracy (see the Appendix Figure 8 for results on our Ru–drug example), but this can, in fact, be increased with future quantum computing methods.

In general, both the Huzinaga-based projection approach and the bootstrap approach can exploit quantum computational resources in different ways. For a multiconfigurational (strong correlation) problem, for instance, projection embedding would allow FreeQuantum to choose an active orbital space within a quantum core for an exact diagonalization approach, with subsequent treatment of the remaining dynamic correlation, whereas bootstrap embedding would dissect the quantum core into smaller overlapping subsystems, so that all orbitals per subsystem can be subjected to exact

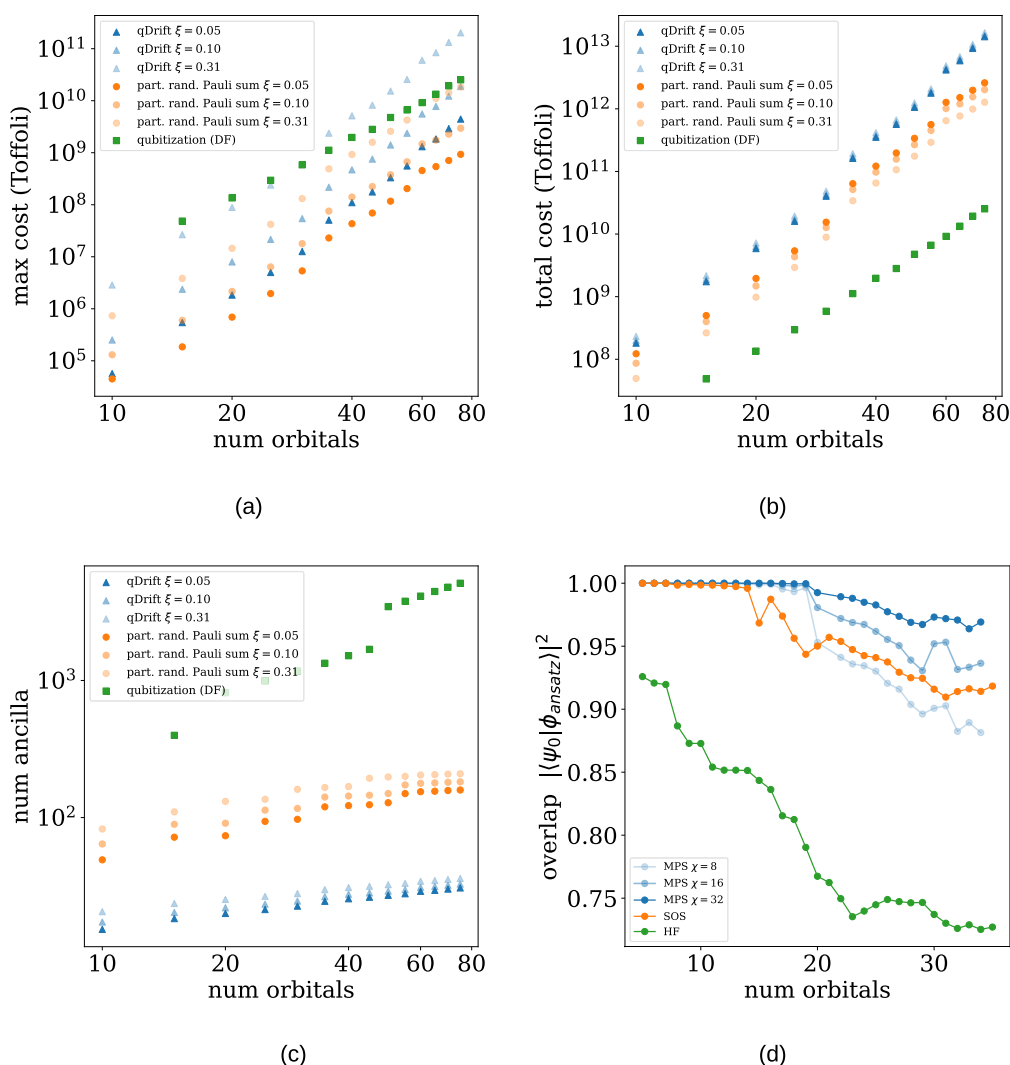


Figure 4. (a) Maximal number of gates per circuit with target accuracy of 0.001 Hartree for active spaces of the ruthenium ligand for partially randomized Trotter (in dependence of overlap) and qubitization. $\xi = \arcsin \frac{1-\eta}{\eta}$, where η is the overlap. (b) Total number of gates per circuit with target accuracy 0.0016 Hartree (4.2 kJ/mol) for active spaces of the ruthenium ligand for partially randomized Trotter (in dependence of overlap) and qubitization. (c) Number of qubits needed in dependence of the above algorithms and overlaps. The product-formula-based methods can avoid all ancilla qubits, at an increased cost of synthesizing single-qubit rotations. (d) Overlap with the ground state for different types of guiding states for the ruthenium system with increasing number of orbitals in the active space. As a proxy for the ground state, we used a converged DMRG state. In orange, the Sum-of-Slater (SOS) overlap with $4 \times N$ Slater determinants, where N is the number of orbitals of the active space. In blue is the overlap of a matrix product state (MPS) obtained by truncating the DMRG state to the indicated bond dimension χ , and in green is the overlap with the Hartree–Fock state (HF). This indicates that good overlap states that can be prepared on a quantum computer exist for our system of interest. (a–d) For details regarding the graphs, see Appendix.

diagonalization (rendering a dynamic-correlation treatment unnecessary).

2.5. Example Application of the Pipeline: Ruthenium-Based Anticancer Drug

We have run the full FreeQuantum pipeline on the binding of a small-molecule drug to a protein. The chosen molecule NKP1339 (synonyms are IT-139, KP-1339)^{47,48} causes cell death through apoptosis via several mechanisms, one of which is the inhibition of the Hsp70-family molecular chaperone BiP (binding immunoglobulin protein), which regulates the unfolded protein response.^{49,50} An illustration is shown in Figure 3a.

As a case of molecular recognition of a small molecule by a large biomacromolecule, which always requires sampling of the configuration space accessible to the emerging host–guest

complex and comparing it to the configuration space relevant for the separated molecules in solution, the chosen case represents, as desired, a high-configurational situation. Metal atoms and their coordination spheres are usually hard to describe with classical force fields. The ruthenium transition-metal complex, furthermore, makes the whole system open shell. The electronic ground state is a spin doublet, further highlighting the need for a quantum-mechanical description. The resulting metal–drug protein complex, which belongs to the top left in the biomolecular quantum simulation quadrangle, is therefore an excellent example to demonstrate the developed FreeQuantum pipeline.

The full run predicts binding with a binding free energy of $\Delta G_{\text{binding}}^{\text{MM}} = -19.1 \pm 1.5$ kJ/mol for a pure MM description and $\Delta G_{\text{binding}}^{\text{MM+ML1}} = -17.0 \pm 2.6$ kJ/mol for the first embedding

approach. The QM description in this first embedding hybrid model was based on DFT, with all shortcomings of approximate exchange–correlation density functionals. Therefore, it is an excellent case for the three-layer QM/QM/MM description, for which we obtained $\Delta G_{\text{binding}}^{\text{MM+ML}^2} = -11.3 \pm 2.9$ kJ/mol for our most advanced wave function methods for the core (NEVPT2 on top of CAS configuration interaction, CAS-CI). Figure 3b presents these results together with the data from a QM/MM embedding based on DFT²⁰ (right column). Figure 7 presents the distribution of energies of the embedded cores for the first and second embeddings, as well as their difference, illustrating that the embedding strategy fundamentally changes the energy landscape.

In Figure 3b, we also provide in the middle column corresponding data for a different system involving an organic closed-shell ligand denoted MCL-1/19G.^{20,21} We have used this system, for which reliable classical force field and experimental results are available, in earlier work as a benchmark to demonstrate that our nested embedding approach based on accurate coupled cluster data reproduces this reference. For the ruthenium drug, we here find a coupled cluster result that deviates by about 6–8 kJ/mol from the DFT and MM results, respectively. At the same time, the NEVPT2 result agrees with the coupled cluster data (–10.8 kJ/mol for the unrestricted coupled cluster and –11.3 kJ/mol for NEVPT2), indicating that both electron correlation methods achieve the same high accuracy for this system. Hence, with our best first-principles approaches, we obtain a free energy of binding for the ruthenium drug of about –11 kJ/mol, which we may consider as a prediction to be challenged by experimental work in the future.

The results on MCL-1/19G also suggest that the error inherent to any embedding approach can be kept small enough such that our strategy allows us to take into account the impact of high-accuracy quantum chemical calculations in a biologically relevant case. However, this strategy also has intrinsic limitations due to the curse of dimensionality. First, any wave function-based traditional method lacks guarantees on the energies, which may therefore compromise the accuracy of the final free energy. Second, it will be severely limited in its applicable size. As we will emphasize, quantum computation has the potential to overcome both problems and thus lead to biological quantum advantage. We will provide estimates of the size, quality, and speed of future quantum computers to realize this potential.

3. THE POTENTIAL OF QUANTUM COMPUTING

In the future, the energy computation of the quantum cores may be carried out with quantum algorithms on quantum computers. Indeed, one of the main motivations of the FreeQuantum pipeline is to anticipate the potential of quantum computation for electronic structure computations and their application to biological problems. In this section, we discuss the options and prospects and give concrete resource estimates for the ruthenium drug–protein complex.

The number of high-accuracy energy data points required depends on uncertainty quantification during machine learning. For the chosen complex, this is found to be around 4000. The required energy accuracy is chemical accuracy (i.e., about 1 kJ/mol), but generally depends on the region where the sampling happens and is again determined by the uncertainty quantification of the ML potential.

The algorithmic framework for computing the ground state energy of the quantum core Hamiltonian is quantum phase estimation (QPE). While there exist variational alternatives for ground state energy computation such as the variational quantum eigensolver (VQE), for generating training data the accuracy guarantees of QPE are valuable, and VQE suffers from fundamental computational limitations for larger systems.⁵² QPE requires the preparation of a guiding state $|\psi\rangle$ on the quantum computer with significant ground state overlap and simulation of the Hamiltonian evolution on the quantum computer. For Hamiltonians in biochemistry, finding guiding states with high ground state overlap is possible in many cases,^{31,53,54} and we also confirm this for the described choice of quantum cores for our ruthenium system. This is shown in Figure 4. We find that the Hartree–Fock state (which is trivial to prepare on a quantum computer) provides sufficient overlap even on fairly large active spaces. Higher overlap is obtained by considering guiding states consisting of a linear combination of a small number of Slater determinants. Low bond-dimension matrix product states (MPS), which can easily be prepared,^{53,54} achieve very high overlap.

The main computational cost is derived from the Hamiltonian simulation, which depends on the required precision and the number of orbitals. We consider two regimes: the first is a regime in which the number of qubits is limited (note that $2N$ qubits are required to represent a state in an active space of N spatial orbitals), and the second is a regime where a significant number of qubits can be used as ancilla qubits. In the first regime, we consider single-ancilla QPE in conjunction with partially randomized product formula (Trotter) methods²⁹ and qDRIFT.^{55,29} If there is a high-overlap guiding state, this can be used to reduce the number of gates per circuit (at the expense of running more circuits), allowing for some degree of parallelization and milder noise requirements, which may be valuable for early fault-tolerant hardware. In this example, such states are indeed available, making this a relevant trade-off (but we note that ideally the classical computational cost of finding these states should also be taken into account). In Figure 4, we show the maximal number of gates per circuit, as well as the total cost, needed for the ruthenium system with varying active space sizes. These estimates require knowledge of the Trotter error and are informed by a calculation of the Trotter error on this system carried out on HPC and GPU clusters with exact state vector simulations of up to 36 qubits.³⁰ In the second regime, when a significant amount can be used as ancilla qubits, qubitization methods outperform product formula-based methods when total runtime is considered; resource estimates for qubitization are also shown in Figure 4.

The runtime resulting from the gate counts reported in Figure 4 depends strongly on the hardware, as well as the fault-tolerant architecture. In order to get a rough idea of the concrete quantum resources required, let us see what is required for a single energy computation to be performed within 20 min (which would correspond to two months for 4000 data points of training data of the pipeline). We first consider the case of an active space of 30 spatial orbitals, which has previously been identified as a point where comparison to traditional methods becomes relevant.³⁴ In this case, the Trotter-based algorithm requires at least 60 qubits to represent the state and gate errors below 10^{-7} , as shown in Figure 4. A computation time of 20 min per energy can be reached if the average gate time is below 1×10^{-7} s. Outperforming DMRG

will require larger active spaces; here, we consider 60 spatial orbitals. While product formulas have benefits in terms of the number of logical qubits and the possibility of breaking up into smaller circuits, at this system size the total runtime becomes the dominating factor, and qubitization-based methods are more relevant. With around 1000 logical qubits, gate errors below 10^{-10} , and an average gate time of 10^{-7} s, the qubitization resource estimates reach the 20 min target. Further optimization of qubitization-based methods is likely to reduce this cost by one to two orders of magnitude.^{56,57} In order to fully take into account dynamical correlation, we could either increase the number of orbitals further, thereby avoiding NEVPT2 altogether, or use quantum algorithms to compute such perturbative corrections.⁵⁸

To put these numbers in perspective at the hardware level, current parameters for superconducting qubits are on the order of 5×10^{-8} s gate time and 5×10^{-3} error.⁵⁹ In ion-based platforms, two-qubit gate times of around 10^{-4} s at gate errors 3×10^{-4} have been reported.⁶⁰ Using fault-tolerant techniques to achieve the required accuracy will lead to an increased average time per logical gate, showing that high gate speed, parallelization, and low time-overhead fault-tolerant operation will be crucial for sufficiently fast operation.

While quantum computers based on superconducting qubits and ions are currently the most mature platforms, other architectures exist and are under constant and rapid development. In particular, quantum computers based on neutral atoms in optical tweezers have been shown to scale well to many qubits and offer flexible connectivity.⁶¹ In this context, it is important to note that our pipeline is agnostic to the exact type of quantum computer employed, making it a future-proof and versatile software solution.

Parallel computation of the energies is possible in our current implemented pipeline if several quantum computers are available, allowing us to reduce total wall-clock runtime. We note that, to compare favorably to experimental efforts, it would be required to execute the entire pipeline within a time frame of 24 h, which could reasonably be achieved by parallelization. If active learning is used in the ML2 training, parallelization of the energy data may be limited. Additionally, as detailed above, for each energy and sufficiently large ground state overlap, one can parallelize the estimate (especially when using randomized product formulas).

4. CONCLUSIONS

An accurate, quantum-mechanically informed assessment of the interactions between molecules in molecular recognition processes can deliver a sound basis for (i) a detailed physics-grounded understanding, (ii) an accurate estimate of the binding free energy, and (iii) a flexibility that guarantees the transferability of a methodology to arbitrary molecule classes. However, a complete computational pipeline capable of routine application has been missing. In this work, we have addressed the feasibility challenges in a software pipeline by combination of various innovations: (i) a two-fold embedding strategy allows us to carve out quantum cores of variable size from a large quantum region so that highly accurate quantum calculations can be carried out for them in order to obtain high-accuracy data for those regions. (ii) These quantum calculations can be performed with traditional methodologies (e.g., based on coupled cluster or multiconfigurational approaches) or with future fault-tolerant quantum computers. (iii) To be able to deal with scarce high-accuracy data, we

exploit an efficient transfer machine learning approach that threads the quantum data into the first level of embedding (that is, into the quantum-classical hybrid model). (iv) An element-agnostic machine learning strategy allows us to sample the free energy of binding at all levels (quantum-classical and quantum-in-quantum embedded in classical) of embedding. (v) For the innermost quantum-in-quantum embedding, we investigated two different embedding strategies to highlight the modularity and flexibility of the overall approach. The first embedding scheme is based on Huzinaga-type projection embedding, and the other one on bootstrap embedding. Hence, a high degree of versatility is guaranteed for the connection of structures from classical configuration space sampling to accurate quantum calculations for the quantum cores of individual configurations.

The entire pipeline enables us to tailor the quantum core size (measured in terms of atoms and/or orbitals) to the capabilities of the available accurate quantum engines. Moreover, our pipeline exhibits a high degree of autonomy, and it is highly modular (i.e., different modules for specified tasks can easily be exchanged) so that future improvements of the technology can be easily incorporated. When the required quantum computational resources as outlined become available, they can be directly applied within the FreeQuantum pipeline. We have done concrete resource estimates for the ruthenium drug–protein complex, which is located at the top left of the biomolecular simulation quadrangle. We chose this complex because standard force fields struggle with open-shell spin doublets, but it was still meaningful to demonstrate the pipeline with traditional HPC calculations. It would have been impossible to attempt a case in the top right of the quadrangle. It should be noted, however, that the resource estimates for quantum phase estimation are dominated by the Hamiltonian simulation, which is determined by the size of the problem and not the ground state correlation structure. Although finding and preparing good guiding states will become more challenging (as they involve traditional quantum chemical methods), we expect our resource estimates for the ruthenium drug–protein complex to be on a similar order of magnitude for cases in the top right of the quadrangle.

From previous work on quantum algorithms for electronic structure computations in realistic cases, one might have concluded that applying quantum computers to free energy computations is unrealistic for the foreseeable future due to the extensive configurational sampling. By reducing the size of the quantum region through a multilayer quantum embedding strategy paired with efficient use of quantum data through machine learning, we change this perspective. We conclude that there is a realistic prospect of obtaining useful results not only in strongly correlated regimes but also where the correlations are not so strong, the vast majority of biological cases. Furthermore, although our work dealt with molecular recognition, the methodology is general and can be used for any atomistic simulation problem in biochemistry, replacing less accurate classical force fields with the ML1 and ML2 potentials in our hybrid model approach.

Our main contribution in this work is therefore to provide a realistic and concrete way to use quantum computers for free energy computations and thereby open a research direction leading to future quantum advantage in biology. When this will exactly happen, in terms of more precise estimates of free energies, will depend on the break-even point with high-accuracy traditional electronic structure approaches in the case

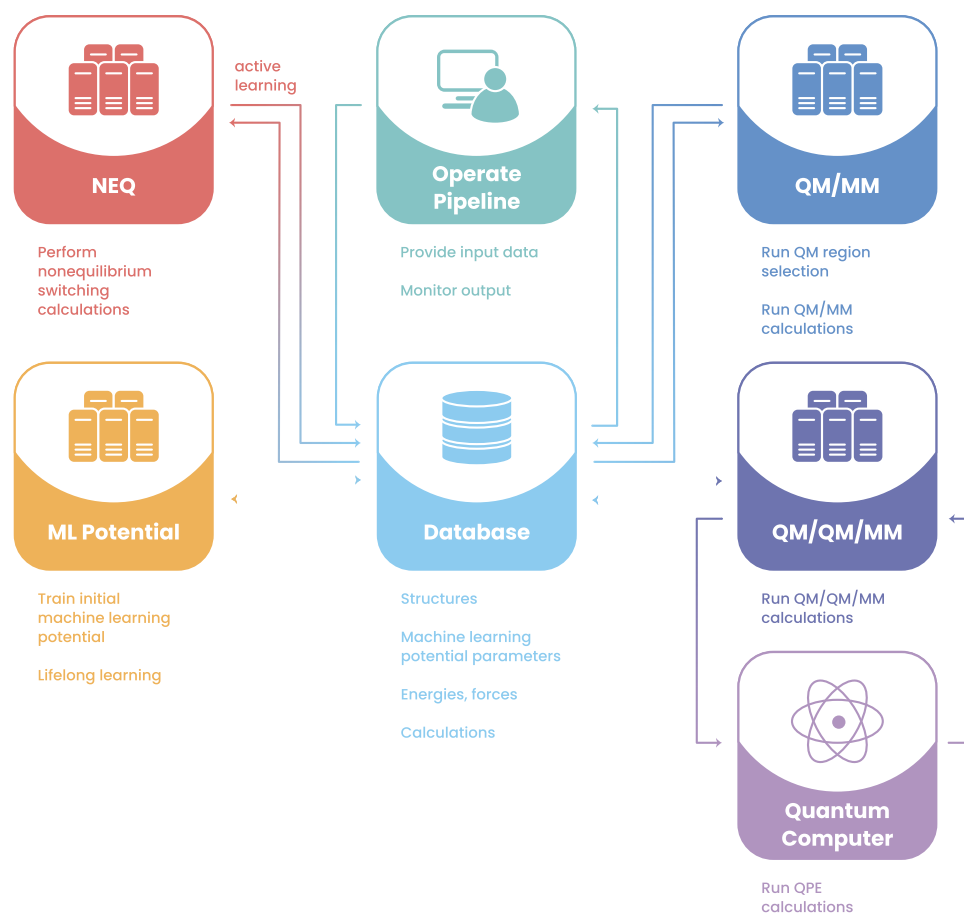


Figure 5. Data flow in the FreeQuantum pipeline. The database (depicted in the center) facilitates the exchange of data between the individual modules of the pipeline, which are shown on the left- and right-hand sides. The QM/MM and QM/QM/MM modules enable the calculation of accurate reference data. The ML potential module creates an ML potential from these reference data, which is then used by the NEQ module to calculate the binding free energy. The pipeline is controlled by a human operator (top center).

of strongly correlated systems or other wave function-based methods (e.g., coupled cluster) in the case of weaker correlations. What is clear, however, is that quantum computing has the potential to change the foundation of free energy calculations by delivering high-quality useful training data within the pipeline. FreeQuantum is open source and free of charge, welcoming both development contributions and deployment for the benefit of biochemistry and the pharmaceutical sciences.

APPENDIX

In the following, we provide further information on the workflow of the FreeQuantum computational pipeline, computational aspects thereof, as well as the quantum computational resource estimates.

Workflow of FreeQuantum

A schematic representation of the data flow within the FreeQuantum pipeline is shown in Figure 5. All steps requiring significant human time are fully automated (with the exception of the initial structure preparation step). However, the individual steps can also be triggered manually via simple Python scripts. Such manual operation can allow close monitoring of the individual steps.

All relevant data are stored in a central MongoDB database.⁶² The software modules which collectively form the pipeline communicate through this database. This

architecture allows for geographically distributed deployment of the individual pipeline modules: literally, any module can run on any computer in the world, as long as it is able to communicate with the database.

After the first, purely MM-based FEP calculation, snapshots (typically on the order of 2000) from the resulting trajectory are extracted via one of the pipeline operation scripts and stored in the database. Optional calculations to automatically determine the QM region based on a few (typically about 100) of these snapshots can be set up if no intuitive QM region selection is available. These calculations are then also stored in the database (i.e., all necessary input to carry out these calculations is stored). Then, these calculations are carried out by the QM/MM calculations module. Once the QM region has been determined, QM/MM calculations are set up for all snapshots stored in the database and subsequently carried out by the QM/MM calculations module. The resulting energies and forces are stored in the database.

These energies and forces are then read from the database by the ML potential module and used as input data to train an ML potential (ML1). The resulting parameters are written back to the database so that an NEQ calculation to refine the potential energy surface with the ML force field can be carried out (this is done by the NEQ calculations module). During the molecular dynamics simulation necessary to carry out the NEQ calculation, structures might be encountered for which the ML force field indicates large uncertainty. If this should be the case,

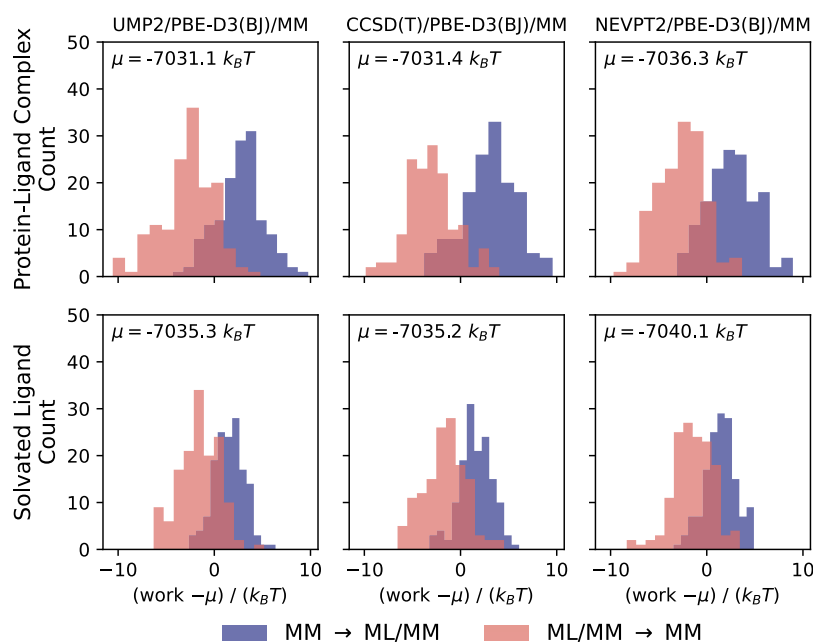


Figure 6. Work distribution for the first (out of six) NEQ simulations for the GRP78-NKP1339 protein–ligand complex, the solvated NKP1339 ligand, and each Huzinaga embedding variant (UMP2/PBE-D3/MM, UCCSD(T)/PBE-D3/MM, and NEVPT2/PBE-D3/MM). The distributions were shifted by their mean μ for clarity.

these structures are reported so that additional QM/MM calculations can be carried out to provide the necessary training data to improve the ML1 potential.

The QM/QM/MM module works along the same concept as the QM/MM module. It retrieves calculations, which are to be carried out, from the database and writes back the resulting energies and, in principle, forces (which will be used to train the ML2 potential). We note that here, for the training of ML2, the presented calculations did not use forces, but could easily be adapted to do so, if the corresponding high-accuracy quantum data become available.

Moreover, we note that the QM/QM/MM module can also interface to a quantum computer engine in order to carry out QPE calculations on the quantum cores. This is because, naturally, a quantum computer cannot directly communicate with the database, but instead has to rely on a classical computer for this. Instead of creating a separate communications module, we opted to integrate this within the QM/QM/MM module. Besides simplifying the overall architecture, this design has the advantage that the input needed to carry out a QPE calculation does not have to be stored in the database. It is worth pointing out that a second-quantized Hamiltonian needs all corresponding one- and two-electron integrals to be fully defined; these integrals can quickly demand a lot of storage space.

Due to the concept of communicating exclusively through a central database, the only requirement that a module must fulfill to be compatible with our pipeline is that it must be able to communicate with the database. We provide a wrapper⁶³ to simplify this communication. Therefore, one has a lot of freedom when building a module, enabling one to construct it exactly according to the specific needs of this module. For example, the QM/MM and QM/QM/MM modules need a large software stack to carry out their calculations, which can significantly complicate deployment. Therefore, we realized these modules as Apptainer (formerly Singularity) images, which are straightforward to share among research groups and

to launch on high-performance computing environments. The actual workflow to carry out QM/MM and QM/QM/MM calculations is implemented in SCINE Puffin,⁶⁴ which is a framework to carry out complex computational tasks, as recently described elsewhere.⁶⁵ The complete source code of the FreeQuantum pipeline is available free of charge and open source on GitHub.⁹⁴

Computational Aspects of FreeQuantum

The model construction for the Ru-based host–guest complex alongside initial molecular mechanics and quantum-classical hybrid data and the foundational machine learning approach has been reported in ref 20. Ref 21 described the principles of our transfer learning approach for the Huzinaga-type quantum-in-quantum embedding, albeit for a different protein–guest complex that is not as challenging as the Ru-based one considered in this work.

The molecular structures of the protein–ligand complex of the Ru-based anticancer compound NKP1339^{47,66} bonded to the heat shock protein GRP78⁶⁷ (PDB entry 3LDO) were taken from the MM FEP end-point trajectories and the active learning structures in ref 20. For each structure, we calculated the electronic energy with Huzinaga-type embedding based on the ansatz described in ref 21. Huzinaga-type embedding⁶⁸ is a variant of projection-based embedding⁶⁹ that allows the combination of nearly arbitrary electronic structure methods by partitioning the occupied orbital space. Huzinaga-type embedding is formally exact within the framework of Kohn–Sham DFT, making the approach robust even if the system partitioning cuts covalent bonds.

In our Huzinaga-type multilevel approach, we combined unrestricted second-order Møller–Plesset perturbation theory (UMP2), the strongly contracted NEVPT2,^{70,71} or unrestricted CCSD(T) (UCCSD(T)) with Kohn–Sham DFT using Perdew–Burke–Ernzerhof’s exchange–correlation functional (PBE)⁷² with Grimme’s D3 dispersion correction⁷³ and Becke–Johnson damping⁷⁴ using the def2-SVP basis set⁷⁵ and

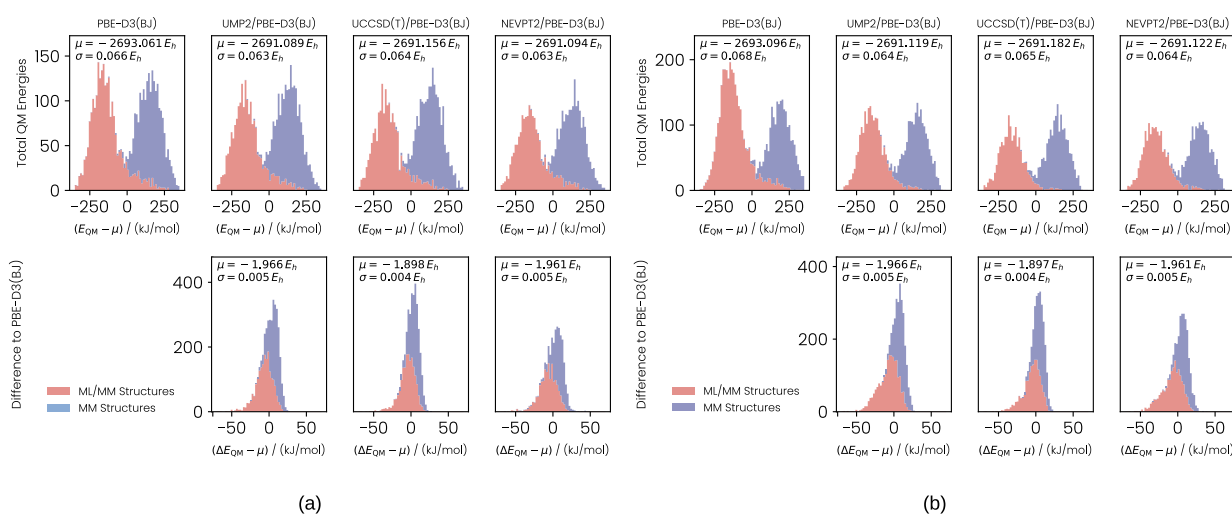


Figure 7. Energy distributions obtained for the Huzinaga-based QM/QM/MM embedding: (a) GRP78/NKP1339 protein–ligand complex and (b) solvated NKP1339 ligand. The distributions were shifted by their mean μ , and the standard deviation σ is provided for clarity. The distributions for MM and ML/MM structures are stacked, i.e., nonoverlapping.

the corresponding effective core potential for Ru.⁷⁶ The MM environment was parameterized using the Amber ff99sb*-ILDN force field.^{77,78}

The MP2, NEVPT2, and unrestricted CCSD(T) calculations were performed with the program PySCF,⁷⁹ taking the valence orbital coefficients and one-particle Hamiltonian from Serenity^{80,81} as input. The actual embedding was performed with Serenity. In addition to the orbital partitioning from ref 21, we relaxed the orbitals in the quantum core in an embedded self-consistent field calculation. Relaxing and canonicalizing the orbitals was required by the UMP2 and NEVPT2 implementations in PySCF.

NEVPT2 corrections were built on top of the DMRG configuration interaction (DMRG-CI) calculations, which were performed with the QCMAQUIS program. The active space was selected based on the single-orbital entropies using autoCAS,^{82,83} and a bond dimension of 250 was used for the underlying DMRG-CI^{84,45} calculation. In order to obtain a consistent active space among different sampled geometries, relaxed orbitals were localized using the intrinsic bond orbital (IBO) scheme⁸⁵ within PySCF. An active space consisting of 5 ruthenium *d*-orbitals was found to be appropriate for this system. The full quantum core contained 43/42 occupied and 310/311 virtual alpha/beta orbitals. The DMRG-CI calculation on the selected active space was performed without bond dimension truncation, and the optimized wavefunction was used to generate the *n*-RDMs ($n = 1-4$) needed for the NEVPT2 correction. Inactive orbitals were canonicalized prior to the NEVPT2 calculation using the generalized Fock operator constructed with the DMRG-CI 1-RDMs.

After transfer learning of the QM/QM/MM (UMP2/PBE-D3/MM, NEVPT2/PBE-D3/MM, or UCCSD(T)/PBE-D3/MM) energies, six NEQ switching simulations were performed for the solvated NKP1339 ligand and the solvated GRP78-NKP1339 protein-ligand complex, using the protocol described in ref 21. These NEQ simulations provided a distribution of 36 $\Delta G_{\text{binding}}^{\text{ML/MM}}$ values. The final value reported in the main article was calculated as their average. The work distributions for the first NEQ simulations of the solvated ligand and protein–ligand complex, based on the three Huzinaga embedding variants UMP2/PBE-D3/MM,

UCCSD(T)/PBE-D3/MM, and NEVPT2/PBE-D3/MM, are shown in Figure 6. The distributions are narrow and show significant overlap between switches from the MM to the ML/MM potential energy surface and back. There is no qualitative difference between the distributions obtained with UMP2/PBE-D3/MM, UCCSD(T)/PBE-D3/MM, and NEVPT2/PBE-D3/MM. The narrow distributions and significant overlap show that the relatively short switching times of 10 ps are sufficient to arrive at converged free energy estimates.

The QM energy distributions for all three Huzinaga-based embedding approaches (UMP2/PBE-D3/MM, UCCSD(T)/PBE-D3/MM, and NEVPT2/PBE-D3/MM) and the ML1 energy model (PBE-D3/MM) are shown in Figure 7a and 7b for the GRP78/NKP1339 protein–ligand complex and the NKP1339 solvated ligand, respectively. Furthermore, the differences between QM/QM/MM and PBE-D3/MM energies are shown. This difference must be effectively learned by the transfer learning approach, elevating the ML1 to the ML2 machine learning potential. The distributions of QM/MM and QM/QM/MM energies are bimodal because structures from the initial MM force field (which are included in this set) have too short Cl–H distances, leading to the right peak (indigo), whereas the DFT computed quantum energies lead to the left peak (red), as discussed in ref 20. The difference between QM/QM/MM and QM/MM energies is unimodal and narrow because QM/MM energies already capture a significant portion of the energy.

For the alternative bootstrap embedding-based inner quantum-in-quantum layer, QM/QM/MM calculations were performed using Q_uEmb.⁸⁶ The methodology has been described in detail in previous works.^{87–90} Bootstrap embedding utilizes the mean-field-level description of the full system to generate a set of small-fragment Hamiltonians that are subsequently evaluated using high-level wavefunction methods. The method forms an ensemble of small fragments, based on atomic connectivity, that overlap and span the entire chemical system. Bootstrap embedding uniquely treats the entire quantum region on an equal footing and converges to the full system result with increasing fragment size. This systematic improvability is especially important for near-term quantum devices, since bootstrap embedding provides a fine

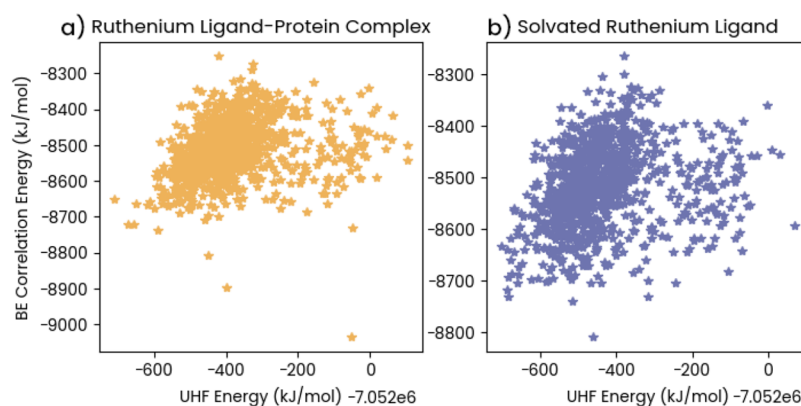


Figure 8. We show the correlation between the bootstrap embedding correlation energies and the UHF energies for each ruthenium ligand–protein complex (a) and solvated ruthenium ligand structure (b). We show that the correlation energies cannot be trivially predicted from the UHF results for either system.

controllable knob between the necessary number of qubits and target accuracy. As such, it can accommodate the resource limitations of near-term devices and also provides a clear pathway for more accurate results with hardware improvements. The method has been validated on a wide variety of chemical systems in minimal and dense basis sets, demonstrating systematic convergence to full-system high-level method energies at reduced cost.^{89,41,91} Recent extensions to perform bootstrap embedding within the QM/MM framework allow the QM/QM/MM-type embedding calculations necessary for large systems in biochemical simulations.⁴¹

QuEmb interfaces with PySCF⁷⁹ to perform the necessary quantum chemistry routines. For bootstrap embedding calculations, mean-field references were obtained with unrestricted Hartree–Fock (UHF) calculations using the def2-SVP basis set⁷⁵ and the def2-SVP effective core potential for ruthenium.⁷⁶ Löwdin localization was used to construct the atom-centered set of orbitals. As introduced in previous work, we generate the entangled interacting bath for each fragment in both the α and β densities, using unrestricted CCSD to solve each fragment.⁹² We omitted chemical potential or density matching, which may lead to further improvements in this work. Likewise, we used BE(1) fragmentation, which assigns each nonhydrogen atom and its attached hydrogen to separate fragments.

Raw data generated for this work will be made available on ERDA.⁹³ The source code for the FreeQuantum pipeline is available open source and free of charge on GitHub.⁹⁴

Quantum Computing Resource Estimates

In the following, we give details of how the quantum computing resource estimates shown in Figure 4 have been obtained. They are essentially based on five different calculations. First, the active spaces of the quantum core located on the ruthenium drug molecule bound to the target protein need to be set up. This defines the Hamiltonian terms, which enter the resource estimates for quantum phase estimation. Here, we use three methods (qDRIFT, partially randomized, and qubitization). Finally, guiding state overlaps need to be computed.

Active spaces were selected from the set of valence orbitals on the basis of the largest single-orbital entropies obtained using autoCAS.^{82,83} Single-orbital entropies were estimated from a low bond dimension ($\chi = 250$) DMRG calculation using QCMAquis.⁹⁵ For each active space, orbitals were

ordered using the Fiedler vector of the mutual information,⁹⁶ and a final DMRG calculation was performed to obtain a proxy of the ground state. A bond dimension $\chi = 1024$ was found to be sufficient in all cases. Subsequently, we performed a symmetry shift for particle number symmetry, as described in ref 29. This lowers the sum of weights $\lambda = \sum_{i=1}^L |h_i|$ for the active space Hamiltonians in Pauli representation, $H = \sum_{i=1}^L h_i P_i$ (i.e., after the fermion-to-qubit mapping).

The resource estimation for qDRIFT and the partially randomized method essentially follows the cost analysis done in ref 29. Hamming weight phasing is used for both methods, which for the randomized method requires an estimate of the expected length of commuting terms when sampling. We used an estimated length of about six for all active spaces. Together with the values of ϵ , λ , and ξ , this is sufficient to arrive at resource estimates for qDRIFT. The partially randomized method is based on the second-order Trotter formula and a representation of the Hamiltonian as a weighted sum of Pauli strings. The partitioning of the Hamiltonian terms into deterministic and randomized ones was chosen such that the total cost (in terms of Toffolis) is minimized. Additionally, the cost of the partially randomized method depends on the Trotter error, $\epsilon_{\text{Trotter}} = C\delta^2$, where δ is the Trotter step size. Computing C exactly is as hard as computing the ground state energy, so for larger system sizes, an approximation scheme is needed. In ref 30, exact Trotter errors were computed for the same active spaces of the ruthenium ligand for up to 16 spatial orbitals. We performed a power-law fit of the Trotter error estimates against the λ value of the active space Hamiltonians, arriving at $C^{1/2} \approx 1.08 \times 10^{-4} \lambda^{1.25}$. Since λ is easy to compute even for the largest active spaces studied here, this fit allowed us to estimate C even for larger systems. The resource estimates for the qubitization approach refer to the ‘double-factorized’ variant³³ and were computed using the OPENFERMION python package.⁹⁷

The overlaps were computed by first converging DMRG calculations to have a meaningful proxy for the ground state. This required going up to a bond dimension of 1024. MPS thus obtained was used to construct the guiding states. Lower bond dimension guiding states were obtained by truncating the bond dimension using singular value decomposition of the density matrix, which maximizes the overlap of the truncated and initial MPS.⁹⁸ Sum-of-Slater guiding states were obtained by sampling the Slater determinants corresponding to the largest coefficients using the sampling-reconstruction of the




Figure 9. Enlarged version of Figure 3a to highlight structural details of the protein Ru-drug complex.

complete active space (SR-CAS).⁹⁹ Guiding states were normalized in all cases.

Enlarged Figure of the Protein–Drug Structure

For the sake of better readability, we provide an enlarged figure of the structure of the chaperone BiP with the ruthenium transition-metal complex NKP1339 in Figure 9.

AUTHOR INFORMATION

Corresponding Authors

Aram W. Harrow – Center for Theoretical Physics – a Lineweber Institute, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, United States; Email: aram@mit.edu

Anders Krogh – Department of Computer Science, Quantum for Life Centre, University of Copenhagen, Copenhagen DK-2100, Denmark; Department of Public Health, University of Copenhagen, DK-2200 Copenhagen, Denmark; Email: akrogh@di.ku.dk

Troy Van Voorhis – Department of Chemistry, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, United States; orcid.org/0000-0001-7111-0176; Email: tvvan@mit.edu

Kresten Lindorff-Larsen – Department of Biology, Linderstrøm-Lang Centre for Protein Science, University of Copenhagen, Copenhagen DK-2200, Denmark; orcid.org/0000-0002-4750-6039; Email: lindorff@bio.ku.dk

Gemma Solomon – Department of Chemistry and Nano-Science Center, University of Copenhagen, Copenhagen DK-2100, Denmark; NNF Quantum Computing Programme, Niels Bohr Institute, University of Copenhagen, Copenhagen DK-2100, Denmark; orcid.org/0000-0002-2018-1529; Email: solomon@chem.ku.dk

Markus Reiher – Department of Chemistry and Applied Biosciences, ETH Zurich, Zurich 8093, Switzerland; orcid.org/0000-0002-9508-1565; Email: mreiher@ethz.ch

Matthias Christandl – Department of Mathematical Sciences, Quantum for Life Centre, University of Copenhagen, Copenhagen DK-2100, Denmark; Email: christandl@math.ku.dk

Authors

Jakob Günther – Department of Mathematical Sciences, Quantum for Life Centre, University of Copenhagen, Copenhagen DK-2100, Denmark

Thomas Weymuth – Department of Chemistry and Applied Biosciences, ETH Zurich, Zurich 8093, Switzerland; orcid.org/0000-0001-7102-7022

Moritz Bensberg – Department of Chemistry and Applied Biosciences, ETH Zurich, Zurich 8093, Switzerland

Freek Witteveen – Department of Mathematical Sciences, Quantum for Life Centre, University of Copenhagen, Copenhagen DK-2100, Denmark; QuSoft, Centrum voor Wiskunde en Informatica, Amsterdam 1098 XG, The Netherlands

Matthew S. Teynor – Department of Chemistry and Nano-Science Center, University of Copenhagen, Copenhagen DK-2100, Denmark; NNF Quantum Computing Programme, Niels Bohr Institute, University of Copenhagen, Copenhagen DK-2100, Denmark; orcid.org/0000-0002-6981-4809

F. Emil Thomasen – Department of Biology, Linderstrøm-Lang Centre for Protein Science, University of Copenhagen, Copenhagen DK-2200, Denmark; orcid.org/0000-0002-2096-4873

Valentina Sora – Department of Computer Science, Quantum for Life Centre, University of Copenhagen, Copenhagen DK-2100, Denmark; Department of Public Health, University of Copenhagen, DK-2200 Copenhagen, Denmark

William Bro-Jørgensen – Department of Chemistry and Nano-Science Center, University of Copenhagen, Copenhagen DK-2100, Denmark; orcid.org/0000-0001-8171-6374

Raphael T. Husistein – Department of Chemistry and Applied Biosciences, ETH Zurich, Zurich 8093, Switzerland; orcid.org/0009-0005-2696-4574

Mihael Erakovic – Department of Chemistry and Applied Biosciences, ETH Zurich, Zurich 8093, Switzerland

Marek Miller – Department of Mathematical Sciences, Quantum for Life Centre, University of Copenhagen, Copenhagen DK-2100, Denmark

Leah Weisburn – Department of Chemistry, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, United States; orcid.org/0009-0001-4026-4946

Minsik Cho – Department of Chemistry, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, United States; orcid.org/0000-0002-9307-8549

Marco Eckhoff – Department of Chemistry and Applied Biosciences, ETH Zurich, Zurich 8093, Switzerland; orcid.org/0000-0002-5581-789X

Complete contact information is available at: <https://pubs.acs.org/10.1021/acs.jctc.5c02088>

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

We acknowledge funding from the Research Project “Molecular Recognition from Quantum Computing”. Work on “Molecular Recognition from Quantum Computing” was supported by Wellcome Leap as part of the Quantum for Bio (Q4Bio) Program. Furthermore, M.C., J.G., T.W., F.W., M.B., M.Ec., R.T.H., M.M., M.R., V.S., M.Er., and A.H. acknowledge support from the Novo Nordisk Foundation (Grant No. NNF20OC0059939 ‘Quantum for Life’). M.M. acknowledges support from the novoSTAR Programme by Novo Nordisk A/S. G.C.S. and W.B.J. acknowledge support from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement no. 865870). G.C.S. and M.S.T. acknowledge support from the Novo Nordisk Foundation, grant number NNF22SA0081175, NNF Quantum Computing Programme and NNF20OC0060019, SolidQ. A.K. acknowledges support from the Novo Nordisk Foundation NNF20OC0062606 and NNF20OC0063268. T.W., R.T.H., and M.R. acknowledge support from the NCCR Catalysis (grant number 180544), a National Centre of Competence in Research funded by the Swiss National Science Foundation. M.R. also acknowledges financial support from the Swiss National Science Foundation through Grant No. 200021_219616. M.Ec. acknowledges an ETH Zurich Postdoctoral Fellowship. The work is cofunded by the European Union (ERC, DynaPLIX, SyG-2022 101071843, to K.L.-L.). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

ABBREVIATIONS

CAS: complete active space
 CCSD: coupled cluster including singly and doubly excited determinants
 CI: configuration interaction
 DFT: density functional theory
 DMRG: density matrix renormalization group

FCI: full configuration interaction
 FEP: free energy perturbation
 ML: machine learning
 ML1: machine learning-based force field to describe (outer) QM region
 ML2: machine learning-based force field to describe quantum cores
 MM: molecular mechanics
 MPS: matrix product state
 NEQ: nonequilibrium switching
 QM: quantum mechanical
 QM/MM: embedding of a quantum mechanically treated part into a (larger) part treated with MM methods
 QM/QM/MM: nested embedding of a quantum mechanically treated part (called quantum cores) into another, also quantum mechanically treated part, which in turn is embedded into a part treated with MM methods
 QPE: quantum phase estimation
 VQE: variational quantum eigensolver

REFERENCES

- (1) Wang, L.; et al. Accurate and reliable prediction of relative ligand binding potency in prospective drug discovery by way of a modern free-energy calculation protocol and force field. *J. Am. Chem. Soc.* **2015**, *137*, 2695–2703.
- (2) Mobley, D. L.; Gilson, M. K. Predicting Binding Free Energies: Frontiers and Benchmarks. *Annu. Rev. Biophys.* **2017**, *46*, 531–558.
- (3) Abel, R.; Wang, L.; Harder, E. D.; Berne, B.; Friesner, R. A. Advancing Drug Discovery through Enhanced Free Energy Calculations. *Acc. Chem. Res.* **2017**, *50*, 1625–1632.
- (4) Courmia, Z.; Chipot, C.; Roux, B.; York, D. M.; Sherman, W. *Free Energy Methods in Drug Discovery: Current State and Future Directions*; ACS Publications, 2021; pp 1–38.
- (5) Wade, A. D.; Bhati, A. P.; Wan, S.; Coveney, P. V. Alchemical Free Energy Estimators and Molecular Dynamics Engines: Accuracy, Precision, and Reproducibility. *J. Chem. Theory Comput.* **2022**, *18*, 3972–3987.
- (6) Lelièvre, T.; Rousset, M.; Stoltz, G. *Free Energy Computations: A Mathematical Perspective*; Imperial College Press: London, UK, 2010.
- (7) York, D. M. Modern Alchemical Free Energy Methods for Drug Discovery Explained. *ACS Phys. Chem. Au* **2023**, *3*, 478–491.
- (8) Sabanés Zariquiey, F.; Farr, S. E.; Doerr, S.; De Fabritiis, G. QuantumBind-RBFE: Accurate Relative Binding Free Energy Calculations Using Neural Network Potentials. *J. Chem. Inf. Model.* **2025**, *65*, 4081–4089.
- (9) Crha, R.; Poliak, P.; Gillhofer, M.; Oostenbrink, C. Alchemical Free-Energy Calculations at Quantum-Chemical Precision. *J. Phys. Chem. Lett.* **2025**, *16*, 863–869.
- (10) Semelak, J. A.; Ignacio, P.; Huddleston, K. K.; Olmos, J.; Grassano, J. S.; Clemente, C.; Drusin, S. I.; Marti, M.; Gonzalez Lebrero, M. C.; Roitberg, A. E.; Estrin, D. A. Advancing Multiscale Molecular Modeling with Machine Learning-Derived Electrostatics. *J. Chem. Theory Comput.* **2025**, *21*, 5194–5207.
- (11) Wang, X.; Wu, X.; Brooks, B. R.; Wang, J. Accurate Free Energy Calculation via Multiscale Simulations Driven by Hybrid Machine Learning and Molecular Mechanics Potentials. *J. Chem. Theory Comput.* **2025**, *21*, 6979–6987.
- (12) Wang, Y.; et al. On the design space between molecular mechanics and machine learning force fields. *Appl. Phys. Rev.* **2025**, *12*, No. 021304.
- (13) Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys.* **1982**, *21*, 467–488.
- (14) Lloyd, S. Universal Quantum Simulators. *Science* **1996**, *273*, 1073–1078.
- (15) Baiardi, A.; Christandl, M.; Reiher, M. Quantum Computing for Molecular Biology. *ChemBioChem* **2023**, *24*, No. e202300120.

- (16) Reiher, M.; Wiebe, N.; Svore, K. M.; Wecker, D.; Troyer, M. Elucidating reaction mechanisms on quantum computers. *Proc. Natl. Acad. Sci. U. S. A.* **2017**, *114*, 7555–7560.
- (17) Finkelmann, A. R.; Stiebritz, M. T.; Reiher, M. Inaccessibility of the μ -hydride species in [FeFe] hydrogenases. *Chem. Sci.* **2014**, *5*, 215–221.
- (18) Ghosh, S.; Das, C. K.; Uddin, S.; Stripp, S. T.; Engelbrecht, V.; Winkler, M.; Leimkühler, S.; Brocks, C.; Duan, J.; Schäfer, L. V.; Happe, T. Protein dynamics affect O₂-stability of Group B [FeFe]-hydrogenase from *Thermosidiminibacter oceani*. *J. Am. Chem. Soc.* **2025**, *147*, 15170–15180.
- (19) Stiebritz, M. T.; Reiher, M. Hydrogenases and oxygen. *Chem. Sci.* **2012**, *3*, 1739–1751.
- (20) Bensberg, M.; Eckhoff, M.; Thomasen, F. E.; Bro-Jørgensen, W.; Teynor, M. S.; Sora, V.; Weymuth, T.; Husstein, R. T.; Knudsen, F. E.; Krogh, A.; Lindorff-Larsen, K.; Reiher, M.; Solomon, G. C. Machine Learning Enhanced Calculation of Quantum-Classical Binding Free Energies. *J. Chem. Theory Comput.* **2025**, *21*, 8182–8198.
- (21) Bensberg, M.; Eckhoff, M.; Husstein, R. T.; Teynor, M. S.; Sora, V.; Bro-Jørgensen, W.; Thomasen, F. E.; Krogh, A.; Lindorff-Larsen, K.; Solomon, G. C.; Weymuth, T.; Reiher, M. Hierarchical quantum embedding by machine learning for large molecular assemblies. *J. Chem. Theory Comput.* **2025**, *21*, 7662–7674.
- (22) Bösel, L.; Thürlmann, M.; Riniker, S. Machine Learning in QM/MM Molecular Dynamics Simulations of Condensed-Phase Systems. *J. Chem. Theory Comput.* **2021**, *17*, 2641–2658.
- (23) Pan, X.; Yang, J.; Van, R.; Epifanovsky, E.; Ho, J.; Huang, J.; Pu, J.; Mei, Y.; Nam, K.; Shao, Y. Machine-Learning-Assisted Free Energy Simulation of Solution-Phase and Enzyme Reactions. *J. Chem. Theory Comput.* **2021**, *17*, 5745–5758.
- (24) Gastegger, M.; Schütt, K. T.; Müller, K.-R. Machine learning of solvent effects on molecular spectra and reactions. *Chem. Sci.* **2021**, *12*, 11473–11483.
- (25) Hudson, P. S.; Woodcock, H. L.; Boresch, S. Use of Nonequilibrium Work Methods to Compute Free Energy Differences Between Molecular Mechanical and Quantum Mechanical Representations of Molecular Systems. *J. Phys. Chem. Lett.* **2015**, *6*, 4850–4856.
- (26) Rufa, D. A.; Bruce Macdonald, H. E.; Fass, J.; Wieder, M.; Grinaway, P. B.; Roitberg, A. E.; Isayev, O.; Chodera, J. D. Towards chemical accuracy for alchemical free energy calculations with hybrid physics-based machine learning/molecular mechanics potentials. **2020**, *bioRxiv*.
- (27) Tkaczyk, S.; Karwounopoulos, J.; Schöller, A.; Woodcock, H. L.; Langer, T.; Boresch, S.; Wieder, M. Reweighting from Molecular Mechanics Force Fields to the ANI-2x Neural Network Potential. *J. Chem. Theory Comput.* **2024**, *20*, 2719–2728.
- (28) Giese, T. J.; Zeng, J.; Lerew, L.; McCarthy, E.; Tao, Y.; Ekesan, S.; York, D. M. Software Infrastructure for Next-Generation QM/MM- Δ MPLP Force Fields. *J. Phys. Chem. B* **2024**, *128*, 6257–6271.
- (29) Günther, J.; Witteveen, F.; Schmidhuber, A.; Miller, M.; Christandl, M.; Harrow, A. Phase estimation with partially randomized time evolution. **2025**, *arXiv:2503.05647*.
- (30) Miller, M.; Günther, J.; Witteveen, F.; Teynor, M. S.; Erakovic, M.; Reiher, M.; Solomon, G. C.; Christandl, M. phase2: Full-State Vector Simulation of Quantum Time Evolution at Scale. **2025**, *arXiv:2504.17881v1*.
- (31) Erakovic, M.; Witteveen, F.; Harley, D.; Günther, J.; Bensberg, M.; Meitei, O. R.; Cho, M.; Van Voorhis, T.; Reiher, M.; Christandl, M. High Ground State Overlap via Quantum Embedding Methods. *PRX Life* **2025**, *3*, No. 013003.
- (32) Low, G. H.; Chuang, I. L. Hamiltonian Simulation by Qubitization. *Quantum* **2019**, *3*, 163.
- (33) von Burg, V.; Low, G. H.; Häner, T.; Steiger, D. S.; Reiher, M.; Roetteler, M.; Troyer, M. Quantum computing enhanced computational catalysis. *Phys. Rev. Res.* **2021**, *3*, No. 033055.
- (34) Goings, J. J.; White, A.; Lee, J.; Tautermann, C. S.; Degroote, M.; Gidney, C.; Shiozaki, T.; Babbush, R.; Rubin, N. C. Reliably assessing the electronic structure of cytochrome P450 on today's classical computers and tomorrow's quantum computers. *Proc. Natl. Acad. Sci. U. S. A.* **2022**, *119*, No. e2203533119.
- (35) Zwanzig, R. W. High-Temperature Equation of State by a Perturbation Method. I. *Nonpolar Gases*. *J. Chem. Phys.* **1954**, *22*, 1420–1426.
- (36) Mey, A. S.; Allen, B. K.; Bruce Macdonald, H. E.; Chodera, J. D.; Hahn, D. F.; Kuhn, M.; Michel, J.; Mobley, D. L.; Naden, L. N.; Prasad, S.; Rizzi, A.; Scheen, J.; Shirts, M. R.; Tresadern, G.; Xu, H. Best Practices for Alchemical Free Energy Calculations [Article v1.0]. *Living J. Comp. Mol. Sci.* **2020**, *2*, 18378.
- (37) Song, L. F.; Merz, K. M. J. Evolution of Alchemical Free Energy Methods in Drug Discovery. *J. Chem. Inf. Model.* **2020**, *60*, 5308–5318.
- (38) Shirts, M. R.; Chodera, J. D. Statistically optimal analysis of samples from multiple equilibrium states. *J. Chem. Phys.* **2008**, *129*, 124105.
- (39) Smith, J. S.; Nebgen, B. T.; Zubatyuk, R.; Lubbers, N.; Devereux, C.; Barros, K.; Tretiak, S.; Isayev, O.; Roitberg, A. E. Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning. *Nat. Commun.* **2019**, *10*, 2903.
- (40) Brunken, C.; Reiher, M. Automated construction of quantum-classical hybrid models. *J. Chem. Theory Comput.* **2021**, *17*, 3797–3813.
- (41) Weisburn, L. P.; Cho, M.; Bensberg, M.; Meitei, O. R.; Reiher, M.; Van Voorhis, T. Multiscale Embedding for Quantum Computing. *J. Chem. Theory Comput.* **2025**, *21*, 4591–4603.
- (42) Helgaker, T.; Jørgensen, P.; Olsen, J. *Molecular electronic-structure theory*; John Wiley & Sons, 2013.
- (43) Stein, C. J.; Reiher, M. Automated selection of active orbital spaces. *J. Chem. Theory Comput.* **2016**, *12*, 1760–1771.
- (44) Wardzala, J.; Hennefarth, M.; Agarawal, V.; Jangid, B.; Seal, A.; Hermes, M.; King, D.; Gagliardi, L. The Modern Era of Multi-reference Methods for Chemistry and Materials Science. **2025**, *ChemRxiv*.
- (45) Baiardi, A.; Reiher, M. The density matrix renormalization group in chemistry and molecular physics: Recent developments and new challenges. *J. Chem. Phys.* **2020**, *152*, No. 040903.
- (46) Liu, Y.; Meitei, O. R.; Chin, Z. E.; Dutt, A.; Tao, M.; Chuang, I. L.; Van Voorhis, T. Bootstrap Embedding on a Quantum Computer. *J. Chem. Theory Comput.* **2023**, *19*, 2230–2247.
- (47) Trondl, R.; Heffeter, P.; Kowol, C. R.; Jakupec, M. A.; Berger, W.; Keppler, B. K. NKP-1339, the first ruthenium-based anticancer drug on the edge to clinical application. *Chem. Sci.* **2014**, *5*, 2925–2932.
- (48) Flocke, L. S.; Trondl, R.; Jakupec, M. A.; Keppler, B. K. Molecular mode of action of NKP-1339 – a clinically investigated ruthenium-based drug – involves ER- and ROS-related effects in colon carcinoma cell lines. *Invest. New Drugs* **2016**, *34*, 261–268.
- (49) Bakewell, S. J.; Rangel, D. F.; Ha, D. P.; Sethuraman, J.; Crouse, R.; Hadley, E.; Costich, T. L.; Zhou, X.; Nichols, P.; Lee, A. S. Suppression of stress induction of the 78-kilodalton glucose regulated protein (GRP78) in cancer by IT-139, an anti-tumor ruthenium small molecule inhibitor. *Oncotarget* **2018**, *9*, 29698.
- (50) Bertolotti, A.; Zhang, Y.; Hendershot, L. M.; Harding, H. P.; Ron, D. Dynamic interaction of BiP and ER stress transducers in the unfolded-protein response. *Nat. Cell Biol.* **2000**, *2*, 326–332.
- (51) Friberg, A.; Vigil, D.; Zhao, B.; Daniels, R. N.; Burke, J. P.; Garcia-Barrantes, P. M.; Camper, D.; Chauder, B. A.; Lee, T.; Olejniczak, E. T.; Fesik, S. W. Discovery of Potent Myeloid Cell Leukemia 1 (Mcl-1) Inhibitors Using Fragment-Based Methods and Structure-Based Design. *J. Med. Chem.* **2013**, *56*, 15–30.
- (52) Larocca, M.; Thanasilp, S.; Wang, S.; Sharma, K.; Biamonte, J.; Coles, P. J.; Cincio, L.; McClean, J. R.; Holmes, Z.; Cerezo, M. Barren plateaus in variational quantum computing. *Nat. Rev. Phys.* **2025**, *7*, 174–189.
- (53) Fomichev, S.; Hejazi, K.; Zini, M. S.; Kiser, M.; Morales, J. F.; Casares, P. A. M.; Delgado, A.; Huh, J.; Voigt, A.-C.; Mueller, J. E.;

Arazola, J. M. Initial state preparation for quantum chemistry on quantum computers. *PRX Quantum* **2024**, *5*, No. 040339.

(54) Berry, D. W.; Tong, Y.; Khattar, T.; White, A.; Kim, T. I.; Low, G. H.; Boixo, S.; Ding, Z.; Lin, L.; Lee, S.; Chan, G. K.-L.; Babbush, R.; Rubin, N. C. Rapid Initial-State Preparation for the Quantum Simulation of Strongly Correlated Molecules. *PRX Quantum* **2025**, *6*, No. 020327.

(55) Campbell, E. Random Compiler for Fast Hamiltonian Simulation. *Phys. Rev. Lett.* **2019**, *123*, No. 070503.

(56) Caesura, A.; Cortes, C. L.; Poll, W.; Sim, S.; Steudtner, M.; Anselmetti, G.-L. R.; Degroote, M.; Moll, N.; Santagati, R.; Streif, M.; Tautermann, C. S. Faster quantum chemistry simulations on a quantum computer with improved tensor factorization and active volume compilation. *PRX Quantum* **2025**, *6*, No. 030337.

(57) Low, G. H.; King, R.; Berry, D. W.; Han, Q.; DePrince, A. E., III; White, A.; Babbush, R.; Somma, R. D.; Rubin, N. C. Fast quantum simulation of electronic structure by spectrum amplification. *Phys. Rev. X* **2025**, *15*, No. 041016.

(58) Günther, J.; Baiardi, A.; Reiher, M.; Christandl, M. More quantum chemistry with fewer qubits. *Phys. Rev. Res.* **2024**, *6*, No. 043021.

(59) Kjaergaard, M.; Schwartz, M. E.; Braumüller, J.; Krantz, P.; Wang, J. I.-J.; Gustavsson, S.; Oliver, W. D. Superconducting Qubits: Current State of Play. *Annu. Rev. Condens. Matter Phys.* **2020**, *11*, 369–395.

(60) Löschnauer, C. M.; Toba, J. M.; Hughes, A. C.; King, S. A.; Weber, M. A.; Srinivas, R.; Matt, R.; Nourshargh, R.; Allcock, D. T. C.; Ballance, C. J.; Matthiesen, C.; Malinowski, M.; Harty, T. P. Scalable, high-fidelity all-electronic control of trapped-ion qubits. *PRX Quantum* **2025**, *6*, No. 040313.

(61) Bluvstein, D.; et al. Logical quantum processor based on reconfigurable atom arrays. *Nature* **2024**, *626*, 58–65.

(62) MongoDB Inc. *MongoDB*. 2024. <https://www.mongodb.com>.

(63) Bensberg, M.; Grimmel, S. A.; Sobez, J.-G.; Steiner, M.; Türtscher, P. L.; Unsleber, J. P.; Reiher, M. *SCINE Database: Release 1.3.0*; Zenodo, 2023.

(64) Bensberg, M.; Brunken, C.; Csizi, K.-S.; Grimmel, S.; Gugler, S.; Sobez, J.-G.; Steiner, M.; Türtscher, P. L.; Unsleber, J. P.; Weymuth, T.; Reiher, M. *SCINE Puffin: Release 1.3.0*; Zenodo, 2023.

(65) Weymuth, T.; Unsleber, J. P.; Türtscher, P. L.; Steiner, M.; Sobez, J.-G.; Müller, C. H.; Mörchen, M.; Klasovita, V.; Grimmel, S. A.; Eckhoff, M.; Csizi, K.-S.; Borgia, F.; Bensberg, M.; Reiher, M. *SCINE — Software for Chemical Interaction Networks*. *J. Chem. Phys.* **2024**, *160*, 222501.

(66) Peti, W.; Pieper, T.; Sommer, M.; Keppler, B. K.; Giester, G. Synthesis of Tumor-Inhibiting Complex Salts Containing the Aniontrans-Tetrachlorobis(indazole)ruthenate(III) and Crystal Structure of the Tetraphenylphosphonium Salt. *Eur. J. Inorg. Chem.* **1999**, *1999*, 1551–1555.

(67) Macias, A. T.; et al. Adenosine-Derived Inhibitors of 78 kDa Glucose Regulated Protein (Grp78) ATPase: Insights into Isoform Selectivity. *J. Med. Chem.* **2011**, *54*, 4034–4041.

(68) Hégyely, B.; Nagy, P. R.; Ferenczy, G. G.; Kállay, M. Exact density functional and wave function embedding schemes based on orbital localization. *J. Chem. Phys.* **2016**, *145*, No. 064107.

(69) Manby, F. R.; Stella, M.; Goodpaster, J. D.; Miller, T. F. A Simple, Exact Density-Functional-Theory Embedding Scheme. *J. Chem. Theory Comput.* **2012**, *8*, 2564–2568.

(70) Angeli, C.; Cimiraglia, R.; Evangelisti, S.; Leininger, T.; Malrieu, J.-P. Introduction of n-electron valence states for multi-reference perturbation theory. *J. Chem. Phys.* **2001**, *114*, 10252–10264.

(71) Angeli, C.; Cimiraglia, R.; Malrieu, J.-P. n-electron valence state perturbation theory: A spinless formulation and an efficient implementation of the strongly contracted and of the partially contracted variants. *J. Chem. Phys.* **2002**, *117*, 9138–9153.

(72) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Phys. Rev. Lett.* **1996**, *77*, 3865.

(73) Grimme, S.; Antony, J.; Ehrlich, S.; Krieg, H. A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu. *J. Chem. Phys.* **2010**, *132*, 154104.

(74) Grimme, S.; Ehrlich, S.; Goerigk, L. Effect of the damping function in dispersion corrected density functional theory. *J. Comput. Chem.* **2011**, *32*, 1456–1465.

(75) Weigend, F.; Ahlrichs, R. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy. *Phys. Chem. Chem. Phys.* **2005**, *7*, 3297–3305.

(76) Andrae, D.; Häußermann, U.; Dolg, M.; Stoll, H.; Preuß, H. Energy-adjusted ab initio pseudopotentials for the second and third row transition elements. *Theor. Chim. Acta* **1990**, *77*, 123–141.

(77) Best, R. B.; Hummer, G. Optimized Molecular Dynamics Force Fields Applied to the Helix-Coil Transition of Polypeptides. *J. Phys. Chem. B* **2009**, *113*, 9004–9015.

(78) Lindorff-Larsen, K.; Piana, S.; Palmo, K.; Maragakis, P.; Klepeis, J. L.; Dror, R. O.; Shaw, D. E. Improved side-chain torsion potentials for the Amber ff99SB protein force field. *Proteins: Struct., Funct., Bioinf.* **2010**, *78*, 1950–1958.

(79) Sun, Q.; Berkelbach, T. C.; Blunt, N. S.; Booth, G. H.; Guo, S.; Li, Z.; Liu, J.; McClain, J. D.; Sayfutyarova, E. R.; Sharma, S.; Wouters, S.; Chan, G. K.-L. PySCF: the Python-based simulations of chemistry framework. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2018**, *8*, No. e1340.

(80) Unsleber, J. P.; Dresselhaus, T.; Klahr, K.; Schnieders, D.; Böckers, M.; Barton, D.; Neugebauer, J. Serenity: A Subsystem Quantum Chemistry Program. *J. Comput. Chem.* **2018**, *39*, 788–798.

(81) Niemeyer, N.; Eschenbach, P.; Bensberg, M.; Tölle, J.; Hellmann, L.; Lampe, L.; Massolle, A.; Rikus, A.; Schnieders, D.; Unsleber, J. P.; Neugebauer, J. The subsystem quantum chemistry program Serenity. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2022**, *13*, No. e1647.

(82) Stein, C. J.; Reiher, M. autoCAS: A Program for Fully Automated Multiconfigurational Calculations. *J. Comput. Chem.* **2019**, *40*, 2216–2226.

(83) Bensberg, M.; Mörchen, M.; Stein, C. J.; Unsleber, J. P.; Weymuth, T.; Reiher, M. *qscine/autocas: Release 2.3.0*; Zenodo, 2024.

(84) White, S. R. Density Matrix Formulation for Quantum Renormalization Groups. *Phys. Rev. Lett.* **1992**, *69*, 2863–2866.

(85) Knizia, G. Intrinsic Atomic Orbitals: An Unbiased Bridge between Quantum Theory and Chemical Concepts. *J. Chem. Theory Comput.* **2013**, *9*, 4834–4843.

(86) Cho, M.; Meitei, O.; Weisburn, L. P.; Weser, O.; Weatherly, S.; Tran, H.; Ye, H.-Z.; Alexiu, A.; Hanscam, B.; Van Voorhis, T. QuEmb: A Toolbox for Bootstrap Embedding Calculations of Molecular and Periodic Systems. *J. Phys. Chem. A* **2025**, *129*, 6538–6551.

(87) Ye, H.-Z.; Van Voorhis, T. Atom-Based Bootstrap Embedding For Molecules. *J. Phys. Chem. Lett.* **2019**, *10*, 6368–6374.

(88) Ye, H.-Z.; Ricke, N. D.; Tran, H. K.; Van Voorhis, T. Bootstrap Embedding for Molecules. *J. Chem. Theory Comput.* **2019**, *15*, 4497–4506.

(89) Ye, H.-Z.; Tran, H. K.; Van Voorhis, T. Bootstrap Embedding For Large Molecular Systems. *J. Chem. Theory Comput.* **2020**, *16*, 5035–5046.

(90) Welborn, M.; Tsuchimochi, T.; Voorhis, T. V. Bootstrap embedding: An internally consistent fragment-based method. *J. Chem. Phys.* **2016**, *145*, No. 074102.

(91) Tran, H. K.; Weisburn, L. P.; Cho, M.; Weatherly, S.; Ye, H.-Z.; Van Voorhis, T. Bootstrap Embedding for Molecules in Extended Basis Sets. *J. Chem. Theory Comput.* **2024**, *20*, 10912–10921.

(92) Tran, H. K.; Ye, H.-Z.; Voorhis, T. V. Bootstrap embedding with an unrestricted mean-field bath. *J. Chem. Phys.* **2020**, *153*, 214101.

(93) Günther, J. et al. Data Archive for the Journal Article “How to use quantum computers for biomolecular free energies”. **2026**, <https://sid.erda.dk/sharelink/eUNBmbdB9w>.

(94) *The FreeQuantum Pipeline*. 2026; <https://github.com/Quantum-for-Life/pipeline>.

(95) Keller, S.; Dolfi, M.; Troyer, M.; Reiher, M. An efficient matrix product operator representation of the quantum chemical Hamiltonian. *J. Chem. Phys.* **2015**, *143*, 244118.

(96) Barcza, G.; Legeza, Ö.; Marti, K. H.; Reiher, M. Quantum-information analysis of electronic states of different molecular structures. *Phys. Rev. A* **2011**, *83*, No. 012508.

(97) McClean, J. R.; et al. OpenFermion: the electronic structure package for quantum computers. *Quantum Sci. Technol.* **2020**, *5*, No. 034014.

(98) Chan, G. K.-L.; Head-Gordon, M. Highly correlated calculations with a polynomial cost algorithm: A study of the density matrix renormalization group. *J. Chem. Phys.* **2002**, *116*, 4462–4476.

(99) Boguslawski, K.; Marti, K. H.; Reiher, M. Construction of CASCI-type wave functions for very large active spaces. *J. Chem. Phys.* **2011**, *134*, 224101.



CAS BIOFINDER DISCOVERY PLATFORM™

**PRECISION DATA
FOR FASTER
DRUG
DISCOVERY**

CAS BioFinder helps you identify
targets, biomarkers, and pathways

Unlock insights

CAS
A division of the
American Chemical Society