Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

AFTER THE QUANTUM ADVANTAGE CLAIM: WHICH DRUG-DISCOVERY PROBLEMS ARE ACTUALLY READY FOR QUANTUM COMPUTATION?

Lucas Meyer[1]*, Anna Schmid[2], Stefan Braun[1], Laura Keller[3]

1. ***Department of Quantum Computation Readiness in Drug Discovery, Faculty of Pharmacy, ETH Zurich, Zurich, Switzerland.***
2. ***Department of Problem-Specific Quantum Applicability, Faculty of Pharmaceutical Sciences, University of Bern, Bern, Switzerland.***
3. ***Department of Beyond-Quantum-Advantage Problem Selection, Faculty of Pharmacy, EPFL Lausanne, Lausanne, Switzerland.***

ARTICLE INFO

Received:
23 April 2026
Received in revised form:
18 August 2026
Accepted:
21 August 2026
Available online:
28 August 2026

Keywords: Quantum computing, Drug discovery, Quantum chemistry, Quantum machine learning, Molecular simulation, Pharmaceutical benchmarking

ABSTRACT

Quantum computation is increasingly presented as a potential response to difficult molecular-simulation, optimization, and machine-learning problems in drug discovery. However, the existence of computationally demanding pharmaceutical tasks does not establish that quantum methods can solve them more accurately, efficiently, or usefully than contemporary classical approaches. This horizon-scanning review evaluates which problem classes currently justify bounded quantum research programs and which remain dependent on unresolved hardware, algorithmic, benchmarking, data, or translational assumptions. The review applies an evidence-gated assessment that separates technical execution, benchmark performance, plausible computational utility, experimental confirmation, pharmaceutical decision impact, and routine readiness. The synthesis distinguishes near-term variational and hybrid approaches from fault-tolerant molecular simulation, quantum-assisted classical workflows, quantum machine learning, and optimization-oriented proposals. Current evidence supports the feasibility of selected small-system calculations, restricted molecular-property models, and hybrid workflow components, but it does not support routine pharmaceutical readiness for any broad drug-discovery problem class. The most defensible long-horizon opportunities concern chemically consequential electronic-structure and dynamical problems in which strong correlation limits classical approximations. Nearer-term value is more plausibly sought through tightly scoped hybrid experiments, application-structured benchmarking, and prospective evaluations that isolate the contribution of the quantum component. The article proposes a quantum readiness horizon map and a problem-level assessment framework rather than a universal technology maturity ranking. These conceptual contributions are intended to organize evidence and research priorities; they are not validated prediction tools. Readiness judgments remain conditional on continuing improvements in classical computation, hardware architecture, error control, resource estimation, data suitability, experimental validation, and integration into pharmaceutical decision workflows.

To Cite This Article: Meyer L, Schmid A, Braun S, Keller L. After the Quantum Advantage Claim: Which Drug-Discovery Problems Are Actually Ready for Quantum Computation? Pharmacophore. 2026;17(4):157-68. https://doi.org/10.51847/ytVqQIwlCE

Introduction

Quantum computation has entered pharmaceutical discourse through claims that molecular electronic structure, conformational exploration, chemical optimization, and data-intensive prediction may contain computational structures that are difficult for conventional methods. Yet quantum chemistry is a problem-dependent field in which algorithmic promise, computational assumptions, and pharmaceutical resource demands must be evaluated together [1-3]. A formal reduction in asymptotic complexity, an executable circuit, or a chemically recognizable demonstration is not equivalent to an end-to-end improvement in a discovery workflow. State preparation, basis and active-space selection, circuit construction, measurement, error control, classical preprocessing, and output interpretation may each determine whether a nominally favorable algorithm remains scientifically useful.

Corresponding Author: Lucas Meyer; Department of Quantum Computation Readiness in Drug Discovery, Faculty of Pharmacy, ETH Zurich, Zurich, Switzerland. E-mail: lucas.meyer@ethz.ch

Drug discovery also cannot be treated as a single computational problem. It contains heterogeneous simulation, learning, search, and optimization tasks whose readiness must be assessed separately [4]. Estimating the electronic energy of an isolated molecular structure, modeling a reaction mechanism, predicting an absorption or toxicity label, generating candidate compounds, optimizing a synthetic route, and supporting a medicinal-chemistry decision require different inputs, outputs, accuracy thresholds, validation strategies, and turnaround times. A quantum approach that is theoretically compatible with one task may be irrelevant to another. Readiness must therefore be assigned to a defined problem under a defined hardware regime and comparator, not to "quantum drug discovery" as an undifferentiated field.

This distinction is particularly important for quantum machine learning. Pharmaceutical quantum-machine-learning studies increasingly demonstrate that quantum kernels, parameterized circuits, or hybrid neural architectures can be applied to selected molecular datasets. Nevertheless, the available evidence remains dominated by small-scale, simulated, or hybrid demonstrations rather than routine workflow utility [5]. Reported predictive performance may depend on descriptor engineering, dimensionality reduction, restricted data access, favorable partitioning, simulator conditions, or weak classical comparators. Even a reproducible performance difference would not by itself establish mechanism, calibrated uncertainty, prospective generalization, experimental success, or value to a pharmaceutical decision.

The purpose of this horizon-scanning review is therefore not to decide whether quantum computing is generally promising, but to identify which drug-discovery problems are sufficiently well specified, evidence supported, and constraint aware to justify particular readiness claims. The central contribution is an evaluative synthesis that separates current evidence, emerging signals, plausible medium-term opportunities, and speculative long-term possibilities. It assesses each opportunity through problem definition, algorithm–hardware compatibility, best-classical comparison, resource scaling, validation setting, pharmaceutical relevance, and the possibility of disconfirmation. The resulting framework is proposed as an organizing interpretation rather than an empirically validated maturity model.

Horizon-Scanning Scope, Sources, and Assessment Framework

A horizon scan is designed to detect and assess emerging developments before their consequences are fully established. For quantum drug discovery, this requires a transparent definition of information sources, signal-capture procedures, prioritization rules, assessment horizons, and updating logic [6]. The present review therefore draws from complementary evidence classes: quantum-computing and quantum-chemistry reviews; algorithmic and resource-estimation studies; hardware demonstrations; pharmaceutical quantum-machine-learning investigations; hybrid computational workflows; chemically or experimentally motivated case studies; and methodological literature on benchmarking and practical advantage. Sources were considered relevant only when they could support a bounded claim about capability, limitation, validation, readiness, or research design.

Signal selection was based on more than technical novelty. Emerging-technology scanning should distinguish novelty, scientific plausibility, expected impact, implementation timing, dependency, and uncertainty rather than ranking signals by visibility alone [7]. In this review, a signal was treated as stronger when it addressed a pharmaceutical problem that is consequential, difficult for strong classical methods, representable on the proposed quantum architecture, and assessable through a realistic validation pathway. Signals were weakened when they depended on unspecified data-loading assumptions, unreported resource costs, unusually small or favorable datasets, incomparable baselines, idealized noise models, or the unsupported assumption that circuit execution would translate directly into pharmaceutical value.

Eligibility and synthesis decisions were organized to remain traceable even though the review is horizon oriented rather than meta-analytic. Transparent review reporting requires explicit eligibility criteria, screening decisions, source flow, and synthesis procedures [8]. Evidence was therefore coded by application domain, computational paradigm, hardware regime, comparator quality, validation setting, maturity, translational limitation, and risk of overstatement. "Not clearly reported" was preferred where a source did not provide enough information to classify a feature. Conflicting findings were interpreted in relation to differences in molecular representation, circuit design, noise assumptions, classical preprocessing, data access, benchmark construction, and the definition of advantage.

The proposed assessment framework separates six questions. First, is the pharmaceutical problem scientifically meaningful and sufficiently defined? Second, is a quantum method theoretically compatible with the structure of that problem? Third, has the method been executed or simulated under conditions that reveal rather than conceal its resource requirements? Fourth, has it been compared with an appropriate best-available classical approach? Fifth, has its output been validated prospectively, experimentally, or within a consequential workflow? Sixth, does the evidence support technical feasibility, benchmark competitiveness, plausible utility, translational value, or routine readiness? **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop horizon-scanning scope, sources, and assessment framework within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for Horizon-scanning scope, sources, and assessment framework

Evidence domain	Eligible evidence or method	Reported capability or claim	Strength of support	Reproducibility or bias concern	Unresolved gap
Horizon-scanning methodology	Systematic or scoping reviews of horizon-scanning	Emerging signals can be identified, prioritized,	Strong for methodological transparency;	Methodological heterogeneity; differences in	No universally validated horizon-scanning protocol

	methods; transparent review-reporting guidance	classified, and periodically reassessed through explicit procedures	indirect for quantum pharmaceutical readiness	terminology, prioritization, and stakeholder perspective	for quantum drug discovery
Quantum-computing paradigms	Technical reviews, algorithm studies, hardware experiments, and resource analyses	Quantum algorithms can represent selected simulation, optimization, or learning problems	Strong for theoretical compatibility; variable for practical execution	Idealized assumptions, small systems, architecture dependence, and selective reporting	Demonstration that theoretical compatibility produces end-to-end pharmaceutical value
Molecular and chemical simulation	Electronic-structure, reaction-mechanism, quantum-dynamics, and hybrid simulation studies	Quantum methods may address chemically difficult states, observables, or dynamics	Moderate for long-horizon scientific plausibility; limited for current workflow readiness	Active-space selection, state-preparation cost, measurement burden, and uncertain scaling	Prospective comparison on chemically consequential problems beyond reliable classical reach
Pharmaceutical quantum machine learning	Quantum-kernel, variational-circuit, transfer-learning, and hybrid predictive studies	Selected molecular datasets can be processed using quantum or quantum-inspired model components	Moderate for bounded feasibility; weak for general advantage	Small datasets, simulation dependence, data leakage, descriptor dominance, and baseline sensitivity	Generalization under matched data, compute, uncertainty, and prospective validation conditions
Hardware, error, and resource evidence	Device benchmarks, noise studies, error-mitigation research, logical-resource estimates, and measurement analyses	Hardware and algorithmic constraints can be quantified for specified workloads	Strong for identifying bottlenecks; conditional for forecasting readiness	Results depend on architecture, compilation, error model, parallelism, and runtime assumptions	Application-structured resource audits connected to pharmaceutical turnaround and accuracy requirements
Translational and workflow evidence	Hybrid pipelines, experimentally followed candidates, workflow challenges, and decision-oriented evaluations	Quantum components can be embedded within broader computational or experimental processes	Emerging and case dependent	Difficult attribution of value to the quantum component; limited ablation and independent replication	Evidence that the quantum component changes a prospective scientific or development decision
Advantage and readiness claims	Matched quantum–classical comparisons, scaling tests, prospective evaluations, and practical-utility frameworks	A quantum approach may offer a defined computational, scientific, or operational advantage	Currently limited and claim specific	Moving classical baselines, unequal resources, narrow tasks, and conflation of accuracy with utility	Reproducible advantage under realistic resources, validated outputs, and consequential use conditions

Figure 1 presents the quantum readiness horizon map, showing how the article's key components and boundaries are connected within horizon-scanning scope, sources, and assessment framework.

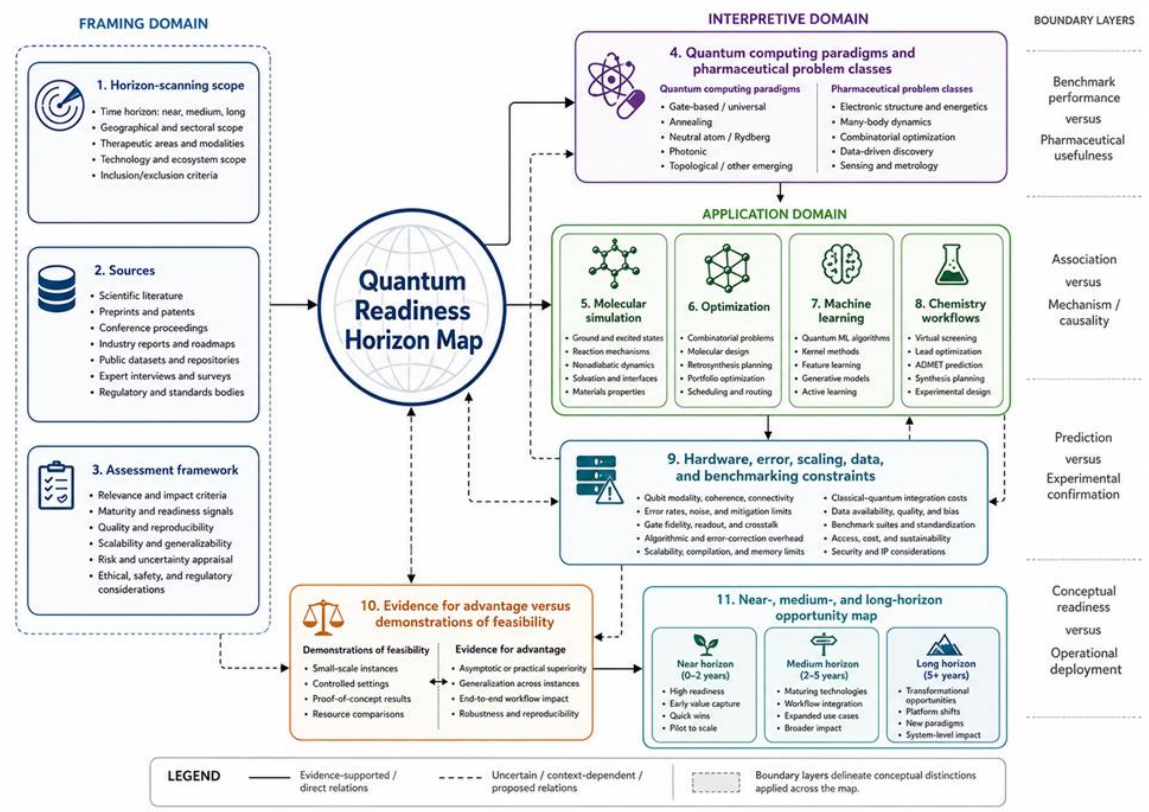

Figure 1. Quantum Readiness Horizon Map

The figure is an original conceptual synthesis that organizes the article's central contribution across horizon-scanning scope, sources, and assessment framework, horizon-scanning scope, assessment framework, quantum computing paradigms and pharmaceutical problem classes, molecular simulation, optimization, machine learning, and chemistry workflows, molecular simulation. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

Quantum Computing Paradigms and Pharmaceutical Problem Classes

Quantum-computing paradigms differ in how they represent a problem, access data, use quantum coherence, interact with classical processors, and tolerate hardware error. Near-term variational approaches reduce circuit-depth demands but inherit coupled optimization, noise, sampling, and measurement constraints [9-11]. In a variational quantum algorithm, a parameterized circuit prepares a state, a quantum processor estimates an objective or observable, and a classical optimizer updates circuit parameters. This hybrid structure enables bounded experiments on noisy devices, but it can relocate rather than remove computational difficulty. Readiness depends on the ansatz, initialization, optimizer, measurement allocation, mitigation method, circuit compilation, hardware connectivity, and scaling behavior as an integrated system.

Fault-tolerant quantum chemistry constitutes a different paradigm. Algorithms based on coherent time evolution, phase estimation, amplitude amplification, or related primitives seek to exploit deep circuits protected by quantum error correction. These approaches may offer more systematic routes to high-accuracy energies, states, or dynamics, but they require hardware capabilities that are qualitatively different from those of present noisy devices. Near-term and fault-tolerant chemistry algorithms consequently answer different scientific questions and should not share a single readiness label [12]. A noisy-device calculation may test circuit construction, embedding, or measurement strategies, whereas a fault-tolerant resource estimate may test whether a future computation is asymptotically and architecturally plausible.

Quantum machine learning forms a third category but contains several distinct access models. Molecular information may be encoded into amplitudes, basis states, rotation angles, kernels, or parameterized circuits, while classical descriptors or learned embeddings frequently perform much of the representational work. Early pharmaceutical studies show that quantum-learning algorithms can be applied to selected drug-discovery datasets, establishing bounded feasibility rather than scalable advantage [13]. Such experiments are scientifically useful when they disclose preprocessing, encoding cost, circuit execution conditions, hyperparameter selection, uncertainty, and strong classical comparisons. They are less informative when the quantum component is evaluated only after aggressive dimensionality reduction or when the classical baseline is not given equivalent tuning and feature access.

The resulting problem–paradigm map is therefore conditional. Fault-tolerant simulation is most plausibly aligned with electronic-structure or dynamical problems in which quantum states, strong correlation, or real-time evolution are central to

the scientific question. Variational algorithms may support near-term studies of restricted active spaces, embeddings, or circuit and measurement design, but their relevance decreases when optimization and sampling costs dominate. Quantum machine learning is most defensible as a hypothesis about a specific representation, kernel, or data-access advantage, not as a substitute for pharmaceutical validation. Quantum optimization proposals require equally explicit mappings between the mathematical objective and the actual constraints of molecular design, synthesis, developability, and experimental testing. Across all paradigms, algorithm availability establishes conceptual compatibility; it does not establish comparative superiority, translational value, or routine readiness.

Molecular Simulation, Optimization, Machine Learning, and Chemistry Workflows

Molecular simulation is the most scientifically defensible domain for long-horizon quantum computation because the underlying quantities are governed by quantum mechanics and because selected chemically consequential systems remain difficult for approximate classical methods. The strongest candidates are not routine small-molecule energy calculations, which are already served by mature classical techniques, but carefully defined problems involving strong electron correlation, transition-metal centers, bond rearrangement, excited states, or competing reaction pathways. Resource analysis of mechanistically important active spaces has shown how fault-tolerant quantum algorithms could, in principle, address electronic structures that challenge conventional approximations [14]. This evidence supports the plausibility of selected reaction-mechanism calculations as long-horizon targets. It does not establish that entire proteins, solvated binding processes, or medicinal-chemistry campaigns can be simulated quantum mechanically, nor does it remove the classical work required to define structures, environments, active spaces, and chemically relevant observables.

Quantum molecular dynamics extends this opportunity from stationary electronic structure to time-dependent phenomena. Nonadiabatic transitions, coupled electronic–nuclear motion, photochemical reactions, charge transfer, and other dynamical processes may contain quantum effects that are expensive to preserve using classical approximations. Quantum algorithms offer conceptually direct representations of real-time evolution, but their pharmaceutical usefulness remains conditional on preparing physically meaningful initial states, propagating them with controlled error, and extracting observables without prohibitive repetition [15]. A technically accurate state evolution would still require connection to experimentally interpretable outputs, environmental effects, conformational ensembles, and competing pathways. Consequently, quantum dynamics is best classified as an exploratory long-horizon opportunity rather than an emerging pharmaceutical service. Its readiness should be tested through bounded molecular problems for which the quantum treatment addresses a known deficiency in classical modeling and produces an observable that can be independently validated.

Nearer-term molecular-simulation research is more likely to use quantum processors as restricted components of larger classical calculations. Subspace expansion, embedding, and related reconstruction strategies seek to obtain more useful information from small quantum registers rather than representing the full molecular problem directly. Virtual quantum subspace expansion has demonstrated that selected effective-basis improvements can be obtained without proportionally increasing the number of physical qubits, although the gain is accompanied by additional measurement and reconstruction requirements [16]. Such methods may be valuable when a compact correlated subsystem can be isolated and when the resulting quantum information improves a complete classical workflow. Their evaluation must therefore include the cost of constructing the subspace, estimating matrix elements, controlling statistical and hardware error, and reintegrating the result into the molecular model. Improvement relative to an uncorrected quantum calculation is not enough; the relevant comparison is against the strongest classical or hybrid method capable of answering the same chemical question.

Machine-learning and optimization workflows occupy a more immediate but less mechanistically secure horizon. Quantum-to-classical transfer learning for toxicity prediction illustrates a bounded hybrid design in which the quantum component contributes representations that are subsequently used by a classical model [17]. Quantum-kernel studies on selected absorption, distribution, metabolism, excretion, and toxicity tasks similarly show that molecular descriptors can be processed within quantum-learning architectures, but simulation-based predictive performance does not demonstrate hardware-level acceleration, prospective generalization, calibrated uncertainty, or improved pharmaceutical decisions [18]. Optimization-oriented applications face an analogous boundary: encoding a docking, generation, or selection problem into a quantum-compatible objective establishes formulation feasibility, not superiority over classical search or medicinal-chemistry judgment. The central near-term question is therefore whether a quantum component contributes reproducible incremental value after matched preprocessing, tuning, compute budgets, and ablation. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop molecular simulation, optimization, machine learning, and chemistry workflows within the article's central argument.

Table 2. Components, relationships, uncertainties, and validation needs within Molecular simulation, optimization, machine learning, and chemistry workflows

Approach or application	Data and task context	Evaluation practice	Evidence maturity	Translational limitation	Review interpretation
Fault-tolerant reaction-mechanism simulation	Strongly correlated active spaces, transition-metal	Resource estimation against chemically meaningful accuracy requirements and	Plausible long-horizon opportunity	Large logical-resource demands, active-space dependence, state	A defensible research target when the unresolved electronic structure

	chemistry, and bond rearrangement	advanced classical methods		preparation, and incomplete environmental representation	is scientifically consequential
Quantum molecular dynamics	Nonadiabatic processes, photochemistry, charge transfer, and coupled electronic–nuclear motion	Validation against model systems, conserved quantities, spectroscopic observables, and reliable classical limits	Exploratory long-horizon evidence	Initial-state preparation, deep time evolution, observable extraction, and integration of environmental effects	Scientifically relevant but not presently ready for pharmaceutical workflow use
Subspace expansion and embedding	Compact correlated regions within larger molecular models	End-to-end comparison including measurement, reconstruction, embedding, and classical postprocessing	Emerging methodological feasibility	Quantum information may be outweighed by measurement and integration costs	Potential niche value must be established against complete classical and hybrid pipelines
Quantum-to-classical transfer learning	Molecular toxicity or property prediction with classical descriptors and limited quantum representations	Matched transfer architectures, ablation, external testing, and uncertainty assessment	Early feasibility	Narrow datasets, descriptor dependence, simulator or device limitations, and unclear quantum attribution	More plausible near term than fully quantum learning, but not evidence of predictive advantage
Quantum-kernel ADME-Tox modeling	Curated molecular-property datasets represented through engineered descriptors	Strong classical kernels, identical data splits, leakage controls, calibration, and prospective validation	Early and simulation-heavy	Small samples, favorable feature spaces, data-loading assumptions, and weak evidence of hardware benefit	Demonstrates model operability on selected tasks rather than pharmaceutical readiness
Quantum optimization and generative chemistry	Docking reductions, molecular selection, constrained generation, and search objectives	Task-complete comparisons including encoding, solution decoding, chemical validity, synthesizability, and search budgets	Exploratory	Mathematical objectives incompletely represent developability, safety, synthesis, and therapeutic value	Formulation readiness is substantially ahead of demonstrated workflow utility
Hybrid chemistry workflows	Classical preprocessing and simulation combined with a restricted quantum subroutine	Component ablation, matched resources, prospective use, and experimental follow-up	Application-relevant but immature	Success is difficult to attribute to the quantum component and may not generalize	The most credible near-term strategy is a bounded hybrid experiment with an explicit stopping rule

Hardware, Error, Scaling, Data, and Benchmarking Constraints

Measurement burden, imperfect error suppression, and noise-driven trainability loss jointly constrain near-term chemistry scaling [19-21]. These limitations interact rather than operate independently. A variational circuit may be shallow enough to execute but still require an impractical number of measurements to estimate its objective. Error mitigation may reduce selected biases but increase sampling demands and depend on assumptions about the device noise. Noise may also flatten an optimization landscape, making additional circuit executions progressively less informative. Consequently, the number of qubits or successful completion of a circuit is an insufficient measure of pharmaceutical readiness. A useful assessment must include circuit depth, connectivity, compilation, state preparation, shot allocation, covariance, optimizer iterations, mitigation overhead, queueing and calibration effects, classical-loop costs, wall-clock time, and the uncertainty of the final chemical quantity.

Trainability is a separate constraint from representational capacity. A highly expressive parameterized circuit may be capable of representing a desired molecular state or decision boundary while remaining practically impossible to optimize. Barren-plateau theory shows that gradients can become exponentially suppressed for broad classes of parameterized circuits, especially under random or insufficiently structured designs [22]. Molecularly informed ansätze, local objectives, careful initialization, layerwise training, and symmetry restrictions may reduce this risk, but they do not create a general guarantee. Pharmaceutical tasks introduce additional difficulties because the objective may be noisy, multiobjective, discontinuous, or only indirectly related to the final experimental endpoint. Claims that a circuit is “expressive enough” should therefore be separated from evidence that it is trainable, reproducible across initializations, robust to hardware variation, and competitive in total optimization effort.

Data and output costs create further boundaries. Quantum machine-learning proposals frequently assume efficient access to quantum states or feature maps, while pharmaceutical data usually arrive as classical molecular graphs, descriptors, images, assay measurements, or text. Encoding these inputs may dominate the proposed speedup, particularly when state preparation must be repeated for every molecule. At the opposite end, a quantum chemistry algorithm must return classical observables, gradients, spectra, ranked candidates, or uncertainty estimates that can support scientific decisions. The required number and precision of outputs may erase an advantage suggested by the internal state evolution. Fault-tolerant resource estimates must therefore report architecture, physical-to-logical overhead, error-correction assumptions, non-Clifford resource supply, parallelism, runtime, and output demands rather than gate counts alone [23]. Resource statements that omit these dependencies should be interpreted as partial algorithmic indicators, not operational forecasts.

Benchmarking must connect hardware capability to application structure without allowing the application label to substitute for validity. Device-level metrics such as fidelity, volume, or circuit throughput are informative only when they predict performance on the circuits required by the target task. Algorithm-level benchmarks must include the full hybrid loop, preprocessing, measurement, mitigation, and postprocessing. Pharmaceutical benchmarks must additionally represent relevant chemical diversity, distribution shift, experimental uncertainty, and decision consequences. The most defensible design is a nested benchmark architecture: device characterization; quantum-subroutine accuracy and cost; matched end-to-end classical comparison; prospective molecular-task validation; experimental confirmation where appropriate; and assessment of whether the output changes a scientific decision. Because hardware and classical algorithms evolve concurrently, benchmark versions, software stacks, calibration conditions, and comparator implementations must remain traceable and periodically updated.

Evidence for Advantage versus Demonstrations of Feasibility

Technical feasibility is the lowest necessary evidence level. Hardware-efficient variational eigensolver experiments established that noisy quantum processors can prepare and optimize bounded molecular Hamiltonians for small systems [24]. This was an important demonstration of operability, but it did not establish superiority over classical quantum chemistry, favorable scaling, or pharmaceutical usefulness. A feasibility claim may therefore state that a mapped problem was executed, that the circuit produced an interpretable estimate, or that a hybrid loop converged under specified conditions. It should not be translated into an advantage claim unless the study also controls for classical simulability, total computational resources, precision, stability, and the strength of alternative classical methods. Small-system execution is an enabling result, not an endpoint.

Benchmark advantage requires a more demanding counterfactual: what would the best feasible classical workflow have achieved on the same scientifically relevant problem under comparable accuracy and resource conditions? Work examining cytochrome P450 electronic structure illustrates why this comparison must include continuing progress in classical multireference chemistry rather than assuming that biochemical complexity automatically creates a quantum opportunity [25]. A molecular system may be biologically important yet remain tractable through improved classical approximations, active-space methods, tensor techniques, or specialized hardware. Conversely, a carefully identified correlated subproblem may remain valuable even when the complete macromolecular environment is treated classically. Credible advantage claims must therefore define the target observable, error tolerance, molecular model, classical comparator, hardware assumptions, and the scientific consequence of the remaining uncertainty.

Workflow feasibility occupies an intermediate position. A hybrid quantum-computing pipeline can demonstrate that quantum-compatible models or optimization components fit within molecular generation, docking, scoring, or screening procedures [26]. This evidence is stronger than an isolated circuit demonstration because it exposes interfaces, preprocessing requirements, data formats, and workflow dependencies. Nevertheless, a drug-related case does not by itself show that the quantum component improved the workflow. The candidate set, descriptors, classical model capacity, search budget, hyperparameter tuning, and post hoc selection may account for the observed result. A workflow claim is credible when the quantum component is clearly defined and ablated, classical counterfactuals receive equivalent resources, unsuccessful cases are disclosed, and performance is assessed on prespecified prospective tasks rather than only on retrospective datasets.

Shared challenges can improve comparability by exposing accuracy, resource, and implementation trade-offs, but a reduced molecular-energy task must not be generalized to the full discovery pipeline [27]. Experimental testing provides a stronger translational signal because it connects computational candidates to physical molecular behavior. The experimentally followed KRAS inhibitor workflow demonstrates that quantum-enhanced computational strategies can participate in a process that produces testable compounds, but experimental activity alone cannot attribute success to the quantum component [28]. Attribution requires matched classical pipelines, component removal studies, independent replication, prospective decision rules, and evidence that the quantum result altered candidate prioritization rather than merely accompanying it. **Table 3** organizes the evidence, constructs, and interpretation boundaries needed to develop evidence for advantage versus demonstrations of feasibility within the article's central argument.

Table 3. Components, relationships, uncertainties, and validation needs within Evidence for advantage versus demonstrations of feasibility

Claim category	Supporting evidence type	Consistency across sources	Quality concern	What can be concluded	What remains uncertain

Technical execution	Hardware or simulator implementation of a defined quantum circuit	Consistent evidence that bounded molecular and learning tasks can be executed	Systems remain small, often classically simulable, and sensitive to noise and sampling	Quantum subroutines are technically operable under specified conditions	Whether execution scales or improves a pharmaceutical workflow
Numerical or benchmark agreement	Agreement with an exact, experimental, or classical reference on a restricted task	Common for selected small problems	Agreement may reflect an easy instance, loose tolerance, or incomplete comparator set	The implementation can reproduce a bounded target quantity	Whether it outperforms strong classical methods in accuracy, cost, or scaling
Quantum computational advantage	Matched quantum–classical comparison with complete resource accounting	Limited and strongly task dependent	Classical baselines may be weak, outdated, or given unequal data and tuning	A defined advantage can be evaluated only relative to an explicit comparator and resource model	Whether any reported difference persists with improved classical methods and larger relevant problems
End-to-end feasibility	Quantum component embedded in a complete chemistry or discovery workflow	Emerging across selected hybrid studies	Classical preprocessing and postprocessing may dominate performance	Integration requirements and interface constraints can be studied empirically	Whether the quantum component contributes net incremental value
Prospective pharmaceutical validation	Prespecified prediction, ranking, or simulation tested on future or held-out chemical problems	Sparse	Dataset shift, assay variability, selection bias, and failed predictions are often underrepresented	Prospective validation can test whether the workflow generalizes beyond retrospective benchmarks	Whether performance changes candidate-selection quality or development efficiency
Experimental confirmation	Synthesized or tested candidates, chemical observables, or mechanistic measurements	Stronger for physical relevance than purely computational evidence	Experimental success may not identify which computational component caused it	Computational outputs can generate experimentally meaningful hypotheses	Whether quantum computation was necessary, superior, or decision changing
Translational or decision utility	Demonstrated influence on a scientific or development decision under realistic constraints	Not established as a general evidence class	Utility depends on turnaround, cost, uncertainty, interpretability, and organizational context	Readiness requires evidence beyond predictive or simulation accuracy	Whether quantum methods improve real pharmaceutical decisions reproducibly
Routine pharmaceutical readiness	Reproducible use across relevant tasks, users, systems, and operating conditions	Not supported by the available evidence	No stable evidence of robustness, scalable advantage, and routine decision impact together	No broad problem class presently warrants a routine-readiness claim	Which narrowly defined tasks may eventually cross all evidence gates

Near-, Medium-, and Long-Horizon Opportunity Map

Near-horizon opportunities should be defined by what can be learned from present or incrementally improved hardware, not by an assumption of imminent pharmaceutical advantage. Suitable activities include application-structured hardware benchmarks, small active-space experiments, circuit and measurement design, error-aware hybrid workflows, and quantum-machine-learning studies with strict ablation and classical controls. Algorithmic work on connectivity-aware electronic-structure simulation shows that circuit structure can be redesigned to reduce depth and better reflect hardware constraints [29]. In the near horizon, the value of such work lies primarily in identifying scaling laws, failure modes, compilation demands, and chemically meaningful test cases. It may support readiness research, but it should not be reported as routine discovery acceleration.

The medium horizon is conditional on improvements that permit larger, more stable quantum subroutines while still relying heavily on classical computation. Low-rank representations and tensor hypercontraction can reduce resource requirements for electronic-structure algorithms by exploiting structure in the molecular Hamiltonian [30]. These advances may move selected calculations closer to feasibility, particularly when a compact correlated region can be separated from its environment. However, reduced asymptotic or constant costs do not eliminate logical-qubit requirements, state preparation, precision demands, measurement or readout costs, and workflow integration. Medium-horizon opportunities therefore include quantum-

assisted embedding, improved estimators, restricted dynamical calculations, and structured learning problems in which the complete access model and classical baseline are favorable and explicitly tested.

The most defensible long-horizon opportunity is fault-tolerant simulation of chemically consequential electronic structure that remains unreliable or prohibitively expensive classically. Computational-catalysis resource studies illustrate the potential scientific value of accurately resolving strongly correlated transition-metal chemistry, while also showing that useful calculations depend on extensive fault-tolerant resources and detailed algorithm engineering [31]. Related pharmaceutical targets may include enzyme active sites, covalent reaction mechanisms, redox processes, photochemistry, and excited-state pathways, but only when the unresolved quantum-mechanical component affects a decision-relevant observable. The long-horizon claim is therefore not that quantum computers will simulate complete drug action. It is that selected electronic or dynamical subproblems may eventually become tractable at a level of reliability that improves a broader multiscale workflow.

Hybrid quantum-assisted classical methods may occupy a bridge between these horizons. Quantum-assisted approaches to fermionic quantum Monte Carlo suggest that quantum information could help reduce a classical bias without moving the entire chemistry calculation onto a quantum processor [32]. Such niches may emerge earlier than end-to-end fault-tolerant simulation, although their general scaling, error sensitivity, and comparative cost remain unresolved. Ultimately, utility-scale quantum computational chemistry requires more than solving an isolated hard instance: it requires reliable throughput, integration with established software and experimental processes, reproducible outputs, useful turnaround, and scientific value that survives comparison with advancing classical methods [33]. The proposed opportunity map is therefore conditional rather than calendar based. Near, medium, and long horizons describe evidence and dependency structures, not promised delivery dates.

Research Standards for Credible Quantum Drug-Discovery Claims

Credible claims should begin with application-structured benchmarking. General device metrics must be supplemented by tests that predict performance on the circuits, measurements, and hybrid loops required by the target pharmaceutical problem. Capability benchmarking research emphasizes that different quantum processors possess different operational strengths and that a single scalar metric cannot fully characterize application performance [34]. A drug-discovery study should therefore report the hardware, topology, calibration context, compiler, circuit family, depth, execution count, mitigation procedure, classical-loop cost, and variability across runs. Where simulators are used, the simulation assumptions and scalability limits must be explicit. Reproducibility requires sufficient information to rerun both the quantum and classical workflows rather than only the final model.

Quantum-machine-learning claims require equally explicit treatment of data. Theoretical and empirical analysis of data in quantum learning shows that apparent advantage depends strongly on how information is accessed, how much data are available, and whether classical learners can recover the relevant structure [35]. Pharmaceutical studies should disclose molecular representation, feature construction, dimensionality reduction, encoding cost, train–validation separation, chemical-series overlap, temporal or scaffold shift, hyperparameter search, uncertainty calibration, and failed runs. Quantum and classical methods must receive matched information and comparable tuning. A result obtained after classical feature extraction may still be useful, but the claimed contribution must be assigned to the quantum component that was actually tested rather than to the entire pipeline.

Speedup claims require a formally specified computational problem, input-access model, target accuracy, success criterion, and defensible classical-hardness assumption [36]. Asymptotic statements should be separated from finite-resource estimates, and both should include state preparation, repetition, output extraction, error correction, and classical preprocessing. Chemistry claims should identify the Hamiltonian, basis, active space, geometry, environmental treatment, observable, and error tolerance. Learning claims should identify whether the claimed advantage concerns training, inference, sample efficiency, representation, or generalization. Optimization claims should report encoding and decoding costs and show that the mathematical solution corresponds to a chemically feasible candidate. Without these specifications, the term “quantum advantage” remains too ambiguous to guide pharmaceutical investment or scientific interpretation.

Practical advantage should be evaluated as a controlled race between quantum and strong classical approaches under matched data, resources, evaluation endpoints, and generalization conditions [37]. Readiness should then be reported through separate evidence gates: scientific relevance, technical execution, matched benchmarking, end-to-end scaling, prospective validation, experimental confirmation, decision impact, and routine utility. Perspectives on practical quantum advantage similarly distinguish programmable quantum simulation and scientific insight from fault-tolerant computation and end-user value [38]. The proposed reporting standard is an original evidence-organizing synthesis and requires future validation; it is not a certified checklist or universal maturity instrument. **Table 4** organizes the evidence, constructs, and interpretation boundaries needed to develop research standards for credible quantum drug-discovery claims within the article’s central argument.

Table 4. Evaluation, implementation, and research priorities arising from After the Quantum Advantage Claim: Which Drug-Discovery Problems Are Actually Ready for Quantum Computation?

Research priority	Evidence gap	Recommended study or reporting design	Validation requirement	Near-term value	Longer-term significance
Establish application-structured	Device metrics and small molecular tasks do not	Versioned benchmark suites spanning circuit execution, complete	Independent replication across hardware and	Reveals present bottlenecks and prevents	Creates a cumulative basis for identifying

pharmaceutical benchmarks	reliably predict discovery-workflow performance	computational tasks, and decision-oriented endpoints	software stacks with strong classical baselines	misleading generalization	genuine scaling or utility
Attribute value to the quantum component	Hybrid workflow success may arise primarily from classical preprocessing, model capacity, or search effort	Prespecified component ablation with matched data, compute, tuning, and candidate budgets	Reproducible incremental benefit after removal or replacement of the quantum component	Improves interpretation of current hybrid studies	Determines whether quantum integration produces durable net value
Strengthen classical counterfactuals	Comparator methods may be weak, outdated, or incompletely optimized	Expert classical reimplementation, living baselines, and transparent resource accounting	Equivalent target accuracy, input information, tuning opportunity, and output requirements	Reduces false advantage claims	Maintains validity as classical methods continue to improve
Move from retrospective to prospective pharmaceutical evaluation	Public benchmark datasets poorly represent temporal, scaffold, assay, and operational shift	Time-split or prospectively registered prediction, ranking, simulation, or candidate-selection studies	Experimental or high-quality computational confirmation using prespecified endpoints	Tests real generalization and exposes failure modes	Establishes whether quantum outputs influence consequential discovery decisions
Report complete resource ledgers	Qubit and gate counts omit measurement, mitigation, compilation, error correction, and wall-clock demands	End-to-end resource reports with uncertainty ranges and architecture-specific assumptions	Cross-checking by independent resource-estimation teams and sensitivity analyses	Supports realistic project selection	Enables credible planning for fault-tolerant pharmaceutical applications
Make data-access and output assumptions explicit	Quantum-learning and simulation proposals may hide expensive state preparation or readout	Formal input/output model, encoding-cost analysis, and matched access conditions	Demonstration that claimed gains persist after complete loading and extraction costs	Clarifies which near-term studies test meaningful hypotheses	Separates formal complexity advantages from operational utility
Define claim-specific readiness language	Feasibility, benchmark accuracy, experimental success, and routine use are frequently conflated	Evidence-gated reporting template with explicit claim category and disconfirming evidence	Agreement among independent reviewers and prospective application to new studies	Improves scientific communication immediately	Supports a longitudinal and updateable readiness taxonomy
Connect quantum outputs to pharmaceutical decisions	Computational accuracy does not automatically imply therapeutic, synthetic, or development value	Workflow studies that record how outputs alter candidate selection, experiments, or mechanistic conclusions	Prospective evidence of decision impact, reproducibility, and acceptable uncertainty	Focuses research on decision-relevant endpoints	Establishes the final bridge from computation to routine pharmaceutical value

Conclusion

The question following a quantum-advantage claim is not whether a quantum algorithm can be associated with a drug-discovery task, but whether the complete evidence supports a defined level of scientific and operational use. The horizon scan indicates that present quantum drug-discovery research is strongest as a program of bounded feasibility testing, resource characterization, hybrid-workflow experimentation, and benchmark design. No broad pharmaceutical problem class currently satisfies the combined requirements of scalable quantum performance, superiority to strong classical methods, prospective validation, experimental confirmation, attributable decision impact, and routine reliability. The most defensible long-horizon targets are selected electronic-structure and dynamical problems in which strong correlation or explicitly quantum evolution limits classical approximation. Nearer-term opportunities are narrower and should emphasize hybrid components, transparent failure analysis, matched classical controls, and prospectively testable hypotheses. The proposed quantum readiness horizon map, opportunity structure, and evidence gates are conceptual tools for organizing claims and priorities rather than validated maturity measures. Their principal boundary is deliberate: execution is not advantage, benchmark advantage is not pharmaceutical translation, and translation is not routine readiness.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Cao Y, Romero J, Olson JP, Degroote M, Johnson PD, Kieferova M, et al. Quantum chemistry in the age of quantum computing. Chem Rev. 2019;119(19):10856-915. doi:10.1021/acs.chemrev.8b00803
2. McArdle S, Endo S, Aspuru-Guzik A, Benjamin SC, Yuan X. Quantum computational chemistry. Rev Mod Phys. 2020;92(1):015003. doi:10.1103/RevModPhys.92.015003
3. Blunt NS, Camps J, Crawford O, Izsak R, Leontica S, Mirani A, et al. Perspective on the current state-of-the-art of quantum computing for drug discovery applications. J Chem Theory Comput. 2022;18(12):7001-23. doi:10.1021/acs.jctc.2c00574
4. Santagati R, Aspuru-Guzik A, Babbush R, Degroote M, Gonzalez L, Kyoseva E, et al. Drug design on quantum computers. Nat Phys. 2024;20(4):549-57. doi:10.1038/s41567-024-02411-5
5. Smaldone AM, Shee Y, Kyro GW, Xu C, Vu NP, Dutta R, et al. Quantum machine learning in drug discovery: Applications in academia and pharmaceutical industries. Chem Rev. 2025;125(12):5436-60. doi:10.1021/acs.chemrev.4c00678
6. Hines P, Yu LH, Guy RH, Brand A, Papaluca-Amati M. Scanning the horizon: A systematic literature review of methodologies. BMJ Open. 2019;9(5). doi:10.1136/bmjopen-2018-026764
7. Garcia Gonzalez-Moral S, Beyer FR, Oyewole AO, Richmond C, Wainwright L, Craig D. Looking at the fringes of MedTech innovation: A mapping review of horizon scanning and foresight methods. BMJ Open. 2023;13(9). doi:10.1136/bmjopen-2023-073730
8. Page MJ, McKenzie JE, Bossuyt PM, Boutron I, Hoffmann TC, Mulrow CD, et al. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ. 2021;372. doi:10.1136/bmj.n71
9. Cerezo M, Arrasmith A, Babbush R, Benjamin SC, Endo S, Fujii K, et al. Variational quantum algorithms. Nat Rev Phys. 2021;3(9):625-44. doi:10.1038/s42254-021-00348-9
10. Bharti K, Cervera-Lierta A, Kyaw TH, Haug T, Alperin-Lea S, Anand A, et al. Noisy intermediate-scale quantum algorithms. Rev Mod Phys. 2022;94(1):015004. doi:10.1103/RevModPhys.94.015004
11. Tilly J, Chen H, Cao S, Picozzi D, Setia K, Li Y, et al. The variational quantum eigensolver: A review of methods and best practices. Phys Rep. 2022;986:1-128. doi:10.1016/j.physrep.2022.08.003
12. Motta M, Rice JE. Emerging quantum computing algorithms for quantum chemistry. WIREs Comput Mol Sci. 2022;12(3). doi:10.1002/wcms.1580
13. Batra K, Zorn KM, Foil DH, Minerali E, Gawriljuk VO, Lane TR, et al. Quantum machine learning algorithms for drug discovery applications. J Chem Inf Model. 2021;61(6):2641-47. doi:10.1021/acs.jcim.1c00166
14. Reiher M, Wiebe N, Svore KM, Wecker D, Troyer M. Elucidating reaction mechanisms on quantum computers. Proc Natl Acad Sci U S A. 2017;114(29):7555-60. doi:10.1073/pnas.1619152114
15. Ollitrault PJ, Miessen A, Tavernelli I. Molecular quantum dynamics: A quantum computing perspective. Acc Chem Res. 2021;54(23):4229-38. doi:10.1021/acs.accounts.1c00514
16. Urbanek M, Camps D, Van Beeumen R, de Jong WA. Chemistry on quantum computers with virtual quantum subspace expansion. J Chem Theory Comput. 2020;16(9):5425-31. doi:10.1021/acs.jctc.0c00447
17. Smaldone AM, Batista VS. Quantum-to-classical neural network transfer learning applied to drug toxicity prediction. J Chem Theory Comput. 2024;20(11):4901-8. doi:10.1021/acs.jctc.4c00432
18. Bhatia AS, Saggi MK, Kais S. Quantum machine learning predicting ADME-Tox properties in drug discovery. J Chem Inf Model. 2023;63(21):6476-86. doi:10.1021/acs.jcim.3c01079
19. Gonthier JF, Radin MD, Buda C, Doskocil EJ, Abuan CM, Romero J. Measurements as a roadblock to near-term practical quantum advantage in chemistry: Resource analysis. Phys Rev Res. 2022;4(3):033154. doi:10.1103/PhysRevResearch.4.033154
20. Cai Z, Babbush R, Benjamin SC, Endo S, Huggins WJ, Li Y, et al. Quantum error mitigation. Rev Mod Phys. 2023;95(4):045005. doi:10.1103/RevModPhys.95.045005
21. Wang S, Fontana E, Cerezo M, Sharma K, Sone A, Cincio L, et al. Noise-induced barren plateaus in variational quantum algorithms. Nat Commun. 2021;12(1):6961. doi:10.1038/s41467-021-27045-6
22. McClean JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H. Barren plateaus in quantum neural network training landscapes. Nat Commun. 2018;9(1):4812. doi:10.1038/s41467-018-07090-4

23. Kim IH, Liu YH, Pallister S, Pol W, Roberts S, Lee E. Fault-tolerant resource estimate for quantum chemical simulations: Case study on Li-ion battery electrolyte molecules. Phys Rev Res. 2022;4(2):023019. doi:10.1103/PhysRevResearch.4.023019
24. Kandala A, Mezzacapo A, Temme K, Takita M, Brink M, Chow JM, et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. Nature. 2017;549(7671):242-6. doi:10.1038/nature23879
25. Goings JJ, White A, Lee J, Tautermann CS, Degroote M, Gidney C, et al. 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(38). doi:10.1073/pnas.2203533119
26. Li W, Yin Z, Li X, Ma D, Yi S, Zhang Z, et al. A hybrid quantum computing pipeline for real world drug discovery. Sci Rep. 2024;14(1):16942. doi:10.1038/s41598-024-67897-8
27. Liang Z, He Z, Sun Y, Herman D, Jiao Q, Zhu Y, et al. Synergizing quantum techniques with machine learning for advancing drug discovery challenge. Sci Rep. 2024;14(1):31216. doi:10.1038/s41598-024-82576-4
28. Ghazi Vakili M, Gorgulla C, Snider J, Nigam AK, Bezrukov D, Varoli D, et al. Quantum-computing-enhanced algorithm unveils potential KRAS inhibitors. Nat Biotechnol. 2025;43(12):1954-9. doi:10.1038/s41587-024-02526-3
29. Kivlichan ID, McClean J, Wiebe N, Gidney C, Aspuru-Guzik A, Chan GKL, et al. Quantum simulation of electronic structure with linear depth and connectivity. Phys Rev Lett. 2018;120(11):110501. doi:10.1103/PhysRevLett.120.110501
30. Lee J, Berry DW, Gidney C, Huggins WJ, McClean JR, Wiebe N, et al. Even more efficient quantum computations of chemistry through tensor hypercontraction. PRX Quantum. 2021;2(3):030305. doi:10.1103/PRXQuantum.2.030305
31. von Burg V, Low GH, Haner T, Steiger DS, Reiher M, Roetteler M, et al. Quantum computing enhanced computational catalysis. Phys Rev Res. 2021;3(3):033055. doi:10.1103/PhysRevResearch.3.033055
32. Huggins WJ, O'Gorman BA, Rubin NC, Reichman DR, Babbush R, Lee J. Unbiasing fermionic quantum Monte Carlo with a quantum computer. Nature. 2022;603(7901):416-20. doi:10.1038/s41586-021-04351-z
33. Castaldo D, Reiher M. Utility-scale quantum computational chemistry. J Phys Chem Lett. 2026;17(29):8140-55. doi:10.1021/acs.jpclett.6c00888
34. Proctor T, Rudinger K, Young K, Nielsen E, Blume-Kohout R. Measuring the capabilities of quantum computers. Nat Phys. 2022;18(1):75-9. doi:10.1038/s41567-021-01409-7
35. Huang HY, Broughton M, Mohseni M, Babbush R, Boixo S, Neven H, et al. Power of data in quantum machine learning. Nat Commun. 2021;12(1):2631. doi:10.1038/s41467-021-22539-9
36. Liu Y, Arunachalam S, Temme K. A rigorous and robust quantum speed-up in supervised machine learning. Nat Phys. 2021;17(9):1013-17. doi:10.1038/s41567-021-01287-z
37. Hibat-Allah M, Mauri M, Carrasquilla J, Perdomo-Ortiz A. A framework for demonstrating practical quantum advantage: Comparing quantum against classical generative models. Commun Phys. 2024;7(1):68. doi:10.1038/s42005-024-01552-6
38. Daley AJ, Bloch I, Kokail C, Flannigan S, Pearson N, Troyer M, et al. Practical quantum advantage in quantum simulation. Nature. 2022;607(7920):667-76. doi:10.1038/s41586-022-04940-6