Pharmacophore

ISSN-2229-5402

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

QUANTUM–CLASSICAL DRUG DISCOVERY AFTER THE ADVANTAGE CLAIM: MATCHING COMPUTATIONAL METHODS TO PROBLEMS THEY CAN CREDIBLY SOLVE

David Thompson[1]*, Sarah Mitchell[1], Rachel Adams[2], James Brown[3]

1. ***Department of Quantum–Classical Drug Discovery and Method Matching, Faculty of Pharmacy, University of Toronto, Toronto, Canada.***
2. ***Department of Credible Problem Solving in Quantum Computing, Faculty of Pharmaceutical Sciences, University of British Columbia, Vancouver, Canada.***
3. ***Department of Beyond-Advantage Claim Computational Methods, Faculty of Pharmacy, McGill University, Montreal, Canada.***

ARTICLE INFO

Received:
07 September 2025
Received in revised form:
15 December 2025
Accepted:
18 December 2025
Available online:
28 December 2025

Keywords: Quantum computing, Drug discovery, Hybrid quantum–classical methods, Quantum machine learning, Computational chemistry, Molecular representation

ABSTRACT

Quantum computing has become increasingly visible in pharmaceutical research, yet the evidentiary meaning of an "advantage" claim remains unsettled. Drug discovery contains heterogeneous scientific decisions, data structures, computational abstractions, validation standards, and translational consequences. Consequently, improvement in circuit execution, molecular-energy estimation, sampling, classification, or candidate generation cannot by itself establish pharmaceutical usefulness. This Original Quantum–Classical Position Article examines how quantum methods should be matched to problems they can credibly solve rather than evaluated through a single performance measure. It develops a proposed Quantum–Classical Credibility Map that links pharmaceutical problem specification, task selection, algorithmic assumptions, hardware conditions, data and encoding requirements, hybrid workflow design, comparator quality, uncertainty, and downstream validation. The synthesis distinguishes computational feasibility from practical advantage, scientific utility, and pharmaceutical value. It further separates electronic-structure calculations, reaction-mechanism analysis, molecular generation, property prediction, and quantum-assisted optimization according to their outputs, resource dependencies, and validation routes. The central proposition is that credibility is conjunctive: a claim cannot be stronger than the least-supported necessary connection between the pharmaceutical decision, computational subproblem, quantum method, classical workflow, and empirical evidence. The proposed framework does not establish quantum advantage, clinical utility, regulatory acceptance, or deployment readiness. Its value lies in organizing research questions, exposing hidden assumptions, improving quantum–classical comparisons, and defining conditions under which a pharmaceutical application should be investigated, reformulated, monitored, or deferred. Prospective benchmarking, independent replication, complete resource accounting, and experimentally connected evaluation are required before the framework itself or the applications organized through it can be considered validated.

To Cite This Article: Thompson D, Mitchell S, Adams R, Brown J. Quantum–Classical Drug Discovery after the Advantage Claim: Matching Computational Methods to Problems They Can Credibly Solve. Pharmacophore. 2025;16(6):120-30. https://doi.org/10.51847/D59nKXqOBj

Introduction

Quantum computing is frequently presented as a response to computational bottlenecks in molecular science, particularly where electronic structure, combinatorial search, high-dimensional sampling, or machine learning may become difficult for conventional methods. The scientific opportunity is real, but it is not uniform across hardware regimes or problem classes. Near-term noisy processors and fault-tolerant quantum computers represent materially different computational regimes, while quantum-chemistry algorithms impose heterogeneous mapping, state-preparation, error-management, measurement, and resource requirements [1-3]. A calculation that is conceptually appropriate for a future error-corrected device may therefore be operationally unsuitable for current hardware. Conversely, a circuit executable on a small processor may address a representation or benchmark chosen for convenience rather than a pharmaceutical bottleneck. The phrase "quantum advantage"

Corresponding Author: David Thompson; Department of Quantum–Classical Drug Discovery and Method Matching, Faculty of Pharmacy, University of Toronto, Toronto, Canada. E-mail: david.thompson@utoronto.ca

obscures these distinctions when it is used without specifying the target problem, comparator, resource model, and scientific consequence of the output.

This ambiguity is especially consequential in drug discovery because the field is not one computational task. It encompasses target-related inference, molecular generation, property prediction, electronic-structure analysis, conformational modeling, interaction assessment, reaction planning, ADME-Tox estimation, synthesis prioritization, and experimental decision-making. Existing analyses of quantum computing for drug discovery show that proposed applications differ substantially in mathematical structure, maturity, hardware dependence, and relationship to established pharmaceutical workflows [4]. It is therefore inadequate to ask whether quantum computing is useful "for drug discovery" in general. A credible evaluation must instead ask which quantum method is proposed, which bounded subproblem it addresses, why that subproblem is consequential, what information enters and leaves the calculation, and how its output could alter a defensible scientific decision.

Quantum machine learning further complicates this assessment because a reported result may depend as much on representation, encoding, preprocessing, data access, model capacity, and classical optimization as on the quantum circuit itself. Current pharmaceutical quantum-machine-learning research includes classification, property prediction, generative modeling, kernel methods, and hybrid architectures, but the evidence remains sensitive to small datasets, simulators, restricted hardware, uncertain scaling, and comparison choices [5]. A result can consequently be valid as a proof of concept while remaining insufficient as evidence of practical advantage. Likewise, synthesis or assay follow-up may support the plausibility of a complete discovery workflow without demonstrating that the quantum component produced an incremental benefit over a carefully matched classical alternative.

This article proposes a problem-first position for quantum–classical drug discovery. Its original contribution is the Quantum–Classical Credibility Map, a conceptual construct that connects pharmaceutical problem structure with task selection, algorithmic assumptions, hardware regime, data requirements, hybrid workflow attribution, fair comparison, uncertainty, and downstream validation. The construct is governed by a weakest-link principle: the credibility of a final claim cannot exceed the least-supported necessary connection in the chain from pharmaceutical question to validated use. This principle rejects the idea that one favorable indicator—such as circuit accuracy, model performance, sampling novelty, energy estimation, or candidate generation—can compensate for an irrelevant target quantity, an unrealistic hardware assumption, an inadequate comparator, or an absent experimental pathway. The framework is proposed rather than validated and is intended to discipline claims, research design, and investment reasoning rather than to certify applications as ready for operational pharmaceutical use.

Separating Quantum Promise from Pharmaceutical Problem Structure

The first requirement is to distinguish the pharmaceutical problem from the computational abstraction chosen to represent it. Pharmaceutical research comprises heterogeneous prediction, generation, prioritization, optimization, and experimental-design activities, and these activities cannot be evaluated through a common success measure. Molecular generation asks whether useful regions of chemical space can be sampled under relevant constraints; property prediction asks whether an endpoint can be estimated for the intended population; medicinal-chemistry design asks whether a candidate can be synthesized, tested, interpreted, and improved within a broader program. These tasks are related, but they are not interchangeable, and their outputs have different evidentiary meanings [6-8]. A generative model can produce syntactically valid structures without producing synthesizable or developable compounds. A property model can discriminate a retrospective dataset without improving prospective prioritization. An electronic-structure calculation can be accurate for a selected Hamiltonian while remaining disconnected from the biological, kinetic, or formulation processes that govern therapeutic performance.

A problem-first formulation should therefore identify at least four layers. The first is the pharmaceutical decision, such as selecting a reaction mechanism for further study, prioritizing compounds for synthesis, or estimating whether an electronic contribution warrants higher-level analysis. The second is the computational target, including an energy, distribution, classifier output, generated structure, or optimization objective. The third is the decision linkage, which explains why improvement in that target should change pharmaceutical action. The fourth is the validation route, which specifies the classical calculation, experimental measurement, or prospective decision study needed to test the claimed value. These layers prevent benchmark design from substituting for problem definition. Comparative studies in molecular learning demonstrate that conclusions about method superiority can change with dataset selection, splitting strategy, endpoint construction, and comparator choice [9]. A quantum model should not receive weaker evaluation merely because the underlying hardware is novel.

The representation used to connect a pharmaceutical object with a computational method is equally important. A molecule may be encoded as a fingerprint, graph, sequence, descriptor vector, Hamiltonian, density representation, conformer ensemble, or learned embedding. Each representation preserves some information, suppresses other information, and imposes assumptions about invariance, locality, geometry, or experimental context. Comparative evidence from molecular property prediction shows that representation choice affects model behavior, predictive performance, and transferability [10]. For quantum methods, this issue extends beyond statistical expressiveness. Encoding may require classical preprocessing, dimensionality reduction, amplitude preparation, qubit mappings, or repeated circuit execution. The claimed benefit must therefore be evaluated after these costs and transformations are included. Data availability is not equivalent to data suitability, and compact encoding is not necessarily information-preserving encoding.

The proposed Quantum–Classical Credibility Map consequently begins with a pharmaceutical problem specification rather than an algorithm. This specification should state the scientific decision, computational abstraction, acceptable uncertainty, current classical limitation, intended output, and downstream validation requirement. Only after these elements are explicit should a quantum or hybrid method be considered. The construct does not imply that every pharmaceutical task requires experimental validation at the same stage; rather, it requires that the appropriate validation pathway be named and that the claim remain bounded until that pathway is completed. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop separating quantum promise from pharmaceutical problem structure within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for Separating quantum promise from pharmaceutical problem structure

Construct or evidence domain	Core question	Scientific basis	Role in the article	Failure or overclaiming risk	Boundary statement
Pharmaceutical decision	What action or scientific judgment is the computation intended to support?	Drug discovery contains distinct design, prediction, prioritization, and validation decisions [6].	Establishes the endpoint against which computational relevance is judged.	A difficult calculation may be mistaken for a consequential pharmaceutical problem.	Computational difficulty alone does not establish decision relevance.
Computational abstraction	Which mathematical object represents the pharmaceutical question?	Molecular-design tasks may be formulated as generation, prediction, simulation, ranking, or optimization [8].	Defines what the method actually calculates.	The abstraction may omit kinetics, biology, synthesis, formulation, or experimental context.	Success on the abstraction applies only to the represented quantity.
Decision linkage	Why should improvement in the computational output alter pharmaceutical action?	Drug-design value depends on connection to medicinal chemistry and experimental constraints [7].	Connects method performance with prospective scientific use.	A favorable output may remain scientifically interesting but operationally irrelevant.	Benchmark improvement does not independently establish pharmaceutical usefulness.
Molecular representation	Which chemical and contextual information is retained or discarded?	Representation choice affects predictive behavior and transferability [10].	Defines the interface between pharmaceutical data and the computational method.	Encoding may remove relevant structure or introduce hidden preprocessing dependence.	Data availability does not establish representation suitability.
Classical comparator	What is the strongest relevant non-quantum method under matched conditions?	Comparative conclusions depend on benchmark and comparator construction [9].	Provides the counterfactual required for an advantage claim.	Weak or outdated baselines can produce artificial superiority.	Novelty of hardware does not justify a weaker comparator.
Resource model	Which costs are included in feasibility or advantage claims?	Quantum workflows may require encoding, compilation, sampling, mitigation, optimization, and verification.	Prevents circuit-only accounting.	Excluded preprocessing or measurement costs may reverse the conclusion.	End-to-end resources must match the scope of the claim.
Uncertainty	Which errors arise from data, modeling, approximation, hardware, and translation?	Pharmaceutical usefulness requires interpretation of multiple uncertainty sources.	Limits confidence in downstream decisions.	Model confidence may be reported as though it were calibrated uncertainty.	Uncertainty must be decomposed and assessed for the intended use.
Validation route	What evidence would confirm that the output is scientifically useful?	Computational outputs require task-appropriate classical, experimental, or prospective confirmation.	Defines how the claim can progress beyond conceptual plausibility.	Molecular plausibility may be reported as mechanistic or therapeutic confirmation.	Prediction, mechanism, experimental confirmation, and clinical utility remain distinct.

Drug-Discovery Tasks Proposed for Quantum Computation

Electronic-structure simulation is the most direct scientific connection between quantum computing and molecular chemistry because the target system is itself quantum mechanical. Proposed applications include ground- and excited-state estimation, treatment of strongly correlated active spaces, reaction-pathway analysis, and calculation of molecular properties derived from electronic structure. These applications are potentially relevant when conventional approximations become unreliable or prohibitively costly for a chemically consequential subsystem. Resource-grounded work on reaction-mechanism elucidation illustrates both the attraction and the difficulty of this task class: a scientifically meaningful mechanism can be mapped to a quantum computation, but the accuracy target, active-space choice, state preparation, error tolerance, and hardware resources must be specified explicitly [11]. Such examples support investigation of quantum simulation; they do not establish that an entire drug-discovery program, protein–ligand system, or biological mechanism can be solved by the same calculation.

Quantum machine learning addresses a different family of tasks. Here, classical molecular or biological information is encoded into a quantum state or parameterized circuit and used for classification, regression, kernel evaluation, or generative sampling. Proof-of-concept studies have applied quantum learning algorithms to drug-related datasets, demonstrating that bounded pharmaceutical prediction tasks can be implemented in quantum or simulated quantum settings [12]. The appropriate interpretation is methodological feasibility rather than generalized superiority. Generative chemistry has similarly been proposed as a near-term application in which quantum distributions, circuit Born machines, or hybrid latent representations contribute to candidate generation [13]. Yet a generated molecular distribution must still be judged by chemical validity, novelty relative to training data, constraint satisfaction, synthesizability, property balance, uncertainty, and performance against contemporary classical generators. Sampling a different distribution is not itself evidence that the distribution is more useful for drug discovery.

Property-prediction studies extend this reasoning to ADME-Tox and related endpoints. A quantum model may produce encouraging benchmark behavior for a selected dataset, but such evidence remains bounded by sample composition, endpoint quality, splitting protocol, representation, tuning, and external validity. Quantum-machine-learning work on ADME-Tox prediction demonstrates a defined application class, not experimental confirmation of absorption, metabolism, toxicity, developability, or clinical safety [14]. More integrated studies can strengthen the evidentiary chain by connecting computational generation with synthesis and biological testing. A hybrid quantum–classical workflow used to identify and evaluate potential KRAS inhibitors illustrates how generated candidates may proceed to experimental follow-up [15]. That progression supports the feasibility of a complete discovery workflow, but it does not independently isolate the causal contribution of the quantum component or prove that it outperforms a matched classical generator.

These task classes should therefore be organized by their output semantics and validation obligations rather than by their shared use of quantum hardware. Electronic-structure tasks output approximations to quantum-mechanical quantities; predictive tasks output endpoint estimates; generative tasks output candidate distributions; optimization tasks output solutions to a specified objective; and quantum-assisted methods output information intended to improve a classical solver. Each class requires a different comparator, uncertainty model, and downstream test. **Figure 1** presents the quantum–classical credibility map, showing how the article's key components and boundaries are connected within drug-discovery tasks proposed for quantum computation.

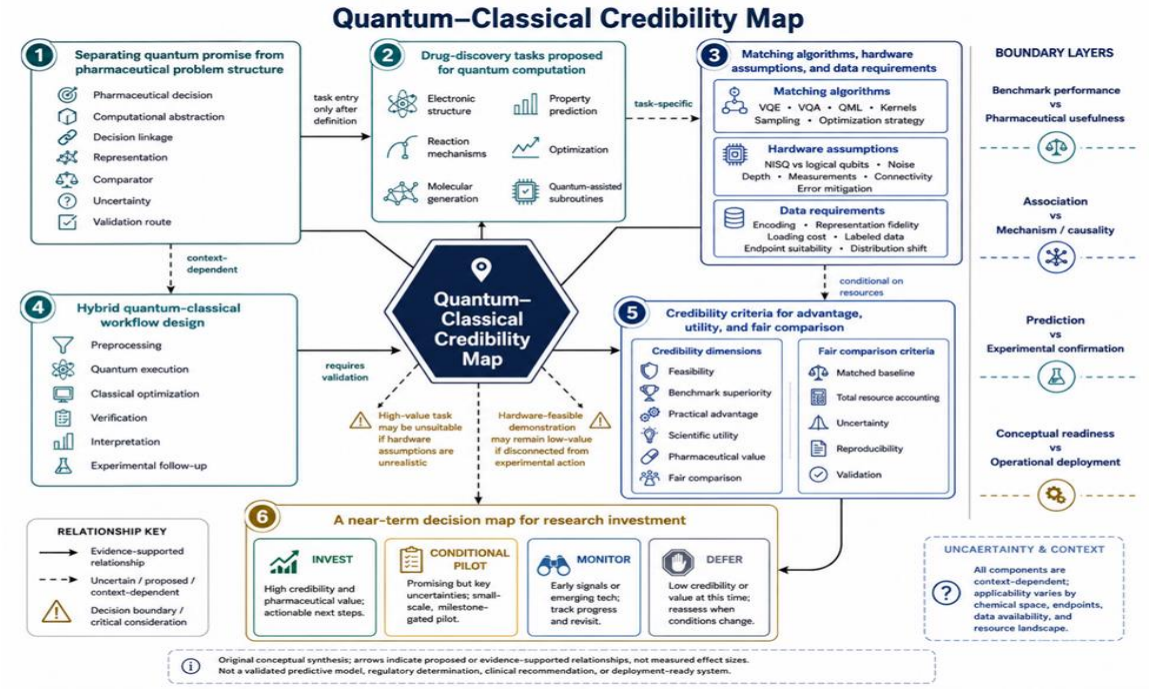

Figure 1. Quantum–Classical Credibility Map.

The figure is an original conceptual synthesis that organizes the article's central contribution across separating quantum promise from pharmaceutical problem structure, drug-discovery tasks proposed for quantum computation, matching algorithms, hardware assumptions, and data requirements, matching algorithms, hardware assumptions, data requirements. 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.

The proposed map treats task selection as a sequence of credibility constraints. A task enters the map only after the pharmaceutical decision and computational abstraction are defined. It then passes through algorithm, data, encoding, hardware, workflow, comparison, and validation layers. The map does not rank electronic structure, generation, prediction, or optimization as intrinsically more valuable. Instead, it asks whether each task possesses a defensible relationship between its mathematical structure and the pharmaceutical decision it is intended to support. A high-value task may remain unsuitable because its hardware requirements are unrealistic, while a hardware-feasible demonstration may remain low value because its output is disconnected from experimental action. This task-level distinction is essential to the article's broader position: quantum–classical drug discovery should advance through credible problem–method matches, not through generalized claims that one computational paradigm has solved pharmaceutical discovery.

Matching Algorithms, Hardware Assumptions, and Data Requirements

Method selection should follow from the mathematical structure of the bounded pharmaceutical subproblem rather than from the availability of a particular circuit or device. Variational quantum algorithms are often proposed for chemistry, optimization, and learning because they divide work between parameterized quantum circuits and classical optimizers. Their suitability, however, depends on the expressiveness of the ansatz, the geometry of the objective function, the number and type of measurements, parameter initialization, trainability, and the noise characteristics of the hardware [16]. A variational method is therefore not a general solution category. It is a family of conditional strategies whose relevance must be justified for the specific Hamiltonian, distribution, classifier, or optimization landscape being addressed.

For electronic-structure calculations, the variational quantum eigensolver illustrates how algorithmic assumptions propagate through a complete computation. Molecular orbitals must be selected, fermionic operators mapped to qubits, an initial state prepared, an ansatz chosen, expectation values measured, and parameters updated through a classical optimizer. Different decisions can alter circuit depth, measurement cost, convergence, chemical accuracy, and susceptibility to noise [17]. Reporting only the final energy or circuit size hides these dependencies. A credible pharmaceutical application should state why the selected active space and observable are relevant to the drug-discovery question, what error is tolerable, and how the result will be checked against accepted classical calculations or experimental evidence.

Hardware feasibility cannot be inferred from qubit count alone. Connectivity, native gates, coherence, control error, readout error, compilation overhead, queueing, and repeated sampling all affect whether a nominally small circuit produces useful information. Noise also interacts with the optimization process rather than merely adding a final perturbation. In variational circuits, noise-induced concentration of cost functions can suppress gradients and impair trainability as the system grows [18]. A method may therefore become practically untrainable before it exceeds the advertised number of available qubits. Hardware assumptions should consequently be reported as part of the scientific hypothesis, not as implementation details that can be supplied after the algorithm has been selected.

Data requirements create an additional matching constraint for quantum machine learning. The scientific meaning of a claimed advantage depends on whether data are classical or quantum, how they are accessed, what encoding is used, how many circuit repetitions are required, and whether the input procedure preserves the computational benefit. Theoretical analyses show that the learning power attributed to quantum models can depend critically on assumptions about data access and representation [19]. Pharmaceutical datasets also contain measurement error, assay heterogeneity, class imbalance, missingness, distribution shift, and endpoint ambiguity that quantum encoding does not remove. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop matching algorithms, hardware assumptions, and data requirements within the article's central argument.

Table 2. Components, relationships, uncertainties, and validation needs within Matching algorithms, hardware assumptions, and data requirements

Component or layer	Inputs	Core function	Expected output	Uncertainty or limitation	Validation requirement
Pharmaceutical task definition	Scientific decision, target quantity, acceptable error, downstream action	Converts a broad discovery objective into a bounded computational problem	Explicit task specification	The selected target may not determine the pharmaceutical decision	Review by pharmaceutical and computational-domain experts
Algorithm family	Hamiltonian, optimization structure, data type, desired observable	Selects the computational strategy	Simulation, prediction, sampling, or optimization result	Suitability may be instance-specific rather than general	Comparison with alternative quantum and classical methods

Ansatz or model architecture	Initial state, symmetry, circuit topology, parameterization	Defines the accessible solution space	Parameterized quantum state or decision function	Limited expressiveness, excessive depth, or optimization difficulty	Sensitivity analysis and chemically meaningful test instances
Molecular representation	Molecular graph, descriptors, orbitals, conformers, Hamiltonian, assay data	Encodes the pharmaceutical object for computation	Quantum state, feature map, operator, or input tensor	Information loss, hidden preprocessing, dimensionality reduction	Representation ablation and domain-validity analysis
State preparation and data loading	Encoded classical or quantum information	Initializes the quantum computation	Prepared input state	Loading cost may dominate the proposed benefit	End-to-end timing and resource accounting
Hardware regime	Physical or logical qubits, connectivity, native gates, noise, correction assumptions	Constrains executable algorithms	Hardware-specific circuit execution	Nominal qubit count may conceal depth, fidelity, or compilation limits	Device-level reporting and cross-platform replication
Classical optimizer	Objective values, gradients or gradient estimates, stopping rules	Updates circuit or model parameters	Candidate parameter set	Local minima, stochasticity, barren plateaus, optimizer dependence	Repeated runs, initialization analysis, and comparator optimizers
Measurement strategy	Observables, shot allocation, grouping, mitigation procedure	Estimates quantities from circuit executions	Expectation values or samples	High variance and sampling overhead	Confidence intervals, convergence analysis, and full shot reporting
Error handling	Calibration data, mitigation assumptions, correction code	Reduces or controls hardware-induced error	Corrected or mitigated estimate	Bias reduction may require substantial additional sampling	Raw-versus-mitigated reporting and independent verification
Classical comparator	Current algorithms, matched data, equivalent tuning and compute budget	Establishes the counterfactual	Relative performance and resource profile	Weak, outdated, or unmatched baselines	Predeclared comparator protocol and periodic baseline refresh
Output verification	Classical calculations, known limits, synthetic controls, experiments	Tests correctness and scientific relevance	Verified or bounded interpretation	Agreement on small instances may not establish scalability	Independent replication and prospective task-specific validation

Hybrid Quantum-Classical Workflow Design

Hybrid design should be understood as an allocation of scientific responsibilities rather than as the mere presence of both quantum and classical code. In the canonical variational loop, the quantum processor prepares a parameterized state and estimates observables, while the classical processor updates parameters and determines subsequent circuit evaluations [20]. Pharmaceutical workflows add further layers: molecular standardization, conformer generation, feature construction, orbital selection, candidate filtering, uncertainty analysis, visualization, expert interpretation, synthesis planning, and experimental testing. Most of these activities remain classical or human-directed. A credible hybrid architecture must identify which output is uniquely attributable to the quantum stage and which outputs arise from preprocessing, postprocessing, or domain constraints.

Error mitigation illustrates why the interface between quantum and classical stages must be explicit. Short-depth mitigation methods estimate or reduce bias through additional circuit executions, noise scaling, probabilistic transformations, or calibrated postprocessing [21]. These procedures may improve an estimate, but they introduce assumptions, variance, and sampling cost. The corrected quantity is therefore a hybrid product rather than an unqualified hardware output. Reporting should preserve raw and mitigated results, calibration conditions, shot counts, uncertainty, and the stability of the conclusion under alternative mitigation choices. Error mitigation cannot be treated as equivalent to fault tolerance, nor should a mitigated benchmark be extrapolated automatically to larger pharmaceutical systems.

Hybridization can also take forms in which the quantum processor assists an established classical solver. Quantum-assisted fermionic Monte Carlo, for example, uses quantum-derived information to reduce bias in a classical computational method [22]. This design is conceptually important because it rejects the false choice between complete quantum replacement and no quantum contribution. In pharmaceutical science, a bounded quantum subroutine may eventually support active-space selection, state preparation, kernel evaluation, distribution sampling, or another narrow operation while the established

classical workflow remains responsible for scale, control, interpretation, and validation. Such assistance should be judged by incremental end-to-end value rather than by whether the quantum calculation appears technically sophisticated in isolation.
A pharmaceutical hybrid workflow should therefore be documented through an explicit workflow contract. The contract should define inputs, ownership of preprocessing, quantum operations, classical optimization, uncertainty propagation, verification, fallback procedures, and the downstream decision. Published hybrid drug-design pipelines demonstrate that quantum and classical components can be arranged within a broader molecular workflow [23]. Their scientific interpretation nevertheless depends on ablation: the complete pipeline should be compared with a classical-only version, a simulator-based version where appropriate, and alternative model components under matched conditions. Without this decomposition, successful candidate ranking or molecular generation cannot be attributed confidently to the quantum stage. The proposed workflow contract is intended to make such attribution testable rather than assumed.

Credibility Criteria for Advantage, Utility, and Fair Comparison

A credible evaluation must distinguish at least four claims: computational feasibility, benchmark superiority, practical advantage, and pharmaceutical utility. Feasibility asks whether a calculation can be executed. Benchmark superiority asks whether it performs better on a specified test under stated conditions. Practical advantage asks whether the complete quantum or hybrid process improves a consequential resource relative to the strongest relevant classical alternative. Pharmaceutical utility asks whether the output improves a scientific or development decision. Chemistry benchmarks, resource analyses of measurement cost, and error-mitigation scholarship jointly show that circuit execution, sampling burden, mitigation overhead, and hardware configuration must be evaluated together [24-26]. None of these dimensions can be inferred from a single energy estimate, classification score, generated molecule, or asymptotic complexity statement.
Fair comparison requires symmetry in the treatment of quantum and classical methods. Data, splits, preprocessing, tuning effort, model capacity, stopping criteria, uncertainty estimation, and resource boundaries should be matched or justified. A recent framework for demonstrating practical quantum advantage through races between quantum and classical generative models emphasizes the importance of matched resources and generalization-oriented assessment [27]. The principle extends beyond generative chemistry. Quantum methods should not be compared with obsolete classical baselines, minimally tuned models, or implementations denied equivalent computational resources. Conversely, the classical comparator should not be granted inaccessible proprietary data or an unlimited optimization budget. The comparison protocol should be predeclared where possible and updated as classical algorithms improve.
Resource accounting should follow the scope of the claim. A circuit-level claim may legitimately report gate count, depth, fidelity, and shots, but it should not be described as an end-to-end pharmaceutical advantage. A workflow-level claim should include data preparation, encoding, compilation, calibration, mitigation, optimization, repeated runs, queueing, postprocessing, verification, and relevant experimental follow-up. Different resources may matter in different settings: wall time, energy use, monetary cost, sample efficiency, memory, human expertise, or access to specialized infrastructure. These dimensions should not be collapsed into one score unless the weighting is explicit and justified. Practical advantage is therefore conditional on a stated resource model and intended user.
Scientific utility also requires evidence beyond faster or more accurate computation. A quantum-assisted result may be useful if it reveals a mechanistic feature, resolves a chemically consequential ambiguity, improves prospective prioritization, or enables a calculation that is otherwise impractical. Utility may exist without a formal speedup, and speedup may exist without utility. Pharmaceutical value demands an additional connection to synthesis, experiments, developability, safety, or another decision-relevant endpoint. Under the proposed framework, claims should progress through a hierarchy: technically feasible, benchmark-supported, practically advantageous, scientifically useful, and pharmaceutically consequential. Movement between levels requires new evidence rather than rhetorical extension of the preceding result.

A Near-Term Decision Map for Research Investment

Near-term research investment should prioritize problem classes with an explicit scientific bottleneck, plausible quantum structure, auditable classical limits, measurable intermediate outputs, and a realistic validation route. Reviews of quantum algorithms for chemistry and materials science indicate that advantage should be sought in selected structures and algorithms rather than assumed for arbitrary simulation or optimization tasks [28]. This supports investment in enabling methods—resource estimation, benchmark construction, active-space design, compilation, mitigation analysis, and verification—even when immediate pharmaceutical advantage is not expected. Such work can reduce uncertainty about which applications deserve larger prospective studies.
A second priority is full-stack co-design. Quantum applications depend on interactions among scientific formulation, algorithms, software, control systems, architecture, and hardware. A systems perspective on quantum computing for scientific discovery argues that these layers should be developed together rather than sequentially [29]. For pharmaceutical research, co-design should also include medicinal chemistry, computational chemistry, pharmacology, data governance, and experimental planning. An application team should be able to explain not only how the circuit operates but also why its output matters, how it will be checked, and what evidence would cause the project to change direction.
Investment decisions should also account for the difference between asymptotic improvement and usable advantage. In fault-tolerant settings, error-correction costs, logical operations, state preparation, and runtime constants may be large enough that modest asymptotic speedups fail to produce a practical benefit [30]. Near-term projects should therefore include explicit stop

rules. A project may be downgraded or deferred when updated classical methods remove the identified bottleneck, resource estimates grow beyond the intended regime, the required data encoding destroys the benefit, or intermediate outputs fail to influence the pharmaceutical decision. Stopping is not evidence that quantum computing lacks scientific value; it is evidence that the particular problem–method match is presently unsupported.

The final category is a conditional pharmaceutical pilot. Such a pilot is justified only when the problem is bounded, the output is consequential, the hardware and algorithm assumptions are explicit, the classical comparator is current, and a prospective validation pathway is available. Resource-aware perspectives on early logical-qubit applications emphasize limited use cases, co-design, and auditable milestones rather than universal claims [31]. The proposed decision map therefore classifies projects as enabling investment, conditional pilot, monitor, or defer. **Table 3** organizes the evidence, constructs, and interpretation boundaries needed to develop a near-term decision map for research investment within the article's central argument.

Table 3. Evaluation, implementation, and research priorities arising from Quantum–Classical Drug Discovery after the Advantage Claim Matching Computational Methods to Problems

Evaluation dimension	What must be tested	Suitable evidence or assessment	Failure signal	Interpretive limitation	Decision relevance
Pharmaceutical problem relevance	Whether the computed quantity changes a defensible scientific decision	Expert-defined decision pathway; prospective prioritization exercise	Output has no effect on candidate selection or scientific interpretation	Decision relevance may differ across programs and stages	Determines whether a project should progress beyond methodological research
Problem–algorithm match	Whether the method exploits structure genuinely present in the task	Complexity analysis; instance scaling; comparison of alternative formulations	Performance depends on an artificial or simplified representation	Small-instance success does not establish scalable suitability	Supports enabling investment or conditional pilot
Hardware feasibility	Whether the required computation fits the stated physical or logical regime	Qubit, depth, connectivity, fidelity, shot, and runtime estimates	Requirements exceed the declared hardware assumptions	Hardware progress may later change the classification	Distinguishes current pilot from monitor or defer
Data and encoding suitability	Whether available data can be loaded and represented without negating the benefit	End-to-end encoding cost; representation sensitivity; provenance review	Preprocessing or loading dominates total cost	Suitable benchmark data may not represent prospective pharmaceutical use	Determines whether QML investigation is scientifically credible
Hybrid attribution	Whether the quantum component adds incremental value	Quantum-off, classical-only, simulator, and component-ablation comparisons	Equivalent output is produced without the quantum stage	Ablation may not capture all interactions among workflow stages	Prevents misallocation of credit and investment
Fair comparison	Whether the comparator is current and resources are matched	Predeclared benchmark; repeated quantum–classical race; baseline refresh	Advantage disappears against a stronger or equally tuned baseline	Best classical practice evolves	Central to any practical-advantage claim
Total resource profile	Whether all relevant computational and operational costs are counted	End-to-end wall time, energy, monetary cost, shots, tuning, mitigation, and expertise	Circuit-only gain is outweighed by workflow overhead	Resource priorities depend on institutional context	Determines whether benchmark superiority has practical meaning
Accuracy and uncertainty	Whether output quality and uncertainty are sufficient for intended use	Reference calculations; calibration; repeated runs; error decomposition	Result is unstable, uncalibrated, or sensitive to implementation choices	Numerical agreement does not establish downstream utility	Defines permissible scientific claims
Reproducibility	Whether another group or platform can reproduce the conclusion	Shared code and settings; independent replication; cross-device testing	Result depends on undisclosed tuning or one transient hardware condition	Reproducibility does not prove generality	Required before investment escalation
Downstream validation	Whether the output survives task-appropriate scientific or	Prospective computation; synthesis; assays;	Benchmark improvement fails to alter or improve	Experimental confirmation of one output does not	Distinguishes conditional pilot from broader research investment

	experimental testing	mechanistic testing where appropriate	downstream evidence	prove general advantage	
Comparator refresh	Whether classical progress has changed the counterfactual	Scheduled reassessment using contemporary classical algorithms	Previously claimed bottleneck is removed classically	Refresh cadence is field-dependent	Supports continuation, redesign, or termination
Research stop rule	Whether unresolved dependencies exceed the expected scientific value	Predefined milestones and termination criteria	Repeated failure at a necessary credibility gate	Stopping one application does not invalidate the wider field	Prevents indefinite continuation of a mismatched project

Limitations and Boundary Conditions

The proposed Quantum–Classical Credibility Map is a conceptual position rather than a validated theory, benchmark standard, or operational decision instrument. Its components were organized from the selected literature and from distinctions necessary to prevent overinterpretation, but the framework has not been tested prospectively across research portfolios. Different pharmaceutical tasks may require different credibility gates, evidence thresholds, and sequencing. The weakest-link principle is deliberately conservative and may understate the value of exploratory work whose downstream relevance cannot yet be specified. It should therefore be used to bound claims and reveal dependencies, not to prohibit hypothesis generation.

The available evidence is also heterogeneous. It includes theoretical analyses, reviews, hardware demonstrations, resource estimates, proof-of-concept pharmaceutical applications, and limited experimentally connected workflows. These evidence types cannot be treated as interchangeable. Results obtained on simulators or small hardware instances may establish implementability without establishing scaling. Resource projections for fault-tolerant computation depend on assumptions about future architectures, correction schemes, compilation, and algorithm development. Pharmaceutical benchmark datasets may not represent prospective chemical space, assay variability, or decision consequences. The framework cannot remove these uncertainties; it can only require that they be stated and linked to the permitted claim.

Several forms of misuse remain possible. The map could be converted improperly into a numerical maturity score, used to rank institutions, or treated as a deterministic investment algorithm. A project may satisfy formal reporting requirements while still relying on a weak pharmaceutical hypothesis or an inappropriate endpoint. Conversely, an early exploratory project may appear weak because it cannot yet supply a complete experimental pathway even though it addresses a scientifically important uncertainty. The framework cannot establish clinical utility, regulatory acceptability, patient benefit, safety, or commercial value. It cannot replace expert review by quantum scientists, computational chemists, medicinal chemists, pharmacologists, data scientists, experimentalists, and research-governance stakeholders.

Research Agenda and Implementation Priorities

The first research priority is prospective task-matching evaluation. Multiple quantum and hybrid projects should be specified before results are known, classified through the proposed framework, and followed through benchmark, replication, and downstream validation stages. Research should test whether independent reviewers agree on the pharmaceutical problem, necessary credibility gates, permissible claim, and invest–monitor–defer classification. Disagreement would reveal where definitions require refinement. Studies should also examine whether use of the framework improves comparator quality, resource reporting, attribution of quantum contribution, and termination of poorly matched projects.

The second priority is shared evaluation infrastructure. Pharmaceutical quantum-computing benchmarks should include chemically meaningful instances, explicit task semantics, current classical implementations, versioned datasets, data-provenance records, hardware metadata, compilation settings, measurement budgets, mitigation details, uncertainty estimates, and reproducible workflow configurations. Benchmark suites should separate NISQ, early logical-qubit, and mature fault-tolerant assumptions rather than combine them. They should also distinguish electronic-structure validation from molecular prediction, generation, optimization, and prospective pharmaceutical decision support. Comparator refresh should be built into governance because classical methods and hardware improve continuously.

The third priority is staged implementation. Initial investment should emphasize reproducible enabling methods, resource estimation, benchmark design, and workflow ablation. Conditional pharmaceutical pilots should follow only when the task is consequential, the method–resource match is plausible, and downstream validation is available. Monitoring is appropriate when one identifiable dependency—such as logical-qubit availability, data-access cost, circuit depth, or measurement burden—prevents a credible pilot. Deferral is appropriate when the computational abstraction is disconnected from the pharmaceutical decision or when a contemporary classical method resolves the proposed bottleneck more effectively. These categories should remain revisable and should document what new evidence would justify movement between them.

Conclusion

Quantum–classical drug discovery should be evaluated after the advantage claim, not organized around it. The proposed Quantum–Classical Credibility Map reframes the field as a task-matching problem in which pharmaceutical relevance,

computational abstraction, algorithmic suitability, hardware regime, data and encoding assumptions, hybrid attribution, fair comparison, uncertainty, and downstream validation are jointly necessary. Its weakest-link principle prevents one favorable performance measure from being extended into unsupported claims of practical advantage or pharmaceutical value. The framework does not establish that quantum computing will or will not transform drug discovery. It instead specifies how individual applications can be investigated without confusing feasibility with usefulness, prediction with confirmation, or conceptual promise with deployment readiness. Progress will depend on problem-first formulation, full-stack co-design, strong classical counterfactuals, complete resource accounting, prospective validation, and willingness to reformulate or defer applications when the match is not credible.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Preskill J. Quantum computing in the NISQ era and beyond. Quantum. 2018;2:79. doi:10.22331/q-2018-08-06-79
2. Cao Y, Romero J, Olson JP, Degroote M, Johnson PD, Kieferová M, et al. Quantum chemistry in the age of quantum computing. Chem Rev. 2019;119(19):10856-915. doi:10.1021/acs.chemrev.8b00803
3. 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
4. Blunt NS, Camps J, Crawford O, Izsák 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
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. Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. Nat Rev Drug Discov. 2019;18(6):463-77. doi:10.1038/s41573-019-0024-5
7. Schneider P, Walters WP, Plowright AT, Sieroka N, Listgarten J, Goodnow RA, et al. Rethinking drug design in the artificial intelligence era. Nat Rev Drug Discov. 2020;19(5):353-64. doi:10.1038/s41573-019-0050-3
8. Walters WP, Barzilay R. Applications of deep learning in molecule generation and molecular property prediction. Acc Chem Res. 2021;54(2):263-70. doi:10.1021/acs.accounts.0c00699
9. Lenselink EB, ten Dijke N, Bongers B, Papadatos G, van Vlijmen HWT, Kowalczyk W, et al. Beyond the hype: Deep neural networks outperform established methods using a ChEMBL bioactivity benchmark set. J Cheminform. 2017;9(1):45. doi:10.1186/s13321-017-0232-0
10. Yang K, Swanson K, Jin W, Coley C, Eiden P, Gao H, et al. Analyzing learned molecular representations for property prediction. J Chem Inf Model. 2019;59(8):3370-88. doi:10.1021/acs.jcim.9b00237
11. 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
12. 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):2641-7. doi:10.1021/acs.jcim.1c00166
13. Pyrkov A, Aliper A, Bezrukov D, Lin YC, Polykovskiy D, Kamya P, et al. Quantum computing for near-term applications in generative chemistry and drug discovery. Drug Discov Today. 2023;28(8):103675. doi:10.1016/j.drudis.2023.103675
14. 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
15. 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
16. 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
17. 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
18. 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
19. 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

20. 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
21. Temme K, Bravyi S, Gambetta JM. Error mitigation for short-depth quantum circuits. Phys Rev Lett. 2017;119(18):180509. doi:10.1103/PhysRevLett.119.180509
22. 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
23. 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
24. McCaskey AJ, Parks ZP, Jakowski J, Moore SV, Morris TD, Humble TS, et al. Quantum chemistry as a benchmark for near-term quantum computers. NPJ Quantum Inf. 2019;5(1):99. doi:10.1038/s41534-019-0209-0
25. 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
26. 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
27. 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
28. Bauer B, Bravyi S, Motta M, Chan GKL. Quantum algorithms for quantum chemistry and quantum materials science. Chem Rev. 2020;120(22):12685-717. doi:10.1021/acs.chemrev.9b00829
29. Alexeev Y, Bacon D, Brown KR, Calderbank R, Carr LD, Chong FT, et al. Quantum computer systems for scientific discovery. PRX Quantum. 2021;2(1):017001. doi:10.1103/PRXQuantum.2.017001
30. Babbush R, McClean JR, Newman M, Gidney C, Boixo S, Neven H. Focus beyond quadratic speedups for error-corrected quantum advantage. PRX Quantum. 2021;2(1):010103. doi:10.1103/PRXQuantum.2.010103
31. Alexeev Y, Batista VS, Bauman N, Bertels L, Claudino D, Dutta R, et al. A perspective on quantum computing applications in quantum chemistry using 25–100 logical qubits. J Chem Theory Comput. 2025;21(22):11335-57. doi:10.1021/acs.jctc.5c01038